

Are You Listening to Me? Fine-Tuning Chatbots for Empathetic Dialogue

Paulo Knob
PUCRS / INCT-SANI

Porto Alegre, RS, Brazil
Email: paulo.knob@edu.pucrs.br

Leonardo Scholler
PUCRS

Porto Alegre, RS, Brazil

Juliano Rigatti
Porto Alegre, RS, Brazil

Soraia Musse
PUCRS
Porto Alegre, RS, Brazil
Email: soraia.musse@pucrs.br

Abstract—Conversational agents have made significant progress since ELIZA, expanding their role across various domains, including healthcare, education, and customer service. As these agents become increasingly integrated into daily human interactions, the need for emotional intelligence, particularly empathetic listening, becomes increasingly essential. In this study, we explore how Large Language Models (LLMs) respond when tasked with generating emotionally rich interactions. We analyzed the emotional progression of the dialogues using both sentiment analysis (via VADER) and expert assessments. While the generated conversations often mirrored the intended emotional structure, human evaluation revealed important differences in the perceived empathy and coherence of the responses. These findings suggest that emotion modeling in dialogues requires not only structural alignment in the expressed emotions but also qualitative depth, highlighting the importance of combining automated and human-centered methods in the development of emotionally competent agents.

I. INTRODUCTION

As technology evolves, it becomes more integrated into people’s daily lives. Many tasks that, in the past, required human intervention can be performed nowadays by specialized computer systems, thanks to Artificial Intelligence (AI) [1]. Since the pioneering work of Joseph Weizenbaum, who proposed the chatbot ELIZA [2], conversational agents have been impacting how humans interact with machines [3]. The advances in technology and, more specifically, in the AI field, led to the development of Natural Generative Language (NGL) models, which include the Generative Pre-trained Transformer (GPT) series [4], released by OpenAI as ChatGPT, and Gemini [5], released by Google, among others. Such systems are based on Large Language Models (LLMs) and can both emulate the conversation between human beings and generate new content, based on the data in which they were trained [6].

Although LLMs, such as ChatGPT and Gemini, can generate coherent information and human-like responses, their ability to identify and describe emotions is uncertain [7]. Emotion is closely related to empathy, especially when we think about its affective attributes, which involve physiological enthusiasm and spontaneous affective responses to someone else’s display of emotions [8]. Therefore, the modeling of an empathetic behavior in large language models is a research problem that is open for discussion and ideas.

The goal of this work is to assess whether or not an LLM-powered chatbot can be fine-tuned, through techniques

of prompt engineering, for empathetic listening, supporting the speaker in understanding feelings and needs for different situations. An experiment was conducted to detect the energy/polarity of the conversation, from the beginning until the end, using two well-known LLM-powered chatbots: ChatGPT, from OpenAI, and Gemini, from Google. Moreover, an expert was asked to evaluate the interactions generated by both chatbots. The results highlight that while lexical sentiment analysis captures general emotional patterns, it is not sufficient to assess whether true empathetic listening occurred, underscoring the importance of human evaluation in this context.

II. RELATED WORK

Arjmand et al. [9] introduce the concept of ‘empathic grounding’ in conversational agents, which includes the listener’s empathy for the speaker’s affective state. The authors argue that models of empathic grounding for embodied agents should be multimodal to facilitate natural and efficient communication. Thus, they describe a multi-modal model that takes as input user speech and facial expression to generate multimodal grounding moves for a listening agent, using a large language model (namely, GPT-3.5). Elyoseph et al. [7] uses the Levels of Emotional Awareness (EA) Scale to analyze ChatGPT’s responses to twenty scenarios, comparing its EA performance with the general population norms, as reported by a previous study. The authors comment that ChatGPT can generate appropriate EA responses, and its performance may improve significantly over time. However, further research is needed in order to understand the potential benefits and risks better, while also refining it to promote mental health. Bill et al. [10] aim to fine-tune an LMM to act in the psychology area. To do so, the authors use a method known as Reinforcement Learning from Human Feedback (RLHF), which is an approach for reinforcement learning that aims to improve the performance of a given LLM by using the feedback of interacting humans. The results achieved presented no significant difference between the pre-trained model and the fine-tuned model, which suggests that the fine-tuned model had little or no difference when compared with the model with no fine-tuning. Chen et al. [11] argue that large language models often tend to provide universal advice when applied in the psychological counseling field. However, in such scenarios, the person is usually looking for more elaborate support, such as

empathy, trust, understanding, and comfort. With that in mind, the authors built a dataset comprised of over 2 million multi-turn empathetic conversations, which was used to fine-tune LLMs, showing that their empathetic response capabilities were significantly enhanced.

III. PROPOSED EXPERIMENT

This work aims to assess whether or not an LLM-powered chatbot can be fine-tuned for empathetic listening, using techniques of prompt engineering, supporting the speaker in understanding feelings and needs for different situations.

A. Datasets Configuration

The goal of this experiment is to evaluate an LLM chatbot that can conduct a conversation with a person to facilitate empathetic listening, thereby creating a sense of comfort in the person. For this experiment, we chose to work with two different chatbots: OpenAI LLM, known as ChatGPT, and Google LLM, known as Gemini. The idea is to utilize techniques of prompt engineering to fine-tune the model and guide it to engage in empathetic listening. For this end, the following steps were conducted:

- (1) Gather examples of conversations curated by an expert, to serve as a base and ground truth;
- (2) Define a benchmark technique to quantify and classify the responses produced by the model;
- (3) Generate the interactions, simulating a real conversation.

For the first step, an expert in communication was assigned to manually create 24 texts with example interactions between two people, where the listener acted as an empathetic listener. A dataset was created with those texts. The conversations present in this dataset portray interactions between a speaker and a listener, where the listener acts as an empathetic listener. Each of the 24 conversations is comprised of 5 interactions, where the speaker (who starts the conversation) has three interactions, and the listener has 2. Table I shows an example of a conversation manually tailored by the expert.

TABLE I
EXAMPLE OF CONVERSATION MANUALLY TAILORED BY THE EXPERT.

Agent	Utterance
Speaker	Yesterday I came home and the sink was once again full of dirty dishes. The guy doesn't care. I've already asked him to wash it as he cooks so it doesn't accumulate, but it's no use! I can not take it anymore.
Listener	I'm listening to you and realizing that there's anger, frustration there, does that make sense?
Speaker	Yes, and tiredness as well.
Listener	And perhaps these feelings are present because you need reciprocity and respect.
Speaker	That's it, I really need it. Of consideration as well.

Then, for the second step, we consider the concept of conversation energy, where negative and positive feelings are evaluated as negative and positive energy, respectively, ranging from -1 (negative) to 1 (positive). To perform this energy evaluation, we utilized the Valence Aware Dictionary and

sEntiment Reasoner (VADER) tool [12], a lexical and heuristic set that generates a numerical component indicating sentiment, ranging from negative to positive, to describe the energy of the text.

Finally, for the third step, the fine-tuned LLM was used to generate several conversations simulating interactions between two people. In this step, four additional datasets were created: the ChatGPT extended dataset, which was created extending the interactions of the dataset manually tailored by the expert, using ChatGPT; the Gemini extended dataset, which was created in the same way, but using Gemini instead; the ChatGPT control dataset, where ChatGPT generated the conversations entirely, with no fine-tuning concerning empathetic listening; and the Gemini control dataset, which was generated in the same manner, but using Gemini instead.

The extended datasets (for both ChatGPT and Gemini) were generated by continuing the conversations of the first dataset (manually tailored by the expert). The idea of extending the data to more points also serves to validate the LLM's coherence and cohesion in maintaining longer conversations until a conclusion. The first points were not changed, so the data analysis is preserved in the first two points. For normalization purposes, all conversations were truncated at 12 interactions, where the speaker (who starts the conversation) counts with 6 interactions, and the listener (role played by the LLM chatbot) counts with another 6. Table II shows an example conversation generated by ChatGPT while extending the conversation customized by the expert, while Table III presents the same example, but generated by Gemini instead.

TABLE II
EXAMPLE OF CONVERSATION GENERATED BY CHATGPT, EXTENDING THE INTERACTION FROM THE EXPERT DATASET.

Agent	Utterance
Speaker	Yesterday I came home and the sink was once again full of dirty dishes. The guy doesn't care. I've already asked him to wash it as he cooks so it doesn't accumulate, but it's no use! I can not take it anymore.
Listener	I'm listening to you and realizing that there's anger, frustration there, does that make sense?
Speaker	Yes, and tiredness as well.
Listener	And perhaps these feelings are present because you need reciprocity and respect.
Speaker	That's it, I really need it. Of consideration as well.
Listener	It sounds like you've been feeling overwhelmed by the situation.
Speaker	Yes, exactly. I feel like I'm the only one taking care of things in the house.
Listener	It must be exhausting to constantly feel like you're carrying all the responsibility.
Speaker	It really is. I just wish he would listen and understand how I'm feeling.
Listener	It sounds like you're seeking understanding and support from him.
Speaker	Yes, that's exactly it. I feel like I'm not being heard.
Listener	It's important to feel heard and validated in your feelings.

In its turn, the control datasets (for both ChatGPT and Gemini) were generated to be conversations which may or may not have empathic components (i.e., there was no specific configuration regarding empathetic listening). Fully generated

TABLE III

EXAMPLE OF CONVERSATION GENERATED BY GEMINI, EXTENDING THE INTERACTION FROM THE EXPERT DATASET.

Agent	Utterance
Speaker	Yesterday I came home and the sink was once again full of dirty dishes. The guy doesn't care. I've already asked him to wash it as he cooks so it doesn't accumulate, but it's no use! I can not take it anymore.
Listener	I'm listening to you and realizing that there's anger, frustration there, does that make sense?
Speaker	Yes, and tiredness as well.
Listener	And perhaps these feelings are present because you need reciprocity and respect.
Speaker	That's it, I really need it. Of consideration as well.
Listener	I can only imagine. It's disheartening when you express a need and it's not met, especially in a shared living space where cooperation is essential. It makes you feel unseen and unheard, right?
Speaker	Exactly! It's like my efforts aren't valued, and I'm carrying the weight of the household chores alone.
Listener	So this isn't just about dirty dishes; it's about feeling undervalued and unsupported in the relationship. That makes perfect sense. It's understandable that you're feeling overwhelmed and resentful.
Speaker	Yes, it's more than just the dishes. It's the feeling that my concerns are constantly dismissed. It's a pattern that's been happening for a while now.
Listener	It sounds like you've been carrying a heavy emotional burden for some time. Have you considered talking to him directly about how this makes you feel, focusing on your feelings rather than blaming?
Speaker	I have tried, but he just says he'll do better and then it happens again. I don't know if it's worth trying again. I'm starting to feel hopeless.
Listener	I hear that hopelessness. It's understandable to feel that way after repeated attempts with little to no change. Have you considered perhaps exploring other ways to address this, maybe with a therapist or counselor who could help you both communicate more effectively? Or perhaps even setting clear boundaries?

through the LMM, these conversations follow the problem and outcome model, similar to empathic listening. It is composed of 30 texts, and, for normalization, only the first 12 interactions of each text were taken into account for the analysis. Table IV shows an example conversation fully generated by ChatGPT with no previous configuration concerning empathetic listening, while Table V presents the same example, but generated by Gemini instead.

In this experiment, we hypothesize that, when analyzing each interaction between the listener and the speaker, it will be possible to observe a tendency of energy going from negative (or lower values) to positive (or higher values), containing these distinct characteristics:

- The conversation begins with the problem, usually negative;
- The empathic agent reiterates the problem to make sure he or she understands, equalizing the energy level of the conversation;
- The conversation progresses toward a positive conclusion, converging the energy of both agents.

In Section IV, we present the analysis of this experiment.

TABLE IV

EXAMPLE OF CONVERSATION FULLY GENERATED BY CHATGPT, WITH NO PREVIOUS CONFIGURATION CONCERNING EMPATHETIC LISTENING.

Agent	Utterance
Speaker	I just finished a really good book, it was so captivating.
Listener	That's awesome! What was the book about?
Speaker	It was a mystery thriller with a plot twist I never saw coming.
Listener	I love when books surprise me like that. It's such a good feeling.
Speaker	Definitely, it keeps you on your toes. Have you read any good books lately?
Listener	Not recently, I've been so busy with work that I haven't had much time for reading.
Speaker	I totally get that, sometimes life just gets in the way of our hobbies.
Listener	Yeah, but it's important to make time for things we enjoy too.
Speaker	You're right, I need to prioritize some me-time soon.
Listener	You deserve it! What's something you like to do for yourself when you have free time?
Speaker	I love going for long walks in nature, it really helps me clear my mind.
Listener	That sounds so relaxing. I should try that sometime.

TABLE V

EXAMPLE OF CONVERSATION FULLY GENERATED BY GEMINI, WITH NO PREVIOUS CONFIGURATION CONCERNING EMPATHETIC LISTENING.

Agent	Utterance
Speaker	My dog, Max, finally stopped chewing on my shoes! It's been a battle for months.
Listener	Oh my gosh, that's amazing! What a relief. Mine still manages to find new and creative ways to destroy things, though.
Speaker	I know, right? It felt like a victory. I think it's finally understood "no" means something. What's your dog getting into these days?
Listener	Mostly the trash. He's a champion garbage diver. I've tried everything, but he's relentless. Speaking of victories, I finally got my car fixed after that fender bender.
Speaker	Oh, good! What happened? And what was wrong with it?
Listener	Some lady backed into me at the grocery store. It was mostly cosmetic damage, but the bumper needed replacing. It took forever to get the insurance sorted out.
Speaker	That's a pain. Insurance companies are the worst. At least it's fixed now, though! Anything fun planned for the weekend?
Listener	Not much, just relaxing mostly. Maybe a movie marathon. How about you?
Speaker	Similar plan, some Netflix and chill. Maybe I'll finally get around to tackling that book I've been meaning to read.
Listener	Sounds nice and relaxing. I'm looking forward to a quiet weekend after this crazy week. Getting some well-deserved rest.
Speaker	Totally agree. We deserve it! I'm already feeling less stressed just thinking about it. Well, I should get going, got a pile of laundry calling my name.
Listener	Haha, same here! Talk to you soon.

B. Expert Questionnaire

In addition to the lexical analysis, we also asked the expert to evaluate the interactions generated for both datasets (i.e., extended and control) by both chatbots (i.e., ChatGPT and Gemini), resulting in a total of four datasets to be evaluated. To do so, the expert completed a form, for each of the four datasets, comprising the questions presented in Table VI, which is an adaptation of the Bartneck questionnaire [13] with one additional question regarding empathetic listening ability. Moreover, we added a free-text field, allowing the expert to

discuss any additional topics not covered by the questions. It is important to clarify that the expert was well aware of the conditions of the experiment, which means he/she knew that an LLM chatbot generated the conversations of those datasets.

TABLE VI
QUESTIONS OF THE EXPERT FORM.

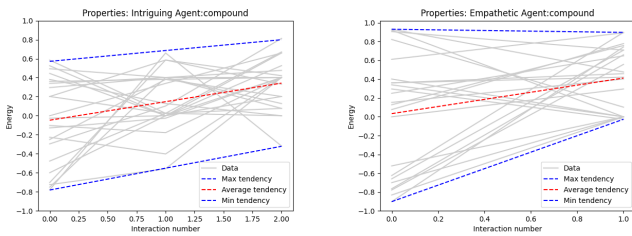
Question	Scale
Looking at how the conversation evolved, rate your impression of the chatbot on these scales, from Unfriendly to Friendly	1-5
Looking at how the conversation evolved, rate your impression of the chatbot on these scales, from Awful to Nice	1-5
Looking at how the conversation evolved, rate your impression of the chatbot on these scales, from Foolish to Sensible	1-5
How do you evaluate the empathetic listening ability of the chatbot, from Very unsatisfied to Very satisfied?	1-5

IV. RESULTS

This section presents the results achieved by this work.

A. Expert Dataset

As commented in Section III-A, a dataset was generated with the conversations that were manually created by the expert, where the listener acted as an empathetic listener. Figure 1 shows the energy tendency of the compound of all 24 interactions, for both the speaker (Intriguing Agent, left) and listener (Empathetic Agent, right). The Y axis shows the energy value of the conversation (as computed by VADER tool), and X axis shows the flow of the interactions (from the beginning to the end). Blue lines show the maximum and minimum energy tendencies, while the red line represents the average energy tendency and grey lines represent each interaction individually. It is also important to notice that all figures start in 0 for X axis, which means that 0 is the first interaction. It is possible to notice a growing pattern of the energy tendency, from lower to higher values, as the conversation flowed, for both agents, which is particularly evident in the red line of the figures (Average tendency).



(a) Energy tendency of the speaker. (b) Energy tendency of the listener.

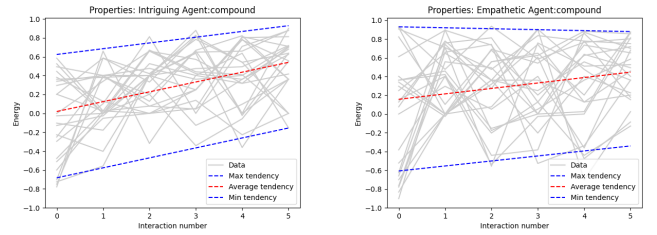
Fig. 1. Energy tendency of the compounded interactions. Y axis shows the energy value (as computed by VADER tool), and X axis shows the flow of the interactions (from the beginning to the end). Blue lines show the maximum and minimum energy tendencies, while the red line represents the average energy tendency. It is possible to notice a growing pattern of the energy tendency, from lower to higher values, as the conversation flowed, which is particularly evident in the red line (Average tendency).

It is interesting to note that this dataset, since the expert created it, can also serve as a model, which means that similar

patterns should arise from other datasets containing empathetic listening. Thus, the growing pattern (i.e., going from lower to higher energy values) of the Average Tendency (i.e., red line of the figures) should arise as well.

B. Extended Dataset

As commented in Section III-A, two datasets were generated by extending the conversations of the expert dataset, using both ChatGPT and Gemini. Figure 2 shows the energy tendency of the compound of all interactions generated by ChatGPT, for both the speaker (Intriguing Agent, left) and listener (Empathetic Agent, right). Since the first two points were preserved, the figure shows only 5 interactions for each speaker and listener. It is possible to notice the continuation of the tendencies generated by the expert dataset (i.e., a maintenance of the growing pattern), which is particularly evident in the red line of the figures (Average tendency).



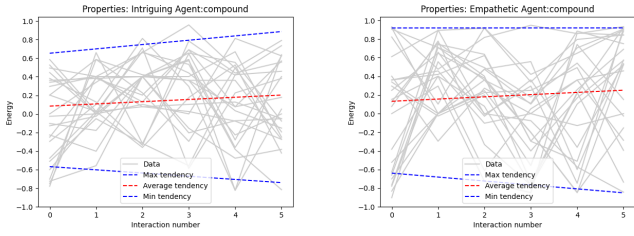
(a) Energy tendency of the speaker. (b) Energy tendency of the listener.

Fig. 2. Energy tendency of the compounded interactions. The Y axis shows the energy value of the conversation (as computed by VADER tool), and X axis shows the flow of the interactions (from the beginning to the end). Blue lines show the maximum and minimum energy tendencies, while the red line represents the average energy tendency. It is possible to notice the continuation of the tendencies generated by the expert dataset, which is particularly evident in the red line of the figures (Average tendency).

In its turn, Figure 3 shows the energy tendency of the compound of all interactions generated by Gemini, for both the speaker (Intriguing Agent, left) and listener (Empathetic Agent, right). As it was observed in the extended dataset generated with ChatGPT, it is possible to notice the continuation of the tendencies generated by the expert dataset, which is particularly evident in the red line of the figures (Average tendency). However, it is also possible to notice a reduction in the minimum tendency of the speaker 3(a) and the listener 3(b), which might indicate a difficulty of Gemini in dealing with an empathetic listener role. This insight would be aligned with the expert evaluation, which is presented later in Section IV-D.

C. Control Dataset

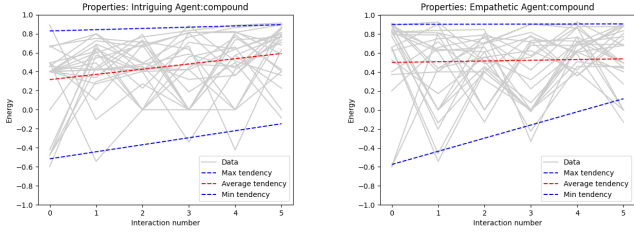
As commented in Section III-A, two datasets were generated to be conversations with no particular configuration regarding empathetic listening, using both ChatGPT and Gemini. Figure 4 shows the energy tendency of the compound of all interactions generated by ChatGPT, for both the speaker (Intriguing Agent, left) and listener (Empathetic Agent, right). It is possible to notice a difference between the two agents,



(a) Energy tendency of the speaker. (b) Energy tendency of the listener.

Fig. 3. Energy tendency of the compounded interactions. Y axis shows the energy value (as computed by VADER tool), and X axis shows the flow of the interactions (from the beginning to the end). It is possible to notice the continuation of the tendencies generated by the expert dataset, which is particularly evident in the red line of the figures (Average tendency). It is also possible to notice a reduction in the minimum tendency of both the speaker(a) and the listener(b).

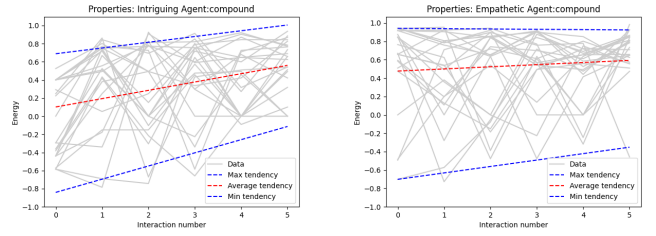
which is particularly evident in the red line of the figures (Average tendency). While the speaker (Intriguing Agent) had a noticeable increase in its energy 4(a), the listener (Empathetic Agent) had a steady energy 4(b), showing that its energy remained almost the same for the whole interaction.



(a) Energy tendency of the speaker. (b) Energy tendency of the listener.

Fig. 4. Energy tendency of the compounded interactions. The Y axis shows the energy value of the conversation (as computed by VADER tool), and X axis shows the flow of the interactions (from the beginning to the end). Blue lines show the maximum and minimum energy tendencies, while the red line represents the average energy tendency. It is possible to notice a difference between both agents, which is particularly evident in the red line of the figures (Average tendency), since the listener (Empathetic Agent(b)) maintained a steady energy during the interactions.

Figure 5 shows the energy tendency of the compound of all interactions generated by Gemini, for both the speaker (Intriguing Agent, left) and listener (Empathetic Agent, right). The Y axis shows the energy value of the conversation (as computed by VADER tool), and X axis shows the flow of the interactions (from the beginning to the end). Blue lines show the maximum and minimum energy tendencies, while the red line represents the average energy tendency. As it was observed in the control dataset generated with ChatGPT, a difference between the two agents is noticeable, which is particularly evident in the red line of the figures (Average tendency). While the speaker had a noticeable increase in its energy 5(a), the listener had a little increase in its energy 5(b), showing that its energy remained almost the same for the whole interaction.



(a) Energy tendency of the speaker. (b) Energy tendency of the listener.

Fig. 5. Energy tendency of the compounded interactions. The Y axis shows the energy value of the conversation (as computed by VADER tool), and X axis shows the flow of the interactions (from the beginning to the end). Blue lines show the maximum and minimum energy tendencies, while the red line represents the average energy tendency. It is possible to notice a difference between the two agents, where the listener's energy remained almost stable.

After collecting data from the datasets, the results seem to validate our hypothesis: when applying the empathetic listening fine-tuning, a tendency of energy going from lower to higher values could be observed, while conversations seem to have evolved to a positive conclusion, converging the energy of both agents. In its turn, when the fine-tuning was not applied (control datasets), the energy of the listener (Empathetic Agent) remained stable. Additionally, when looking at the images, it is possible to notice that, usually, the avg energy of the speaker varied more than the avg energy of the listener. Thus, it is possible that the speaker followed the listener, which could also imply that the behavior of the listener, indeed, impacted on the behavior of the speaker.

D. Expert Evaluation

As described in Section III-B, we asked the expert to evaluate the interactions generated for the datasets. In order to do so, the expert answered a few questions about each dataset, as presented in Table VI. For the first three questions (Friendly, Nice, and Sensible), the expert gave the maximum evaluation (i.e., 5, Very Satisfied) to all four datasets (i.e., Extended and Control, for both ChatGPT and Gemini). However, for the last question (Empathetic listening ability), the expert rated the chatbot with the worst evaluation (i.e., 1, Very Unsatisfied), across all datasets as well. We are aware that having only one expert for this evaluation is a limitation of our work, and should be addressed in future works.

By the end of the questionnaire, we included a free text box, so the expert would be able to freely talk about his/her impressions of the evaluated interaction, and even justify the evaluations given. For the Extended dataset generated by ChatGPT, the expert commented that "Although we can classify the chatbot as 'friendly', 'nice' and 'sensible', his performance from the point of view of empathic listening was disastrous. The length of the conversation proposed by the chatbot is very judgmental and does little to explore the feelings and needs of the interlocutor.". For the Extended dataset generated by Gemini, the expert commented that "The extent of the conversation generated by the chatbot is full of

judgments and is not restricted to investigating feelings and needs, which would be the function of empathic listening.”.

In its turn, for the Control dataset generated by ChatGPT, the expert commented that “In the dialogues created by AI in this case, there is no investigation whatsoever by any of the interlocutor’s feelings and needs.”. In its turn, for the Control dataset generated by Gemini, the expert commented that “There is no interest whatsoever on the part of the interlocutors in the feelings and needs of the other, which would be the main characteristic of empathic listening.”.

The results achieved with the energy tendency reading and with the expert evaluation, for both datasets (Extended and Control) and both chatbots (ChatGPT and Gemini), suggest that, although useful, a lexical analysis should not be used as the only evaluation metric of an empathetic listener chatbot. Even when the energy of the conversation rises over time, the interaction itself might not be beneficial to the user, especially if the chatbot, which should be acting as an empathetic listener, does not fulfill this role. The feedback of the expert suggests that the LLM-powered chatbot failed to act as an empathetic listener, given its judgmental comments and limited emotional expression.

On the other hand, based on the comments, it is possible to notice that the fine-tuned chatbot tried to fill the empathetic listener role, but failed, as the expert commented that it was judgmental and did little to explore the user’s feelings and needs. In its turn, the chatbot with no fine-tuning didn’t even try to fill the role, as the expert commented that there was no investigation about the feelings and needs of the user.

V. FINAL CONSIDERATIONS

This study evaluated whether fine-tuning large language models can improve their ability to engage in empathetic listening dialogue. Our experiments involved ChatGPT and Gemini, both adapted through prompt engineering and assessed using energy trajectory analysis and expert feedback. The obtained results reinforce the relevance of using both human-centered and automated approaches to analyze emotional progressions in dialogue. Beyond that, the study highlights two key take-aways. First, it was encouraging to observe that LLMs can be prompted and guided toward producing more emotionally supportive responses, showing potential for fine-tuning in the direction of empathy. However, the second—and perhaps more critical—finding is that truly empathetic listening goes beyond generating positive or soothing messages. It involves engaging with the speaker’s emotional state, helping them reflect and better understand their own problems—something that current LLMs still struggle to achieve. Finally, while sentiment lexicons, such as VADER, can be helpful tools for analyzing emotional trends, they alone are not sufficient to evaluate the presence or quality of empathetic listening. Empathy in dialogue is a complex, multi-layered process that requires context, intention, and human sensitivity—dimensions that remain challenging for automated systems.

From a human-centered AI perspective, this research reinforces the value of tailoring dialogue systems for affective

sensitivity. Empathetic listening is critical in areas such as mental health support, digital companionship, and training simulations. Our methodology highlights the importance of expert-curated datasets and emotional metrics in evaluating conversational empathy. For future work, we propose extending this investigation to additional LLMs, such as LLaMA or DeepSeek, to assess generalization. Moreover, involving interdisciplinary experts—including psychologists—and conducting user-based studies would enrich both the technical and social understanding of empathetic virtual agents.

ACKNOWLEDGMENT

This study was partly financed by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brazil (CAPES) – Finance Code 001; by the Conselho Nacional de Desenvolvimento Científico e Tecnológico - Brazil (CNPq) - Process Numbers 309228/2021-2; 406463/2022-0; 153641/2024-0.

REFERENCES

- [1] P. Manickam, S. A. Mariappan, S. M. Murugesan, S. Hansda, A. Kaushik, R. Shinde, and S. Thipperudraswamy, “Artificial intelligence (ai) and internet of medical things (iomt) assisted biomedical systems for intelligent healthcare,” *Biosensors*, vol. 12, no. 8, p. 562, 2022.
- [2] J. Weizenbaum, “Eliza—a computer program for the study of natural language communication between man and machine,” *Communications of the ACM*, vol. 9, no. 1, pp. 36–45, 1966.
- [3] A. Talyshinskii, N. Naik, B. Z. Hameed, P. Juliebø-Jones, and B. K. Somani, “Potential of ai-driven chatbots in urology: revolutionizing patient care through artificial intelligence,” *Current Urology Reports*, vol. 25, no. 1, pp. 9–18, 2024.
- [4] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [5] G. Team, R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican *et al.*, “Gemini: a family of highly capable multimodal models,” *arXiv preprint arXiv:2312.11805*, 2023.
- [6] A. Abd-Alrazaq, R. AlSaad, D. Alhuwail, A. Ahmed, P. M. Healy, S. Latifi, S. Aziz, R. Damseh, S. A. Alrazak, J. Sheikh *et al.*, “Large language models in medical education: opportunities, challenges, and future directions,” *JMIR Medical Education*, vol. 9, no. 1, p. e48291, 2023.
- [7] Z. Elyoseph, D. Hadar-Shoval, K. Asraf, and M. Lvovsky, “Chatgpt outperforms humans in emotional awareness evaluations,” *Frontiers in Psychology*, vol. 14, p. 1199058, 2023.
- [8] H. Prendinger, J. Mori, and M. Ishizuka, “Using human physiology to evaluate subtle expressivity of a virtual quizmaster in a mathematical game,” *International journal of human-computer studies*, vol. 62, no. 2, pp. 231–245, 2005.
- [9] M. Arjmand, F. Nouraei, I. Steenstra, and T. Bickmore, “Empathic grounding: Explorations using multimodal interaction and large language models with conversational agents,” in *Proceedings of the ACM International Conference on Intelligent Virtual Agents*, 2024, pp. 1–10.
- [10] D. Bill and T. Eriksson, “Fine-tuning a llm using reinforcement learning from human feedback for a therapy chatbot application,” 2023.
- [11] Y. Chen, X. Xing, J. Lin, Z. Wang, Q. Liu, X. Xu *et al.*, “Soulchat: Improving llms’ empathy, listening, and comfort abilities through fine-tuning with multi-turn empathy conversations,” in *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- [12] C. Hutto and E. Gilbert, “Vader: A parsimonious rule-based model for sentiment analysis of social media text,” in *Proceedings of the international AAAI conference on web and social media*, vol. 8, no. 1, 2014, pp. 216–225.
- [13] C. Bartneck, E. Croft, and D. Kulic, “Measuring the anthropomorphism, animacy, likeability, perceived intelligence and perceived safety of robots,” 2008.