

Rethinking the Alignment of Psychotherapy Dialogue Generation with Motivational Interviewing Strategies

Xin Sun^{*1,2}, Xiao Tang³, Abdallah El Ali^{2,5}, Zhuying Li³,
Pengjie Ren⁴, Jan de Wit⁶, Jiahuan Pei⁷, Jos A. Bosch¹

¹University of Amsterdam, the Netherlands;

²Centrum Wiskunde & Informatica (CWI), the Netherlands;

³Southeast University, China; ⁴Shandong University, China;

⁵Utrecht University, the Netherlands; ⁶Tilburg University, the Netherlands;

⁷Vrije Universiteit Amsterdam, the Netherlands

Abstract

Recent advancements in large language models (LLMs) have shown promise in generating psychotherapeutic dialogues, particularly in the context of motivational interviewing (MI). However, the inherent lack of transparency in LLM outputs presents significant challenges given the sensitive nature of psychotherapy. Applying *MI strategies*, a set of MI skills, to generate more controllable therapeutic-adherent conversations with explainability provides a possible solution. In this work, we explore the alignment of LLMs with MI strategies by first prompting the LLMs to predict the appropriate strategies as reasoning and then utilizing these strategies to guide the subsequent dialogue generation. We seek to investigate whether such alignment leads to more controllable and explainable generations. Multiple experiments including automatic and human evaluations are conducted to validate the effectiveness of MI strategies in aligning psychotherapy dialogue generation. Our findings demonstrate the potential of LLMs in producing strategically aligned dialogues and suggest directions for practical applications in psychotherapeutic settings.

1 Introduction

Motivational interviewing (MI) is a client-centered counseling technique aimed at encouraging individuals to change behaviors via motivating conversations (Miller and Rollnick, 2002a). It can boost self-directed (i.e., instead of instructed or coerced) and autonomous motivation for behavior change through collaboration between therapists and clients, by emphatic conversations that address ambivalence about such change, and by enhancing a sense of self-efficacy (Martins and McNeil, 2009a). The effectiveness of MI as a conversational technique is strongly supported by

empirical evidence that demonstrates clients' adherence to interventions (Alperstein and Sharpe, 2016). Without the use of MI, traditional techniques can potentially cause resistance and disengagement from clients due to their confrontational ways of thinking (Miller and Rollnick, 2002b). Key to ensuring MI's effectiveness are strategic schemes, such as motivational interviewing skill code (MISC) (Miller et al., 2002), for guiding therapeutic conversations and progress in MI sessions.

Most MI chatbots produce dialogues through expert-written scripts and rules to ensure explainability and controllability (Xu and Zhuang, 2020; Park et al., 2019; Zhang et al., 2020a; Sun et al., 2023). However, this results in limited dialogue diversity and high costs for involving domain expertise in dialogue design. With the advent of natural language generation (NLG) (Dong et al., 2022; Gatt and Krahmer, 2018), several MI chatbots focus on rephrasing client utterances and generating MI dialogues with templates (Almusharaf et al., 2020; He et al., 2022; Min et al., 2023). Studies explored how to integrate therapeutic expertise (Welivita and Pu, 2023; Li et al., 2023), such as counseling strategies, into the dialogue generation process (Welivita and Pu, 2023; Tu et al., 2022; Li et al., 2023). However, this is limited by reliance on domain-specific data required by NLG approaches such as pre-training (Devlin et al., 2019) or fine-tuning (Ziegler et al., 2020).

The emergence of Large language models (LLMs) (Naveed et al., 2023) presents new prospects for generating diverse, flexible, and engaging dialogues in data-scarce domains. In addition, in-context learning with few-shot capabilities enables the integration of MI expertise into the generation process with LLMs (Madotto, 2020; Peng et al., 2020). Notwithstanding these promising advantages, LLM-generated dialogues face challenges in controllability and explainability to elicit behavior changes in sensitive contexts such as psy-

*Corresponding author. Email: x.sun2@uva.nl

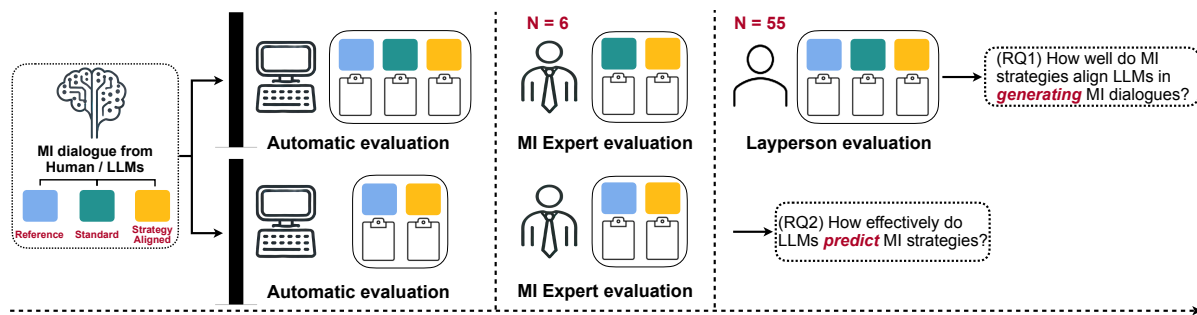


Figure 1: Visual summary of the research questions and corresponding experimental evaluations did in this work.

chotherapy (Sun et al., 2024). Using domain expertise effectively is the key to improving controllability and explainability in the generation process (Welivita and Pu, 2023; Tu et al., 2022; Li et al., 2023). Inspired by the concept of Chain-of-Thoughts (Wei et al., 2023; Wang et al., 2023), we explore utilizing LLMs to predict the next therapist’s MI strategies, i.e., MI skills code (Miller et al., 2002) with its definition as internal reasoning, and generate the therapist’s utterance subsequently strictly following the MI strategies from reasoning. The present study aims to answer the following research questions:

(RQ1) How well do MI strategies align LLMs in generating MI-adherent dialogues?

(RQ2) How effectively do LLMs predict MI strategies aligned with MI principles?

To this end, we conduct extensive experiments to assess the effectiveness of strategy-aligned MI dialogue generation, using both automatic metrics and human evaluation from MI experts and lay evaluators. Our findings demonstrate that MI strategy can effectively instruct LLMs to generate dialogues adherent to MI principles. It enables the controllability and explainability of adopting LLMs to real-world MI applications such as MI chatbots for psychotherapeutic interventions.

The contributions of this work to the current body of knowledge are three-fold: 1) The study is the first to investigate the use of MI strategies to align LLMs for controllable and explainable dialogue generation; 2) We utilize both automatic and human evaluations to validate the effectiveness of such alignment; 3) We combine theoretical analysis and empirical evidence to support the findings.

2 Related Work

2.1 NLG in Motivational Interviewing

Motivational Interviewing (MI) is a therapeutic counseling technique aimed at encouraging self-

directed motivation to change behaviors (Miller and Rollnick, 2002a; Martins and McNeil, 2009b). It emphasizes empathic interaction, respects the client’s autonomy, and has proved to be highly effective in motivating behavior change. As a conversational technique, it can be flexibly incorporated in various therapeutic settings. Experts examine MI-based interactions with strategic schemes such as the Motivational Interviewing Skill Code (MISC) (Miller et al., 2002). This coding scheme is used to assess the effectiveness of MI sessions through the quantification of essential counseling skills and adherence to MI principles.

The role of NLG in MI has revolutionized the way digital therapeutic support is provided. Initially, NLG in MI was limited to replicating ongoing conversations, relying on pre-scripted templates from MI experts (Almusharraf et al., 2020; He et al., 2022; Welivita and Pu, 2023). The rapid advancements in LLMs now enable these models to rephrase what clients say, reflecting their words or even emotions, in ways that appear empathetic (Dieter et al., 2019; Rose et al., 2022; Shen et al., 2020). This development shows promising potential in enhancing client engagement and adherence to therapeutic goals. Despite these benefits, integrating NLG capability and applying LLMs into MI applications presents challenges, particularly in ensuring the generated content adheres to MI principles while being controllable, explainable, and free from harm. There are also ethical considerations, including the impact of generative AI on therapist-client relationships and the risks of LLMs generating inappropriate responses.

2.2 Instructed Dialogue Generation

The instructed dialogue generation (Gupta et al., 2022; Kwak et al., 2023) marks a significant shift in NLG, moving from focusing on linguistic fluency to incorporating specific dialogue objectives. This

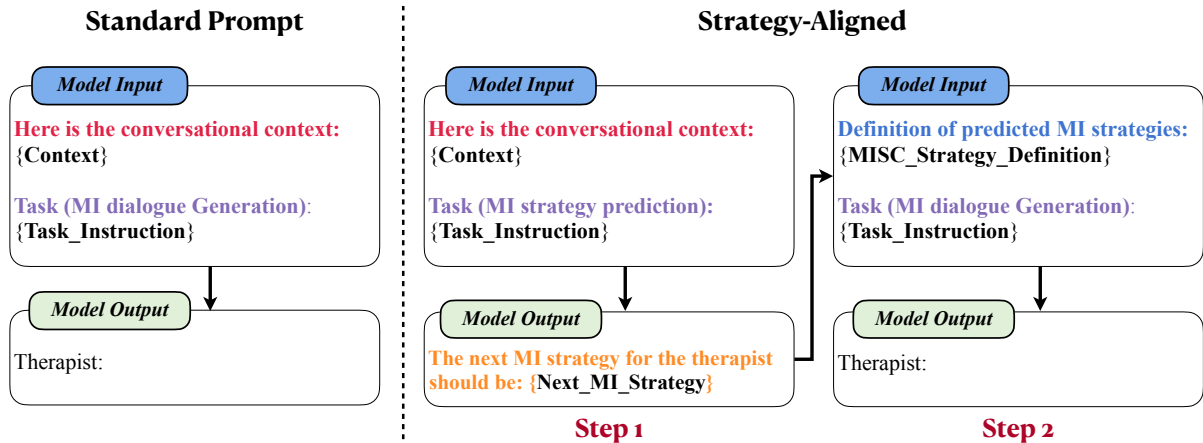


Figure 2: The “strategy-aligned” dialogue generation with the strategy prediction as step 1 and the subsequent therapist’s utterances generation as step 2 in the context of Motivational Interviewing.

transition is evident in the progression from rule-based systems, which rely on static dialogue scripts, to generative models that adapt to dialogue contexts and specific instructions (Gupta et al., 2022), intent (Deng et al., 2023a) or strategy (Welivita and Pu, 2023), and being able to engage in “mixed-initiative” dialogues (Tu et al., 2022; Deng et al., 2023b) where models and users can both lead conversations with specific dialogue objectives. In the realm of therapeutic dialogue generation, instructed dialogue generation efforts, such as those by (Yang et al., 2024; Shah et al., 2022; Rashkin et al., 2019; Gao et al., 2023), highlight the importance of embedding psychological and empathetic principles into response generation, aiming for alignment with therapeutic goals while maintaining dialogue engagement (Sharma et al., 2023). These works inspire our current exploration of enhancing the applications of LLMs in MI. By aligning LLMs with specific MI strategies, we expect these can improve (i.e., strategically aligned and MI-adherent) dialogues. This approach has the dual benefit of not only controlling LLMs with MI principles but also providing more transparent generations with explicitly stated strategic objectives. Additionally, advanced reasoning capabilities of LLMs, such as the Chain-of-Thought (CoT) concept (Wei et al., 2023), further pave the way for strategic alignment by enabling LLMs to predict the next MI strategy as internal reasoning (Cao et al., 2019). Besides, in-context learning (Brown et al., 2020) has emerged as a critical approach for effectively utilizing LLMs by selecting contextual prompts (Fu et al., 2023). These advancements enhance LLMs’ ability to generate dialogues that are well-grounded and contextually coherent (Wang

et al., 2023).

3 Aligning LLMs with MI Strategy

3.1 Approach

Inspired by (Li et al., 2023), we propose an approach we denote “strategy-aligned”, to generate MI dialogues with LLMs guided by domain-specific MI strategies (i.e., the MISC (Miller et al., 2002)), thereby making LLMs align to such strategies with controllability and explainability. We employ the concept of Chain-of-Thought (CoT) (Wei et al., 2023; Zhang et al., 2022), enabling LLMs to internally reason the next MI strategies as the prediction based on dialogue context. Subsequently, LLMs generate dialogues aligned with the predicted MI strategy. For example, LLMs can generate empathetic utterances to promote therapeutic support guided by MI strategies “Reflection”. Figure 2 demonstrates the approach “strategy-aligned”, which combines two steps: MI strategy prediction and strategy-aligned MI dialogue generation.

3.2 Prompt Design

We design two prompting methods for our experimental purpose, including the Standard Prompt and Strategy-aligned Prompt as illustrated in Figure 2. Detailed prompt design is attached in Appendix C.

3.2.1 Standard Prompt

For the standard prompt, we only include the dialogue context and task instruction to generate the next therapist’s utterances. The objective is defined as below. “c” represents the MI dialogue context; “u” represents the next utterance of the therapist we expect the LLM to generate; “k” is the number of

dialogue sizes in the context and we choose it as 5:

$$\mathcal{M} : c_{i-k,i-1} \rightarrow u_i \quad (1)$$

3.2.2 Strategy-aligned Prompt

Compared with the standard prompt, the strategy-aligned prompt is dynamic to give more specific information to LLMs. Specifically, we first request LLM to predict the next MISC strategy for therapists as the internal reasoning and use the reasoning output as the internal state in the subsequent MI dialogue generation process. Strategy-aligned prompting has three elements: 1) the MI dialogue context; 2) the definition of MI strategies as based on the MISC (Miller et al., 2002). 3) the LLM-predicted MI strategies of the therapist’s next utterance (i.e., the type of MI strategy we expect the LLM to predict). We define the objective as follows, in which “s” stands for MI strategies; “d” stands for the definition of MI strategies:

$$\mathcal{M} : c_{i-k,i-1}, s_{i-k,i-1}, d_{str} \rightarrow s_i \rightarrow u_i \quad (2)$$

4 Experiments

4.1 Task Definition

We conduct comparative experiments across benchmark LLMs using different prompting types to assess their performance in generating dialogues adherent to MI principles. Specifically, we define two experimental tasks to address the research questions accordingly: 1) MI dialogue “generation” and 2) MI strategy “prediction”. The procedure of the experiments is illustrated in Figure 1.

4.2 Datasets

Two MI datasets are used with MI strategy annotated by scheme called MISC (Miller et al., 2002). The first dataset is “AnnoMI” (Zixiu et al., 2022; Wu et al., 2023) which has MI conversations and a single coarse-grained MI strategy per utterance. The second dataset is “BiMISC” (Sun et al., 2024) with MI conversations and multiple fine-grained MI strategies per utterance. Table 1 is an example of MI dialogues and MI strategies in these two datasets. Detailed MI strategies are attached in Appendix D.

4.3 Benchmark LLMs

We benchmark several prominent LLMs, focusing on LLMs renowned for their size, perfor-

Therapist	Strategy	
	BiMISC	AnnoMI
That is a good example you give.	AFF	
The sense of smoking is not only motivated by the need [...]	GI	RF
You say if you don’t see it [...]	CR	

Table 1: MI strategies in BiMISC and AnnoMI datasets. **AnnoMI**: utterance is annotated with a single MI strategy, such as “RF” (reflection). **BiMISC**: utterance is annotated with multiple strategies, like “AFF” (affirm), “GI” (give information), and “CR” (complex reflection).

mance, and open-source nature. We select six open-sourced LLMs: **Flan-t5-xxl**, **Vicuna-13B**, **Qwen-14B**, **Qwen2-7B**, **Llama-2-13B**, and **Llama-3-8B**. All these open-sourced LLMs are recognized for their capability to align closely with human instructions (Ouyang et al., 2022), particularly in dialogue interactions. Additionally, we choose **GPT-4** as a commercial benchmark, noted for its superior performance in dialogue generation scenarios.

4.4 Automatic Evaluation Metrics

To objectively evaluate the quality of generations, we apply following automatic evaluation metrics.

- **BLEU & ROUGE** (Papineni et al., 2002; Lin, 2004) assesses the overlap of n-grams between the generation and reference in terms of precision and recall, respectively. We measure $n = 1$.
- **METEOR** (Banerjee and Lavie, 2005) evaluates semantic and syntactic accuracy, including synonym and paraphrase use for linguistic precision.
- **BERTScore** (Zhang et al., 2020b) assesses semantic similarity by BERT embeddings, measuring contextual relevance of generations.
- **Entropy** (Wikipedia, 2024) quantifies the unpredictability and assesses the effectiveness of strategy in controlling generation. Lower entropy indicates more aligned responses.
- **Belief** (Wei et al., 2022) is for hypothesis testing updates posterior probabilities of generations under both hypotheses (H_0 : MI strategies are effective, and H_1 : they are not). The mathematic derivation is in Appendix B.

4.5 Human Evaluation

We employ a two-phase human evaluation: one with MI experts to capture the alignment and adherence to MI principles and another with laypeople focusing on client-perceived quality.

We conduct expert evaluation to assess two as-

pects of alignment: 1) how effective the LLM-generated dialogues align with the MI strategies and adhere to MI principles, and 2) how well the LLM-predicted MI strategies adhere to therapeutic MI principles. We select 100 MI dialogue contexts from datasets for evaluation. The expert assessment involving six experts explores the nuanced effectiveness of strategy-aligned MI dialogue generation by LLMs.

The expert assessment focuses on six criteria (EC): (EC1) how effectively the MI strategy guides the generation of utterance; (EC2) how independent the generated utterance is with the MI strategy; (EC3) how well the generated utterance aligns with the dialogue context; (EC4) how well the generated utterance aligns with the MI principles; (EC5) how the quality of the generated utterance compares to that of a human therapist; (EC6) how well the MI strategy aligns with the dialogue context and MI principles. The first five criteria (EC1-EC5) assess the strategic alignment of generated utterances with the MI strategy. The last criterion (EC6) assesses the performance of MI strategy prediction with LLMs adherent to MI principles. Detailed assessing statements of these criteria are in Table 2.

Criteria	Assessing Statement
EC 1	MI strategy effectively guides therapist’s utterance generation
EC 2	MI strategies impact the therapist’s utterance generation
EC 3	generated therapist’s utterance aligns with the dialogue context
EC 4	generated therapist’s utterance aligns with the MI principles
EC 5	quality of generated utterances is comparable to therapists
EC 6	MI strategy aligns with dialogue context and MI principles

Table 2: Assessing items of expert criteria (EC1-EC6).

Moreover, we are interested in client perceptions. We select 200 MI contexts from datasets for assessment. 55 lay evaluators assess generated utterances and references from human therapists. Each evaluator assesses 14 MI contexts with three criteria: appropriateness, coherence, and relevance (Braggaar et al., 2023; Celikyilmaz et al., 2021).

5 Outcomes

5.1 Empirical Analysis on Automatic Metrics

Table 3 demonstrates the standard prompt yields the lowest scores in automatic metrics. This outcome shows that strategy-aligned generation with MI strategy can effectively instruct LLMs to generate dialogue following specific MI principles. From the model perspective, the commercial GPT-4 model consistently achieves the highest scores across metrics. However, notable is the performance of open-sourced LLMs Flan-T5, Vicuna-

13B and Qwen2, which closely rival that of GPT-4. This highlights the significant advancements in open-sourced LLMs for MI dialogue generation.

To comprehensively evaluate the effectiveness of MI strategies for MI-adherent dialogue generation, we applied Bayesian inference (Wei et al., 2022), which allows to update the belief in a hypothesis based on multiple automatic metrics listed in Table 3 (i.e., BLEU, ROUGE, BERTScore, and Entropy). We calculated the likelihood of each generation for two hypotheses: H_0 (MI strategy is effective for dialogue generation) and H_1 (MI strategy is not effective). Our findings show that all LLMs exhibited a higher probability for H_0 , indicating that MI strategies are generally effective for guiding dialogue generation. Notably, open-source models Qwen 14B and Llama2-13B demonstrated the highest belief in H_0 , surpassing GPT-4. This highlights the potential of open-source LLMs in strategic alignment tasks, offering promising alternatives to proprietary models.

Moreover, to understand how LLMs utilize MI strategies in dialogue generation, Figure 3 visualizes the attention distribution (Vig, 2019) of LLM generations with and without MI strategy. The attention distribution for the strategy-aligned generation shows a significantly denser focus on the MI strategy compared to the standard prompting components (i.e., task instruction and dialogue context only). This emphasizes LLM’s consideration of MI strategies in strategy-aligned dialogue generation.

5.2 Aligning LLM with MI Strategy for Dialogue Generation

In expert evaluations, MI experts assess the alignment and quality of LLM-generated MI dialogues with either “Standard” or “Strategy-aligned” prompting methods in Section 3.2, focusing on determining how well the generated utterances align with the MI strategies and dialogue context. Figure 4 and paired-samples t-test (Ross and Willson, 2017) show that “Strategy-aligned” is significantly more effective in guiding the generation of utterances than standard prompt by criteria EC1 on both datasets ($p < .01$) and the generated utterances are significantly more dependent on MI strategy (EC2) on AnnoMI ($p < .05$) and BiMISC ($p < .01$), proving the effectiveness of MI strategy in aligning dialogue generation with MI principles. Further analyses revealed that the quality of generated utterances with “strategy-aligned” are significantly higher than the generations with the standard

Model	Length		BLEU \uparrow		ROUGE \uparrow		METEOR \uparrow		BERTScore \uparrow		Entropy \downarrow		Belief \uparrow
	/wo	/w	/wo	/w	/wo	/w	/wo	/w	/wo	/w	/wo	w	H_0
AnnoMI													
Flan-T5-XXL-11B	13.3	10.9	10.1	10.9	8.2	8.4	10.7	11.6	84.8	85.6	2.9	2.6	0.62
Vicuna-13B	40.5	30.1	14.3	14.7	12.1	12.3	17.7	17.3	85.1	85.5	4.9	4.5	0.65
Qwen-14B	38.8	37.7	7.8	12.5	6.5	10.4	10.7	15.1	62.4	84.4	3.8	4.5	0.82
Qwen-2-7B	59.7	30.7	11.9	13.8	10.2	10.5	17.3	16.0	84.2	85.2	5.5	4.6	0.64
Llama-2-13B	36.2	44.9	7.5	14.5	7.7	11.8	12.0	18.2	79.8	84.4	4.2	5.0	0.75
Llama-3-8B	57.6	61.1	8.7	8.7	7.6	8.3	13.8	14.2	81.1	80.7	5.1	5.1	0.58
GPT-4	54.3	23.1	13.6	14.3	11.2	11.2	18.7	18.9	84.1	85.5	5.3	4.3	0.66
BiMISC													
Flan-T5-XXL-11B	31.3	26.5	9.5	10.8	9.7	9.7	11.6	12.2	82.7	83.8	2.7	2.9	0.70
Vicuna-13B	51.9	38.0	8.4	12.1	8.0	10.0	13.1	16.7	82.7	84.3	4.9	4.7	0.51
Qwen-14B	41.2	42.6	7.7	11.1	6.4	9.3	10.7	14.9	61.7	83.9	3.9	4.8	0.78
Qwen-2-7B	64.5	40.8	9.1	10.9	7.7	8.3	14.5	14.2	82.1	84.1	5.5	4.8	0.63
Llama-2-13B	6.1	20.7	1.6	6.6	4.7	5.2	2.0	7.8	20.3	82.3	4.0	3.9	0.74
Llama-3-8B	61.2	61.1	7.2	7.2	7.1	7.4	11.5	11.9	81.2	81.0	5.1	5.0	0.65
GPT-4	60.3	36.3	10.9	13.7	9.8	10.0	16.0	16.9	83.4	84.5	5.4	4.6	0.64

Table 3: Results from the automatic evaluation on two datasets with seven benchmark LLMs and two different types of prompt: standard without strategy (/wo) vs. strategy-GT with strategy (/w). Belief is the Bayesian post probabilities for hypothesis H_0 that “MI strategies are effective for MI-adherent dialogue generation”.

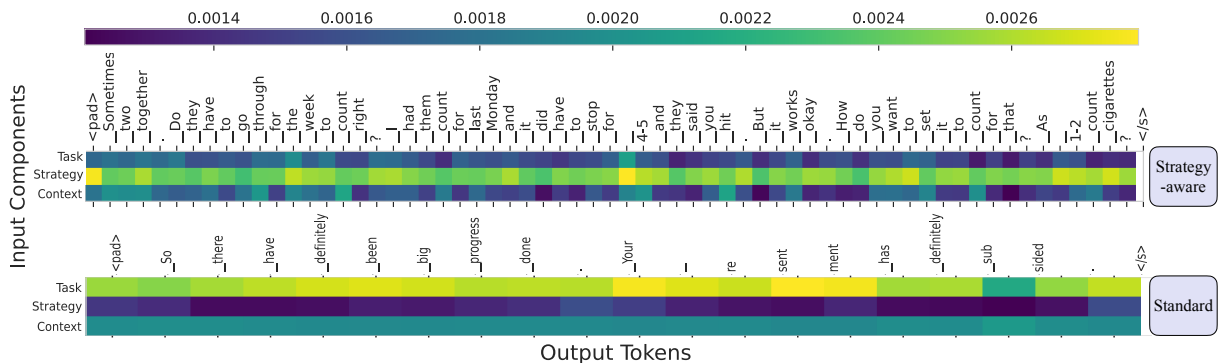


Figure 3: Comparison of attention score distributions from LLM (i.e., Flan-T5 in Appendix A with Encoder-Decoder architecture, last layer, and averaged across all heads) for dialogue generation, with (“strategy-aligned”) and without (“standard”) the MI strategy. The input conversational context and strategy match the ‘Case Study’ example for consistency. Attention to input tokens is aggregated into three prompting components for better comparison.

prompt (EC3 & EC4) ($p < .01$): both prompting approaches achieving above-average scores compared to human therapist’s utterances (EC5), indicating the potential of LLMs in generating therapeutic dialogues in MI. Thus, expert evaluation solves our first research question that MI strategy can guide LLMs to generate dialogues that are strictly aligned with MI principles and are comparable to those of human therapists.

Lay people’s evaluations offer valuable insights from the client’s perspective, complementing expert assessments. As shown in Table 4, Vicuna model using “strategy-aligned” prompts scored significantly higher than both standard prompt and reference utterances across all three assessment dimensions: “Appropriateness”, “Coherence”, and “Empathy” with $p < .01$ on both datasets, confirmed by post hoc Tukey HSD tests (Keselman and Rogan, 1977). This underscores the effective-

ness of the MI strategy in aligning Vicuna with MI principles. However, GPT-4 showed a different trend, with a significant improvement in “Empathy” ($p < .05$) of the standard prompt, but not in “Appropriateness” ($p = .07$) or “Coherence” ($p = .19$). This suggests that while experts confirm the effectiveness of MI strategy aligning GPT-4 with MI principles (results from Figure 4), it might affect the empathetic conversational nature from the client’s perspective (Syed et al., 2024). A further “Case Study” provides concrete examples demonstrating a nuanced balance between MI strategic alignment and client perceptions.

5.3 Predicting MI Strategy as Reasoning

We explore how effectively LLMs predict MI strategies aligned with MI principles to address RQ2. Building on work (Cao et al., 2019), we first benchmark LLMs for the MI strategy prediction task. As

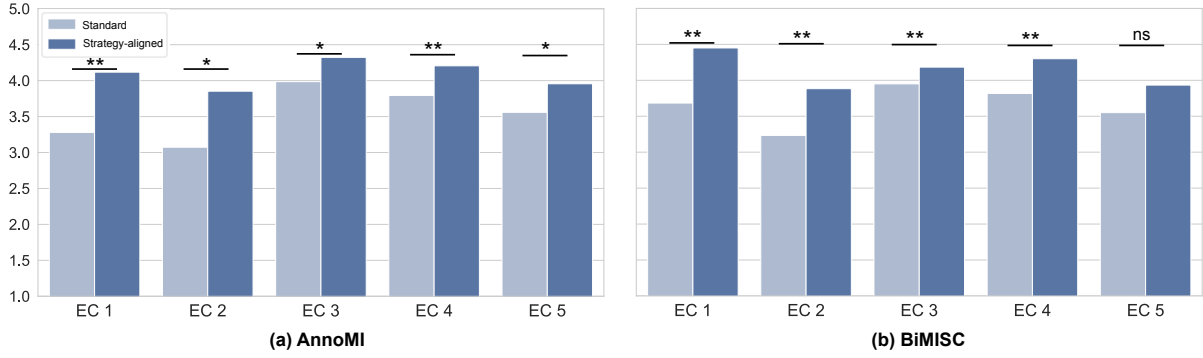


Figure 4: Experts evaluation on two datasets based on assessing criteria (“EC1-EC5”) in Table 2. It assesses the alignment between MI strategy and utterances generated by “Standard” and “Strategy-aligned” prompts. The y-axis denotes the average ratings ranging from 1 (Strongly disagree) to 5 (Strongly agree). (** $p < .01$, * $p < .05$, “ns” $p < .1$)

Dataset	AnnoMI			BiMISC		
	App	Coh	Emp	App	Coh	Emp
Reference	2.77	2.93	2.92	2.71	2.65	2.87
Vicuna-Std	3.75	3.69	3.70	3.62	3.57	3.59
Vicuna-Aligned	3.77	3.75	3.77	3.81	3.79	3.85
GPT-4-Std	3.91	3.88	3.94	4.00	3.94	3.99
GPT-4-Aligned	3.71	3.72	3.68	3.80	3.79	3.87

Table 4: Mean of laypeople assessment on generated utterances by either standard (-Std) or “strategy-aligned” (-Aligned) prompts from Vicuna and GPT-4 models, and reference utterances from therapists. “App”: Appropriateness; “Coh”: Coherence; “Emp”: Empathy.

shown in Table 5, GPT-4 achieves the highest accuracy. Specifically, GPT-4 is the pre-trained model with zero-shot setup, while GPT-4o (FT) is the model fine-tuned on two datasets. Fine-tuning significantly improves LLM performance for MI strategy prediction. Additionally, the accuracy drops in BiMISC with multiple fine-grained strategies compared to AnnoMI with a single coarse-grained strategy (as shown in Table 1), indicating higher complexity of multi-label prediction in MI context.

Dataset	AnnoMI		BiMISC	
	Acc	F1	Acc	F1
Flan-T5	46.2	77.6	19.1	17.4
Vicuna-13B	44.7	76.2	10.5	18.8
GPT-4	50.0	78.9	33.6	27.9
GPT-4o (FT)	63.6	81.4	47.2	36.9

Table 5: The next MI strategy prediction for single-strategy in AnnoMI vs. multiple-strategy in BiMISC.

Although the performance of prediction in Table 5 is modest and the strategy is incorrectly predicted compared to ground truth, the LLM-predicted MI strategies could still be appropriately aligned with context and MI principle, because the MI strategies in MI sessions are usually not uniquely appropriate. “Case Study” further

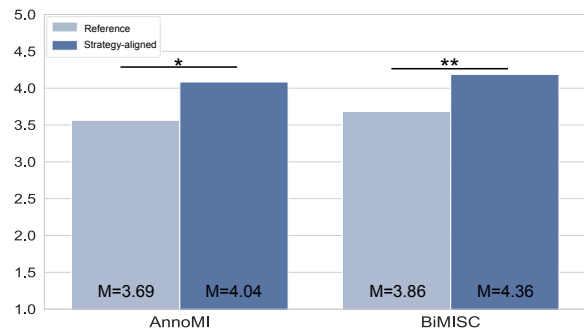


Figure 5: Experts assess the MI strategy prediction by GPT-4 using criteria (EC6): “MI strategy aligns with dialogue context and MI principles”. (** $p < .01$, * $p < .05$)

gives an example demonstrating such nature of MI. Therefore, we conduct the expert evaluation to assess the alignment of the predicted MI strategy with dialogue context and MI principles. According to expert criteria (EC6), Figure 5 shows that “strategy-aligned” demonstrates significantly superior contextual alignment and adherence to MI principles compared to “standard” (i.e., the LLM-generated utterances by Standard prompt and the ground truth MI strategy annotated in the datasets) (with $p < .05$ in AnnoMI and $p = .001$ in BiMISC) by paired-samples t-test. This highlights the significant effectiveness of LLMs for MI strategy prediction.

6 Case Study

This case study aims to shed light on: 1) LLM-predicted strategies can still align with MI principles and context, although prediction differs from ground-truth strategy; 2) balance between therapeutic strategic alignment and client perception.

Figure 6 shows LLM-predicted MI strategies “affirm; closed question” in “strategy-aligned” receives a higher alignment expert score (EC6) (5.0) than ground-truth strategy (3.0) as only “closed

MI Conversational Context	Predicted Next MI Strategy Generated Therapist's Utterance	Expert Rating (EC1-5 / EC6)	Layperson Rating (App / Coh / Emp)
Therapist: Oh at the end of January it was okay. But you don't smoke that much anymore. Client: No, five cigarettes. I smoked one today.	Reference: 'closed question' "Is that at certain moments or do you say well I already know that I am going to smoke a cigarette at twelve o'clock?"	—	4.0 / 3.0 / 4.0
Therapist: How does that go if you go through those five cigarettes? Client: Good.	Standard: 'closed question' (strategy not used for generation) "That's good to hear. It sounds like you've found a rhythm that works for you and helps you manage your intake. If your goal is to cut down even further or maintain this level, keep doing what works best for you."	3.0 / 4.0	5.0 / 4.0 / 5.0
Therapist: So it is actually just right so to speak. Client: Yes, that is going well. I just don't smoke for about four or five behind each other, and then I put a cigarette up again. [...]	Strategy-aligned: 'affirm: closed question' "Sounds like you're managing it quite thoughtfully, which is commendable. How do you feel about your current smoking pattern, are there any changes you'd like to make in the future?"	5.0 / 5.0	4.0 / 3.0 / 4.0

Figure 6: An example of MI dialogue generated by two prompting methods as well as a reference from a human therapist, assessed by MI experts and laypeople.

question”, indicating more effectiveness and empathy, and illustrating the potential of LLM for predicting MI strategy. Therefore, expert evaluations in Figure 5 argue that while the prediction accuracy in Table 5 is modest, some inaccurately predicted strategies are still appropriate and adhere to MI principles within the specific dialogue context, due to the non-uniqueness nature in MI context.

Moreover, experts give higher scores for alignment between generated utterances and strategy (EC1-5) in “strategy-aligned” compared to standard prompt, indicating that strategy-aligned generations are more aligned with MI principles which are “affirmative” and “questional” for controlling the dialogue flow. However, lay evaluators rate generations of standard prompt higher, especially in dimension of “empathy”. The utterance “How do you feel about your current smoking pattern, are there any changes you’d like to make in the future?” generated by “strategy-aligned” is strictly aligned with MI strategy “closed question”, whereas the standard prompt generates “That’s good to hear. [...], every small step is progress.” reflects more empathy as it is MI strategy of “reflection” instead of the “closed question”, which can not lead dialogue to next topic in MI session. Laypeople may perceive dialogues incorporating strict MI strategies as rigid or lacking emotional nuances, suggesting a gap between strategic alignment and client perceptions. This underscores the complexity of translating MI adherence into dialogues perceived as empathetic, highlighting the need for LLMs to balance MI strategies with greater naturalness in psychotherapeutic conversations.

7 Discussion

7.1 Strategy-Aligned MI Dialogue Generation

Current MI applications rely heavily on expert pre-scripted dialogues, which are time and resource-

intensive and require significant domain expertise for conversational design. These pre-scripted dialogues often lack dialogue diversity, which is crucial in MI. LLMs present a promising solution by generating diverse and coherent dialogues, introducing a greater variety of responses, and enhancing the adaptive nature of MI dialogues, in addition to reducing dependency on experts.

The uncontrollable nature of LLMs, however, poses risks in sensitive areas like psychotherapy. LLM-generated dialogues can deviate from MI principles and therapeutic goals, potentially leading to inappropriate outputs. This is where strategic alignment with MI principles becomes critical. Ensuring LLM-generated content adheres to MI guidelines safeguards against deviations that could undermine the therapeutic process. Integrating MI strategies into dialogue generation helps LLMs produce utterances that are both relevant and consistent with MI principles. Additionally, strategic alignment provides explainability in MI dialogue generation, which is essential for applying LLMs in sensitive contexts like psychotherapy. It enhances the safety and perceived trustworthiness of LLM-generated outputs, making the integration of LLMs into MI applications more acceptable.

Extensive evaluations in this work indicate the great potential of strategic alignment in generating MI dialogue with controllability and explainability by ensuring that LLMs follow MI strategies, bridging the gap between automated dialogue generation and therapeutic nuances of effective MI.

7.2 Challenges of Applying LLMs in MI

Applying LLMs in MI presents distinct challenges. This work indicates that strategy can effectively guide LLM-generated dialogues adherent to MI principles. While strategy-aligned utterances are technically precise within therapeutic goals, they may lack naturalness and empathy, which laypeople-

ple prioritize (Syed et al., 2024). This highlights the need for balances between empathetically engaging utterances and those aligned with MI principles. Achieving such balance is crucial for the success of LLM-assisted psychotherapeutic conversation tools, which must meet professional standards while resonating with people, ensuring both strategic accuracy and genuine human connection.

Moreover, the way LLM-generated utterances influence dialogue flow is pivotal. For instance, the generated question utterances could alter the expected course of therapist-led interactions, affecting the ongoing dialogue flow in MI. MI strategy prediction by LLMs provides the potential to ensure that dialogues are contextually relevant and adhere to MI principles. It makes the generative and mixed-initiative systems (Tu et al., 2022) more controllable and aligned with domain expertise. However, inappropriately predicted strategies could lead dialogues in unintended directions, detracting from the therapeutic goals. Enhancing LLMs' ability to accurately understand and predict MI strategies at each reasoning step is key to ensuring dialogues remain on course. Future advancements in LLM reasoning enhanced by domain adaption and further fine-tuning are crucial to overcome these challenges and maximizing the potential of LLM-generated therapeutic dialogues.

8 Conclusion

This work tackles the challenge of utilizing LLMs in the sensitive domain of psychotherapy, confronted by the uncontrollable nature of LLMs. We rethink the way to align LLMs with MI strategy for safer MI dialogue generation. We conduct extensive experiments and analyses with automatic and human evaluations to validate that the MI strategy can effectively align LLMs to generate dialogues adherent to MI principles. The findings also highlight the need for balancing strategic alignment with empathetic engagement in psychotherapeutic interactions. We provide a controllable and explainable solution for MI dialogue generation by LLMs, setting a foundation for future research to enhance the efficacy of LLMs in psychotherapy.

Limitations

While this work provides valuable insights, several limitations should be acknowledged.

First, the generalizability of our findings is limited by the specific data and demographics used.

Expanding the dataset to include a more diverse range of MI scenarios and client demographics could improve the generalizability of the findings.

Second, while “strategy-aligned” offers the potential to guide LLMs in MI dialogue generation, it may not fully capture the dynamic nature of human therapeutic communication and emotional nuances of client and therapist interaction. Enhancing the models with advanced techniques like domain adaption, fine-tuning, or grounding with domain-specific knowledge bases to better capture these dynamic and emotional nuances could help.

Third, the reliance on subjective human evaluations and traditional automatic metrics might not capture the full mental resonance and therapeutic effectiveness of the dialogues. To develop and incorporate more nuanced and comprehensive evaluation metrics could help.

Fourth, while the study assesses strategic alignment with MI principles, it does not measure the impact of generated dialogues on actual therapeutic outcomes like client motivation or behavior change. The practical application of LLM-empowered MI in real-world settings remains untested, and their effectiveness in live sessions needs empirical validation. Additionally, deploying LLMs in therapeutic contexts raises further ethical concerns, including handling sensitive information and potential biases in the generated content. Future research could conduct longitudinal studies to evaluate long-term effects, pilot studies or controlled trials to test real-world effectiveness with comprehensive ethical considerations.

Future work will focus on refining LLM capabilities to capture nuanced human conversational interactions and conducting empirical studies to validate the effectiveness and ethical deployment of LLMs in live therapeutic settings.

Ethical Statement

This research adheres to rigorous ethical standards to ensure the responsible use of AI and the protection of participant data. Large Language Models (LLMs) were carefully monitored to prevent harmful or biased content, particularly in the sensitive context of MI. Informed consent was obtained from all human evaluators, ensuring their anonymity and right to withdraw at any time without consequence (see Appendix E for details). Ethical considerations included addressing potential risks, such as handling sensitive client information and mitigating

biases in generated content. This work complies with institutional ethical requirements, and future researchers are expected to uphold these standards to ensure the responsible use of data and the advancement of knowledge in the field.

Data Anonymization and Privacy

Data privacy is a priority. Part of the data utilized in this work originates from MI counseling sessions and thus contains sensitive information. To protect the privacy of the individuals involved, we implement rigorous data anonymization procedures for the human-involved evaluation. All identifiable information, including names, addresses, and specific personal details, are meticulously removed to ensure confidentiality and anonymity. Recognizing the potential implications of AI technology in therapeutic settings, we advocate for ongoing dialogue and collaboration between AI experts, ethicists, and psychotherapeutic professionals to guide the future development of LLM-assisted therapeutic systems and applications.

Use of AI Tools

We only use the AI tools (i.e., ChatGPT and Grammarly) for checking the grammatical errors.

Acknowledgment

This work is funded by the European Commission in the Horizon H2020 scheme, awarded to Jos A.Bosch (TIMELY Grant agreement ID: 101017424). Additionally, Zhuying Li is supported by the National Natural Science Foundation of China (No. 62302094) Pengjie Ren is supported by the Natural Science Foundation of China (No. 62472261, 62102234), the Key R&D Program of Shandong Province (grant: 2024CXGC010108) All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- Fahad Almusharraf, Jonathan Rose, and Peter Selby. 2020. [Engaging unmotivated smokers to move toward quitting: Design of motivational interviewing-based chatbot through iterative interactions](#). *J Med Internet Res*, 63(6):776–777.
- Dion Alperstein and Louise Sharpe. 2016. [The efficacy of motivational interviewing in adults with chronic pain: A meta-analysis and systematic review](#). *The Journal of Pain*, 17(4):393–403.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Anouck Braggaar, Christine Liebrecht, Emiel van Miltenburg, and Emiel Kraemer. 2023. [Evaluating task-oriented dialogue systems: A systematic review of measures, constructs and their operationalisations](#). *Preprint*, arXiv:2312.13871.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Jie Cao, Michael Tanana, Zac Imel, Eric Poitras, David Atkins, and Vivek Srikumar. 2019. [Observing dialogue in therapy: Categorizing and forecasting behavioral codes](#). pages 5599–5611.
- Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2021. [Evaluation of text generation: A survey](#). *Preprint*, arXiv:2006.14799.
- Wentao Deng, Jiahuan Pei, Zhaochun Ren, Zhumin Chen, and Pengjie Ren. 2023a. [Intent-calibrated self-training for answer selection in open-domain dialogues](#). *Transactions of the Association for Computational Linguistics*, 11:1232–1249.
- Yang Deng, Wenxuan Zhang, Yifei Yuan, and Wai Lam. 2023b. [Knowledge-enhanced mixed-initiative dialogue system for emotional support conversations](#). *Preprint*, arXiv:2305.10172.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *Preprint*, arXiv:1810.04805.

- Justin Dieter, Tian Wang, Arun Tejasvi Chaganty, Gabor Angeli, and Angel X. Chang. 2019. [Mimic and rephrase: Reflective listening in open-ended dialogue](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 393–403, Hong Kong, China. Association for Computational Linguistics.
- Chenhe Dong, Yinghui Li, Haifan Gong, Miaoxin Chen, Junxin Li, Ying Shen, and Min Yang. 2022. [A survey of natural language generation](#). *ACM Computing Surveys*, 55(8):1–38.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023. [Complexity-based prompting for multi-step reasoning](#). *Preprint*, arXiv:2210.00720.
- Pan Gao, Donghong Han, Rui Zhou, Xuejiao Zhang, and Zikun Wang. 2023. [Cab: Empathetic dialogue generation with cognition, affection and behavior](#). *Preprint*, arXiv:2302.01935.
- Albert Gatt and Emiel Kraemer. 2018. Survey of the state of the art in natural language generation: core tasks, applications and evaluation. *J. Artif. Int. Res.*, 61(1):65–170.
- Prakhar Gupta, Cathy Jiao, Yi-Ting Yeh, Shikib Mehri, Maxine Eskenazi, and Jeffrey P. Bigham. 2022. [Instrctdial: Improving zero and few-shot generalization in dialogue through instruction tuning](#). *Preprint*, arXiv:2205.12673.
- Linwei He, Erkan Başar, Reinout Wiers, Marjolijn Antheunis, and Emiel Kraemer. 2022. [Can chatbots help to motivate smoking cessation? a study on the effectiveness of motivational interviewing on engagement and therapeutic alliance](#). *BMC Public Health*.
- H J Keselman and Joanne C Rogan. 1977. The tukey multiple comparison test: 1953–1976. *Psychol. Bull.*, 84(5):1050–1056.
- Jin Myung Kwak, Minseon Kim, and Sung Ju Hwang. 2023. [Context-dependent instruction tuning for dialogue response generation](#). *Preprint*, arXiv:2311.07006.
- Anqi Li, Lizhi Ma, Yaling Mei, Hongliang He, Shuai Zhang, Huachuan Qiu, and Zhenzhong Lan. 2023. [Understanding client reactions in online mental health counseling](#). *Preprint*, arXiv:2306.15334.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Andrea Madotto. 2020. Language models as few-shot learner for task-oriented dialogue systems.
- Renata Martins and Daniel McNeil. 2009a. [Review of motivational interviewing in promoting health behaviors](#). *Clinical psychology review*, 29:283–93.
- Renata Martins and Daniel McNeil. 2009b. [Review of motivational interviewing in promoting health behaviors](#). *Clinical psychology review*, 29:283–93.
- William Miller and Stephen Rollnick. 2002a. [Motivational interviewing: Preparing people for change](#), 2nd ed. *Journal For Healthcare Quality*, 25:46.
- William Miller and Stephen Rollnick. 2002b. [Motivational interviewing: Preparing people for change](#), 2nd ed. *Journal For Healthcare Quality*, 25:46.
- William R Miller, Theresa B Moyers, Denise Ernst, and Paul Amrhein. 2002. Manual for the motivational interviewing skill code (misc). *Unpublished manuscript*. Albuquerque: Center on Alcoholism, Substance Abuse and Addictions, University of New Mexico.
- Do Min, Veronica Perez-Rosas, Ken Resnicow, and Rada Mihalcea. 2023. [Verve: Template-based reflective rewriting for motivational interviewing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10289–10302, Singapore. Association for Computational Linguistics.
- Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Nick Barnes, and Ajmal Mian. 2023. [A comprehensive overview of large language models](#). *arXiv preprint arXiv:2307.06435*.
- Long Ouyang, Jeff Wu, Xu Jiang, Carroll L. Wainwright, Diogo Almeida, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- SoHyun Park, Jeewon Choi, Sungwoo Lee, Changhoon Oh, Changdai Kim, Soohyun La, Joonhwan Lee, and Bongwon Suh. 2019. [Designing a chatbot for a brief motivational interview on stress management: Qualitative case study](#). *J Med Internet Res*, 21(4):e12231.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. [Few-shot natural language generation for task-oriented dialog](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 172–182, Online. Association for Computational Linguistics.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. [Towards empathetic open-domain conversation models: A new benchmark and](#)

- dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.
- Jonathan Rose, Imtihan Ahmed, Eric Keilty, Carolynne Cooper, and Peter Selby. 2022. [Generation and Classification of Motivational-Interviewing-Style Reflections for Smoking Behaviour Change Using Few-Shot Learning with Transformers](#).
- Amanda Ross and Victor L. Willson. 2017. *Paired Samples T-Test*, pages 17–19. SensePublishers, Rotterdam.
- Raj Sanjay Shah, Faye Holt, Shirley Anugrah Hayati, Aastha Agarwal, Yi-Chia Wang, Robert E. Kraut, and Diyi Yang. 2022. [Modeling motivational interviewing strategies on an online peer-to-peer counseling platform](#). *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).
- Ashish Sharma, Inna W. Lin, Adam S. Miner, David C. Atkins, and Tim Althoff. 2023. [Human-ai collaboration enables more empathic conversations in text-based peer-to-peer mental health support](#). *Nature Machine Intelligence*, 5(1):46–57.
- Siqi Shen, Charles Welch, Rada Mihalcea, and Verónica Pérez-Rosas. 2020. [Counseling-style reflection generation using generative pretrained transformers with augmented context](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 10–20, 1st virtual meeting. Association for Computational Linguistics.
- Xin Sun, Dimosthenis Casula, Arathy Navaratnam, Anna Popp, Franziska Knopp, Giovanni Busini, Jan Wesołowski, Marie Van Reeth, Elke Reich, Reinout Wiers, and Jos A. Bosch. 2023. [Virtual support for real-world movement: Using chatbots to overcome barriers to physical activity](#). In *HAI 2023: Augmenting Human Intellect*, volume 368 of *Frontiers in Artificial Intelligence and Applications*, pages 201–214.
- Xin Sun, Jiahuan Pei, Jan de Wit, Mohammad Aliannejadi, Emiel Kraemer, Jos TP Dobber, and Jos A Bosch. 2024. [Eliciting motivational interviewing skill codes in psychotherapy with llms: A bilingual dataset and analytical study](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5609–5621.
- Sara Syed, Zainab Iftikhar, Amy Wei Xiao, and Jeff Huang. 2024. [Machine and human understanding of empathy in online peer support: A cognitive behavioral approach](#). In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.
- Quan Tu, Yanran Li, Jianwei Cui, Bin Wang, Ji-Rong Wen, and Rui Yan. 2022. [Misc: A mixed strategy-aware model integrating comet for emotional support conversation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 308–319, Dublin, Ireland. Association for Computational Linguistics.
- Jesse Vig. 2019. [A multiscale visualization of attention in the transformer model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 37–42, Florence, Italy. Association for Computational Linguistics.
- Hongru Wang, Rui Wang, Fei Mi, Yang Deng, Zezhong Wang, Bin Liang, Ruifeng Xu, and Kam-Fai Wong. 2023. [Cue-CoT: Chain-of-thought prompting for responding to in-depth dialogue questions with LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12047–12064, Singapore. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Zhengxiao Wei, Aijun Yang, Leno Rocha, Michelle F Miranda, and Farouk S Nathoo. 2022. [A review of bayesian hypothesis testing and its practical implementations](#). *Entropy (Basel)*, 24(2).
- Anuradha Welivita and Pearl Pu. 2023. [Boosting distress support dialogue responses with motivational interviewing strategy](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Wikipedia. 2024. [Entropy \(information theory\)](#).
- Zixiu Wu, Simone Balloccu, Vivek Kumar, Rim Helaoui, Diego Reforgiato Recupero, and Daniele Riboni. 2023. [Creation, analysis and evaluation of annomi, a dataset of expert-annotated counselling dialogues](#). *Future Internet*, 15(3).
- Bei Xu and Ziyuan Zhuang. 2020. [Survey on psychotherapy chatbots](#). *Concurrency and Computation: Practice and Experience*.
- Wenjie Yang, Anna Fang, Raj Sanjay Shah, Yash Mathur, Diyi Yang, Haiyi Zhu, and Robert E. Kraut. 2024. [What makes digital support effective? how therapeutic skills affect clinical well-being](#). *Proc. ACM Hum.-Comput. Interact.*, 8(CSCW1).
- Jingwen Zhang, Yoo Jung Oh, P. Lange, Zhou Yu, and Yoshimi Fukuoka. 2020a. [Artificial intelligence chatbot behavior change model for designing artificial intelligence chatbots to promote physical activity and a healthy diet: Viewpoint](#). *Journal of Medical Internet Research*, 22.

- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. [Automatic chain of thought prompting in large language models](#). *Preprint*, arXiv:2210.03493.
- Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2020. [Fine-tuning language models from human preferences](#). *Preprint*, arXiv:1909.08593.
- Wu Zixiu, Balloccu Simone, Kumar Vivek, Helaoui Rim, Reiter Ehud, Reforgiato Recupero Diego, and Riboni Daniele. 2022. [Anno-mi: A dataset of expert-annotated counselling dialogues](#). In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6177–6181.

Appendix A Details of LLMs in the experiments

We benchmark several prominent LLMs to evaluate their performance, focusing on LLMs renowned for their size, performance, and open-source nature. We select six open-sourced LLMs: Flan-T5-XXL ¹, Vicuna-13B ², Qwen-14B-Chat ³, Qwen2-7B-Instruct ⁴, Llama-2 ⁵, and Llama-3 ⁶. All these open-sourced LLMs are recognized for their capability to align closely with human instructions (Ouyang et al., 2022), particularly in dialogue interactions. We use the “Transformers” package from the HuggingFace ⁷ to do generations by all these open-sourced models.

Additionally, we choose GPT-4 ⁸ as the commercial benchmark, noted for its superior performance in various NLP tasks and especially in the dialogue scenarios. We use openai Python library and requests library to send requests to the OpenAI API ⁹ and to do the generations with GPT-4.

The models were used in compliance with their respective licenses and terms at the time of the study.

¹<https://huggingface.co/google/flan-t5-xxl>

²<https://huggingface.co/lmsys/vicuna-13b-v1.5>

³<https://huggingface.co/Qwen/Qwen-14B-Chat>

⁴<https://huggingface.co/Qwen/Qwen2-7B-Instruct>

⁵<https://huggingface.co/meta-llama/Llama-2-13b-hf>

⁶<https://huggingface.co/meta-llama/Meta-Llama-3-8B>

⁷https://huggingface.co/docs/transformers/en/main_classes/text_generation

⁸<https://platform.openai.com/docs/models/gpt-4-and-gpt-4-turbo>

⁹<https://platform.openai.com/docs/overview>

Appendix B Bayesian inference to validate the effectiveness of MI strategy

Introduction

Bayesian inference provides a way for us to evaluate the effectiveness of the Motivational Interviewing (MI) strategy in guiding the LLM generations. It updates the beliefs about hypotheses given new evidence (i.e., new generation in our case). More importantly, this approach allows us to evaluate the overall effectiveness of the MI strategy based on a set of automatic metrics as discussed in Section 4.4.

Hypotheses

- H_0 : MI strategy can effectively guide LLM for MI-adherent dialogue generation.
- H_1 : MI strategy can NOT effectively guide LLM for MI-adherent dialogue generation.

Step 1: Initial Setup

Initial Data Collection:

Collect all the generated dialogues and compute the mean (μ) and standard deviation (σ) for BLEU, ROUGE, METEOR, BERTScore, and Entropy for both dialogues generated with (“strategy-GT”) and without (“standard”) the MI strategy.

Initial Statistics:

Mean (μ) and Standard Deviation (σ) Value for generations with MI strategy (“strategy-GT”):

$$\mu_{\text{BLEU,with}}, \mu_{\text{ROUGE,with}}, \mu_{\text{METEOR,with}}, \mu_{\text{BERTScore,with}}, \mu_{\text{Entropy,with}}$$
$$\sigma_{\text{BLEU,with}}, \sigma_{\text{ROUGE,with}}, \sigma_{\text{METEOR,with}}, \sigma_{\text{BERTScore,with}}, \sigma_{\text{Entropy,with}}$$

Mean (μ) and Standard Deviation (σ) for generations without MI strategy (“standard”):

$$\mu_{\text{BLEU,without}}, \mu_{\text{ROUGE,without}}, \mu_{\text{METEOR,without}}, \mu_{\text{BERTScore,without}}, \mu_{\text{Entropy,without}}$$
$$\sigma_{\text{BLEU,without}}, \sigma_{\text{ROUGE,without}}, \sigma_{\text{METEOR,without}}, \sigma_{\text{BERTScore,without}}, \sigma_{\text{Entropy,without}}$$

Step 2: Bayesian Calculation

Prior Probabilities:

Initially, set $P(H_0) = P(H_1) = 0.5$, reflecting no prior preference for either hypothesis.

Likelihood Calculation:

For automatic evaluation metrics of all observed generations D :

$$P(D|H_0) = \prod_{i=1}^N \prod_{\text{metric}} \frac{1}{\sqrt{2\pi\sigma_{\text{metric,with}}^2}} \exp\left(-\frac{(D_{i,\text{metric}} - \mu_{\text{metric,with}})^2}{2\sigma_{\text{metric,with}}^2}\right)$$
$$P(D|H_1) = \prod_{i=1}^N \prod_{\text{metric}} \frac{1}{\sqrt{2\pi\sigma_{\text{metric,without}}^2}} \exp\left(-\frac{(D_{i,\text{metric}} - \mu_{\text{metric,without}})^2}{2\sigma_{\text{metric,without}}^2}\right)$$

Posterior Probability Calculation:

Using Bayes’ Theorem, calculate the posterior probabilities:

$$P(H_0|D) = \frac{P(D|H_0) \cdot P(H_0)}{P(D)}$$

$$P(D) = P(D|H_0) \cdot P(H_0) + P(D|H_1) \cdot P(H_1)$$

The posterior for H_1 is similarly calculated:

$$P(H_1|D) = \frac{P(D|H_1) \cdot P(H_1)}{P(D)}$$

Step 3: Final Bayesian Belief Update

The final belief in the effectiveness of the MI strategy is the **posterior probability** $P(H_0|D)$ after considering all generations from LLM in one condition. It provides a comprehensive assessment of the effectiveness of MI strategy-aware dialogue generation using Bayesian inference, grounded in empirical generated dialogue data from the experiment. The automatic metrics (BLEU, ROUGE, METEOR, BERTScore, and Entropy) provide a robust assessment of dialogue quality. This method ensures that the belief in the effectiveness of the MI strategy is based on a thorough analysis of all empirical data, resulting in a comprehensive calculation.

$$\text{Final Belief for } H_0 = P(H_0|D)$$

Example Calculation

Considering we have the following metrics for dialogues generated with and without the MI strategy:

$$\begin{aligned} \mu_{\text{BLEU,with}} &= 3, & \sigma_{\text{BLEU,with}} &= 0.5 \\ \mu_{\text{BLEU,without}} &= 2.5, & \sigma_{\text{BLEU,without}} &= 0.5 \\ \mu_{\text{ROUGE,with}} &= 5, & \sigma_{\text{ROUGE,with}} &= 1 \\ \mu_{\text{ROUGE,without}} &= 4.5, & \sigma_{\text{ROUGE,without}} &= 1 \\ & & \vdots & \end{aligned}$$

For a new generation of MI dialogue with BLEU = 3.2, ROUGE = 5.5, ...

We can calculate the likelihoods $P(D|H_0)$ and $P(D|H_1)$:

$$P(D|H_0) = \left(\frac{1}{\sqrt{2\pi \cdot 0.5^2}} \exp\left(-\frac{(3.2 - 3)^2}{2 \cdot 0.5^2}\right) \right) \cdot \left(\frac{1}{\sqrt{2\pi \cdot 1^2}} \exp\left(-\frac{(5.5 - 5)^2}{2 \cdot 1^2}\right) \right) \cdot \dots$$

$$P(D|H_1) = \left(\frac{1}{\sqrt{2\pi \cdot 0.5^2}} \exp\left(-\frac{(3.2 - 2.5)^2}{2 \cdot 0.5^2}\right) \right) \cdot \left(\frac{1}{\sqrt{2\pi \cdot 1^2}} \exp\left(-\frac{(5.5 - 4.5)^2}{2 \cdot 1^2}\right) \right) \cdot \dots$$

Then, using these likelihoods, compute the posterior probabilities.

$$P(D) = P(D|H_0) \cdot 0.5 + P(D|H_1) \cdot 0.5$$

$$P(H_0|D) = \frac{P(D|H_0) \cdot 0.5}{P(D)}$$

$$P(H_1|D) = \frac{P(D|H_1) \cdot 0.5}{P(D)}$$

This calculation yields a final belief in the effectiveness of MI strategy in guiding LLM for dialogue generation based on empirical data from experiments considering chosen automatic evaluation metrics.

Appendix C Example of the prompt template

Component	Content
Conversational Context	<p>Conversational context: [We give historical conversations]</p> <p>Therapist: Yes, those were not really your moments, they were not really your smoking moments, that was a bit literally and figuratively, especially at the end of the day. [...] Client: Yes. Therapist: Yes, okay, so you say I am actually satisfied with the current state of affairs and ... Client: Yes I, I already said that, I like that with losing weight, I have a striving that I am between 85 and 90, that I still want to throw smoking out all the way, it is better anyway And cheaper.</p>
For “Strategy-aligned” Only	<p>The next MI strategy is: the next MISC code(s) and the definition of the MISC code(s) as well as the MISC examples (optional).</p>
Next MISC Code(s) (optional)	<p>The next MISC code(s) for therapist should be: [We give the next MI strategy from the prediction]</p> <p>The next MISC code(s) for the therapist should be: [the Next MI code(s)]</p>
MISC Definition (optional)	<p>The definition of the next MISC code(s) for therapist: [We give the definition of MISC according to the MISC code(s)]</p> <p>’reflection’: reflection is a statement made by the therapist that captures and mirrors back the essence of what the client has said or expressed. [...]</p> <p>’question’: question is made by the therapist to gain more clarity or to explore the client’s perspective, feelings, thoughts, or experiences. [...]</p> <p>’therapist_input’: therapist_input is any other therapist utterance that is not codable as ’question’ or ’reflection’. [...]</p>
MISC Examples (optional)	<p>Examples of each code in MISC: [We can optionally give examples of MISC code(s) according to the next MISC code(s)]</p> <p>’reflection’: Example 1: Client: ’I’m scared of the consequences if I don’t stop smoking.’ Therapist: ’You’re expressing fear about the potential effects of continued smoking.’ [...]</p> <p>’question’: Example 1: Client: ’I think I need to stop smoking.’ Therapist: ’Have you tried quitting before?’ [...]</p>
Task Instruction	<p>Task: [We give instruction to explain the dialogue generation task]</p> <p>Given the conversational context, please generate the next therapist’s utterance that strictly follows the next MISC strategy and its definition. Please only generate an utterance and do not generate question statement if the next MISC strategy is neither “open question” nor “closed question”.</p> <p>The next therapist’s utterance is:</p>

Table 1: An example of the prompt template used in the experiments for the “standard” or “strategy-aligned” MI dialogue generation.

Appendix D MI strategy: MISC codes and their definitions

Therapist Code	Description (abbreviated version)	Example
Question (QUS)	Asking questions for a wide range of answers.	What are your main goals for our sessions?
Reflection (RF)	Utterance that mirrors the client's thoughts, feelings, or experiences to show understanding.	It sounds like you're feeling overwhelmed with your current schedule.
Therapist Input (TI)	Any therapist's utterances are not coded as "question" or "reflection".	Okay. Keep going.

Client Code	Description	Example
Change Talk (CT)	Utterance indicate a desire, ability, reasons, or need for change.	I really want to quit smoking because I want to be healthier.
Sustain Talk (ST)	Utterance argue against change or express a desire to maintain the status.	I've tried quitting smoking before, but it's just too hard for me.
Neutral Talk (NT)	Utterance are neither for nor against change.	I've been smoking since I was a teenager. It's been a part of my life for a long time.

Table 2: The complete coarse-grained MISC codes in the AnnoMI dataset.

Therapist Code	Description (abbreviated version)	Example
Open question (OQ)	Asking questions for a wide range of answers.	Can you tell me more about your drinking habits?
Closed question (CQ)	Asking questions for concise answers: "Yes" or "no", a number.	Did you use heroin this week?
Simple reflection (SR)	Conveying shallow understanding without additional information.	You don't want to do that.
Complex reflection (CR)	Conveying deep understanding with additional information.	That's where you drew the line.
Advice (ADV)	Providing suggestions or recommendations.	Consider starting with small, manageable changes like taking a short walk daily.
Affirm (AFF)	Conveying positive or complimentary information.	You did well by seeking help.
Direct (DIR)	Offering an imperative order, command, or direction.	You've got to stop drinking.
Emphasize control (EC)	Emphasizing client's freedom of choice.	It's up to you to decide whether to drink.
Facilitate (FA)	Encouraging the client to keep sharing.	Tell me more about that.
Filler (FIL)	Filler utterances are not related to behavior change.	Good Morning!
Giving information (GI)	Offering relevant information, explanations, or feedback.	There are several treatment options available for managing stress.
Support (SP)	Offering encouragement and reassurance	I'm here to support you through your recovery journey.
Structure (STR)	Offering a treatment process during the client's journey.	First, let's discuss your drinking, and then we can explore other issues.
Warn (WAR)	Offering a warning or negative consequences.	You could go blind if you don't manage your blood sugar levels.
Permission seeking (PS)	Asking for consent before providing information or advice.	May I suggest a few stress management techniques?
Opinion (OP)	Expressing a viewpoint or judgment	In my opinion, addressing your stress can help reduce your drinking.

Client Code	Description	Example
Follow/Neutral (FN)	No indication of client inclination toward or away from change.	Yeah.
Ask (ASK)	Asking for clarification or information.	What treatment options are available?
Commitment (CM+/CM-)	An agreement, intention, or obligation regarding future change.	I will try to reduce my drinking.
Taking step (TS+/TS-)	Concrete steps the client has recently taken to make a change.	I threw away all of my cigarettes.
Reason (R+/R-)	Rationale, basis, justification, or motive to make a change.	It would be so good for my kids.
Other (O+/O-)	Other statements clearly reflect intention of change.	My family doesn't believe I can quit.

Table 3: The complete fine-grained MISC codes in the BiMISC dataset. The symbols "+" and "-" represent the client's desire to change (+) or not change (-) their behaviors with CM, TS, R or O intention.

Appendix E Information letter and consent form

Expert Evaluation	
Information Letter	<p>Dear participant,</p> <p>Thank you for contributing to our study. This letter aims to provide you with essential information about the study's background and objectives.</p> <p>Background of the study Our research focuses on evaluating Artificial Intelligence (AI) generated dialogues within the context of Motivational Interviewing (MI) counseling.</p> <p>What is this survey about? You are invited to participate in an online survey that will take approximately 45-60 minutes. The survey involves evaluating AI-generated dialogue responses in health counseling settings. Your insights are invaluable to understanding the effectiveness of these AI models. There are no right or wrong answers, so please respond based on your personal perspective. If you encounter any issues or have questions, please contact us at your convenience.</p> <p>Participation Your participation is entirely voluntary. During the experiment, you are free to stop participating at any moment without giving a reason for doing so. No personal identifying information will be collected, and all demographic data will remain confidential. Only the research team will have access to the data, and any published results will be anonymized. Still, do not hesitate if you have any concerns, and please communicate with the responsible researcher if anything happens.</p> <p>Further information Should you have questions about this study at any given moment, please contact the responsible researcher(s): [anonymous]. Formal complaints about this study can be addressed to the Ethics Review Board; [anonymous]. For questions and formal complaints about the protection of your personal information, please contact the Data Protection Officer: [anonymous]</p> <p>Thank you,</p>
Consent Form	<p>In this form, we refer to the information letter describing the research in which you participate. By signing this form, you declare that you understand the nature and methods of this study as described in the information letter.</p> <p>Should you have questions about this study at any given moment, please contact the responsible researchers: [anonymous]. Formal complaints about this study can be addressed to the Ethics Review Board: [anonymous]. For questions and formal complaints about the protection of your personal information, please contact the Data Protection Officer: [anonymous]</p> <p>By selecting "Agre" you confirm that:</p> <ul style="list-style-type: none">• I am 18 years or older.• I have read and understood the information letter.• I agree to participate in the study and use the data obtained with it.• I understand that I can withdraw the participation from the study at any moment without providing any reason.

Table 4: The information letter and consent form used for "Expert Evaluation".

Laypeople Evaluation

Information Letter

Dear participant,

Thank you for showing interest in our study. This letter aims to provide you with essential information about the study's background and objectives.

Background of the study

Our research focuses on evaluating Artificial Intelligence (AI) generated dialogues within the context of Motivational Interviewing (MI) counseling. We aim to assess the performance of various AI models developed based on real-life therapeutic dialogues, with a particular emphasis on their application in health counseling scenarios, such as encouraging healthy behaviors.

What is this survey about?

You are invited to participate in an online survey that will take approximately 30 minutes. The survey involves evaluating AI-generated dialogue responses in health counseling settings. Your insights are invaluable to understanding the effectiveness of these AI models. We recommend completing the survey in a quiet setting and focusing on the instructions and questions provided. There are no right or wrong answers, so please respond based on your personal perspective. If you encounter any issues or have questions, please contact us at your convenience.

Participation

Your participation is entirely voluntary. To ensure the quality of our data, responses completed in under 15 minutes will be considered invalid. As a token of our appreciation, you will receive a compensation of 4 Euros upon completion.

Discomfort, risks, and insurance

We will try to minimize all potential risks during the lab procedures. During the experiment, you are free to stop participating at any moment without giving a reason for doing so. Still, do not hesitate if you have any concerns, and please communicate with the responsible researcher if anything happens.

Your privacy is guaranteed

We assure you that your privacy is of utmost importance. No personal identifying information will be collected, and all demographic data will remain confidential. Only the research team will have access to the data, and any published results will be anonymized.

Further information

Should you have questions about this study at any given moment, please contact the responsible researcher(s): [anonymous]

Thank you.

Consent Form

In this form, we refer to the information letter describing the research in which you participate. By signing this form, you declare that you understand the nature and methods of this study as described in the information letter.

Should you have questions about this study at any given moment, please contact the responsible researchers: [anonymous]. Formal complaints about this study can be addressed to the Ethics Review Board: [anonymous]. For questions and formal complaints about the protection of your personal information, please contact the Data Protection Officer: [anonymous]

By selecting "Agre" you confirm that:

- I am 18 years or older.
- I have read and understood the information letter.
- I agree to participate in the study and use the data obtained with it.
- I understand that I can withdraw the participation from the study at any moment without providing any reason.

Table 5: The information letter and consent form used for "Laypeople Evaluation".