

Aspect Based Emotion Analysis for Dialogue Understanding

Author: Zied Mustapha
Co-supervisor: Kalpani Anuradha Welivita
Supervisor: Pearl Pu
HCI Group, EPFL, Switzerland

Abstract

We present an innovative approach for aspect-based emotion analysis in distress-related conversations utilizing large language models (LLMs), specifically the GPT-3.5 model. The focus of this study is on the design and implementation of advanced prompt techniques, including zero-shot, few-shot, and chain-of-thought prompting, to generate structured responses for further analysis. The method was tested on two distinct datasets: CounselChat, a collection of therapy dialogues, and the Reddit Empathetic Dialogues (RED) dataset, a collection of peer support dialogues from Reddit dataset. By leveraging the power of LLMs, this study automates the task of emotion identification and corresponding cause categorization, and the extraction of Motivational Interviewing Treatment Integrity (MITI) categories from professional responses, bypassing the need for labor-intensive manual labeling. A detailed analysis was performed on the frequency of emotions, their causes, and the corresponding MITI categories, leading to interesting observations about the relationship between emotions, causes, and treatment responses. One possible application is the evaluation of responses in the RED dataset, for which GPT-3.5 can be instructed to use the relationship between emotions, causes and MITI response categories as a reference and assess the responses accordingly (refer to figure 1). The study also introduced a modular and robust code architecture to handle a large number of subsequent requests to the GPT-3.5 model. The results obtained provide compelling evidence for the potential of LLMs in understanding and annotating distress-related conversations, contributing to the development of more empathetic and effective conversational agents for addressing psychological distress.

Keywords: emotion analysis, dialogue understanding, aspect-based, sentiment analysis, machine learning

1. Introduction

In the evolving field of Natural Language Processing (NLP), one critical application lies in the realm of aspect-based sentiment and emotion analysis. This refers to the process of identifying and categorizing sentiments or emotions expressed towards specific aspects or entities within a text. In distress-related conversations, recognizing these sentiments or emotions becomes even more crucial. By uncovering the underlying emotions expressed in a request for help, we gain a deeper understanding of the individual's mental state, which can aid in offering more empathetic, accurate, and effective responses.

Conversational agents, particularly chatbots, have been at the forefront of technology aimed at providing assistance and comfort to individuals in distress. Properly designed, these chatbots have the potential to be a non-intrusive and constant source of support. However, their effectiveness hinges on their ability to understand and respond appropriately to a person's emotional state. Herein lies the importance of aspect-based sentiment and emotion analysis.

Motivational interviewing (MI) is a counseling method that involves the counselor engaging with the client in a manner that inspires them to change behavior. The success

of this approach relies heavily on the counselor's ability to accurately identify and address the client's emotions. The Motivational Interviewing Treatment Integrity (MITI) coding system was developed to provide a means of evaluating the fidelity of motivational interviewing. It provides an operational definition of MI, presenting a set of categories (e.g., giving information, affirming, reflecting) used to classify a counselor's verbal responses. By identifying and studying the correlations between the expressed emotions and aspects in distress-related conversations and the MITI categories used to address them, we can gain invaluable insights into the design of more effective chatbots.

In this project, we utilize the GPT-3.5-turbo model for emotion analysis of messages seeking help and a MITI category analysis of the professional responses to these messages. The overarching goal is to improve our understanding of the interplay between emotions, their causes, and the MI strategies, and to leverage this understanding to improve chatbot design for better support of individuals in distress.

Machine learning techniques have greatly advanced the field of emotion and sentiment analysis. However, these traditional methods often hinge upon the availability of large amounts of large human-annotated datasets. To train these models, it is required to have a vast corpus of text, each labeled with the correct emotion or sentiment. This labelling process is

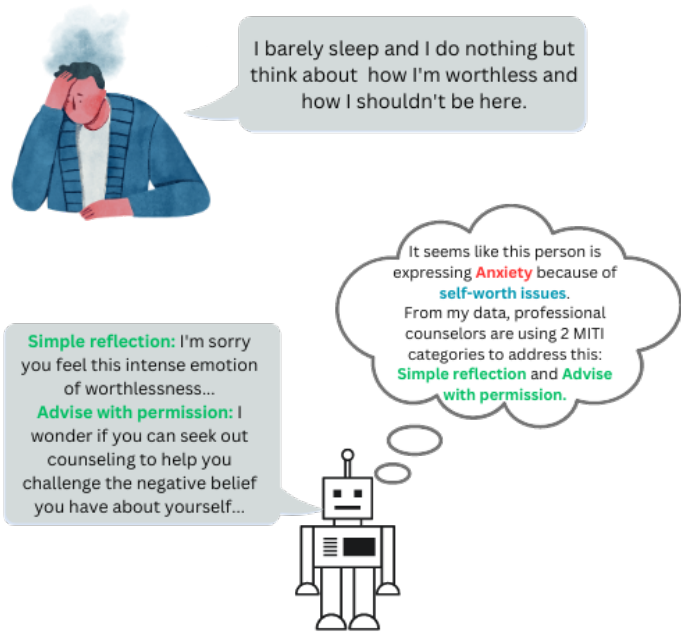


Figure 1: Visualisation of possible application.

a time-consuming and labor-intensive task, often requiring expertise in psychology and linguistics to accurately identify and annotate emotional states or aspects. Furthermore, the context and subtleties of language make this task even more complex and challenging. In distress-related conversations, where precision of emotional recognition is paramount, these challenges are magnified. This often leads to a compromise in the scalability and adaptability of these traditional machine learning models. It raises the need for more advanced and efficient models capable of understanding and analysing emotion and sentiment in a more nuanced manner.

With the advent of Language Learning Models (LLMs) like GPT-3.5-turbo, a paradigm shift has occurred in the field of emotion and sentiment analysis. These models are pre-trained on a massive scale and have remarkable capabilities in understanding and generating human-like text. This pretraining phase essentially equips these models with a comprehensive understanding of human language, context, and even emotions. Moreover, these LLMs can be fine-tuned or adapted to act as classifiers using zero-shot and few-shot prompting. This eliminates the necessity for human-annotated labels, which makes these models highly scalable and efficient compared to traditional machine learning models.

Zero-shot and few-shot learning refers to the ability of a model to understand and perform a task without any or with very few example prompts. These abilities significantly reduce the time and resources required to train the model. Furthermore, they make LLMs capable of adapting to new tasks swiftly, a trait that is particularly useful in the rapidly

evolving landscape of human conversation and sentiment. This advantage is especially important in distress-related conversations, where prompt and accurate understanding and response are crucial. Thus, the use of LLMs for emotion and sentiment analysis offers great potential for developing more effective and responsive conversational agents.

In the context of this research, we have utilized two different datasets: CounselChat and RED (Reddit Emotion Dataset), each of which provides a unique perspective on distress-related conversations. CounselChat is a dataset of real anonymized online counseling conversations, while RED is a dataset that comprises Reddit conversations labeled with various emotions.

Our approach was to experiment with different prompt designs to get the Language Learning Model, specifically GPT-3.5-turbo, to annotate the utterances in these datasets with emotions, the corresponding causes that trigger these emotions and the MITI response types that are used to address these emotions. The idea was to transform the problem into a format that the model has been trained on. By reframing the task as a language completion problem, we leveraged the vast knowledge encapsulated in the LLM. This allowed us to overcome the limitations associated with traditional machine learning methods, such as the need for large amounts of labeled data.

Prompt engineering, which involves creatively designing the input prompts to the model, was a critical part of our methodology. Through iterative experimentation and careful analysis of the model's responses, we fine-tuned the prompts to optimize the model's performance in emotion and sentiment analysis.

Our work demonstrates the use of the GPT-3.5-turbo model for aspect-based emotion analysis, eliminating the need for extensive labeled data. Through prompt engineering and linking identified emotions and causes to corresponding MITI categories, we present a novel approach for designing more empathetic conversational agents. This work paves the way for advanced mental health support systems and promotes further research in this field.

2. Related Work

The aspect-based sentiment and emotion analysis have been the focus of numerous studies in the past. Traditional approaches to this problem were primarily rule-based, heavily relying on the manual definition of rules and patterns for detecting sentiments and emotions in text (Mayur Wankhade, 2022). Over time, these methods evolved into more sophisticated approaches utilizing neural network-based models, such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and their variations (Yu and Sun, 2020). Despite their effectiveness, these methods require extensive amounts of labeled data, making the model training process

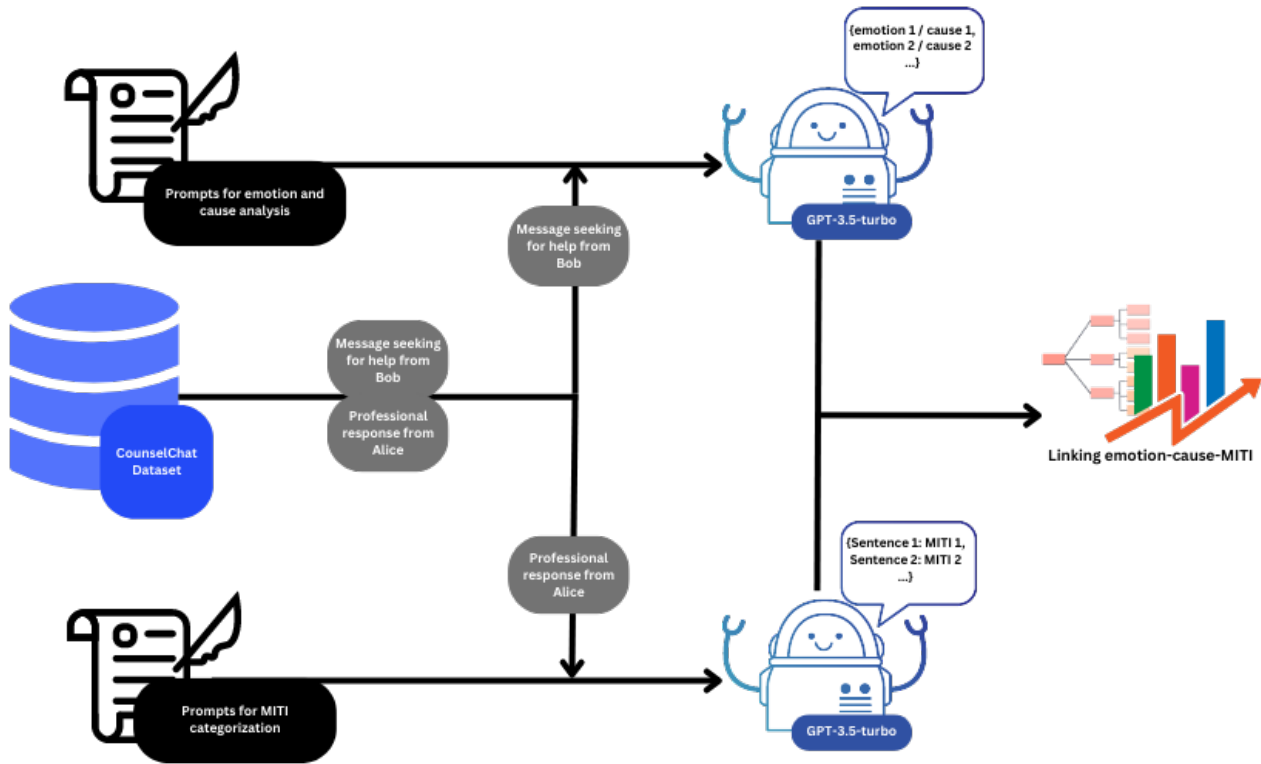


Figure 2: Process for prompting and analysis.

labor-intensive and time-consuming.

In parallel to this, there has been an emergence of large language models (LLMs) such as GPT-3 (Brown et al., 2020), GPT-4, Llama, OPT, and PaLM (Ni et al., 2020). These models, pretrained on massive text corpora, have showcased their potential in a variety of applications, ranging from text generation to machine translation, summarization, and even to tasks requiring reasoning and understanding of the context (Moghaddam and Honey, 2023).

More recently, there has been an increased interest in leveraging these LLMs for classification tasks. The advantages include the fact that they don't require a vast amount of labeled data, unlike the traditional methods, and they can be adapted to act as classifiers using zero-shot and few-shot prompting techniques (Suzgun, 2022). As such, they offer a promising direction for aspect-based sentiment and emotion analysis, as we explore in our work.

3. Datasets

In this work, two different datasets are used to conduct the analysis: CounselChat and Reddit Empathetic Dialogues (RED).

3.1. CounselChat

CounselChat is a dataset that consists of text from online counseling sessions, which are typically structured as question-

answer pairs. The conversations in this dataset cover a wide range of mental health issues, providing valuable insight into the emotions and concerns of individuals seeking help. The dataset consists of 2,129 post-response pairs that span across 31 distress related topics, the most frequent topics being depression, relationships, and intimacy. Due to the professional nature of the responses, the dataset is very useful, offering a high level of authenticity and practical application for mental health support.

3.2. Reddit Empathetic Dialogues (RED)

The RED dataset, derived from Reddit, a popular online discussion forum, contains a significant amount of conversational data from various subreddits focused on support and empathy. The dataset encompasses approximately 100,000 conversations, embodying over a million utterances from individuals sharing personal experiences and receiving emotional support from others. The conversations' average length in this dataset is roughly 10 turns.

These conversations, typically informal in nature, capture the natural language and diverse emotional expressions of everyday people. Despite this informality, the RED dataset serves as a robust platform for studying how empathy and support are communicated in online dialogues. It includes an expansive range of 32 different emotional categories, offering nuanced insights into emotional expression and response dynamics in conversations.

These two datasets provide a comprehensive set of conversational data for understanding emotions and their causes, as well as the responses to these emotions, making them particularly suitable for this study.

4. Methodology

The methodology employed in this study was divided into three distinct parts (refer to figure 2 for illustration):

4.1. Emotion and Aspect Analysis

The first stage of our methodology was an in-depth emotion analysis, carried out using the OpenAI GPT-3.5 model. We engaged the model in the task of extracting and identifying a range of emotions and their causative factors present within the CounselChat dataset. To ensure a structured and manageable output, we formulated specific prompts that would elicit responses in an easily parsable format. The prompts follow the chain-of-thought principle by decomposing his task in sentences as simple as possible to understand:

Assume the role of a specialist in text sentiment analysis.

Your task involves scrutinizing a text provided to you, written by an individual.

Your objective is to pinpoint each emotion expressed by the writer.

The only emotions possible are: [...]

Do not use other emotions.

Subsequently, link every emotion to a specific cause.

The only possible causes are:[...]

These are the only causes that should be utilized in your analysis, write them as they are, do not add details.

—[DATA to insert]—

When delivering your analysis, ensure it's structured as follows:

Emotion: 'specified emotion' / Cause: 'identified cause'

Emotion: 'specified emotion' / Cause: 'identified cause' [...]

Your findings should solely pertain to the emotions articulated by the author.

Only output the structure described, do not add anything else.

Again the only emotions that you can use are: [...]

And the only causes that you can use are: [...]

Recognizing the complexity of human emotions, we imposed certain constraints to guide the focus of our analysis. Primarily, we restricted our scope to a selected set of emotions, identified through a data-driven approach. Our selection was based on the most frequent emotions within the dataset, a finding established through an exploratory statistical analysis of a representative sample of 300 messages. This allowed us to condense the wide spectrum of emotions into a manageable subset (refer to figure 3 and see Appendix A for the final list of emotions). In terms of identifying and categorizing the causes of these emotions, we initially attempted to have the GPT-3.5 model undertake this task. However, we encountered

issues with this approach, particularly with the model providing emotions as categories. To overcome this, we utilized k-means clustering to group the causes into 10 categories, each represented by 10 words. We then engaged the GPT-3.5 model to discern commonalities among these words. The final categorization, which was fine-tuned based on interpretability and common sense, led us to establish five distinct categories: Love Relationship Issues, Mental Issues, Physical Health, Personal Struggles, and Family Conflict.

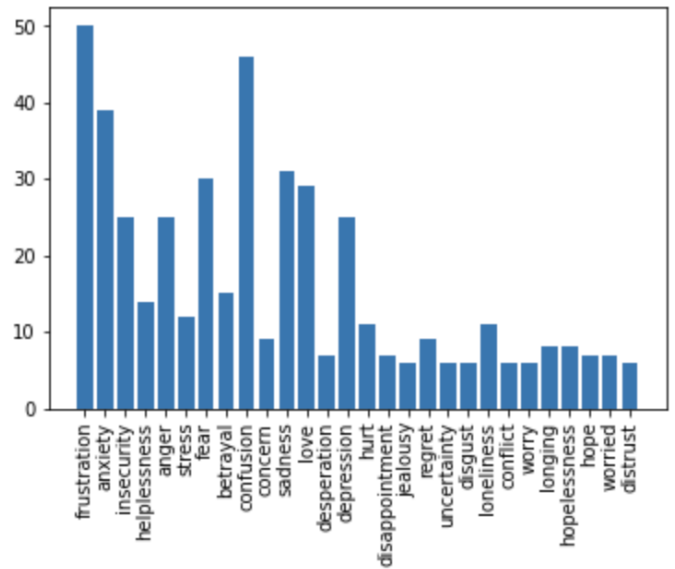


Figure 3: Emotion frequencies.

4.2. MITI Categories Analysis

The second phase concentrated on analyzing the Motivational Interviewing Treatment Integrity (MITI) categories. We tasked GPT-3.5 with identifying the MITI categories employed in each sentence of the response messages from professional counsellors in the CounselChat dataset. As with the previous phase, this was executed using specific prompts that ensured the output was in a pre-designated format for easier future use (refer to the Appendix B). In this step, we used few-shot prompting, providing examples of sentences for each category to improve the accuracy of the categorization:

You have been provided a professional counselor's response to a help-seeking message:

—[DATA to insert]—

Each sentence should be classified under one or more of the following categories: [...]

Your categorizations need to look like this:

Closed question, Open question if the first sentence is Closed question and the second sentence is Open question.

Output strictly this, nothing else.

These categories are part of a framework known as the Motivational Interviewing Treatment Integrity (MITI) system.

To enhance your understanding of these categories, here is some additional information:

Emotion to Cause to MITI Category Sankey Diagram

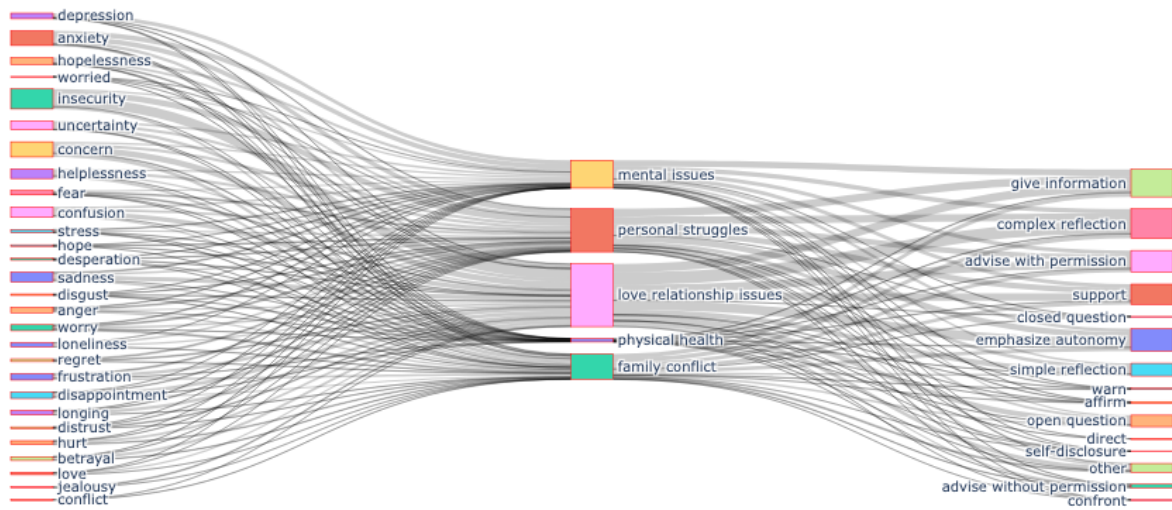


Figure 4: Emotion to Cause to MITI category Sankey plot diagram

1. Closed Question:

“Questions that can be answered with an yes/no response or a very restricted range of answers. “ Example: “Do you think this is an advantage? “

2. Open Question:

“Questions that allow a wide range of possible answers. “ Example: “What is your take on that? “ [...] Again, output only what was asked (MITI categories), do not respond anything else.

4.3. Emotion-Cause-MITI Link Representation

The third phase involved representing the connections between the identified emotions/causes and the corresponding MITI categories used to address them. This was represented using a specific structure, a variation of the JSON format. This representation allowed for the creation of a Sankey diagram that effectively illustrated these connections (see Figure 4).

Note: refer to Appendix B for full prompts.

5. Results

Considerable experimentation was undertaken to hone the prompt design, seeking results that closely aligned with our expectations. Notably, manual inspection of the results showed that the application of few-shot learning greatly enhanced the efficacy of the MITI categorisation process. In contrast to the results gleaned from zero-shot learning, which demonstrated

limited diversity in categorisation, the introduction of few-shot learning significantly improved the scope and accuracy of categories identified by GPT-3.5.

Our analysis of the CounselChat and RED datasets unveiled a variety of recurrent emotions. The most frequently detected emotions included frustration, anxiety, insecurity, helplessness, anger, stress, fear, betrayal, confusion, concern, sadness, love, desperation, depression, hurt, disappointment, jealousy, regret, uncertainty, disgust, loneliness, conflict, worry, longing, hopelessness, hope, worry, and distrust (refer to Figure 2 for detailed distributions).

When we compare this list to the categories identified by pre-existing models such as Ekman’s basic emotions, Plutchik’s emotion wheel, the 32 emotions classified in EmpathicDialogues, and the GoEmotions model, we observe that GPT-3.5 can identify a broader range of emotions. Particularly, emotions like ‘betrayal’, ‘concern’, ‘desperation’, ‘regret’, ‘uncertainty’, ‘conflict’, ‘longing’, and ‘worry’ present a departure from the more generic categories often captured in existing models.

The Sankey plot depicted in Figure 4 provides a visual representation of the connections between emotions, their underlying causes, and the corresponding MITI categories used to address them. Each link in the plot signifies that within the CounselChat dataset, a professional counselor has used a particular MITI category to address a specific emotion resulting from a certain cause.

The following key statistics emerge from the data:

- Emotion with most total links: insecurity, number of links: 3132, percentage of total: 12.39%
- Cause with most total links: love relationship issues, number of links: 9774, percentage of total: 38.65%
- MITI category with most total links: complex reflection, number of links: 4606, percentage of total: 18.21%
- Emotion with most unique links: anxiety, number of links: 5
- Cause with most unique links: personal struggles, number of links: 15
- MITI category with most unique links: give information, number of links: 5
- Most frequent emotion-cause pair: ('insecurity', 'love relationship issues'), count: 1539
- Least frequent emotion-cause pair: ('betrayal', 'mental issues'), count: 2
- Most frequent cause-MITI pair: ('love relationship issues', 'complex reflection'), count: 1942
- Least frequent cause-MITI pair: ('physical health', 'affirm'), count: 2

Note that a "unique link" refers to a distinct connection between two nodes. It is a specific emotion leading to a specific cause. And "total links" includes multiple instances of the same connection between two nodes.

6. Possible Application

The final phase of our methodology involved the application of the emotion-cause and corresponding MITI category recognition to evaluate responses in the Reddit Empathetic Dialogues (RED) dataset. For this, GPT-3.5-turbo was instructed to use the Sankey diagram as a reference and assess the responses accordingly.

We are here giving only the first steps to this process. We first reformatted the Sankey diagram so that every link is represented as follows:

""The emotion frustration was caused 281 times by family conflict: This was addressed 635 times by using give information. This was addressed 689 times by using complex reflection. This was addressed 519 times by using advise with permission [...]
The emotion confusion was caused 700 times by love relationship issues: This was addressed 1331 times by using give information. This was addressed 1942 times by using complex reflection [...]"

Subsequently, we instructed GPT as follows:

I need you to evaluate therapeutic responses to messages seeking for help. I will teach you how to do it by showing you how professional counselor addresses particular emotion that was caused by a specific aspect: ... Evaluate the therapeutic responses in 2 sentences.

When trying to evaluate some instances of the RED dataset, we get GPT-3.5 to respond as such as: ""The response is professional and therapeutic as it challenges the negative beliefs expressed in the original message and provides a counter-example to the idea that no one cares about the person. The response also acknowledges the person's feelings of confusion and inability to handle life. The use of personal examples in the response aligns with the MITI category of "giving information" and "providing feedback" which are effective in promoting change and increasing motivation. The response also aligns with the MITI category of "emphasizing autonomy" by encouraging the person to challenge their negative beliefs and consider the possibility that there are people who care about them. Overall, the response is effective in promoting change and increasing motivation by challenging negative beliefs and providing a counter-example to the idea that no one cares about the person.""

Note that this application can be greatly enhanced by better prompting and more experimentation.

7. Discussion

This study presented a novel approach to aspect-based emotion analysis, leveraging the capabilities of the GPT-3.5 language model. We applied a variety of prompts, including zero-shot, few-shot, and chain-of-thought, to annotate the emotions and their associated causes in distress-related conversations extracted from the CounselChat and RED datasets. Additionally, we utilized the Motivational Interviewing Treatment Integrity (MITI) categories to examine the responses from professional counselors, which unveiled significant insights into the correlations between these categories, specific emotions, and their causes.

The evaluation of our results illustrates that our methodology was largely effective. It provided a comprehensive emotion and aspect analysis, and the incorporation of the MITI categories in analyzing counselor responses enriched the depth of our investigation. Our findings are particularly evident when compared with existing emotion models such as Ekman's basic emotions, Plutchik's emotion wheel, the 32 emotions in EmpatheticDialogues, and GoEmotions. In many instances, GPT-3.5's performance aligned with these well-established frameworks, attesting to its accuracy in emotion identification and aspect association (Rashkin et al., 2019).

However, despite our methodological success, we also recognized certain limitations. The performance of GPT-3.5 could fluctuate depending on the complexity and subtleties of the dialogues analyzed. A notable challenge was experienced during the categorization of causes, leading us to revise our

approach and employ k-means clustering. The task of applying MITI response types to address emotions and aspects posed a complex challenge. Although few-shot learning increased the diversity of categorization, some variations remained, suggesting the need for further refinement in prompt strategies.

Our research also resulted in the creation of a robust, modular codebase. This system has been designed to handle large volumes of subsequent requests efficiently, ensuring outputs are structured appropriately for downstream analysis. Notably, the code includes mechanisms for recovery from errors using a threading model and an exception catching architecture. This feature facilitates continuity in data analysis and mitigates disruptions to the workflow. The versatility of the codebase lends itself to reuse and adaptation in future research applications and analysis. The entire code can be found at: [https://github.com/Laniakea1999/master_proj_2023]

The Sankey plot used to illustrate the relationships between emotions, causes, and MITI categories proved invaluable. However, its interpretation heavily depends on context and data quality, implying it might be less effective with high variance or inconsistent datasets.

In summary, the applied approach and methodology have shown promising potential for future research, particularly in aspect-based emotion analysis in dialogue systems. However, continuous refinement and adaptation to specific data are imperative in all machine learning applications. Based on our findings, we recommend careful consideration of prompt design in future research, customizing it to align with the specific requirements of the task at hand.

8. Future Work

Looking forward, there is ample opportunity to expand upon the work conducted in this study. Given that GPT models can generate human-like text, one promising direction is to explore the role of these models in the evaluation process. Manual annotation and evaluation are labor-intensive tasks that often constitute bottlenecks in Natural Language Processing (NLP) pipelines. By investigating ways to employ GPT models in these processes, we may be able to streamline the pipeline and expedite the development of sophisticated dialogue systems.

Moreover, with the evolution of GPT models and the recent development of GPT-4, it is worth exploring its potential to yield more accurate or nuanced results. This could potentially improve both emotion identification and aspect categorization.

Additionally, our study made significant strides in prompt design, but further work can be done to refine these prompts and improve the accuracy of our emotion and aspect categorizations. Similarly, while we identified and focused on five key cause categories, additional categories could potentially be extracted with a more fine-grained analysis, or with the

application of different NLP techniques.

Finally, It's important to underscore that the application (evaluation of therapeutic responses by GPT) performance can significantly be improved through enhanced prompting and additional experimental iterations.

As the field of Large Language Models continues to evolve, new strategies and techniques will undoubtedly emerge, providing even more tools for improving the performance of aspect-based emotion analysis. Our study is just a stepping stone in the promising path of understanding human emotions in dialogues and designing better conversational agents for addressing psychological distress.

Acknowledgements

I am deeply grateful for the guidance, support, and mentorship provided by my supervisor, Kalpani Anuradha Welivita, throughout this project. Your expertise and insights have been invaluable and have greatly enriched this research. I would also like to extend my sincerest thanks to Professor Pearl Pu for the generous provision of resources which significantly contributed to the completion of this project. The conducive environment for learning and research that you have fostered is truly commendable.

Appendix A.

- Causes:
 - Love Relationship Issues
 - Mental Issues
 - Physical Health
 - Personal Struggles
 - Family Conflict
- Emotions:
 - frustration
 - anxiety
 - insecurity
 - helplessness
 - anger
 - stress
 - fear
 - betrayal
 - confusion
 - concern
 - sadness
 - love
 - desperation
 - depression
 - hurt
 - disappointment
 - jealousy
 - regret
 - uncertainty
 - disgust
 - loneliness
 - conflict
 - worry
 - longing
 - hopelessness
 - hope
 - worried
 - distrust

Appendix B.

Assume the role of a specialist in text sentiment analysis. Your task involves scrutinizing a text provided to you, written by an individual. Your objective is to pinpoint each emotion expressed by the writer. The only emotions possible are: - frustration - anxiety - insecurity - helplessness - anger - stress - fear - betrayal - confusion - concern - sadness - love - desperation - depression - hurt - disappointment - jealousy - regret - uncertainty - disgust - loneliness - conflict - worry - longing - hopelessness - hope - worried - distrust Do not use other emotions. Subsequently, link every emotion to a specific cause. The only possible causes are: - Love Relationship Issues - Mental Issues - Physical Health - Personal Struggles - Family Conflict

These are the only causes that should be utilized in your analysis, write them as they are, do not add details.

—[DATA to insert]—

When delivering your analysis, ensure it's structured as follows: Emotion: 'specified emotion' / Cause: 'identified cause' Emotion: 'specified emotion' / Cause: 'identified cause' ... Your findings should solely pertain to the emotions articulated by the author. Only output the structure described, do not add anything else. Again the only emotions that you can use are: - frustration - anxiety - insecurity - helplessness - anger - stress - fear - betrayal - confusion - concern - sadness - love - desperation - depression - hurt - disappointment - jealousy - regret - uncertainty - disgust - loneliness - conflict - worry - longing - hopelessness - hope - worried - distrust

And the only causes that you can use are: - Love Relationship Issues - Mental Issues - Physical Health - Personal Struggles - Family Conflict

You have been provided a professional counselor's response to a help-seeking message:

—[DATA to insert]—

Each sentence should be classified under one or more of the following categories: - Closed question - Open question - Simple reflection - Complex reflection - Give information - Advise with permission - Affirm - Emphasize Autonomy - Support - Advise without permission - Confront - Direct - Warn - Self-disclosure - Other

Your categorizations need to look like this: Closed question, Open question if the first sentence is Closed question and the second sentence is Open question. Output strictly this, nothing else.

These categories are part of a framework known as the Motivational Interviewing Treatment Integrity (MITI) system. To enhance your understanding of these categories, here is some additional information:

1. Closed Question: "Questions that can be answered with an yes/no response or a very restricted range of answers. " Example: "Do you think this is an advantage? "
2. Open Question: "Questions that allow a wide range of possible answers. " Example: "What is your take on that? "
3. Simple Reflection: "Repetition, rephrasing, or paraphrasing of speaker's previous statement. " Example: It sounds like you're feeling

worried. 4. Complex Reflection: “Repeating or rephrasing the previous statement of the speaker but adding substantial meaning/emphasis to it. “ Example: “Speaker: Mostly, I would change for future generations. Listener: It sounds like you have a strong feeling of responsibility. “ 5. Give Information: “Educating, providing feedback, or giving an opinion without advising. Example: “Logging your cravings is important as cravings often lead to relapses. “ MI Adherent Behaviour Codes: 6. Advise with Permission: “Advising when the speaker asks directly for advice. Indirect forms of permission can also occur, such as when the listener says to disregard the advice as appropriate.” Example: “If you agree with it, we could try to brainstorm some ideas that might help. “ 7. Affirm: “Encouraging the speaker by saying something positive or complimentary. “ Example: You should be proud of yourself for your past’s efforts. 8. Emphasize Autonomy: “Emphasizing the speaker’s control, freedom of choice, autonomy, and ability to decide. “ Example: “It is really up to you to decide. “ 9. Support: “Statements of compassion or sympathy. “ Example: “I know it’s really hard to stop drinking. “ MI Non-Adherent Behaviour Codes: 10. Advise without Permission: “Making suggestions, offering solutions or possible actions without first obtaining permission from the speaker. “ Example: “You should simply scribble a note that reminds you to take a break. “ 11. Confront: “Directly and unambiguously disagreeing, arguing, blaming, criticizing, or questioning the speaker’s honesty. “ Example: Yes, you are an alcoholic. You might not think so, but you are. 12. Direct: “Giving orders, commands, or imperatives. “ Example: “Don’t do that! “ 13. Warn: “A statement or event that warns of something or that serves as a cautionary example. “ Example: “Be careful, DO NOT stop taking meds without discussing with your doctor. “ Other: 14. Self-Disclose: “The listener discloses his/her personal information or experiences. “ Example: “I used to be similar where I get obsessed about how people look. “ 15. Other: “Statements that are not classified under the above codes.” Example: “Good morning, Hi there.”

Again, output only what was asked (MITI categories), do not respond anything else.

I need you to evaluate therapeutic responses to messages seeking for help. The Motivational Interviewing Treatment Integrity (MITI) are types of sentences used for therapeutic responses. I will teach you how to do it by showing you statistics on the MITI categories of sentences that were used by professional counselor to address particular emotion that was caused by a specific cause: —[DATA to insert]—

Grade how professional and therapeutic the response is and justify it using the statistics you learned.

References

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A. et al. (2020), ‘Language models are few-shot learners’, *arXiv preprint arXiv:2005.14165* .
- Dataonomy (2023), ‘What is ai prompt engineering: Examples and how to’, <https://dataonomy.com/2023/01/27/what-is-ai-prompt-engineering-examples-how/> .
- Innodata (2023), ‘How to enhance chatgpt with prompt engineering’, <https://innodata.com/how-to-enhance-chatgpt-with-prompt-engineering/> .
- Jules White, Quchen Fu, S. H. M. S. C. O. H. G. A. E. J. S.-S. and Schmidt, D. C. (2023), ‘A prompt pattern catalog to enhance prompt engineering with chatgpt’.
- Mayur Wankhade, A. C. S. R. . C. K. (2022), ‘A survey on sentiment analysis methods, applications, and challenges’, <https://link.springer.com/article/10.1007/s10462-022-10144-1> .
- Moghaddam, S. R. and Honey, C. J. (2023), ‘Boosting theory-of-mind performance in large language models via prompting’, <https://arxiv.org/abs/2304.11490> .
- Moyers, T. B., Manuel, J. K. and Ernst, D. (2015), Motivational interviewing treatment integrity coding manual 4.2.1.
- Ni, K., Doran, C., Driggs-Campbell, K. and Kochenderfer, M. J. (2020), Planning and execution with flexible sketch plans, in ‘2020 IEEE International Conference on Robotics and Automation (ICRA)’, IEEE, pp. 9486–9492.
- Rashkin, H., Smith, E. M., Li, M. and Boureau, Y.-L. (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’, Association for Computational Linguistics, pp. 5370–5381.
- Suzgun, M. (2022), ‘Prompt-and-rerank: A method for zero-shot and few-shot arbitrary textual style transfer with small language models’, <https://aclanthology.org/2022.emnlp-main.141.pdf> .
- Welivita, A. and Pu, P. (2022), Curating a large-scale motivational interviewing dataset using peer support forums.
- X, M. (2021), ‘Building an aspect-based sentiment analysis pipeline using gpt-3’, <https://medium.com/@xmikex83/building-an-aspect-based-sentiment-analysis-pipeline-using-gpt-3-67e26f2a8a5b>. Accessed: yyyy-mm-dd.
- Yu, D. and Sun, S. (2020), ‘A systematic exploration of deep neural networks for eda-based emotion recognition’, <https://www.mdpi.com/2078-2489/11/4/212> .