

Measuring What Matters!! Assessing Therapeutic Principles in Mental-Health Conversation

Abdullah Mazhar[†], Het Riteshkumar Shah[†], Aseem Srivastava^{†,*},
Smriti Joshi[#], Md. Shad Akhtar[†]

[†]IIIT Delhi; [‡]MBZUAI; [#]Wysa

{abdullahm, het22213, shad.akhtar}@iiitd.ac.in,
smriti@touchkin.com, aseem.srivastava@mbzuai.ac.ae

Abstract

The increasing use of large language models in mental health applications calls for principled evaluation frameworks that assess alignment with psychotherapeutic best practices beyond surface-level fluency. While recent systems exhibit conversational competence, they lack structured mechanisms to evaluate adherence to core therapeutic principles. In this paper, we study the problem of evaluating *AI-generated therapist-like responses* for clinically grounded appropriateness and effectiveness. We assess each therapists utterance along six therapeutic principles: *non-judgmental acceptance*, *warmth*, *respect for autonomy*, *active listening*, *reflective understanding*, and *situational appropriateness* using a fine-grained ordinal scale. We introduce FAITH-M, a benchmark annotated with expert-assigned ordinal ratings, and propose CARE, a multi-stage evaluation framework that integrates intra-dialogue context, contrastive exemplar retrieval, and knowledge-distilled chain-of-thought reasoning. Experiments show that CARE achieves an F-1 score of 63.34 versus the strong baseline Qwen3 F-1 score of 38.56 which is a 64.26% improvement, which also serves as its backbone, indicating that gains arise from structured reasoning and contextual modeling rather than backbone capacity alone. Expert assessment and external dataset evaluations further demonstrate robustness under domain shift, while highlighting challenges in modeling implicit clinical nuance. Overall, CARE provides a clinically grounded framework for evaluating therapeutic fidelity in AI mental health systems.

1 Introduction

The integration of Artificial Intelligence (AI) in mental health support has gained considerable traction, driven by its potential to deliver accessible (Babu and Joseph, 2024), scalable (AIMakinah et al., 2024), and cost-effective care (Cross

*Work done while at IIIT Delhi.

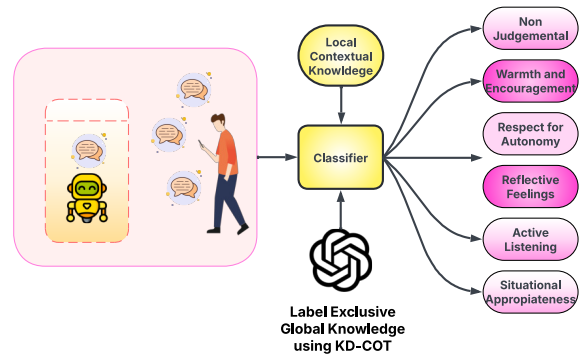


Figure 1: Our problem statement involves the ordinal classification of mental health conversation focusing on six therapeutic principles. We employ retrieval context generation along with knowledge distillation using a chain-of-thought in our proposed method, CARE, ordinal classification of therapist response.

et al., 2024). Conversational systems ranging from simple rule-based bots to large language models have been deployed in numerous support applications (Song et al., 2024). Recent studies show that 41% of support seekers have turned to conversational agents, with over 80% relying on LLMs such as ChatGPT rather than clinically validated therapeutic tools¹. Prior work shows that lay participants often rate ChatGPT-generated therapeutic responses as comparable to or even more favorable than those written by trained clinicians (Hatch et al., 2025). This suggests that standard human evaluations may be driven by surface-level qualities such as fluency or positivity, rather than clinically grounded appropriateness, thereby highlighting limitations in current evaluation practices (Chang et al., 2025).

Although LLMs demonstrate strong fluency and contextual coherence, their application in mental health raises concerns about clinical alignment. Effective therapy depends not only on linguistic quality but also on adherence to core therapeutic

¹<https://thehemingwayreport.beehiiv.com/p/research-report-mental-health-in-an-ai-world>

tic principles, such as non-judgmental acceptance, reflective listening, emotional validation, and respect for client autonomy, which are essential for trust and safety. Despite their sophistication, LLM-generated responses often fall short of professional standards. *Therapeutic principles* (Sutton and Stewart, 2017) formalize these guidelines for appropriate, context-sensitive counseling.

Recent works have explored empathetic and emotionally aware agents (Goel et al., 2021; Zaranis et al., 2021), but do not explicitly address therapeutic alignment. Most approaches rely on surface-level cues or rule-based templates, which fail to capture the structured, principle-driven engagement. Critically, little work evaluates conversational models against core therapeutic principles using grounded, interpretable frameworks.



To support structured evaluation of therapeutic alignment, we build on the publicly available HOPE dataset (Malhotra et al., 2022), which contains therapist client conversations across diverse counseling scenarios. We extend HOPE by introducing fine-grained expert annotations for therapeutic alignment, resulting in FAITH-M, a benchmark comprising 10,172 therapist utterances across 167 sessions. Each utterance is annotated with *six independent labels*, one for each core therapeutic dimension, using a 5-point ordinal scale ranging from *strong negative* (-2) to *strong positive* ($+2$). The six dimensions, derived from established counseling theory (Sutton and Stewart, 2017), include: (a) *non-judgmental acceptance*, (b) *warmth and encouragement*, (c) *respect for autonomy*, (d) *active listening*, (e) *reflective understanding*, and (f) *situational appropriateness*.

We further propose CARE², a clinically grounded evaluation framework that emulates expert reasoning by integrating conversational context and exemplar-driven therapeutic knowledge. CARE encodes conversational history through a *relevant context* module and leverages exemplar interactions and structured reasoning via a *knowledge distillation-cum-chain-of-thought* (KD-CoT) module. Our evaluation across four categories of baselines demonstrates the effectiveness of CARE, achieving a weighted F1-score of 63.34 and outperforming the strongest baseline (Qwen3, also used as the backbone within CARE) by +64.26%. We further support our findings with extensive quantitative, qualitative, and expert analyses, including

evaluations on external expert-guided datasets to assess generalization. Figure 1 illustrates a high-level overview of our approach.

Problem Formulation. We formulate the task as a multi-label ordinal classification problem over therapist utterances in a patient–therapist conversation $\mathcal{C} = [p_1, u_1, p_2, u_2, \dots, p_n, u_n]$, where p_t and u_t denote the patient’s and therapist’s utterances, respectively. Each therapist utterance u_t is evaluated along six therapeutic dimensions, with labels $y_i \in \{-2, -1, 0, +1, +2\}$ reflecting ordinal alignment. Given a local context window, the model predicts these labels while optimizing a loss that jointly captures ordinal structure and classification fidelity. Our key contributions are as follows:

- We introduce a novel task for evaluating therapist responses against six core therapeutic principles.
- We release **FAITH-M**, an expert-guided benchmark with fine-grained ordinal labels for principled evaluation of clinical dialogue.
- We propose **CARE**, a clinically grounded evaluation framework that integrates local conversational context and exemplar-based reasoning for therapeutic assessment, and demonstrate that its advantages persist under cross-dataset evaluation on trauma-focused (Suhas et al., 2025) and CBT based³ benchmarks.

Reproducibility. Code and dataset is open-sourced on [GitHub](#)  [GitHub](#) .

2 Related Work

Mental health has emerged as a critical global concern, motivating interdisciplinary research on AI-powered diagnostic tools and conversational support systems (Cruz-Gonzalez et al., 2025). While significant progress has been made in detection, assessment, and automated support delivery, comparatively little work examines whether AI-generated responses adhere to established therapeutic standards an omission that is particularly consequential in emotionally sensitive settings.

AI Applications in Mental Health. A substantial body of work focuses on using AI to detect mental health issues from textual data. Studies have used social media and clinical transcripts to identify depressive or anxiety-related symptoms (Mazhar et al., 2025) using supervised learning and lexicon-based techniques (Yadav et al., 2020;

²CARE: Clinical Alignment and Reasoning Evaluation

³<https://cheeseburgertherapy.org/>

#	Therapeutic Principles
1	<p>Non-Judgmental Acceptance: This principle refers to the therapist’s use of language that accepts the patient’s emotions and experiences without criticism, blame, or dismissal. <i>Pos Indicator:</i> The therapist uses accepting language, demonstrating empathy toward the patient’s feelings, while <i>Neg Indicator:</i> the therapist uses dismissive language.</p> <p>→ Pos Example: “It’s okay to feel that way [SP]”. “You have the right to your feelings”.</p> <p>→ Neg Example: “That’s not a healthy way to feel [SN]”. “You shouldn’t think like that [SN]”.</p>
2	<p>Warmth & Encouragement: This principle captures the therapist’s expression of care, support, and Pos regard toward the patient. <i>Pos Indicator:</i> The therapist uses encouraging language to convey support and empathym while <i>Neg Indicator:</i> the therapist uses emotionally distant language.</p> <p>→ Pos Example: “I’m glad you shared that [MP]”. “You’re doing great by opening up. [SP]”.</p> <p>→ Neg Example: “I don’t think discussing this will help you right now.”[SN]”.</p>
3	<p>Respect for Autonomy: This principle reflects the therapist’s acknowledgment of the patient’s right to make independent decisions. <i>Pos Indicator:</i> The therapist uses language that supports the patient’s choices, while <i>Neg Indicator:</i> the therapist uses directive language that overrides the patient’s preferences.</p> <p>→ Pos Example: “What do you think would be best for you? [SP]”.</p> <p>→ Neg Example: “You need to do this. [SN]”. “This is the only solution. [SN]”.</p>
4	<p>Active Listening: This principle reflects the therapist’s attentive engagement with the patient’s verbal and emotional cues. <i>Pos Indicator:</i> The therapist uses reflective responses that acknowledge the patient’s statements, while <i>Neg Indicator:</i> the therapist uses inattentive responses that ignore or redirect the patient’s concerns.</p> <p>→ Pos Example: “It sounds like you’re feeling frustrated. [SP]”.</p> <p>→ Neg Example: “Let’s focus on something else. [MN]”.</p>
5	<p>Reflective Feelings: This principle captures the therapist’s ability to identify and reflect the patient’s emotional state. <i>Pos Indicator:</i> The therapist uses language that accurately acknowledges the patient’s emotions, while <i>Neg Indicator:</i> the therapist uses responses that ignore or misinterpret the patient’s feelings.</p> <p>→ Pos Example: “You seem really upset about that experience. [SP]”.</p> <p>→ Neg Example: “So you quit your job last week. [MN]”. “What’s your plan now? [MN]”.</p>
6	<p>Situational Appropriateness: This principle captures the therapist’s ability to respond in alignment with the conversational context and the patient’s emotional needs. <i>Pos Indicator:</i> The therapist uses context-aware responses aligned with the situation, while <i>Neg Indicator:</i> the therapist uses responses that are irrelevant or disconnected from the context.</p> <p>→ Pos Example: “Fabulous. Okay. So would you mind if we spend a few minutes going through that now?[MP]”.</p> <p>→ Neg Example: “Okay. Did you have any other questions about what you read on that shape? [MN]”.</p>

Table 1: Definitions of the six therapeutic principles used in FAITH-M, accompanied by representative Pos and Neg therapist utterances. All examples are annotated using a five-point ordinal scale, ranging from Strong Neg to Strong Pos, with Mild Neg and Mild Pos indicating intermediate levels of therapeutic alignment.

Cha et al., 2022; William and Suhartono, 2021). The DAIC corpus was also introduced, incorporating verbal and non-verbal signals to support multimodal depression detection. While impactful for early screening, these systems are not designed for interactive therapeutic use and often lack contextual reasoning or dialogic modeling (Gratch et al., 2014).

Conversational Systems in Mental Health. Conversational agents are increasingly explored for supportive roles in mental health contexts (Srivastava et al., 2025c). The coherence of LLM responses in therapeutic dialogues has been examined (Song et al., 2024), and ChatGPT responses have been evaluated against therapeutic principles (Hatch et al., 2025). Other efforts focus on counseling summarization (Srivastava et al., 2022; Chen and Verma, 2006; Srivastava et al., 2024), dialogue understanding (Malhotra et al., 2022; Saha et al., 2021), and response generation (Srivastava et al., 2023; Lippe et al., 2020). End-to-end systems such as CareBot (Craeto et al., 2021) aim to of-

fer structured support. However, most evaluations emphasize fluency, empathy, or informativeness, relying on general-purpose metrics or subjective judgments rather than clinically grounded assessment.

Empathy in Mental Health Interactions. Empathy is widely recognized as central to therapeutic communication. Prior work has explored computational empathy in online forums (Sharma et al., 2021), peer-support conversations (Srivastava et al., 2025a), and dialogue systems (Cai et al., 2024; Srivastava et al., 2025b). Approaches range from affective classification to empathy-driven generation using transformer models (Goel et al., 2021; Zaranis et al., 2021). However, empathy alone does not guarantee therapeutic safety or professional adherence, as it may overlook essential principles such as respect for autonomy or situational appropriateness.

3 Dataset: FAITH-M

In this section, we discuss the dataset construction pipeline. FAITH-M is derived from the publicly available HOPE dataset (Malhotra et al., 2022), with each therapist utterance annotated using six ordinal labels based on therapeutic principles from *Learning to Counsel* (Sutton and Stewart, 2017). To ensure reliability and clinical validity, expert clinicians developed the annotation guidelines and validated labels at multiple stages. The following subsections detail the dataset construction process.

3.1 Annotation Guidelines

We adopt a structured annotation process to evaluate therapist responses in FAITH-M across six core therapeutic principles. In collaboration with domain experts, we develop annotation guidelines to reflect clinically grounded counseling practices. We assign an ordinal score for each therapist’s utterance, indicating the degree of alignment with a given therapeutic principle, ranging from *strong negative* (−2) to *strong positive* (+2). A score of (0) denotes a *neutral* case in which the principle is neither expressed nor contextually required. Positive scores (+1, +2) indicate a mild to strong presence of the principle, while negative scores (−1, −2) are assigned when the utterance contradicts an expected therapeutic principle. We define the undertaken therapeutic principles as follows and list a few examples in Table 1.

3.2 Annotation Process

We conduct annotation in a structured, phase-wise manner to ensure reliability and consistency. Three annotators participated in the process, including two primary annotators and one senior supervising annotator. In addition, an expert clinical psychologist provided domain-specific guidance and validated the annotations.

Phase-1 (Training). Annotators were trained over approximately three weeks. They were first introduced to the annotation guidelines and then participated in multiple rounds of annotation and discussion on a subset of 100 therapeutic conversations. During these sessions, the senior moderator and domain expert resolved disagreements and clarified ambiguities. Cohen’s kappa score (Cohen, 1960) was computed after each round, and training concluded upon achieving an overall kappa of 0.69 across all therapeutic dimensions (*NJ*: 0.76, *W&E*: 0.70, *RF*: 0.64, *AL*: 0.65, *RA*: 0.73, and *SA*: 0.70).

Phase-2 (Annotation). Annotators labeled the full dataset over a period of 3 months. Their progress was periodically reviewed by peer annotators to ensure consistency with the guidelines. Final labels were consolidated by the experts, who resolved any tie-breaks between annotators.

Phase-3 (Validation). Following the annotation, the expert counselor validated the labels to ensure clinical accuracy and consistency with therapeutic principles and expert-level interpretations.

3.3 Data Analysis

FAITH-M comprises 10,172 utterances from 167 dialogues and is annotated with six ordinal labels corresponding to core therapeutic principles. The dataset is split session-wise into *train* (6,902 utterances from 116 sessions), *validation* (949 utterances from 17 sessions), and *test* (2,321 utterances from 34 sessions) sets in a 70:10:20 ratio. The detailed distribution for each therapeutic dimension is presented in Table 2.

Therapeutic Dimension	SN	MN	Neu	MP	SP
Non-Judgmental Acceptance	107	413	624	1133	771
Warmth and Encouragement	42	217	733	1134	921
Respect for Autonomy	112	208	502	1401	825
Active Listening	54	89	546	1458	900
Reflective Feelings	84	155	116	1056	589
Situational Appropriateness	79	151	290	1279	1248

Table 2: Label-wise distribution of utterances in FAITH-M. Labels range from *strong negative* (SN: -2) to *neutral* (neu: 0) to *strong positive* (SP: +2).

Therapeutic Adherence. The distribution of ordinal labels across therapeutic dimensions is shown in Figure 2. It reflects the strong adherence of therapists to therapeutic principles. The dataset reveals that negative labels (SN and MN) are rare, while the majority of responses are neutral or positive, indicating consistent alignment with counseling best practices. This highlights the dataset’s reliability in evaluating therapist-client interactions.

4 Methodology

Therapeutic assessment depends on conversational context, as utterance meaning is shaped by prior dialogue. Accordingly, CARE encodes local conversation history through a *relevant context module* and incorporates clinical reasoning via *chain-of-thought-based knowledge distillation*. Figure 3 illustrates the overall architecture.

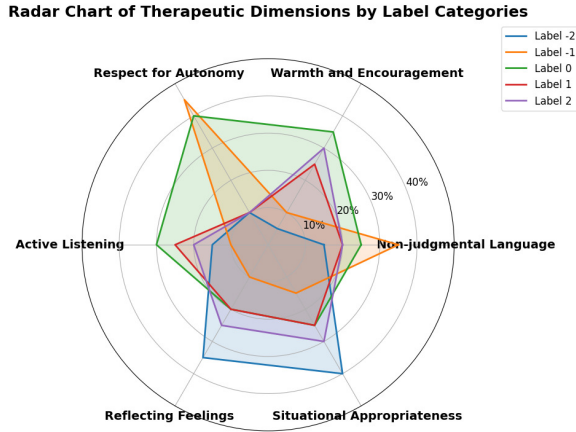


Figure 2: Distribution of ordinal labels across therapeutic dimensions in the FAITH-M dataset.

4.1 Relevant Context Module

We construct a local context window \mathbf{C} for each therapist utterance u_t , containing the utterance itself and k preceding conversational turns: $\{p_{t-k}, u_{t-k}, \dots, p_t, u_t\}$, where p_i and u_i denote the patient’s and therapist’s utterance at turn i , respectively. This window provides immediate conversational context, enabling the encoder’s self-attention (Vaswani et al., 2017) to capture dependencies between prior therapist interventions and the patient’s evolving state. The resulting representation \mathbf{R}_{ctx} serves as one of the input streams to the final prediction layer, capturing the situational grounding necessary for accurate therapeutic alignment assessment.

4.2 KD-CoT

To explicitly embed clinical reasoning and therapeutic knowledge into our model, we utilize a knowledge distillation approach based on Chain-of-Thought prompting. This module enables the model to learn not only from raw examples but also from structured rationales that emulate expert reflection. The KD-CoT pipeline comprises the following phases:

Label-Exclusive Exemplar Sets. To distinguish positive and negative therapeutic alignments for each principle, we construct label-exclusive exemplar sets from the *training split* of FAITH-M. We retain only therapist utterances labeled as strongly positive or strongly negative, excluding neutral and borderline cases to reduce ambiguity and capture canonical patterns of adherence and violation. For each principle, we curate two exemplar pools of patient–therapist pairs, where the therapist response immediately follows the patient utterance to en-

sure contextual coherence. Utterances from the validation and test splits are strictly excluded from exemplar construction and are used only for model selection and final evaluation.

Reference Retrieval. We embed each patient-therapist pair using a Sentence Transformer (Reimers and Gurevych, 2019). For a given test instance, we compute cosine similarity between its embedding and those of the reference sets. We then retrieve the top two most semantically similar reference pairs per dimension and polarity. The aim is to keep the selected examples both contextually relevant and therapeutically aligned. The retrieved exemplars from the previous step are passed to GPT-4o along with a structured prompt⁴. GPT-4o is used to generate dimension-specific CoT explanations grounded in therapeutic norms Figure 6 (c.f. Appendix).

Knowledge Encoding. The retrieved exemplar pairs are provided to a LLM via structured, few-shot prompts to elicit dimension-specific explanations grounded in therapeutic norms. For each therapeutic principle, the model generates concise rationales explaining why a therapist’s response demonstrates alignment or constitutes a violation. These explanations act as expert-like reasoning traces that capture clinically salient cues beyond surface lexical patterns. In our primary setup, we use GPT-4o as the reasoning teacher for this knowledge distillation step and report its results in the main experiments. The generated rationales are concatenated with the current therapist utterance and encoded using a Qwen3 (Yang et al., 2025) model, producing a distilled knowledge representation \mathbf{R}_{KD} , which is integrated with other model streams for final ordinal prediction.

4.3 Ordinal Classification Block

We integrate the context-aware representation R_{ctx} , the distilled knowledge representation R_{KD} , and the utterance embedding r_t using a cross-attention-based fusion mechanism. Specifically, r_t is used as the query, while R_{ctx} and R_{KD} serve as key-value pairs, enabling the model to attend over contextual and knowledge signals. We then pass the resulting fused representation to an ordinal classification head to predict alignment scores across five levels. To jointly model ordinal structure and categorical

⁴Refer to Appendix to view prompts (cf Figure 7).

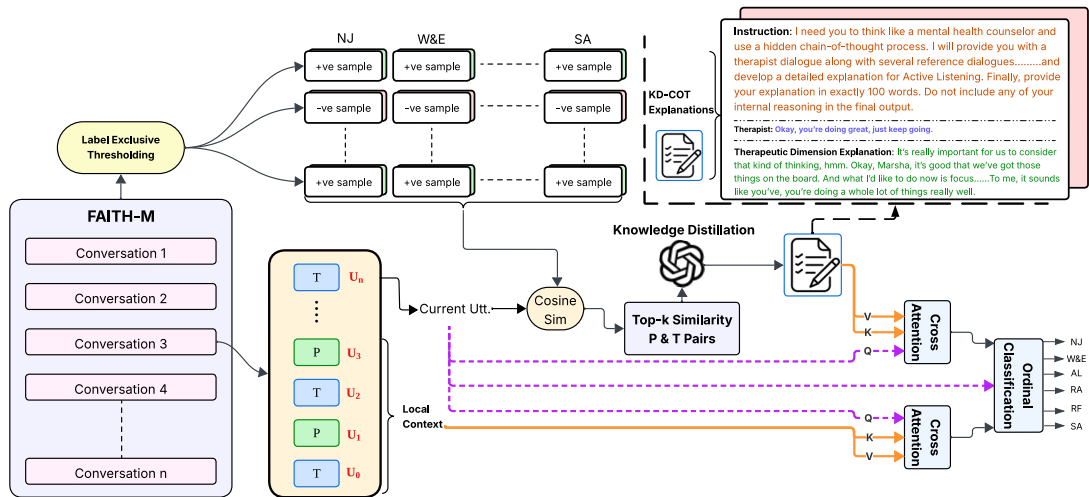


Figure 3: Proposed framework CARE: Therapist utterances are contextualized using intra-conversational self-attention, followed by expert-informed reasoning through knowledge-distilled chain-of-thought explanations. A fusion module integrates contextual, semantic, and distilled knowledge representations, enabling ordinal classification across six core therapeutic principles.

accuracy, we use a hybrid loss:

$$\mathcal{L} = \alpha \cdot \text{MSE}(\hat{y}, y) + \beta \cdot \text{CE}(\hat{y}, y), \quad (1)$$

where MSE captures ordinal distance and CE models class-wise accuracy. We tune $\alpha, \beta \in 0.3, 0.5, 0.7$ on the validation set and observe stable performance, with $\alpha = \beta = 0.5$ performing best. We use this setting in all experiments to balance ordinal consistency and classification accuracy.

5 Experiments and Results

Here, we compare the CARE performance against baseline methods and present a detailed ablation study to analyze the contribution of individual components. We further provide quantitative and qualitative analyses, including error analysis and human expert evaluations, to assess performance and alignment with core therapeutic principles.

Baselines. We benchmark CARE against a diverse set of 15 baselines, comprising 4 prompt-based and 11 trainable models, evaluated under a unified experimental setup. The baseline methods can be viewed in Table 3, and we include more details pertaining to the baseline checkpoints in Appendix A. All baselines use the same local context window ($k = 2$) and, where applicable, are trained with the same ordinal-aware loss function (Equation 1).

Performance Comparison. Table 3 reports performance on the FAITH-M benchmark. Encoder and encoder-decoder models capture coarse therapeutic signals but struggle with fine-grained ordi-

nal distinctions, particularly in mid-spectrum categories. Decoder-only LLMs perform better, but their predictions often rely on surface-level empathy cues, limiting robustness. CARE achieves substantial gains over the strongest decoder-only baseline, Qwen3, which also serves as the backbone of the CARE architecture. These improvements are most pronounced in weighted F1, reflecting better handling of both frequent and infrequent ordinal labels. Both Qwen3 and LLaMA 3.2 are included due to their competitive baseline performance, highlighting that CARE consistently improves across different LLM backbones. To enable a fairer comparison, we also fine-tune MentalLLaMA using the same KD-CoT supervision employed within CARE, resulting in a marked improvement over its zero-shot counterpart. Label-wise performance is reported in (Appendix Table 11, with the strongest results observed for Non-Judgmental Language and comparatively lower scores for Warmth and Encouragement and Reflecting Feelings. Overall, CARE continues to outperform all baselines, indicating that CARE’s gains arise from the joint integration of contrastive exemplar-based reasoning and explicit local contextual encoding.

Ablation Study. We conduct an ablation study to assess the contribution of core CARE components, including local contextual encoding and the KD-CoT module. Here, GPT-4o and MentalLLaMA are used as reasoning teachers to generate Chain-of-Thought rationales within the KD-CoT frame-

	Models	Acc	Pre	Re	F1 _w
ZS	Mental Llama	27.94	8.69	27.74	12.97
	GPT-4o	31.09	36.19	31.09	30.49
FS	Mental Llama	27.74	8.62	27.74	12.97
	GPT-4o	30.44	32.28	30.44	22.40
Enc	MentalBERT	23.85	30.05	23.85	25.73
	RoBERTa	30.09	33.08	30.09	30.72
	ALBERT	42.32	41.58	42.32	34.49
	DeBERTa	33.79	35.32	33.79	34.52
En-De	MentalBART	41.26	46.08	41.27	27.30
	BART	33.34	34.61	33.34	33.96
Dec	Gemma	36.56	37.23	36.56	36.88
	Phi4	36.89	37.91	35.81	37.37
	LLaMA 3.1	43.36	42.81	43.46	37.73
	LLaMA 3.2	44.91	44.78	44.91	37.90
	CARE-LLaMA 3.2	62.07	64.11	62.07	63.07
	Qwen3	45.47	45.10	45.38	38.56
CARE-Qwen3	63.30	64.05	62.65	63.34	
	$\Delta_{Baseline}(\%)$	$\uparrow 39.21\%$	$\uparrow 42.03\%$	$\uparrow 38.05\%$	$\uparrow 64.26\%$
	CARE-Qwen3 – (LC)	57.08	58.06	57.02	57.2
	CARE-Qwen3 – (LE)	53.81	54.04	53.02	53.08

Table 3: Comparative analysis across multiple model classes. LLaMA-3.2 and Qwen-3 represent the strongest-performing baseline backbones, and CARE yields consistent and substantial improvements over both, indicating that gains are not specific to a single base model. LC: label-context; LE: label-exclusive; ZS: zero-shot; FS: few-shot; Enc: encoder; Dec: decoder.

work, while the final predictions are produced by the student model. Table 4 reports performance as we progressively add contextual grounding, label-exclusive exemplar retrieval, and contrastive KD-CoT reasoning, showing consistent improvements with each component. The largest gains arise from combining context with polarity-aware exemplar reasoning, indicating that CARE’s effectiveness stems from structured reasoning rather than dialogue context or prompt supervision alone. Figure 5 visualizes these trends, where moderate context sizes ($k = 2-3$) yield peak performance and larger windows show diminishing returns. Additional ablation analyses and implementation details are provided in (Appendix C).

6 Analysis

Here, we discuss analysis, expert evaluation, and generalizability.

Generalizability. In addition to the CARE’s benchmarking on FAITH-M (c.f. Table 3), we extended the CARE’s inference on two cross-dataset evaluations: PTSD (Suh et al., 2025) and Cheeseburger⁵. Table 5 shows detailed results and high-

⁵<https://cheeseburgertherapy.org/>

#	Param	Acc.	Macro			Weighted		
			F1	Prec.	Rec.	F1	Prec.	Rec.
GPT-4o	$k = 1$	55.66	53.77	54.68	52.89	56.03	56.42	55.65
	$k = 2$	55.35	53.26	53.56	52.98	56.04	56.79	55.32
	$k = 3^{\ominus}$	56.37	55.18	56.27	54.15	57.55	58.84	56.32
	$k = 4$	54.99	51.38	52.15	50.64	55.04	55.84	54.97
	$k = 5$	48.86	43.41	44.12	42.74	49.20	49.56	48.85
Mental Llama	$k = 1$	45.36	43.25	44.51	42.07	46.56	47.85	45.34
	$k = 2$	45.01	43.30	44.68	42.01	46.36	47.74	44.97
	$k = 3$	45.97	43.76	45.19	42.32	47.21	48.64	45.88
	$k = 4$	43.67	43.11	44.57	41.76	45.18	46.87	43.64
	$k = 5$	43.18	40.92	41.23	40.63	44.78	46.54	43.15

Table 4: Ablation Study for relevant context (k) and KD-CoT Modules (with GPT-4o and MentalLlama) on combined positive (+ve) and negative (-ve) indicators. Results with more indicators and local context are included in Appendix (c.f. Table 10).

Model	Acc.	Macro			Weighted		
		F1	Prec.	Rec.	F1	Prec.	Rec.
<i>Dataset: PTSD</i>							
GPT-4o	38.26	25.53	28.79	28.49	35.06	37.09	38.26
Qwen	38.00	16.80	19.10	25.30	31.10	34.90	38.00
Llama 3.2	29.80	8.80	8.60	18.90	17.50	14.50	29.80
CARE	52.34	46.80	48.80	46.50	50.90	51.30	52.40
<i>Dataset: CheeseBurger</i>							
GPT-4o	26.33	12.62	19.61	10.14	31.50	42.22	26.33
Qwen	28.05	11.25	10.71	16.58	26.94	29.35	28.05
Llama 3.2	23.15	9.01	16.98	17.89	18.95	39.13	23.15
CARE	48.96	20.46	28.28	20.07	50.30	65.25	48.96

Table 5: Zero-shot performance comparison. The CARE method consistently outperforms baselines on both PTSD and CheeseBurger datasets.

lights that the advantage persists under cross-dataset evaluation. On the PTSD dataset, which contains trauma-focused Prolonged Exposure therapy dialogues, and the CBT-based Cheeseburger dataset, CARE consistently outperforms prompt-based and fine-tuned LLM baselines. Although absolute performance decreases relative to in-domain evaluation, as expected with considerable domain-shift, the relative improvements remain stable, suggesting that CARE generalizes beyond the distribution on which it was developed.

Quantitative Analysis. We analyze the confusion matrices to understand error patterns across therapeutic dimensions using the updated model configuration in Figure 4. For the *Reflecting Feelings* dimension, the model shows reliable recognition of extreme categories, with accuracies of 66% for Strong Negative and 49% for Strong Positive. As expected in ordinal classification, most errors occur between adjacent labels, particularly between Mild Positive and Strong Positive (30%), reflecting the difficulty of distinguishing subtle emotional af-

Type	Utterance	NJL	Warmth	Autonomy	Listening	Reflect	SA
P	Well, I was raised Catholic.						
T	So you consider yourself Catholic. But you feel like there's a lot of rules.	N / N	N / MP	MP / MP	N / SP	N / MP	MP / SP
P	Yeah. I mean, there's a lot of do's and don'ts. Our church was very like, hell oriented. Like, if you do this, you're gonna go to hell...						
T	How do you think those spiritual values influence your daily life?	MP / MP	SP / MP	N / SP	MP / MP	N / N	SP / SP

Table 6: Examples of *mispredictions* by CARE. These errors typically arise from subtle pragmatic mismatches or an overestimation of empathic alignment. N: Neutral, SP: Strong Positive, MP: Mild Positive, SN: Strong Negative, MN: Mild Negative.

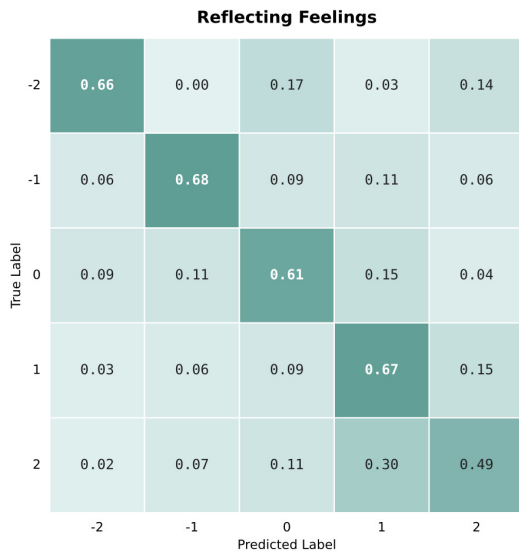


Figure 4: Confusion matrix for *Reflecting Feelings*. Refer to Figure 10 for other dimensions' matrices.

firmation from strong empathic reflection. Similar trends are observed across the remaining therapeutic principles (Appendix, Figure 10). The model consistently maintains strong diagonal mass for extreme labels, achieving approximately 65–70% accuracy for Strong Negative and Mild Negative classes across dimensions such as *Non-Judgmental Language*, *Active Listening*, and *Situational Appropriateness*. In contrast, mid-spectrum categories remain challenging, with *neutral* labels often confused with adjacent *mild positive* or *mild negative* classes (15–20% misclassification, depending on the dimension). Overall, these confusion patterns indicate that the model captures ordinal structure effectively, with errors largely constrained to neighboring categories rather than cross-polarity confusions. This behavior suggests sensitivity to graded therapeutic alignment while highlighting the inherent ambiguity of mid-spectrum judgments.

Error Analysis. Table 6 presents cases where CARE diverges from expert annotations. Errors typically arise from misinterpreting subtle therapeutic cues, such as emotionally grounded paraphrasing,

which the model sometimes labels as overly positive or neutral, consistent with patterns observed in earlier analyses. Across 6 utterances, the model deviated from gold labels in 7 out of 18 dimension-level predictions, usually by a margin of ± 1 . For example, the utterance “*So you consider yourself Catholic. But you feel like there’s a lot of rules.*” was annotated as *neutral* for *non-judgmental acceptance*, while the model predicted *mild-positive*, reflecting a tendency to overestimate surface-level reflectiveness as affirming. Errors were most frequent for *situational appropriateness* and *reflective feelings* (3/6 cases each), whereas *non-judgmental acceptance* and *respect for autonomy* showed higher agreement (4/6 correct each). The model occasionally carries forward prior sentiment across topic shifts, leading to misalignment in *situational appropriateness* and highlighting challenges in modeling mid-range ordinal labels.

Expert Assessment. We consult a licensed clinical psychologist to independently review model predictions and assess alignment with human therapeutic judgments across all dimensions. The expert agreed with the model’s predictions in a substantial proportion of cases. Agreement was highest for dimensions with explicit emotional and supportive cues, including *warmth and encouragement* (84.0%) and *non-judgmental acceptance* (81.6%). In contrast, agreement was lower for more interpretive dimensions such as *respect for autonomy* (70.3%) and *reflecting feelings* (66.7%), which require deeper contextual and emotional understanding. Qualitative feedback further indicates that while CARE reliably captures overt emotional signals, it occasionally overestimates superficially supportive responses and underestimates subtle acknowledgements. The model also shows limitations when conversations shift abruptly, affecting its handling of *situational appropriateness*. Detailed agreement statistics are reported in Table 7.

#	In-Context	Dim	Macro			Weighted		
			F1	Prec.	Rec.	F1	Prec.	Rec.
GPT-4o	Zero Shot	NJ	23.67	25.81	25.75	23.67	25.81	25.75
		W&E	20.44	33.60	24.55	20.44	33.60	24.55
		RA	30.99	31.32	34.73	30.99	31.32	34.73
		AL	37.20	42.19	33.53	37.20	42.19	33.53
		RF	38.94	45.01	37.42	38.94	45.01	37.42
		SA	31.68	39.18	30.53	31.68	39.18	30.53
		Avg.	30.49	36.18	31.08	30.49	36.18	31.08
	Few-Shot	NJ	13.01	19.14	25.74	13.01	19.14	25.74
		W&E	39.22	44.65	44.91	39.22	44.65	44.91
		RA	18.52	45.83	27.24	18.52	45.83	27.24
		AL	24.40	38.45	31.13	24.40	38.45	31.13
		RF	18.27	27.95	23.65	18.27	27.95	23.65
		SA	25.31	33.95	37.42	25.31	33.95	37.42
		Avg.	23.12	34.99	31.69	23.12	34.99	31.69

Table 8: Per-dimension macro and weighted scores for GPT-4o under zero-shot and few-shot settings. Short labels: NJ = Non-Judgmental, W&E = Warmth & Encouragement, RA = Respect for Autonomy, AL = Active Listening, RF = Reflecting Feelings, SA = Situational Appropriateness.

NJL	Warmth	Autonomy	Listening	Reflect	SA
81.60%	84.00%	70.30%	70.10%	66.70%	69.20%

Table 7: Agreement score: Expert vs CARE.

Discussion We analyze GPT-4o under zero-shot and few-shot prompting to examine the limitations of prompt-based therapeutic evaluation. While few-shot prompting yields a modest gain in average accuracy, it leads to a substantial drop in weighted F1, indicating poorer calibration across ordinal categories. This suggests that in-context examples improve format adherence without enhancing sensitivity to ordinal severity. Per-dimension results in Table 8 show stronger performance on *Warmth and Encouragement* and *Active Listening*, where alignment is often conveyed through surface-level cues, and weaker performance on *Non-Judgmental Language* and *Respect for Autonomy*, which require reasoning about implicit judgment and relational intent, as evident under zero-shot prompting (Figure 11), GPT-4o frequently collapses negative and neutral instances into the Neutral class. Few-shot prompting (Figure 12) shifts predictions toward *mild* and *strong positive* labels without recovering *negative* or *neutral* categories. Figure 9 (c.f. Appendix) illustrates GPT-4o’s tendency to replicate exemplar patterns rather than adapt to the target context, highlighting structural limits of prompt-only evaluation for fine-grained therapeutic assessment.

7 Conclusion

We presented CARE, a context-aware framework for evaluating therapist responses across six therapeutic principles using fine-grained ordinal judgments. By integrating local conversational context with knowledge-distilled chain-of-thought reasoning, CARE enables more reliable assessment of therapeutic alignment beyond surface-level fluency. We also introduced FAITH-M, an expert-guided benchmark to support principled evaluation of clinical dialogue. Our findings highlight the limitations of prompt-only methods in ordinal calibration and nuanced therapeutic reasoning, particularly for interpretive dimensions. In contrast, CARE provides a structured and extensible evaluation approach. Future work will explore multi-turn assessment and patient-specific and cross-cultural settings.

8 Ethics and Limitations.

While CARE demonstrates strong performance in utterance-level therapeutic evaluation, several limitations remain. The six principles considered capture only a subset of clinical competencies; aspects such as cultural competence, trauma-informed care, crisis assessment, and intervention timing are not modeled. This scope was intentionally limited to communicative alignment rather than holistic clinical judgment. CARE operates on individual utterances with limited local context and does not capture long-term therapeutic dynamics such as alliance building across sessions. Moreover, therapeutic appropriateness is inherently context- and culture-dependent, and interpretations of warmth or support may vary across settings. Cultural and contextual adaptation was not a primary focus and remains an important direction for future work.

Finally, CARE is designed solely for evaluating AI-generated therapist-like responses and does not provide clinical interventions or replace professional care. All data are anonymized, annotations were conducted under expert supervision, and any real-world use should include clinician oversight and appropriate safeguards.

9 Acknowledgment

The authors acknowledge the support of the Infosys Foundation through CAI at IIIT-Delhi.

References

- Marah Abdin, Sahaj Agarwal, Sahaj Agarwal, Ahmed Awadallah, Vidhisha Balachandran, Harkirat Behl, Lingjiao Chen, Gustavo de Rosa, Suriya Gunasekar, Mojan Javaheripi, Neel Joshi, Piero Kauffmann, Yash Lara, Caio César Teodoro Mendes, Arindam Mitra, Besmira Nushi, Dimitris Papailiopoulos, Olli Saarikivi, Shital Shah, Vaishnavi Shrivastava, Vibhav Vineet, Yue Wu, Safoora Yousefi, and Guoqing Zheng. 2025. [Phi-4-reasoning technical report](#).
- Rawan AlMakinah, Andrea Norcini-Pala, Lindsey Disney, and M Abdullah Canbaz. 2024. Enhancing mental health support through human-ai collaboration: Toward secure and empathetic ai-enabled chatbots. *arXiv preprint arXiv:2410.02783*.
- Anithamol Babu and Akhil P Joseph. 2024. Artificial intelligence in mental healthcare: transformative potential vs. the necessity of human interaction. *Frontiers in Psychology*, 15:1378904.
- Mingxiu Cai, Daling Wang, Shi Feng, and Yifei Zhang. 2024. Empcrl: Controllable empathetic response generation via in-context commonsense reasoning and reinforcement learning. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5734–5746.
- J. Cha, S. Kim, and E. Park. 2022. [A lexicon-based approach to examine depression detection in social media: the case of twitter and university community](#). *Humanities and Social Sciences Communications*, 9:325.
- Crystal T Chang, Hodan Farah, Haiwen Gui, Shawheen Justin Rezaei, Charbel Bou-Khalil, Ye-Jean Park, Akshay Swaminathan, Jesutofunmi A Omiye, Akaash Kolluri, Akash Chaurasia, et al. 2025. Red teaming chatpt in medicine to yield real-world insights on model behavior. *npj Digital Medicine*, 8(1):149.
- Ping Chen and Rakesh Verma. 2006. A query-based medical information summarization system using ontology knowledge. In *19th IEEE Symposium on Computer-Based Medical Systems (CBMS'06)*, pages 37–42. IEEE.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Reuben Crasto, Lance Dias, Dominic Miranda, and Deepali Kayande. 2021. Carebot: a mental health chatbot. In *2021 2nd international conference for emerging technology (INCET)*, pages 1–5. IEEE.
- Shane Cross, Imogen Bell, Jennifer Nicholas, Lee Valentine, Shaminka Mangelsdorf, Simon Baker, Nick Titov, Mario Alvarez-Jimenez, et al. 2024. Use of ai in mental health care: community and mental health professionals survey. *JMIR Mental Health*, 11(1):e60589.
- Pablo Cruz-Gonzalez, Aaron Wan-Jia He, Elly PoPo Lam, Ingrid Man Ching Ng, Mandy Wingman Li, Rangchun Hou, Jackie Ngai-Man Chan, Yuvraj Sahni, Nestor Vinas Guasch, Tiev Miller, et al. 2025. Artificial intelligence in mental health care: a systematic review of diagnosis, monitoring, and intervention applications. *Psychological Medicine*, 55:e18.
- Raman Goel, Seba Susan, Sachin Vashisht, and Armaan Dhanda. 2021. [Emotion-aware transformer encoder for empathetic dialogue generation](#). In *2021 9th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, page 1–6. IEEE.
- Jonathan Gratch, Ron Artstein, Gale Lucas, Giota Stroutou, Stefan Scherer, Angela Nazarian, Rachel Wood, Jill Boberg, David DeVault, Stacy Marsella, David Traum, Skip Rizzo, and Louis-Philippe Morency. 2014. [The distress analysis interview corpus of human and computer interviews](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3123–3128, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- S. Gabe Hatch, Zachary T. Goodman, Laura Vowels, H. Dorian Hatch, Alyssa L. Brown, Shayna Guttman, Yunying Le, Benjamin Bailey, Russell J. Bailey, Charlotte R. Esplin, Steven M. Harris, D. Payton Holt, Jr., Merranda McLaughlin, Patrick O'Connell, Karen Rothman, Lane Ritchie, D. Nicholas Top, Jr., and Scott R. Braithwaite. 2025. [When eliza meets therapists: A turing test for the heart and mind](#). *PLOS Mental Health*, 2(2):1–16.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Raisa Islam and Owana Marzia Moushi. 2024. Gpt-4o: The cutting-edge advancement in multimodal llm. *Authorea Preprints*.
- Shaoyong Ji, Tianlin Zhang, Luna Ansari, Jie Fu, Prayag Tiwari, and Erik Cambria. 2021. [Mentalbert: Publicly available pretrained language models for mental healthcare](#).
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [ALBERT: A lite BERT for self-supervised learning of language representations](#). *CoRR*, abs/1909.11942.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. [BART: denoising sequence-to-sequence pre-training](#)

- for natural language generation, translation, and comprehension. *CoRR*, abs/1910.13461.
- Phillip Lippe, Pengjie Ren, Hinda Haned, Bart Voorn, and Maarten de Rijke. 2020. Diversifying task-oriented dialogue response generation with prototype guided paraphrasing. *arXiv preprint arXiv:2008.03391*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Ganeshan Malhotra, Abdul Waheed, Aseem Srivastava, Md Shad Akhtar, and Tanmoy Chakraborty. 2022. Speaker and time-aware joint contextual learning for dialogue-act classification in counselling conversations. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, pages 735–745.
- Abdullah Mazhar, Zuhair Hasan Shaik, Aseem Srivastava, Polly Ruhnke, Lavanya Vaddavalli, Sri Keshav Katragadda, Shweta Yadav, and Md Shad Akhtar. 2025. [Figurative-cum-commonsense knowledge infusion for multimodal mental health meme classification](#). In *Proceedings of the ACM on Web Conference 2025*, WWW '25, page 637–648, New York, NY, USA. Association for Computing Machinery.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Tulika Saha, Dhawal Gupta, Sriparna Saha, and Pushpak Bhattacharyya. 2021. Emotion aided dialogue act classification for task-independent conversations in a multi-modal framework. *Cognitive Computation*, 13(2):277–289.
- Ashish Sharma, Inna W Lin, Adam S Miner, David C Atkins, and Tim Althoff. 2021. Towards facilitating empathic conversations in online mental health support: A reinforcement learning approach. In *Proceedings of the Web Conference 2021*, pages 194–205.
- Inhwa Song, Sachin R. Pendse, Neha Kumar, and Munmun De Choudhury. 2024. [The typing cure: Experiences with large language model chatbots for mental health support](#).
- Aseem Srivastava, Tanya Gupta, Alison Cerezo, Sarah Peregrine Lord, Md Shad Akhtar, and Tanmoy Chakraborty. 2025a. Critical behavioral traits foster peer engagement in online mental health communities. *PloS one*, 20(1):e0316906.
- Aseem Srivastava, Smriti Joshi, Tanmoy Chakraborty, and Md Shad Akhtar. 2024. [Knowledge planning in large language models for domain-aligned counseling summarization](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17775–17789, Miami, Florida, USA. Association for Computational Linguistics.
- Aseem Srivastava, Gauri Naik, Alison Cerezo, Tanmoy Chakraborty, and Md. Shad Akhtar. 2025b. [Sentiment-guided commonsense-aware response generation for mental health counseling](#).
- Aseem Srivastava, Ishan Pandey, Md Shad Akhtar, and Tanmoy Chakraborty. 2023. Response-act guided reinforced dialogue generation for mental health counseling. In *Proceedings of the ACM Web Conference 2023*, pages 1118–1129.
- Aseem Srivastava, Zuhair Hasan Shaik, Tanmoy Chakraborty, and Md Shad Akhtar. 2025c. [Trust modeling in counseling conversations: A benchmark study](#).
- Aseem Srivastava, Tharun Suresh, Sarah P. Lord, Md Shad Akhtar, and Tanmoy Chakraborty. 2022. [Counseling summarization using mental health knowledge guided utterance filtering](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, page 3920–3930, New York, NY, USA. Association for Computing Machinery.
- BN Suhas, Andrew M Sherrill, Rosa I Arriaga, Christopher Wiese, and Saeed Abdullah. 2025. Thousand voices of trauma: A large-scale synthetic dataset for modeling prolonged exposure therapy conversations. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Jan Sutton and William Stewart. 2017. *Learning to Counsel: How to Develop the Skills, Insight and Knowledge to Counsel Others*, 4 edition. Little, Brown Book Group, London.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshiev, Francesco Visin, Gabriel Rasskin, Gary Wei,

- Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahlteinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024. [Gemma 2: Improving open language models at a practical size.](#)
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- David William and Derwin Suhartono. 2021. [Text-based depression detection on social media posts: A systematic literature review.](#) *Procedia Computer Science*, 179:582–589.
- Shweta Yadav, Jainish Chauhan, Joy Prakash Sain, Krishnaprasad Thirunarayan, Amit Sheth, and Jeremiah Schumm. 2020. [Identifying depressive symptoms from tweets: Figurative language enabled multitask learning framework.](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 696–709, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- An Yang, Anpeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, and Sophia Ananiadou. 2023. [Mentallama: Interpretable mental health analysis on social media with large language models.](#)
- Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, Jimin Huang, and Sophia Ananiadou. 2024. [Mentallama: Interpretable mental health analysis on social media with large language models.](#) In *Proceedings of the ACM Web Conference 2024, WWW '24*, page 4489–4500. ACM.
- Emmanouil Zaranis, Georgios Paraskevopoulos, Athanasios Katsamanis, and Alexandros Potamianos. 2021. [Empbot: A t5-based empathetic chatbot focusing on sentiments.](#)

Appendix

A Baseline Methods

We benchmark CARE against a diverse set of prompt-based and trainable baselines under a unified experimental setup. All baselines use the same local context window ($k = 2$) and, where applicable, are trained with the same ordinal-aware loss function (Equation 1).

(a) Zero-shot LLMs: GPT-4o (Islam and Moushi, 2024) and MentalLLaMA (Yang et al., 2024) are evaluated without task-specific fine-tuning to gauge the upper bound of out-of-the-box performance. **(b) Few-shot LLMs:** We further extend this by evaluating GPT-4o under few-shot prompting 8, using carefully constructed in-context examples (Appendix 6) to examine whether performance improves with light in-context supervision. A detailed analysis of both zero-shot and few-shot performance including prompt formats, per-dimension scores, and failure analysis is presented in (Discussion 6). **(c) Transformer encoders:** MentalBERT (Ji et al., 2021), ALBERT (Lan et al., 2019), DeBERTa (He et al., 2020), and RoBERTa (Liu et al., 2019) test how well discriminative, parameter-efficient encoders capture therapeutic nuances when fine-tuned on our data; **(d) Encoder-decoders:** BART (Lewis et al., 2019) and its domain-adapted variant MentalBART (Yang et al., 2023) assess the utility of sequence-to-sequence pretraining for dialogue-level classification tasks; **(e) Decoder-only LLMs:** Qwen3 (Yang et al., 2025), LLaMA 3.1/3.2 (Grattafiori et al., 2024), Phi-4 (Abdin et al., 2025), and Gemma (Team et al., 2024) probe the limits of autoregressive generators when repurposed for structured ordinal prediction.

B Experimental Setup

All experiments were conducted on 2 NVIDIA A100 GPUs. Models were trained for 10 epochs using the AdamW optimizer with a learning rate of 1×10^{-5} , weight decay of 0.01, dropout of 0.2, and gradient clipping at 1.0. Training was performed with a batch size of 16, and input sequences were truncated to the maximum length supported by each backbone. A sliding window of size 6 was used to handle longer dialogues. Parameter-efficient fine-tuning was implemented using LoRA with rank $r = 16$ and scaling factor $\alpha = 32$. All experiments were run with a fixed random seed of 42. A com-

Hyperparameter	Value
Learning Rate (lr)	1×10^{-5}
Adam Beta1	0.9
Adam Beta2	0.999
Dropout	0.2
Weight Decay	0.01
Batch Size	16
Max Sequence Length	4096
Sliding Window Size	6
Gradient Clipping	1.0
Learning Rate Scheduler	Constant
Seed	42
LoRA Rank (r)	32
LoRA Alpha (α)	32

Table 9: Detailed hyperparameters used in the experiments.




Figure 5: Performance trends across KD-CoT configurations, showing the variation in Weighted F1 score with different context window sizes and exemplar retrieval.

plete list of hyperparameters is provided in Table 9. Inference time varies with the choice of the KD-CoT reasoning teacher. Using open-source models (e.g., Qwen) for rationale generation yields approximately 1 second per instance with a weighted F1 of 57.75, whereas using GPT-4o increases inference time to around 5 seconds but improves weighted F1 to 63.34. This highlights a clear trade-off between efficiency and evaluation accuracy.

C Additional Ablation Details.

Table 10 provides a comprehensive breakdown of the ablation experiments investigating the impact of context window size (k) and the label polarity settings on the performance of the Relevant Context and KD-CoT modules. Consistent with observations discussed earlier, the Relevant Context Module achieves optimal performance at $k = 2$,

#	Param	Acc.	Macro			Weighted		
			F1	Prec.	Rec.	F1	Prec.	Rec.
Rel. Context	$k = 1$	49.68	46.80	47.05	46.56	50.41	51.24	49.62
	$k = 2$	53.12	47.60	48.34	46.92	54.02	55.24	52.87
	$k = 3$	50.58	47.07	47.86	46.32	51.26	52.01	50.55
	$k = 4$	48.95	46.21	46.67	45.77	50.17	51.52	48.89
	$k = 5$	47.73	43.98	44.67	43.32	48.43	49.24	47.65
KD-COT Module (GPT-4o)	<i>+ve & -ve labels</i>							
	$k = 1$	55.66	53.77	54.68	52.89	56.03	56.42	55.65
	$k = 2$	55.35	53.26	53.56	52.98	56.04	56.79	55.32
	$k = 3^{\ominus}$	56.37	55.18	56.27	54.15	57.55	58.84	56.32
	$k = 4$	54.99	51.38	52.15	50.64	55.04	55.84	54.97
	$k = 5$	48.86	43.41	44.12	42.74	49.20	49.56	48.85
	<i>+ve Labels Only</i>							
	$k = 1$	46.20	42.70	42.05	43.39	48.09	50.17	46.19
	$k = 2$	49.68	47.23	47.53	46.94	51.02	52.47	49.66
	$k = 3$	50.22	47.22	47.81	46.65	51.73	53.36	50.20
	$k = 4$	50.12	47.12	48.43	45.88	51.14	52.24	50.09
	$k = 5$	48.81	45.66	46.39	44.97	49.13	51.77	48.76
	<i>-ve Labels Only</i>							
	$k = 1$	45.23	41.58	40.89	42.31	44.26	43.36	45.21
	$k = 2$	45.41	43.99	44.59	43.41	45.84	46.31	45.38
$k = 3$	47.98	46.35	47.69	45.10	48.72	49.52	47.96	
$k = 4$	47.57	45.42	46.51	44.39	48.22	48.96	47.52	
$k = 5$	46.86	42.82	42.21	43.45	46.95	47.08	46.84	
KD-COT Module (Mental Llama)	<i>+ve and -ve Labels</i>							
	$k = 1$	45.36	43.25	44.51	42.07	46.56	47.85	45.34
	$k = 2$	45.01	43.30	44.68	42.01	46.36	47.74	44.97
	$k = 3$	45.97	43.76	45.19	42.32	47.21	48.64	45.88
	$k = 4$	43.67	43.11	44.57	41.76	45.18	46.87	43.64
	$k = 5$	43.18	40.92	41.23	40.63	44.78	46.54	43.15
	<i>+ve Labels Only</i>							
	$k = 1$	43.72	43.50	44.67	42.39	44.93	46.34	43.63
	$k = 2$	45.02	44.34	45.21	43.52	45.66	46.71	44.77
	$k = 3$	45.11	44.56	45.62	43.56	45.42	45.98	44.89
	$k = 4$	44.87	44.14	45.28	43.16	45.00	45.68	44.35
	$k = 5$	43.17	43.04	44.52	41.67	44.33	45.85	42.91
	<i>-ve Labels Only</i>							
	$k = 1$	42.96	42.35	43.56	41.22	44.30	45.94	42.78
	$k = 2$	42.37	42.27	43.67	40.96	43.53	45.31	41.89
$k = 3$	43.05	42.87	44.02	41.78	43.38	46.05	42.84	
$k = 4$	42.75	41.99	42.66	41.36	43.94	45.34	42.63	
$k = 5$	42.63	41.67	42.83	40.59	43.81	45.13	42.57	

Table 10: Ablation Study for Relevant Context and KD-COT Modules. The best-performing configuration is highlighted.

balancing sufficient conversational history without introducing noise. For the KD-CoT module, evaluations across different configurations utilizing both positive and negative labels, positive only, and negative-only exemplars demonstrate that contrastive reasoning with both polarities significantly enhances accuracy and F1 scores. Notably, GPT-4o outperforms Mental LLaMA by a substantial margin across all metrics and configurations, reinforcing the critical role of a high-quality reasoning backbone. The detailed metrics, including accuracy, macro and weighted F1, precision, and recall, further validate these findings, emphasizing the robustness of the proposed framework under varying settings.

Therapeutic Dimension	Prec	Rec	F1
Non-Judgmental Language	65.80	64.45	65.10
Warmth and Encouragement	62.10	61.70	61.90
Active Listening	64.20	63.40	63.80
Reflecting Feelings	62.40	62.00	62.20
Situational Appropriateness	63.90	63.30	63.60
Respect for Autonomy	63.70	63.20	63.44

Table 11: Label-wise precision, recall, and F1 scores (%) for each therapeutic dimension on FAITH-M using CARE. Highest F1 is observed for Non-Judgmental Language; lowest for Warmth and Reflecting Feelings.

D Prompting

D.1 Few-shot vs Zero-shot

The few-shot setting provides multiple in-context examples with detailed explanations, enabling the model to learn implicit scoring patterns and align its judgments accordingly. In contrast, the zero-shot prompt relies solely on task instructions and the immediate dialogue context, without prior demonstrations. This distinction highlights how contextual grounding through examples can influence model reasoning, consistency, and interpretability in therapeutic assessment tasks. Overall, Figure 6 emphasizes the trade-off between guided learning in few-shot prompting and the generalization capability of zero-shot evaluation.

D.2 Prompting with Label-Guided Reasoning

To obtain precise and contextually grounded explanations of therapeutic principles, we designed a structured prompt that incorporates *label-exclusive* utterances dialogue snippets characteristic of a specific therapeutic dimension to guide the model’s reasoning. The prompt instructs the language model to act as a mental health counselor, engaging in internal chain-of-thought reasoning before generating a final explanation.

Specifically, the model is provided with a set of reference utterances and a target therapist statement. It is asked to internally synthesize insights from the references and produce a 100-word explanation focused on a single therapeutic principle such as *Active Listening* while omitting any visible reasoning steps. This approach encourages the model to generate concise yet rich rationales grounded in domain-relevant context.

An example of the full prompt template used to elicit such explanations from GPT-4o is shown in Figure 7. The retrieved exemplars from the previous stage (i.e., label-exclusive utterances) are incor-

Few-Shot vs Zero-Shot Prompt Design for GPT-4o Evaluation

Few-Shot Prompt:

Instruction: You are a clinical NLP evaluator. Score a therapist response from **-2 to +2** across six therapeutic principles. You will be shown few-shot examples and then asked to score a new interaction.

Example

Patient: <previous utterance>

Therapist: <therapist response>

Explanations:

- Non-Judgmental Language: ...
- Warmth and Encouragement: ...
- Respect for Autonomy: ...
- Active Listening: ...
- Reflecting Feelings: ...
- Situational Appropriateness: ...

Scores: NJ: X, W: X, RA: X, AL: X, RF: X, SA: X

Repeat for multiple examples.

Now evaluate this response:

Patient: <current patient utterance>

Therapist: <current therapist response>

Explanations: ...

Return scores in same format.

Zero-Shot Prompt:

Instruction: You are a mental health evaluation assistant trained in psychotherapy communication techniques. Please assess the quality of the **therapist's response** using the following six therapeutic dimensions. Use the preceding **patient's statement** to help you understand the context and appropriateness of the therapist's message.

Patient said: "<patient utterance>"

Therapist responded: "<therapist utterance>"

Rate the therapist's response from **-2 to +2** for each dimension: -2 = strongly violates the principle -0 = neutral or unclear +2 = strongly demonstrates the principle

Output Format (verbatim):

Non-Judgmental: X, Warmth: X, Respect: X, Active Listening: X, Reflecting Feelings: X, Situational Appropriateness: X

Figure 6: Comparison of the few-shot (top) and zero-shot (bottom) prompt formats used to evaluate GPT-4o and Mental LLaMA responses. The few-shot prompt includes in-context examples and explanations, while the zero-shot prompt relies only on task instruction and current dialogue context.

Prompt Design for Dimension Specific COT Explanation

Instruction: I need you to think like a mental health counselor and use a hidden chain-of-thought process. I will provide you with a therapist dialogue along with several reference dialogues. First, internally analyze and understand the reference dialogues to build expert knowledge. Then, using your hidden chain-of-thought reasoning, analyze the therapist dialogue and develop a detailed explanation for Active Listening. Finally, provide your explanation in exactly 100 words. Do not include any of your internal reasoning in the final output.

Label Exclusive Utterances: It's really important for us to consider that kind of thinking, hmm. Okay, Marsha, it's good that we've got those things on the board. And what I'd like to do now is focus, stepping out a bit from your situation and looking, looking at it from another point of view. To me, it sounds like you've, you're doing a whole lot of things really well. Can you list all of the things that you're doing really well?

Therapist Current Utterance: Okay, you're doing great, just keep going.

Figure 7: Prompt used to extract dimension-specific Chain-of-Thought explanations from GPT-4o. The retrieved exemplars from the previous step are passed to the model along with the structured instruction, guiding the explanation of a therapeutic principle such as Active Listening.

Model-generated Rationale for Active Listening

Therapist Response:

“I appreciate you answering these questions. For me, this has been helpful.”

Model-Generated Explanation (Active Listening):

Active listening is a crucial therapeutic skill that involves fully engaging with a client’s words, emotions, and underlying meanings. In the therapist dialogue, the phrase “I appreciate you answering these questions” demonstrates validation, showing the client their input is valued. By stating, “This has been helpful,” the therapist reinforces the significance of the dialogue, encouraging further sharing. This approach fosters a safe space for clients to express themselves openly. Active listening not only enhances understanding but also builds rapport, allowing clients to feel heard and supported in their journey towards healing and self-discovery.

Figure 8: Representative example of a model-generated rationale used within the KD-CoT pipeline for the *Active Listening* therapeutic principle.

porated into the prompt, which guides the model in producing dimension specific justifications aligned with therapeutic principle.

E Human-Level Rational Assessment

CARE incorporates expert-level reasoning through the KD-CoT module, which relies on large language models to generate fine-grained, expert-standard explanations that guide ordinal classification. A natural concern arises regarding the credibility of these generated rationales, particularly in light of the strong agreement observed between expert annotations and CARE across multiple dimensions in Table 7.

To address this concern, we conduct a human-level explanation assessment to evaluate the quality of rationales generated by GPT-4o and Mental-LLaMA, which are subsequently utilized within the CARE framework. We randomly selected 15 instances, and for each instance, three clinically relevant dimensions were chosen. For each dimension, both LLMs generated an explanation.

These explanations were independently evaluated by 25 human evaluators. Each explanation was rated on a 5-point Likert scale (1 = poor, 5 = excellent), focusing on clarity, contextual relevance, and clinical appropriateness.

Our analysis indicates that GPT-4o explanations are consistently rated higher, with approximately 80% receiving a score of 4 or 5, reflecting strong alignment with expert reasoning and contextual grounding. In contrast, Mental-LLaMA explanations are predominantly rated between 3 and 4, suggesting moderate quality with occasional limitations in specificity or contextual depth. These findings suggest that GPT-4o produces more reliable and clinically coherent rationales, making

it a strong choice for guiding CARE through the KD-CoT mechanism.

Rationale Examples Figure 8 presents a representative example of a model-generated rationale for the *Active Listening* dimension. The explanation highlights clinically meaningful cues such as validation, responsiveness, and conversational engagement. These examples are intended to demonstrate the clarity, contextual relevance, and clinical plausibility of the generated rationales, rather than to serve as ground-truth therapeutic justifications.

F Behavioral Differences Between CARE and GPT-4o

Figure 10 presents normalized confusion matrices across the six therapeutic dimensions, comparing CARE with GPT-4o under zero-shot and few-shot prompting. Rows denote gold labels and columns predicted labels under the ordinal scale: SN, MN, N, MP, and SP.

Across dimensions, CARE exhibits a coherent ordinal error structure, with most misclassifications occurring between adjacent categories (e.g., Neutral vs. Mild Positive). This behavior indicates sensitivity to ordinal severity and reliable discrimination at the extremes (SN and SP), reflecting robust handling of clearly misaligned and strongly aligned responses.

In contrast, GPT-4o under zero-shot prompting (Figure 11) shows a strong bias toward Mild Positive and Strong Positive predictions, frequently collapsing negative and neutral instances into positive categories. This effect is particularly evident for principles such as *Non-Judgmental Language* and *Respect for Autonomy*, suggesting reliance on surface-level linguistic cues rather than relational or pragmatic intent.

Full Few-Shot In-Context Example Used for GPT-4o Evaluation

Instruction: You are an expert in mental health counselling. Evaluate each therapist's response on a scale from -2 to +2 based on the following six therapeutic principles:

Non-Judgmental Language; Warmth and Encouragement; Respect for Autonomy; Active Listening; Reflecting Feelings; Situational Appropriateness.

Use the reference examples below to guide your judgment. Each example includes a therapist response, detailed explanations for each principle, and final scores. Follow the same reasoning pattern when evaluating the new response.

Example A

Therapist Response: "So you have difficulty falling asleep. But once you fall asleep."

Explanations: **Non-Judgmental Language:** Non-judgmental language in therapy is essential for creating a safe and supportive environment where clients feel comfortable sharing their experiences without fear of criticism. **Warmth and Encouragement:** Warmth and encouragement in therapy create a safe space for clients to explore their feelings and challenges. **Respect for Autonomy:** Respect for autonomy emphasizes recognizing the client's ability to make their own decisions regarding their mental health. **Active Listening:** Active listening involves fully concentrating on the client's words and demonstrating understanding. **Reflecting Feelings:** Reflecting feelings acknowledges the client's emotional experience by summarizing their struggle with sleep difficulties. **Situational Appropriateness:** The therapist response appropriately addresses a common clinical concern raised by the client.

Scores: Non-Judgmental: 2, Warmth: 1, Respect: 1, Active: 2, Reflecting: 1, Situational: 2

Example B

Therapist Response: "What County are we in?"

Explanations: **Non-Judgmental Language:** The response does not express judgment and maintains neutrality. **Warmth and Encouragement:** The question provides limited warmth and encouragement. **Respect for Autonomy:** The response does not restrict the client's autonomy. **Active Listening:** The response demonstrates minimal engagement with the client's expressed concerns. **Reflecting Feelings:** The question serves as a grounding technique rather than emotional reflection. **Situational Appropriateness:** The response is weakly appropriate as it diverts attention from the client's emotional needs.

Scores: Non-Judgmental: 1, Warmth: 0, Respect: 1, Active: 1, Reflecting: 1, Situational: 1

Now evaluate this response:

Therapist Response: "Okay, good. I'll see you next week. And we'll see how you did with those goals."

Explanations: **Non-Judgmental Language:** Non-judgmental language fosters a supportive therapeutic environment. The phrase "we'll see how you did with those goals" avoids explicit criticism but implies evaluation. **Warmth and Encouragement:** The response conveys mild encouragement through continuity and reassurance. **Respect for Autonomy:** The therapist acknowledges the client's responsibility for goal progress without imposing directives. **Active Listening:** The response reflects acknowledgment of prior discussion about goals and future follow-up. **Reflecting Feelings:** Emotional reflection is limited, as the response focuses on task completion rather than affect. **Situational Appropriateness:** The response is appropriate as a session-closing statement that reinforces accountability and continuity.

Figure 9: Complete few-shot in-context example used for GPT-4o evaluation. The prompt includes prior therapist responses with full principle-wise explanations and ordinal scores, followed by a new therapist response evaluated using the same structure.

Few-shot prompting improves output consistency but does not correct this bias. As shown in Figure 12, predictions remain skewed toward positive labels, with limited recovery of negative and neutral classes. Compared to CARE, GPT-4o's errors are less ordinaly localized, highlighting the limitations of prompt-based evaluation for nuanced therapeutic assessment.

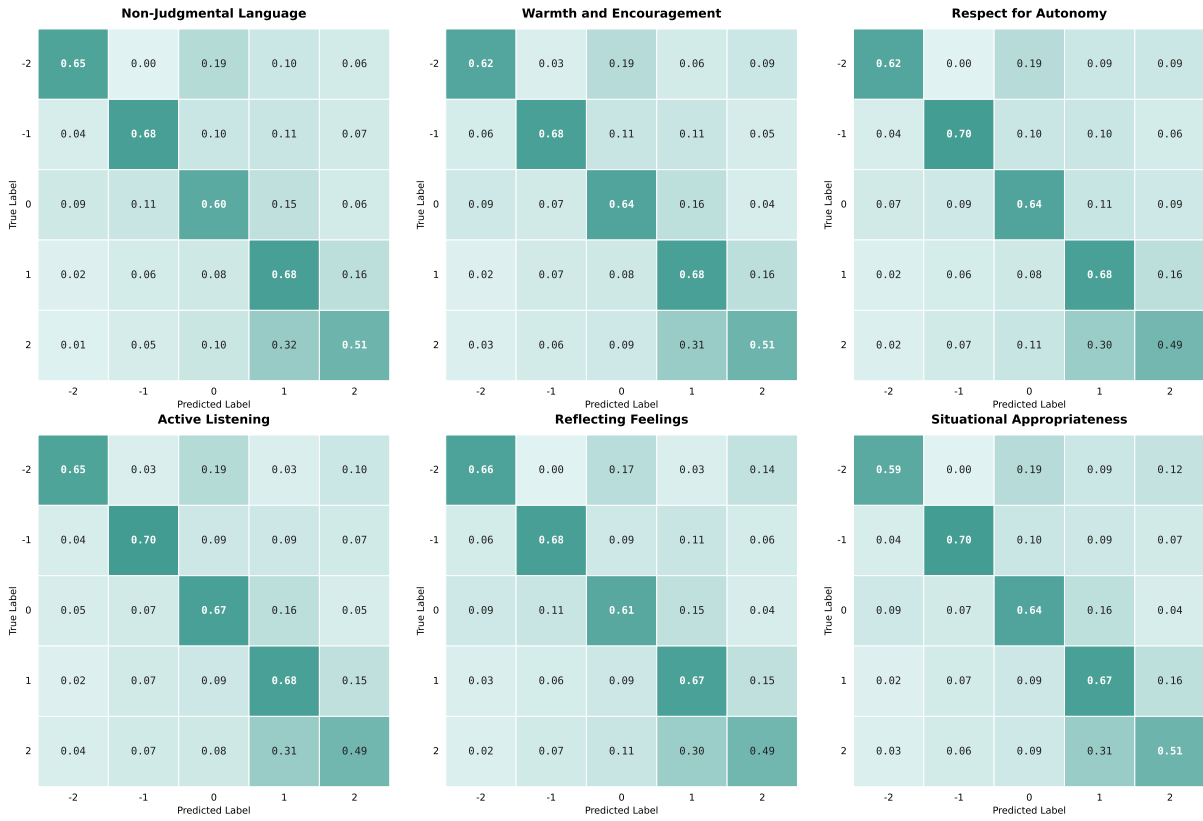


Figure 10: Confusion matrices computed on our proposed dataset, FAITH-M using our model CARE, across six therapeutic principles.

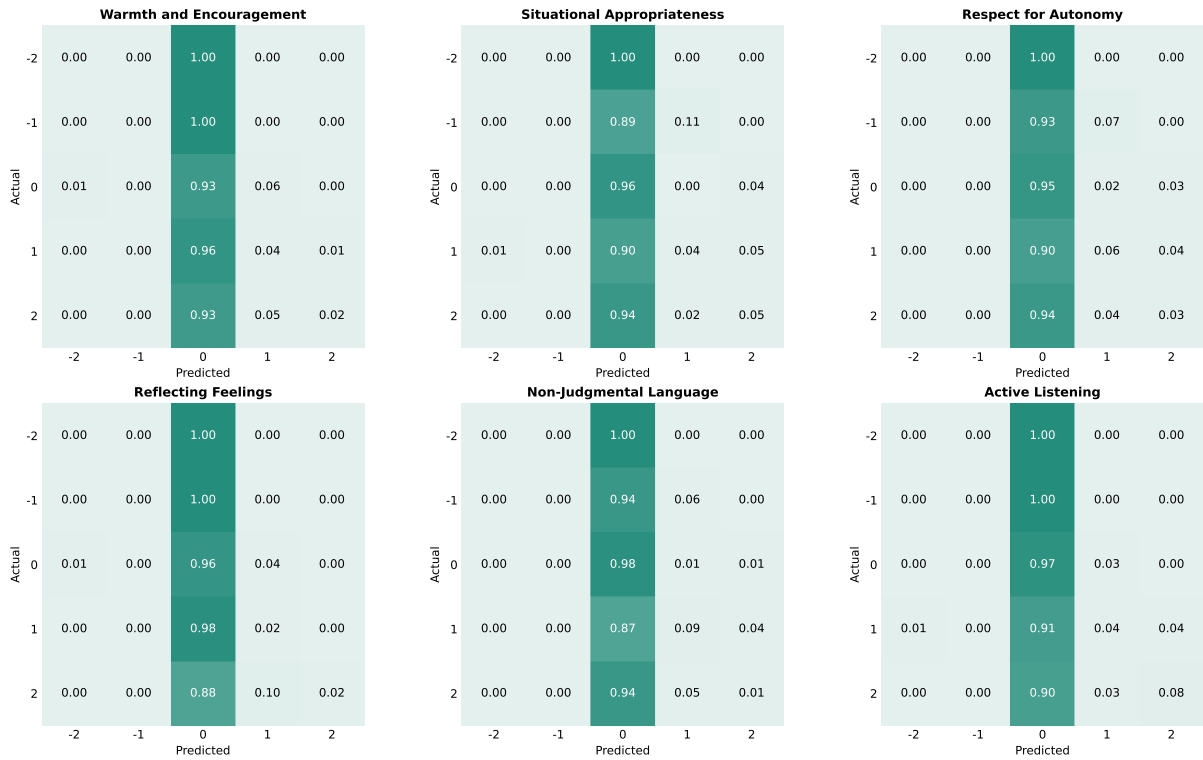


Figure 11: Confusion matrices computed on FAITH-M using GPT-4o under zero-shot prompting across six therapeutic principles. The matrices reveal a strong bias toward positive predictions, with negative and neutral categories frequently collapsed into Mild Positive or Strong Positive classes.

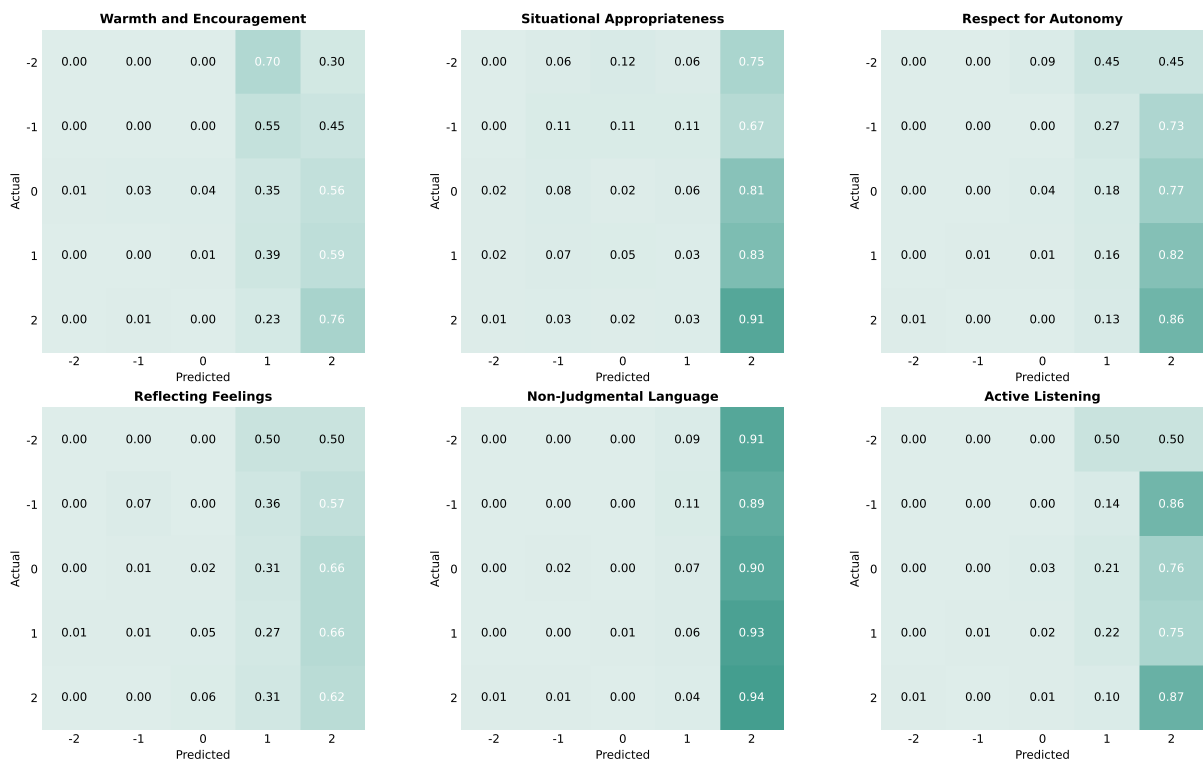


Figure 12: Confusion matrices computed on FAITH-M using GPT-4o under few-shot prompting across six therapeutic principles. While few-shot prompting sharpens confidence around positive labels, it does not correct the systematic under-recognition of negative and neutral categories.