

Chain-of-Interaction: Enhancing Large Language Models for Psychiatric Behavior Understanding by Dyadic Contexts

Guangzeng Han*, Weisi Liu*, Xiaolei Huang*, Brian Borsari^{†‡}

*Department of Computer Science, University of Memphis, Memphis, United States

{ghan,wliu9,xiaolei.huang}@memphis.edu

†San Francisco Veterans Affairs Health Care System, San Francisco, United States

brian.borsari@va.gov

‡Department of Psychiatry and Behavioral Sciences, University of California, San Francisco, United States

Abstract—Automatic coding patient behaviors is essential to support decision making for psychotherapists during the motivational interviewing (MI), a collaborative communication intervention approach to address psychiatric issues, such as alcohol and drug addiction. While the behavior coding task has rapidly adapted language models to predict patient states during the MI sessions, lacking of domain-specific knowledge and overlooking patient-therapist interactions are major challenges in developing and deploying those models in real practice. To encounter those challenges, we introduce the Chain-of-Interaction (CoI) prompting method aiming to contextualize large language models (LLMs) for psychiatric decision support by the dyadic interactions. The CoI prompting approach systematically breaks down the coding task into three key reasoning steps, extract patient engagement, learn therapist question strategies, and integrates dyadic interactions between patients and therapists. This approach enables large language models to leverage the coding scheme, patient state, and domain knowledge for patient behavioral coding. Experiments on real-world datasets can prove the effectiveness and flexibility of our prompting method with multiple state-of-the-art LLMs over existing prompting baselines. We have conducted extensive ablation analysis and demonstrate the critical role of dyadic interactions in applying LLMs for psychotherapy behavior understanding.¹

Index Terms—Behavioral coding, large language models, psychotherapy, motivational interview

I. INTRODUCTION

Motivational interviewing (MI) [1] is defined as “particular way of talking with people about change and growth to strengthen their own motivation and commitment.” MI can facilitate change in behaviors associated with mental health issues, which are among the most complex health problems impacting over one billion people worldwide [2]. The Motivational Interviewing Skill Code (MISC) [3] is a coding system that provides comprehensive and structured coding schema for examining therapist and patient behaviors (utterances) and assessing different aspects of the intervention (e.g., therapist empathy). However, MI session assessments require significant time commitment, labor costs, and professional training [4], which may not meet immediate training or clinical needs. An

alternative solution is to build machine learning models to automate coding process in Fig. 1.

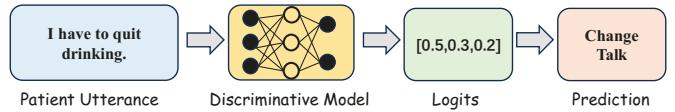


Fig. 1. Diagram illustrating automatic behavioral coding methods based on discriminative models.

Fig. 1 presents current standard process of automatically coding MI sessions that first vectorize patient utterances into neural representations and then feed the representations to discriminative classifiers for final predictions [5]. For example, recent studies have developed MI classifiers by recurrent neural networks (e.g., LSTM or GRU) [6], [7] and BERT [8], [9] for automatically predicting MISC codes at utterance and session levels, which have achieved promising performance. However, a major issue of the existing ML classifiers is missing incorporating the psychiatric domain knowledge, including MISC coding manuals, which guide and train human professionals to annotate patient behaviors. The existing methods require costly and time-consuming human annotations and may not integrate domain-specific knowledge from the MISC coding schema by only feeding utterances into classifiers for predictions, such as MISC concepts and definitions used by human experts in behavioral coding. Thus, a concrete question is: *how can we instruct models to annotate MI sessions following the MISC schema like psychiatric professionals?*

Large language models (LLMs), such as GPT [10], [11] and Llama2 [12], has undergone a rapid evolution in recent years and have demonstrated a significant potential transform mental health and psychotherapy interventions, such as depression detection in social media [13], [14] and psychotherapy chatbot [15]. Effective prompting design for LLMs is the key to achieve precise diagnosis and understanding of patient behaviors. Prompting methods involve how we put together and format questions or commands for the model to instruct generating specific responses. For example, prompting meth-

¹Code available at <https://github.com/trust-nlp/CoI-Psychotherapy>

TABLE I

SAMPLES OF MISC-CODED DATA. THE SAME SUB-CODES MAY BE ENCODED AS DIFFERENT PATIENT CODES DUE TO THEIR DIFFERENT VALENCE (+/-)

Patient code	Sub-code	Description	Therapist utterance (Code)	Patient utterance
Change Talk	Commitment+	Patient makes a changing commitment.	Do you want to change? (Closed Question)	I am going to stop smoking.
	Ability+	Assess client's ability to change.	You've been through a lot. (Support)	I can do it.
	Desire+	A desire to alter the target behavior.	I want to talk about your motivation. (Structure)	I want to quit doing drugs.
	Reason+	Reasons for changing the behavior	What are some reasons for quitting? (Open Question)	I'm killing myself.
Follow/Neutral	N/A	Reporting information.	How often do you drink? (Closed Question)	Usually 4-5 days.
	N/A	Follow therapist.	Do you have children? (Closed Question)	I have three daughters.
	N/A	Ask question.	The support group meets from 4 to 5 pm. (Give Information)	Do you know when they meets?
Sustain Talk	Commitment-	Patient makes a maintaining commitment.	Speeding will cost you your license. (Warn)	I will never slow down when I drive.
	Ability-	Assess client's ability to change.	Only you know what's best for you. (Emphasize Control)	I cannot stop overeating even when I try.
	Desire-	A desire to maintain the behavior.	If you get bored you'll use drugs. (Warn)	I want to keep getting high.
	Reason-	Reasons for maintaining the behavior.	Put your health first! (Direct)	At least I get relaxed when I drink.

ods like Chain of Thought (CoT) [16], Thread of Thought [17], and Retrieval-Augmented Generation (RAG) [18], [19] can effectively solve common sense reasoning tasks — by splitting complex questions into several easier and key parts of general questions. However, these kinds of prompting techniques do not consider the dyadic contexts of MI sessions, interactions between patients and therapists. The dyadic interaction between psychotherapists and patients is a natural characteristic of MI sessions, which is commonly overlooked in the existing prompt approaches. As the interactive context can provide key insights for accurately inferring patient behaviors, the challenges call for new methods to integrate psychiatric knowledge, understand dyadic contexts, and equip with strong reasoning capability.

To address these specific challenges, we propose the *Chain-of-Interaction* prompting method aiming to incorporate domain knowledge from MISC coding schema and in-session interactions. It utilizes a series of prompt stages to enhance the reasoning abilities of LLMs in the task of coding patient utterance, drawing upon psychological domain knowledge and the style of interactions between patients and therapists from three consecutive stages: Interaction Definition, Involvement Assessment, and Valence Analysis. The Interaction Definition stage aims to help LLMs define the therapist-patient interactions by categorizing the unique MISC codes of each utterance. Next, the Involvement Assessment evaluates patient engagement by assessing willingness for self-exploration and emotional expression. Finally, Valence Analysis stage is to integrate the general sentiment of patient and the clues from previous two stages, we expect LLMs can utilize the integrated information to mimic the reasoning process of human experts.

We conduct extensive experiments on datasets derived from real-world MI sessions addressing alcohol usage disorder. We examine three prompting baselines and experiment with four state-of-the-art auto-regressive LLMs, including Llama2 [12], Falcon [20], Mistral [21], and ChatGPT [22]. Our experiments demonstrate the effectiveness of the Chain-of-Interaction (CoI) prompting method, which can outperform multiple baselines by a large margin. The consistent improvements in Macro-F1 and Micro-F1 scores can indicate that CoI may enhance LLMs under data and label imbalance. Furthermore, our ablation study reveals the critical contribution of domain knowledge and dyadic interactions in LLM performance.

II. DATA

We used sessions from two independent randomized clinical trials [23], which conducted MI sessions with 249 university students who had been mandated for alcohol-related infractions (e.g., possession, vandalism). Each MI session lasts between 45 to 60 minutes and were audio recorded. To protect user privacy and identities, these recorded sessions were transcribed into text, de-identified, and coded using the Motivational Interviewing Skills Code (MISC version 2.0) [24]. MISC encodes patient utterances into three mutually exclusive overarching categories: Follow/Neutral, Change Talk, and Sustain Talk. “Follow/Neutral” encompasses responses where the patient aligns with, remains neutral, or seeks information from the therapy without directly addressing behavior change; “Change Talk” involves statements indicating an intention to change the behavior; and “Sustain Talk” includes utterances that suggest maintaining the current behavior. Change language can be conceptualized as being on a continuum with a positive valence indicating Change Talk, a negative valence representing Sustain Talk, and a neutral valence corresponding to Follow/Neutral.

To control data quality, this study did not consider therapy records that had annotation errors or incomplete annotations. Per recording session, we lower-cased utterances and extracted data segments consisting of every ten non-overlapping utterances as one data entry, aiming to provide dyadic interactions between patients and therapists. For each data entry, we utilized the MISC code of the last patient utterance of the entry as the ground truth label. This method ensures that the model’s focus is on the patient’s most recent utterance, accurately reflecting their behavioral state at that particular moment in the therapy session. We summarize data statistics in Table II for raw and processed data to allow for replication.

TABLE II

DATA STATISTICS. \bar{L} REFERS TO THE AVERAGE NUMBER OF TOKENS PER UTTERANCE. \bar{U} IS THE AVERAGE UTTERANCE COUNT PER THERAPY SESSION.

Utterance	Change Talk	Follow/Neutral	Sustain Talk	\bar{L}
48,529	13,298	29,025	6,026	30.2
Session	Female	Male	Therapist	\bar{U}
249	95	154	14	194.9

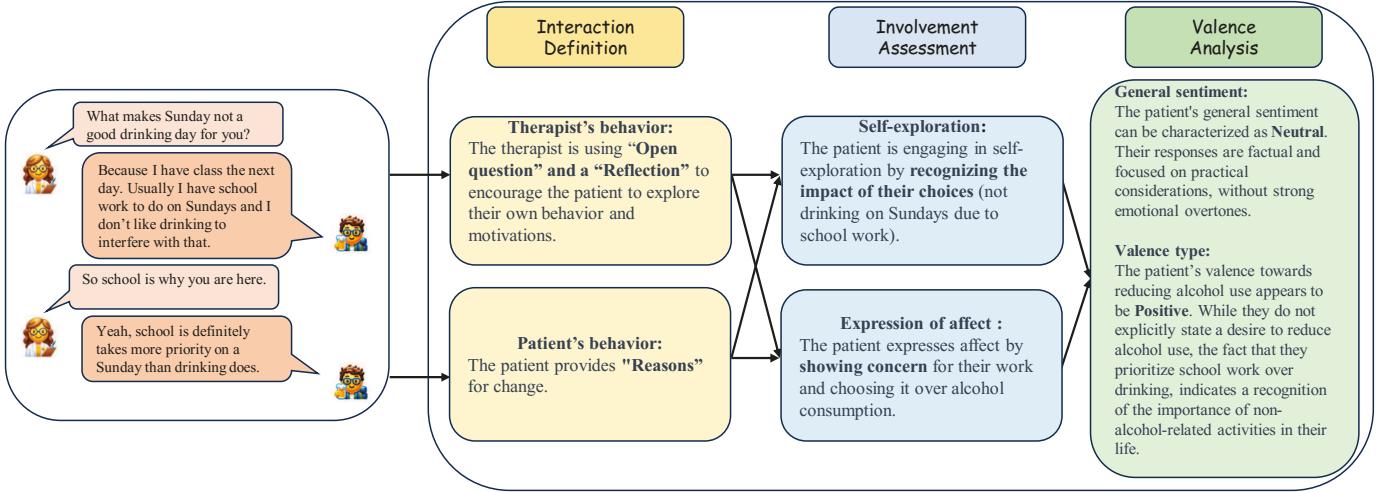


Fig. 2. Overview of the “Chain-of-Interaction” prompting method, we use three colors to denote different stages.

In this study, we utilize the three MISC client language categories for the auto behavior coding task, and we present examples of the MISC codes and utterances in Table I. The examples show interactive and contextual effects exist between patients and therapists that patient utterances can shape therapeutic strategies for interventionists and in turn the conversation engagement will impact on patient behaviors. While past studies [7], [9] have shown concatenating patient utterance and its previous utterances can improve prediction accuracy of the MISC coding task, how the patient-therapist interactions can inform model predictions is underexplored. To effectively utilize this information, we proposed the Chain-of-Interaction (CoI) prompting, which helps large language models in automatically coding motivational interview patient behavior by leveraging the interactions between patients and therapists, as well as the engagement of patients during counseling sessions.

III. THE CHAIN-OF-INTERACTION PROMPTING

MI efficacy theory [25], [26] suggests that dyadic interactions play a critical role in therapist skills, patient behaviors, and outcomes — guide and inspire the development of our new prompting method. Two aspects of this dyadic interaction are linked to MI efficacy. First, there is a technical component in which therapist MI Consistent (MICO) utterances can selectively elicit and strengthen client change language which is predictive of subsequent behavior change. There is also a relational component which incorporates the interpersonal aspects of the therapist and client relationship such as MI Spirit and empathy. We propose the *Chain-of-Interaction* prompting, aiming to incorporate the psychological domain knowledge and the two aspects of therapist-patient interaction to promote the contextual awareness of large language models in the behavioral coding task. Fig. 2 illustrates our prompting framework of Chain-of-Interaction with three continuous stages: 1) Interaction Definition, 2) Involvement Assessment, 3) Valence Analysis. Through this structured prompting approach, we

imbue LLMs with a new level of understanding and analytical depth that parallels professional psychological expertise.

First, the Interaction Definition stage enables the LLMs learning MI domain knowledge and objectively understanding the interaction patterns. Next, the Involvement Assessment stage inquiries LLMs to examine emotional and self-expression clues from the patient utterances to infer the patient’s engagement level, aiming to bridge correlations between patient mental states and their behavior outcomes. Finally, in the Valence Analysis stage, LLMs predict the patient’s utterance valence by integrating information about the patient’s emotional state and interaction pattern. Through this structured approach, we imbue LLMs with a level of understanding and analytical depth that parallels professional psychological expertise. More details about the Chain-of-Interaction prompting methods can be found in Fig. 3.

A. Interaction Definition

Our goal is to equip predictive models with strong reasoning capabilities as human professionals, while concatenating dyadic interactions and patient utterances as features is the major approach for existing end-to-end MISC classifiers [6], [7], [9], [27], [28]. Unlike the feature integration approach, human professionals ideally utilize MICO strategies to interact with patients and the MISC coding schema can be used to annotate patient behaviors, which is a fundamental difference. To encounter the wide-existing yet unsolved issue, we introduce the *Interaction Definition* stage, which allows Large Language Models (LLMs) to understand and interpret the interactive contexts of therapists and patients by the guidance of MISC behavior coding schema and definitions.

We design instructions to guide LLMs understanding interactions between therapists and patients and mental states of patients via two major parts of prompt modules, 1) definitions of MISC codes and coping with 2) dyadic contexts. The instructional prompts aim to inspire LLMs to capture MI domain knowledge and master the MISC coding schema. To

Code Descriptions [C]	Utterance History [X]	Prompt Template T + Code Descriptions [C] + Utterance History [X]
<p>Affirm: Acknowledge admirable behavior or quality of the patient.</p> <p>.....</p> <p>.....</p> <p>Support: Sympathetic behavior.</p>	<p>Therapist utterance[1] Patient utterance [1]</p> <p>.....</p> <p>.....</p> <p>Therapist utterance[5] Patient utterance [5]</p>	<p><u>Task Definition:</u> Based on the Motivational Interview transcript between a psychotherapist and a patient with alcohol abuse issues, please identify the patient's valence about changing their behavior (i.e. reducing alcohol abuse) as either neutral, positive, or negative.</p> <p>Here are the motivational interviewing behavior Code Descriptions [C]</p> <p><u>Test Instance:</u> Utterance History[X]</p> <p><u>Chain-of-Interaction Prompt:</u> Based on the Motivational Interview record, answer the following Chain-of-Interaction questions.</p> <p>Stage1: What are the prime motivation interview behaviors of the doctor and patient?</p> <p>Stage2: How is the involvement of patient in this interview, Specifically does the patient self-explore, and does the patient have any emotional expression?</p> <p>Stage3: What is the general sentiment of patient language? And What is the patient's valence about reducing the alcohol use?</p> <p><u>Output:</u></p>

Fig. 3. Details of the prompt design for the “Chain-of-Interaction” method.

achieve this, we start with feeding LLMs with extracted MI definitions from MISC coding manuals [1], [24]. After the initial instruction, we then ask LLMs to review a patient utterance with its dyadic contexts. The process is to enable LLMs behave like human professionals that apply domain knowledge by their in-context learning ability. We request LLMs to attempt annotating utterances of the dyadic context by the MISC therapist MICO categories (e.g. open-ended questions, affirmation, reflections) and the patient utterance subcategories shown in Table I. We illustrate examples of the process and prompt examples in Fig. 2 and 3. For example, Fig. 2 shows that LLMs encode the patient and therapist utterances as “Reason” and “Open question”, respectively, which will help LLMs understand patient dynamic mental states and provide more information for the following stages. The stage meets our initial goal by instructing LLMs by the behavior coding task, incorporating domain knowledge as instructions, and equipping LLMs with domain-specific reasoning capabilities.

B. Involvement Assessment

MI coding teams objectively code MISC patient behavior while also subjectively rating the relational aspects of the session [23] (e.g., engagement). Prior research [29] has also demonstrated that patients who are highly engaged in treatment have more Change Talk, and less Sustain Talk, therefore assessing patient engagement will help the model better understand patient tendencies regarding target behavior. To utilize these auxiliary subjective paths to help the objective behavioral coding, we propose the *Involvement Assessment* stage, which makes LLMs to mimic the rating process of the MI coders by assessing the involvement of patient.

In this stage, we formulate prompts to remind LLMs to revisit the original dyadic contexts and the Interaction Definitions (i.e. outputs of *Interaction Definition stage*), extracting two important cues about engagement from them: self-exploration and emotional expression [30], [31]. These prompts aims to guide LLMs to emulate the rating process professionals use when coding MI patient behavior. To aid in simulating this process, we first require LLMs to consider not only the original dyadic contexts but also the Interaction

Definitions of stage 1 outputs. We then prompt LLMs to assess patient engagement through aspects of self-exploration and emotional expression. The example in Fig. 2 illustrates how the LLM interprets specific methods of self-exploration and emotional expression based on the “reason” behavior of patients from the first stage.

From the clear interaction patterns obtained in the *Interaction Definition* stage to the subjective assessment of engagement in this stage, LLMs can benefit from domain-specific knowledge and emulate the thinking habits of professionals. Furthermore, the emotional expressions identified in this stage serve as a foundation for the analysis in the final stage.

C. Valence Analysis

In MI, valence reflects the client’s utterances regarding changing the target behavior. A positive valence indicates Change Talk, a negative valence represents Sustain Talk, and a neutral valence corresponds to Follow/Neutral. According to their correspondence, we transform the behavioral coding task into valence coding task. Valence is a component of sentiment, focusing specifically on the positive-negative axis of emotional response. Sentiment, however, is a more comprehensive term that captures the full spectrum of emotional states and attitudes, including valence as one of its dimensions. Based on this characterization, we propose the *Valence Analysis* stage, which aims to help LLMs better perform the valence coding task through two sub-stages. The first sub-stage involves utilizing general sentiment analysis to aid LLMs in comprehending sentimental states of patients. The second one integrates outputs obtained from all three stages to make LLMs to more comprehensively utilize the domain-specific knowledge and the interactive contexts.

In this stage, we start with instructing LLMs to perform a generic sentiment analysis based on the original utterances and the emotional expressions obtained from the *Involvement Assessment* stage. The goal is to allow LLMs to learn more about the patient’s sentimental state before reasoning about our final task objective: Valence, through a simpler and more generalized task. After the sentiment analysis, we finally prompt LLMs to review the outputs obtained through the three progressive stages, which include the required objective

criteria in the MISC coding schema, as well as the subjective rating that professionals would refer to when coding, and the sentimental state of the patient’s language. By integrating the output from these three stages, LLMs not only understood the pattern of interaction between the patient and therapist, but also assessed the patient’s level of engagement as well as analyzed the patient’s emotional state. These procedures help LLMs to reason and code more rationally about the patient’s Valence by utilizing domain knowledge and patient-therapist interactions as professional coders do.

IV. EXPERIMENT

To examine the effectiveness of our proposed ‘Chain-of-Interaction’ (CoI) prompting method, we conducted a comparative analysis with state-of-the-art baselines and performed a detailed ablation study. We set our experimental results against three established baseline methods: Zero-Shot prompting [32], [33], Few-Shot prompting [34], [35], and Zero-Shot Chain of Thought (ZeroCoT) [36] prompting. In addition, our work includes an ablation study focusing on the proposed Chain-of-Interaction (CoI) method, where we removed each individual stage within the CoI. By assessing the performance resulting from the removal of each stage, we measured its contribution to the overall effectiveness. We selected four representative auto-regressive LLMs to perform the experiments, including Llama2-13B-Chat [12], Falcon-7B-Instruct [20], Mistral-7B-Instruct [21], and ChatGPT [22]. We used model-specific official prompt templates and followed each model’s default sampling strategies. Our assessment employed Micro-F1 score to gauge overall model accuracy and Macro-F1 score to ensure fairness across different classes, offering a balanced evaluation of the models’ performance in varied environments.

A. Baselines

To demonstrate the effectiveness of our proposed Chain-of-Interaction prompting method, we compare it with three state-of-the-art prompting techniques: 1) Zero-Shot Prompting [32], [33], 2) Few-Shot Prompting [34], [35], 3) Zero-Shot Chain-of-Thought (ZeroCoT) Prompting [36].

1) *Zero-Shot Prompting*: Zero-shot prompting [32], [33], where LLMs receive only a sample and task description without corresponding examples or specially designed prompts, may be insufficient for complex and domain-specific tasks. This approach relies entirely on the LLMs’ inherent knowledge and adaptability, potentially causing it to bypass important reasoning steps and overlook critical information when lacking necessary context. In MI behavioral coding, for instance, Zero-shot prompting might lead to suboptimal performance due to the absence of domain knowledge and contextual understanding. In contrast, our CoI can integrate domain knowledge, enabling LLMs to pay attention to underlying information that are unavailable in the utterances.

2) *Few-Shot Prompting*: LLMs have been proven to have the capability of in-context learning, which makes Few-Shot Prompting [34], [35] more effective than Zero-Shot Prompting. Before performing a specific task, Few-Shot prompting

provides LLMs with a demonstration of examples (i.e., a small number of “shots”) for tuning, including task examples and their corresponding ground truth outputs. These examples serve as a brief learning guide, providing LLMs with some context or insight about the nature of the task. However, Few-Shot prompting does not actively guide the model in breaking down the task into multiple reasoning steps. Therefore, for tasks requiring multi-step reasoning, such as basic mathematical problems [37] and the MI behavioral coding in this study, Few-shot prompting still struggles to perform well. In contrast, our CoI prompting reduces the complexity of the behavioral coding task by breaking it down into three sequential sub-tasks and prompt stages.

3) *Zero-Shot Chain-of-Thought Prompting*: Marking a departure from traditional Chain-of-Thought [16], which typically relies on hand-crafted, detailed Few-shot examples for each task, Zero-shot Chain-of-Thought (ZeroCoT) [36] is a template-based prompting method for LLMs. It focuses on sequential reasoning without the need for specific training examples. By employing a generalized, one-size-fits-all prompt such as “Let’s think step by step”. ZeroCoT effectively guides LLMs to decompose various complex tasks into multiple reasoning steps. Its robustness has been proven across a wide range of applications, from intricate arithmetic and symbolic reasoning problems to other logical reasoning tasks. However, MI behavioral coding is not merely a task of logical reasoning. It also requires the use of domain-specific knowledge and the characteristics of interactions between patients and therapists to assist in coding. For this reason, our CoI method, through three well-designed prompt stages, breaks the MI behavior coding task down into three sequential sub-tasks based on professional guidelines and coding schema, rather than letting LLMs decompose the task on their own as ZeroCoT. This approach aids LLMs in understanding the interactions between patients and therapists and simulates the reasoning process of human coders.

B. Models

We performed experiments on four LLMs, namely Llama2-13B-Chat [12], Falcon-7B-Instruct [20], Mistral-7B-Instruct [21], and ChatGPT [22], which are available by OpenAI API² or Hugging Face³. The reason for selecting these LLMs over their non-fine-tuned base versions (e.g. Llama2-13B-Base, Falcon-7B-Base) is their superior ability to follow dialogue-style instructions.

1) *Llama2-13B-Chat*: Llama2 [12] is an auto-regressive Transformer model pre-trained on publicly available online data. Additionally, it utilizes ghost attention [38] and grouped query attention to enhance consistency in multi-turn dialogues and extend its context length, respectively. In this study, we use the Llama2-Chat variant, which undergoes supervised fine-tuning followed by reinforcement learning from human feedback (RLHF) [11], [39], [40], which includes rejection

²<https://platform.openai.com/docs/guides/gpt/chat-completions-api>

³<https://huggingface.co/>

sampling and proximal policy optimization. While enhancing safety, its performance remains competitive with other open-source LLMs.

2) *Falcon-7B-Instruct*: [41] demonstrated the intrinsic link between the performance of LLMs and the quality of their training data. Consequently, the Falcon-7B-Instruct [20], known for its superior performance, was trained on the meticulously curated RefinedWeb [41] dataset. To further enhance its efficiency and reduce computational overhead, Falcon employs multi-query and flash attention [42] mechanisms, enables it to support sequences up to 2,048 tokens.

3) *Mistral-7B-Instruct*: Mistral-7B-Instruct [21], an instruction-tuned version of the Mistral model checkpoint, leverages the capabilities of Grouped-Query Attention and Sliding Window Attention to expedite inference processes and handle longer sequences. Compared to LLMs with an equivalent number of parameters, Mistral performs better in various domains, including Commonsense Reasoning, Reading Comprehension, and Mathematics.

4) *ChatGPT*: ChatGPT [22] represents an evolution in the Generative Pre-trained Transformer (GPT) series, with a focus on conversational capabilities. It showcases superior language understanding and generation abilities, excelling in benchmarks like SuperGLUE [43] and HumanEval [44], which require deep contextual insight and logical coherence. It also integrates reinforcement learning from human feedback (RLHF) [11], [39], [40], allowing it to generate more contextually appropriate and human-like responses in dialogue scenarios.

C. Sampling Strategy and Experiment Details

To ensure stable and high-quality output, we adhere to the default sampling strategies specific to each LLM. For ChatGPT, we utilize GPT-3.5-Turbo via its *ChatCompletion* API, employing Nucleus Sampling [45] with top-p and temperature parameters set to 1. This method known for enhancing output diversity by selecting the next word from a restricted set of highly probable options, thus facilitating more creative and varied responses. For all other LLMs, we employ Greedy Sampling, a strategy that selects the most probable next word, ensuring deterministic and predictable text generation. Considering the different context length limitations of various LLMs, we established a fair comparison by setting the Few-Shot N value to 1, which means providing one example before each sample. Additionally, the Few-shot examples were randomly chosen from the support set to mitigate the risk of data leakage.

D. Answer Extraction

The open-ended nature of language generation makes it very challenging to evaluate the performance of LLMs on classification tasks. Two generalized evaluation models are to convert the classification task into multiple-choice questions (MCQ) and then extract the probability of the first token to select the label, and the other is to extract the corresponding answer from the complete answer based on a regular expression. Previous

study [46] demonstrated the MCQ pattern may not consistently reflect the final response output due to varying response modes of the models. Therefore, we use regular expressions to extract the first output that conforms to the format in the third stage (Valence Analysis). For example, if a model response is: “The patient’s valence should be coded as neutral or positive,” we extract the first matching token “neutral” as the model prediction. If the models do not follow the instructions, leading to no tokens in the output matching any labels, in this case, we randomly select a label. Furthermore, the models we selected are aligned to ensure their outputs are consistent with human professional guidelines. But, this occasionally triggers the Content Safety Policy when generating content related to psychotherapy. Therefore, we opted to exclude those samples that trigger this policy. These special treatments reduce the errors caused by the uncertainty of the generated model and improve the reliability of the research.

V. RESULTS

Due to the data imbalance shown in Table II, we chose Micro-averaged F1 and Macro-averaged F1 scores for our evaluation. Micro-averaged F1 assesses overall accuracy, while Macro-averaged F1 handles label imbalance by giving equal weight to all classes. These metrics offer a comprehensive view of our Chain-of-Interaction approach’s performance, both for comparing it against baseline methods and for conducting an ablation study.

A. Main Results

Table III presents performance results for different prompting methods applied to the MI behavioral coding task across various large language models (LLMs). The results show that our CoI prompting significantly outperforms the state-of-the-art prompting baselines. For example, our method outperforms the second-best approach, Zero-shot Chain-of-Thought [36], in average performance across four LLMs. Specifically, it shows improvements of 6.3% in Micro-F1 and 1.3% in Macro-F1 score. This superior performance can be attributed to our method’s strategic breakdown of the behavioral coding task into three key stages that leverage domain knowledge. By segmenting the task into multiple interactive stages, our approach facilitates step-by-step reasoning and enables the model to navigate the complexities of the task in a structured and informed way. This progressive reasoning simulates the layered understanding that a psychological professional may employ, leading to more accurate predictions and a more profound interpretation of patient behaviors.

While different baseline prompting methods exhibit significant performance variations across various LLMs, our method consistently achieves state-of-the-art results. For instance, the performance of ZeroCoT may be comparable to our proposed CoI method on LLaMA2 but appears closer to the simpler Zero-shot prompting on Falcon, highlighting the performance instability of baseline methods. Few-shot prompting improves performance over Zero-shot on three LLMs, but its impact on ChatGPT may be minimal, possibly suggesting its training

TABLE III
MAIN RESULTS FOR THE BEHAVIORAL CODING TASK, EXPRESSED IN PERCENTAGES. WE BOLDEN THE BEST PERFORMANCE.

Methods (%)	Llama2		Falcon		Mistral		ChatGPT		Average	
	Micro-F1	Macro-F1								
Zeroshot	42.7	33.6	40.3	31.0	39.6	30.4	52.0	38.0	43.7	33.3
Fewshot	46.5	33.8	43.9	31.1	47.3	33.6	51.9	38.8	47.4	34.3
ZeroCoT	62.6	33.6	40.4	30.9	44.9	31.5	55.6	41.0	50.9	34.3
CoI	63.5	34.5	53.2	32.6	52.0	34.6	60.2	40.5	57.2	35.6

may already include similar tasks. Despite this instability in baseline methods, our approach stands out by consistently outperforming them, demonstrating robustness and adaptability across different LLMs.

B. Ablation Study

To systematically explore the contribution of each stage in our Chain-of-Interaction (CoI) method, we conduct a series of ablation studies and present the results in Table IV. These studies start with a baseline scenario, termed *Zeroshot*, where all CoI stages are omitted. This baseline indicates the performance of models without the benefit of any structured CoI stages. Subsequently, we gradually removed Interaction Definition stage (*w/o ID*), Involvement Assessment stage (*w/o IA*), and Valence Analysis stage (*w/o VA*).

We summarize our results in Table IV. Removing the Interaction Definition stage (*w/o ID*) leads to LLMs beginning their subjective rating process about the patient's emotional expression and self-exploration, without fully understanding the therapist-patient interaction context by the guidance of MISC behavior coding schema. Compared to the full CoI method, Micro-F1 and Macro-F1 decrease by 8.7% and 0.9%, respectively, which reflect the impact of the loss in initial insight into the dynamics of the interaction. When we exclude the Involvement Assessment stage (*w/o IA*), the LLMs no longer mimic the subjective process used by MI professionals to assess patient involvement. In this setting, LLMs analyze the patient's general sentiment and valence directly after reasoning about the specific patterns of interaction between the patient and the therapist. Valence was analyzed without the help of the patient's self-exploration and emotional expression, which are two key indicators of it. The absence results in a decrease of 7.2% in Micro-F1 and 1.7% in Macro-F1 scores compared to the full method. Finally, skipping the general sentiment analysis in the Valence Analysis stage (*w/o VA*) deprives the model of fully indications of the patient's emotional state. This absence can affect the model's ability to accurately code patient behaviors, as it misses out on initial emotional cues that could influence the overall understanding of the patient's attitude towards change. In this scenario, the average performance of the LLMs experiences a 1.5% drop in Micro-F1 and a 0.3% drop in Macro-F1 compared to the full method.

In conclusion, the full CoI method yields the best results, showcasing the method's effectiveness in leveraging domain knowledge to enable the model to perform detailed reasoning. Each stage contributes to a composite understanding, which is reflected in the model performance comparisons, confirming

that the integrated stages of CoI are crucial for a nuanced task like MI behavioral coding.

VI. RELATED WORK

A. Automatic Behavioral Coding

Automatic behavior coding is a critical task to examine the fidelity of MI sessions which, if conducted in real time, can have valuable training and clinical applications. Due to domain-expertise and time-consuming natures of the task, the common strategy is to develop a machine learning classifier to predict and assess patient behaviors in MI sessions, such as substance disorder [47], suicide [48], and alcohol addiction [6]. In recent years, neural models have dominated the automatic behavioral coding task and showed their supremacy performance than non-neural approaches, such as linguistic features [49] and topic models [4]. Existing studies [6], [50]–[52] develop end-to-end classification pipelines by encoding patient utterances into neural feature vectors and utilizing the vectors for predictions by neural classifiers, such as recurrent neural network (RNN). For example, [6], [7], [53] deployed Long-Short Term Memory (LSTM) or Gated Recurrent Unit (GRU) to enrich contextual representations and [50], [51] extended the RNN variants (LSTM and GRU) with attention mechanisms, which identify salient words and patterns in utterances. More recent studies [28], [52] have switched to pre-trained language models (e.g., BERT [8]) as the data feature extractor and neural network classifier, which requires additional fine-tuning steps by the annotated MI corpora.

While the behavioral coding task has achieved promising performance, lacking explicit incorporation of domain knowledge is a major issue of the existing methods, which can lead to unreliable performance comparing to human professionals and get worse when a large amount of training data is not available. In this study, we fill the domain knowledge issue, propose a new prompt learning approach, and employ LLMs for the automatic coding of motivational interviewing, which can enable classification models behave as domain experts and learn the domain knowledge required for MISC coding as instructions.

B. Large Language Models for Mental Health

The field of mental health research is witnessing significant advancements with the integration of LLMs, as illustrated by a series of pioneering studies [54]–[56]. Their comprehensive analysis discusses the potential and limitations of LLMs for mental health, covers technologies ranging from pre-training to instruction tuning and prompt tuning.

TABLE IV
RESULTS OF THE ABLATION ANALYSIS: UTILIZING MICRO-F1, MACRO-F1, EXPRESSED IN PERCENTAGES

Methods (%)	Llama2		Falcon		Mistral		ChatGPT		Average	
	Micro-F1	Macro-F1								
Zeroshot	42.7	33.6	40.3	31.0	39.6	30.4	52.0	38.0	43.7	33.3
w/o ID	48.4	33.5	43.1	30.6	50.0	33.9	52.5	40.8	48.5	34.7
w/o IA	54.6	33.5	45.7	31.1	43.4	31.2	56.2	39.9	50.0	33.9
w/o VA	62.9	33.9	51.4	32.7	50.1	33.8	58.5	40.6	55.7	35.3
CoI	63.5	34.5	53.2	32.6	52.0	34.6	60.2	40.5	57.2	35.6

Consistent with their exploratory research, both instruction tuned and frozen LLMs have been demonstrated to be productive in the Mental health domain. When high-quality data is available, instruction tuning [57], [58] LLMs tends to significantly improve their ability to perform on relevant tasks. For example, ChatCounselor [15] and ChatDoctor [59], as consulting chatbots, leverage instruction-tuned open-source LLMs to achieve performance levels comparable to GPT-4 [10], while requiring significantly fewer computational resources. And [60] conducted instruction tuning on LLMs using real-world psychological Q&A sessions, enabling these models to acquire psychological knowledge and enhance their capability to provide counseling services. In scenarios where high-quality data is lacking, existing studies usually utilize state-of-the-art LLMs to generate synthetic data for fine-tuning their LLMs. For instance [13] used ChatGPT to generate instructions for training open-source LLMs, the instruction-tuned models achieved state-of-the-art performance on mental health detection tasks. Although the instruction tuned LLMs are intuitively more effective than the frozen LLMs, the frozen LLMs with special prompting methods [16], [17], [61], [62] or feature extractors [63] can also perform well with health tasks, such as suicidal risk classification [61] and cognitive distortions detection [62]. Specifically, [61] employed three strategies to assess the performance of LLMs on suicidal risk classification, including Zero-shot prompting, Few-shot prompting and instruction tuning. The results show that both prompting and instruction tuning can improve the reasoning ability of LLMs. Moreover, [62] introduced the Diagnosis of Thought (DoT) prompting framework, which is designed to strategically prompt LLMs to generate diagnosis rationales, with a particular focus on the detection of cognitive distortions.

Likewise, our Chain-of-Interaction method focuses on prompting frozen LLMs. Our approach differs from the existing prompting methods of LLMs for mental health studies, including Zero Chain-of-Thought [36] and [34], [35]. Our approach deconstructs reasoning steps grounded in domain-specific knowledge and the interaction patterns between patient and therapist, while Zero Chain-of-Thought [36] decomposes based on general knowledge and the Few-shot [34], [35] method relies on in-context learning without considering domain knowledge derived from coding schema.

VII. CONCLUSION

In this study, we focus on the task of automatic coding of patient utterances during MI sessions. LLMs have demon-

strated excellent performance on many tasks, and the prompting method is a key factor in their performance. Previous prompting methods have enabled LLMs to excel in tasks like elementary mathematics and other common sense reasoning, but they fall short in domain-specific tasks like coding patient behavior that require specialized knowledge and emphasize the dyadic contexts.

To address these specific challenges, we propose the *Chain-of-Interaction* prompting method, which aims to leverage domain knowledge and patient-therapist interactions to enhance reasoning in LLMs. This method decomposes the task of coding patient utterances in MI into multiple key reasoning steps through three sequential stages. It leverages the interaction characteristics between patients and therapists, allowing the large language model to reason using psychological domain knowledge about behavioral coding without further training.

On a real-world dataset of motivational interviews, we compared our proposed Chain-of-Interaction prompting method with three other popular prompting methods using four advanced large language models. The results indicate that our method achieves state-of-the-art performance. We also conducted an ablation analysis, and the experimental results show that removing any stage of the Chain-of-Interaction leads to a significant decrease in performance, demonstrating the effectiveness of each stage.

VIII. LIMITATION

In this study, we focus on uni-modal large language models that are trained only on text, unlike multi-modal models like GPT-4V [10] which can also process audio. This limitation means that our models cannot leverage audio features from therapy recordings that have not been transcribed. Additionally, the sensitive nature of privacy in psychotherapy presents significant challenges in obtaining experimental data. Therefore, our experiments were only conducted using data from a select group of college students who were enrolled in mandatory alcohol cessation interventions.

IX. ETHICS STATEMENT

We access the raw data in accordance with data agreements and underwent required training. Out of respect for ethical and privacy concerns, we will not release any clinical data that could be linked to individual patients. Instead, we are dedicated to sharing our code and detailed guidelines to facilitate the replication of our results. Our focus lies primarily on computational methodologies, and we do not engage in

direct data collection from human subjects. Furthermore, our institution's review board has verified that our study does not require IRB approval.

X. ACKNOWLEDGMENT

We would like to thank Precious Jones (NSF REU under the project IIS-2245920, PI: Huang) for her contributions in preprocessing the data errors and extracting coding schema from MI coding manuals. The codes analyzed were created in the NIH-NIAAA R01AA017427 project (PI: Borsari). We also appreciate the insightful comments of the ICHI-24 anonymous reviewers.

REFERENCES

- [1] W. R. Miller and S. Rollnick, *Motivational interviewing: Helping people change, 3rd edition*, ser. Applications of motivational interviewing. New York, NY, US: Guilford Press, 2013.
- [2] M. Freeman, "The world mental health report: transforming mental health for all," *World Psychiatry*, vol. 21, pp. 391–392, 10 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/10.1002/wps.21018>
- [3] T. B. Moyers, T. Martin, J. K. Manuel, and W. R. Miller, "The motivational interviewing treatment integrity (miti) code: Version 2.0," 2003.
- [4] D. C. Atkins, M. Steyvers, Z. E. Imel, and P. Smyth, "Scaling up the evaluation of psychotherapy: evaluating motivational interviewing fidelity via statistical text classification," *Implementation Science*, vol. 9, p. 49, 2014.
- [5] M. Tanana, K. A. Hallgren, Z. E. Imel, D. C. Atkins, and V. Srikanth, "A comparison of natural language processing methods for automated coding of motivational interviewing," *Journal of Substance Abuse Treatment*, vol. 65, pp. 43–50, 2016, motivational Interviewing in Substance Use Treatment. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0740547216000222>
- [6] X. Huang, L. Liu, K. Carey, J. Woolley, S. Scherer, and B. Borsari, "Modeling temporality of human intentions by domain adaptation," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, E. Riloff, D. Chiang, J. Hockenmaier, and J. Tsujii, Eds. Brussels, Belgium: Association for Computational Linguistics, Oct.-Nov. 2018, pp. 696–701. [Online]. Available: <https://aclanthology.org/D18-1074>
- [7] J. Cao, M. Tanana, Z. Imel, E. Poitras, D. Atkins, and V. Srikanth, "Observing dialogue in therapy: Categorizing and forecasting behavioral codes," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, A. Korhonen, D. Traum, and L. Márquez, Eds. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 5599–5611. [Online]. Available: <https://aclanthology.org/P19-1563>
- [8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186.
- [9] L. Tavabi *et al.*, "Multimodal automatic coding of client behavior in motivational interviewing," in *Proceedings of the 2020 International Conference on Multimodal Interaction*, ser. ICMI '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 406–413.
- [10] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," 2023. [Online]. Available: <https://arxiv.org/abs/2303.08774>
- [11] Ouyang *et al.*, "Training language models to follow instructions with human feedback," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. New Orleans, Louisiana: Curran Associates, Inc., 2022, pp. 27730–27744. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf
- [12] H. Touvron *et al.*, "Llama 2: Open foundation and fine-tuned chat models," 2023. [Online]. Available: <https://arxiv.org/abs/2307.09288>
- [13] K. Yang, T. Zhang, Z. Kuang, Q. Xie, S. Ananiadou, and J. Huang, "Mentallama: Interpretable mental health analysis on social media with large language models," 2023. [Online]. Available: <https://arxiv.org/abs/2309.13567>
- [14] S. Ji, T. Zhang, L. Ansari, J. Fu, P. Tiwari, and E. Cambria, "MentalBERT: Publicly available pretrained language models for mental healthcare," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, J. Odijk, and S. Piperidis, Eds. Marseille, France: European Language Resources Association, Jun. 2022, pp. 7184–7190.
- [15] J. M. Liu, D. Li, H. Cao, T. Ren, Z. Liao, and J. Wu, "Chatcounselor: A large language models for mental health support," 2023. [Online]. Available: <https://arxiv.org/abs/2309.15461>
- [16] J. Wei *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 24 824–24 837. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf
- [17] Y. Zhou, X. Geng, T. Shen, C. Tao, G. Long, J.-G. Lou, and J. Shen, "Thread of thought unravelling chaotic contexts," 2023. [Online]. Available: <https://arxiv.org/abs/2311.08734>
- [18] Y. Zhang, M. Zhu, Y. Gong, and R. Ding, "Optimizing science question ranking through model and retrieval-augmented generation," *International Journal of Computer Science and Information Technology*, vol. 1, no. 1, pp. 124–130, 2023.
- [19] Z. Jing, Y. Su, Y. Han, B. Yuan, H. Xu, C. Liu, K. Chen, and M. Zhang, "When large language models meet vector databases: A survey," 2024. [Online]. Available: <https://arxiv.org/abs/2402.01763>
- [20] E. Almazrouei *et al.*, "The falcon series of open language models," 2023. [Online]. Available: <https://arxiv.org/abs/2311.16867>
- [21] A. Q. Jiang *et al.*, "Mistral 7b," 2023. [Online]. Available: <https://arxiv.org/abs/2310.06825>
- [22] OpenAI, "Chatgpt: Optimizing language models for dialogue," <https://openai.com/blog/chatgpt/>, 2022, accessed: 2023-07-24.
- [23] B. Borsari *et al.*, "In-session processes of brief motivational interventions in two trials with mandated college students," *Journal of consulting and clinical psychology*, vol. 83, pp. 56–67, 2 2015.
- [24] W. R. Miller, T. B. Moyers, D. Ernst, and P. Amrhein, "Manual for the motivational interviewing skill code (misci) version 2.1," *Substance Abuse and Addiction (CASAA)*, University of New Mexico, vol. 8, pp. 901–4, 2008.
- [25] W. R. Miller and G. S. Rose, "Toward a theory of motivational interviewing," *The American psychologist*, vol. 64, pp. 527–37, 9 2009.
- [26] M. Magill and K. A. Hallgren, "Mechanisms of behavior change in motivational interviewing: do we understand how mi works?" *Current opinion in psychology*, vol. 30, pp. 1–5, 12 2019.
- [27] K. Mishra, P. Priya, M. Burja, and A. Ekbal, "e-THERAPIST: I suggest you to cultivate a mindset of positivity and nurture uplifting thoughts," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023, pp. 13 952–13 967. [Online]. Available: <https://aclanthology.org/2023.emnlp-main.861>
- [28] T. Tran *et al.*, "Multimodal analysis and assessment of therapist empathy in motivational interviews," in *Proceedings of the 25th International Conference on Multimodal Interaction*, ser. ICMI '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 406–415.
- [29] A. Gagnier, "Respiratory syncytial virus: Motivational interviewing: A powerful tool to address vaccine hesitancy," *Canada Communicable Disease Report*, vol. 46, no. 4, p. 93, 2020.
- [30] S. Rollnick, W. R. Miller, C. C. Butler, and M. S. Aloia, "Motivational interviewing in health care: Helping patients change behavior," *COPD: Journal of Chronic Obstructive Pulmonary Disease*, vol. 5, pp. 203–203, 1 2008.
- [31] S. A. Cole, D. Sannidhi, Y. T. Jadotte, and A. Rozanski, "Using motivational interviewing and brief action planning for adopting and maintaining positive health behaviors," *Progress in Cardiovascular Diseases*, vol. 77, pp. 86–94, 2023.

[32] W. Wang, V. W. Zheng, H. Yu, and C. Miao, “A survey of zero-shot learning: Settings, methods, and applications,” *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, jan 2019.

[33] W. Yin, J. Hay, and D. Roth, “Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3914–3923. [Online]. Available: <https://aclanthology.org/D19-1404>

[34] T. B. Brown *et al.*, “Language models are few-shot learners,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS’20. Red Hook, NY, USA: Curran Associates Inc., 2020.

[35] X. Liu *et al.*, “Large language models are few-shot health learners,” 2023. [Online]. Available: <https://arxiv.org/abs/2305.15525>

[36] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, “Large language models are zero-shot reasoners,” 2023. [Online]. Available: <https://arxiv.org/abs/2205.11916>

[37] K. Cobbe *et al.*, “Training verifiers to solve math word problems,” 2021. [Online]. Available: <https://arxiv.org/abs/2110.14168>

[38] Y. Bai *et al.*, “Constitutional ai: Harmlessness from ai feedback,” 2022. [Online]. Available: <https://arxiv.org/abs/2212.08073>

[39] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, “Deep reinforcement learning from human preferences,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2017/file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf

[40] N. Stiennon *et al.*, “Learning to summarize with human feedback,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 3008–3021. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/1f89885d556929e98d3ef9b86448f1951-Paper.pdf

[41] G. Penedo *et al.*, “The refinedweb dataset for falcon llm: Outperforming curated corpora with web data, and web data only,” 2023. [Online]. Available: <https://arxiv.org/abs/2306.01116>

[42] T. Dao, D. Fu, S. Ermon, A. Rudra, and C. Ré, “Flashattention: Fast and memory-efficient exact attention with io-awareness,” in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 16 344–16 359. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/67d57c32e20fd0a7a302cb81d36e40d5-Paper-Conference.pdf

[43] A. Wang *et al.*, “Superglue: A stickier benchmark for general-purpose language understanding systems,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2019/file/4496bf24afe7fab6f046bf4923da8de6-Paper.pdf

[44] M. Chen *et al.*, “Evaluating large language models trained on code,” 2021. [Online]. Available: <https://arxiv.org/abs/2107.03374>

[45] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi, “The curious case of neural text degeneration,” in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=rygGQyrFvH>

[46] X. Wang *et al.*, ““my answer is c”: First-token probabilities do not match text answers in instruction-tuned language models,” 2024. [Online]. Available: <https://arxiv.org/abs/2402.14499>

[47] B. Xiao, C. Huang, Z. E. Imel, D. C. Atkins, P. Georgiou, and S. S. Narayanan, “A technology prototype system for rating therapist empathy from audio recordings in addiction counseling,” *PeerJ Computer Science*, vol. 2, p. e59, 4 2016.

[48] Z. L. Inbar and Elyoseph, “Suicide risk assessments through the eyes of chatgpt-3.5 versus chatgpt-4: Vignette study,” *JMIR Ment Health*, vol. 10, p. e51232, 9 2023.

[49] D. Can, P. G. Georgiou, D. C. Atkins, and S. S. Narayanan, “A case study: detecting counselor reflections in psychotherapy for addictions using linguistic features,” in *Proc. Interspeech 2012*, 2012, pp. 2254–2257.

[50] K. Singla *et al.*, “Using prosodic and lexical information for learning utterance-level behaviors in psychotherapy,” *Interspeech*, vol. 2018, pp. 3413–3417, 9 2018.

[51] J. Gibson, D. Can, P. G. Georgiou, D. C. Atkins, and S. S. Narayanan, “Attention networks for modeling behaviors in addiction counseling,” in *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*, F. Lacerda, Ed. ISCA, 2017, pp. 3251–3255.

[52] J. Gibson, D. C. Atkins, T. A. Creed, Z. Imel, P. Georgiou, and S. Narayanan, “Multi-label multi-task deep learning for behavioral coding,” *IEEE Transactions on Affective Computing*, vol. 13, no. 1, pp. 508–518, 2022.

[53] K. Singla, Z. Chen, D. Atkins, and S. Narayanan, “Towards end-2-end learning for predicting behavior codes from spoken utterances in psychotherapy conversations,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault, Eds. Online: Association for Computational Linguistics, Jul. 2020, pp. 3797–3803. [Online]. Available: <https://aclanthology.org/2020.acl-main.351>

[54] D. Demszky, D. Yang, D. S. Yeager, C. J. Bryan, M. Clapper, S. Chandhok, J. C. Eichstaedt, C. Hecht, J. Jamieson, M. Johnson, M. Jones, D. Krettek-Cobb, L. Lai, N. Jones-Mitchell, D. C. Ong, C. S. Dweck, J. J. Gross, and J. W. Pennebaker, “Using large language models in psychology,” *Nature Reviews Psychology*, vol. 2, no. 11, pp. 688–701, 2023. [Online]. Available: <https://doi.org/10.1038/s44159-023-00241-5>

[55] T. He *et al.*, “Towards a psychological generalist ai: A survey of current applications of large language models and future prospects,” 2023. [Online]. Available: <https://arxiv.org/abs/2312.04578>

[56] S. Nerella, S. Bandyopadhyay, J. Zhang, M. Contreras, S. Siegel, A. Bumin, B. Silva, J. Sena, B. Shickel, A. Bihorac *et al.*, “Transformers in healthcare: A survey,” *arXiv preprint arXiv:2307.00067*, 2023.

[57] R. Lou, K. Zhang, and W. Yin, “A comprehensive survey on instruction following,” 2024. [Online]. Available: <https://arxiv.org/abs/2303.10475>

[58] R. Lou and W. Yin, “Toward zero-shot instruction following,” in *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, N. Falk, S. Papi, and M. Zhang, Eds. St. Julian’s, Malta: Association for Computational Linguistics, Mar. 2024, pp. 50–60. [Online]. Available: <https://aclanthology.org/2024.eacl-srw.5>

[59] Y. Li, L. Zihan, K. Zhang, D. Ruilong, S. Jiang, and Y. Zhang, “Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge,” *Cureus*, vol. 15, no. 6, 2023. [Online]. Available: <https://www.proquest.com/scholarly-journals/chatdoctor-medical-chat-model-fine-tuned-on-large/docview/2844019753/se-2>

[60] T. Lai *et al.*, “Psy-llm: Scaling up global mental health psychological services with ai-based large language models,” 2023. [Online]. Available: <https://arxiv.org/abs/2307.11991>

[61] H. Qi *et al.*, “Supervised learning and large language model benchmarks on mental health datasets: Cognitive distortions and suicidal risks in chinese social media,” 2023. [Online]. Available: <https://arxiv.org/abs/2309.03564>

[62] Z. Chen, Y. Lu, and W. Wang, “Empowering psychotherapy with large language models: Cognitive distortion detection through diagnosis of thought prompting,” in *Findings of the Association for Computational Linguistics: EMNLP 2023*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023, pp. 4295–4304. [Online]. Available: <https://aclanthology.org/2023.findings-emnlp.284>

[63] Z. Lai, J. Wu, S. Chen, Y. Zhou, A. Hovakimyan, and N. Hovakimyan, “Language models are free boosters for biomedical imaging tasks,” *arXiv preprint arXiv:2403.17343*, 2024.