

AuroraCounsel: A Multi-Agent Large Language Model Collaboration Framework for Psychological Counseling

Yuean Zhu^{*†}, Zichao Song^{*}, Haiming Luo

Jiaying University, Meizhou, China

zhuyuean@jyu.edu.cn, songzichao084@gmail.com

 GitHub https://github.com/ananzhu/AuroraCounsel_open

Abstract

While Large Language Models (LLMs) show promise in alleviating mental health resource shortages, existing approaches struggle with complex emotional dynamics and maintaining theory-grounded counseling trajectories. In particular, they fail to adapt interaction rhythms across counseling phases, often resulting in excessive exploration or premature interventions. To address these limitations, we propose AuroraCounsel, an agent-based counseling framework that models professional counseling workflows. The framework integrates an emotion-sensing agent for affect recognition, a reflection agent for structured psychological analysis, and a guidance agent for response generation. Through emotion-driven dynamic routing, AuroraCounsel selectively activates reflective reasoning only when negative emotions are detected, enabling accurate problem tracing while avoiding unnecessary computation. Building on this architecture, we further align counseling behaviors by optimizing the guidance agent with Grouped Relative Policy Optimization (GRPO), allowing it to regulate questioning and supportive strategies across dialogue phases and overcome the limitations of supervised fine-tuning in capturing clinical counseling rhythms, such as those prescribed by REBT. Evaluations using LLM-based evaluators and expert assessments by professional counselors demonstrate consistent improvements over existing approaches in professionalism, empathy, and interaction appropriateness.

1 Introduction

With the rising global prevalence of mental disorders affecting approximately 970 million people (WHO, 2022), the disparity between surging public demand and the scarcity of professional counselors remains critical, particularly in less developed regions (WHO, 2021). Although

Large Language Models (LLMs) exemplified by GPT (Brown et al., 2020), LLaMA (Touvron et al., 2023), and Qwen (Bai et al., 2023) have demonstrated therapeutic potential in systems like PsyDT (Xie et al., 2025) and CPsyCoun (Zhang et al., 2024), they exhibit inherent limitations regarding complex emotional processing, reasoning stability, and therapeutic efficacy. Recent work (Xu et al., 2025a; Ozgun et al., 2025) adopts multi-agent frameworks to decompose counseling into specialized functions, improving interpretability and theoretical grounding. However, these approaches overlook alignment with human counselors' phase-dependent interaction behaviors, which require dynamically adjusting questioning and guidance strategies beyond what supervised fine-tuning can capture.

Based on this observation, we propose **AuroraCounsel**, a multi-agent framework that explicitly models clinical counseling workflows. AuroraCounsel decomposes counseling into emotion sensing, structured psychological reflection, and guided response generation, enabling phase-aware coordination among specialized agents.

The main contributions of this paper are as follows:

1. We propose **AuroraCounsel**, a phase-aware multi-agent counseling framework that models professional counseling workflows and enables behavior-aligned intervention through emotion-driven coordination among specialized agents.
2. In collaboration with psychology experts, we construct a chain-of-thought reasoning dataset and optimize the guidance agent using Grouped Relative Policy Optimization (GRPO) (Shao et al., 2024; Guo et al., 2025) to align phase-dependent counseling behaviors that supervised fine-tuning fails to capture.

^{*}Equal contribution.

[†]Corresponding author.

3. Comprehensive evaluations using LLM-based automated assessments with GPT-4o (OpenAI, 2024) and Claude-Sonnet-4 (Anthropic, 2025), together with expert judgments from professional psychologists, demonstrate consistent improvements over existing methods in professionalism, authenticity, and empathy.

2 Related Work

2.1 LLMs and Multi-Agent Collaboration in Psychological Counseling

Recent studies have explored large language models for mental health counseling and emotional support. Although fine-tuned counseling models have improved empathy and domain relevance, most existing approaches still rely on single-agent generation, which limits process stability, explicit clinical reasoning, and adaptability in complex multi-turn interactions. Recent work has therefore introduced modular and multi-agent designs. AutoCBT decomposes cognitive behavioral therapy into multiple autonomous agents, Trustworthy AI Psychotherapy combines multi-agent counseling with explainable mental disorder diagnosis, Multi-AgentESC structures emotional support generation into dialogue analysis, strategy deliberation, and response generation, and MAGneT applies coordinated multi-agent collaboration to synthesize multi-turn mental health counseling sessions. However, these systems mainly focus on emotional support, diagnostic subtasks, or synthetic data generation, and do not fully model the end-to-end workflow of professional counseling in real interactions. Recent studies have likewise highlighted the lack of standardized dialogue frameworks that capture the sequential logic of therapeutic practice across recognition, formulation, intervention, and risk management. (Xu et al., 2025a; Ozgun et al., 2025; Bunge and Desage, 2025; Ding et al., 2024; Xu et al., 2025b; Mandal et al., 2025)

To bridge this gap, we propose **AuroraCounsel**, a multi-agent counseling framework that organizes the psychotherapy process through the sequential integration of **emotion recognition, psychological reasoning, and counseling feedback**. Unlike prior modular or multi-agent systems that mainly target strategy selection, diagnosis, or synthetic session construction, AuroraCounsel is designed for real multi-turn therapeutic dialogue. Its core **AC-Reflect** module performs structured, theory-guided reasoning, making counselor reasoning more ex-

plicit and interpretable. Together with emotion-driven dynamic routing and phase-aligned response optimization, AuroraCounsel supports both reflection depth control and stage-aware counseling adaptation. In this sense, AuroraCounsel extends existing multi-agent or modular LLM counseling systems by integrating clinical workflow emulation, explicit therapeutic reasoning, adaptive control, and crisis awareness within a unified end-to-end framework. (Xu et al., 2025b; Mandal et al., 2025)

2.2 Challenges of Theoretical Alignment in Psychological Counseling

Aligning LLMs with established psychological theories is essential for professional effectiveness. Significant progress has been made by works like PsyDTLLM (Xie et al., 2025), which aligned model behaviors with Rational Emotive Behavior Therapy (REBT) (Turner, 2016) using domain-specific datasets (Sun et al., 2025). These efforts validate the feasibility of grounding LLMs in specific psychological frameworks through data-driven approaches.

However, Supervised Fine-Tuning (SFT) alone proves insufficient for capturing the dynamic counseling rhythms required by REBT—specifically, the strategic transition from high-frequency questioning in the exploration stage to reflective guidance in later stages (Dryden, 2024; Carona et al., 2021). SFT-based models often struggle to master this adaptive shift, tending towards static or extreme questioning behaviors that compromise professional alignment.

To address this, we introduce reinforcement learning via the GRPO algorithm, incorporating a multi-dimensional reward function that explicitly enforces the REBT-prescribed questioning trajectory. Unlike ToMAP, which utilizes auxiliary modules to enhance reasoning in persuasion tasks, our approach directly embeds counseling theory into the learning objective via a question-rate penalty. This enables the model to internalize the natural rhythm of therapeutic dialogue without relying on external mechanisms, thereby achieving a deeper theoretical alignment beyond mere procedural knowledge.

3 Method

3.1 Framework Overview

As illustrated in Figure 1, the AuroraCounsel framework comprises three integral modules:




Figure 1: Overall structure of the AuroraCounsel framework.

AC-Sense: Functioning as the perceptual unit, this module performs real-time sentiment analysis to classify client inputs into positive or negative states.

AC-Reflect: Operating as the cognitive core, this component activates upon negative detection to construct structured, theoretically grounded analysis chains.

AC-Guide: Serving as the execution unit, this module synthesizes responses by directly handling positive inputs or integrating reasoning chains for negative scenarios.

3.2 AC-Sense

The emotion analysis module, **AC-Sense**, functions as a fine-grained sentiment polarity classifier built upon the pre-trained BERT-Base-Chinese architecture (Devlin et al., 2019). To train this module, we constructed a hybrid sentiment classification dataset integrating the NLPCC 2014 corpus (TCCI, 2014) with over 7,000 psychological counseling dialogues manually annotated by researchers. Following rigorous data cleaning and expert verification, we established a balanced data partition comprising 16,000 training samples and 1,000 testing samples. Architecturally, the system utilizes the dense representation of the [CLS] token as an aggregated semantic embedding, which is subsequently processed by a task-specific Multi-Layer Perceptron (MLP) with a ReLU activation function to generate sentiment probabilities. Through

end-to-end optimization minimizing cross-entropy loss, the model effectively adapts the general semantic features of BERT to the specific decision boundaries required for accurate sentiment analysis. Ultimately, this integration of general and domain-specific training data endows the model with the robustness and generalization capabilities crucial for real-world psychological interventions.

3.3 AC-Reflect

3.3.1 Dataset Construction Principle

To transcend superficial generation and incorporate professional cognitive depth, we introduce the **AC-Reflect** module, a reasoning engine built upon Qwen2.5-7B-Instruct (Qwen et al., 2025). Upon detecting negative emotions, this module executes a structured analysis grounded in theoretical frameworks. Specifically, we collaborated with psychology professors to design the role-playing system prompt shown in Appendix A. Utilizing this prompt to reconstruct the PsyDTCorpus via model-based role-playing, we generated a **reasoning dataset** comprising **2,400 sessions** and **44,000 dialogue turns** that embody professional cognitive logic. By explicitly integrating **pattern recognition**, **psychological mechanisms**, **intervention strategies**, and **crisis warning**, the resulting reasoning chains ensure interpretable responses that strictly adhere to counseling standards.

3.3.2 Core Dimension Design

We redefine the following four dimensions when negative emotions are detected to guide the construction of the reasoning chain:

Pattern Recognition. This dimension involves an integrated analysis of observable behaviors and emotional states to uncover enduring cognitive or relational patterns. The process entails identifying dominant and underlying emotions, clarifying their adaptive functions as defense mechanisms, and mapping these dynamics to long-standing interpersonal schemas or core beliefs.

Psychological Mechanism Analysis. This phase systematically traces the identified patterns back to their deeper psychological mechanisms and developmental origins, such as attachment insecurity and family system dynamics. Consequently, it frames the client’s present difficulties within a broader developmental psychopathological context.

Intervention Structure. The third dimension formulates an actionable intervention framework by incorporating evidence-based strategies like cognitive restructuring and resource activation. This approach aims to support the management of immediate challenges while simultaneously establishing a foundation for long-term psychological adjustment.

Crisis Warning. The detection of potential risk signals necessitates the immediate activation of safety-oriented protocols, ranging from risk assessment to professional crisis referrals. This mechanism ensures the system incorporates robust ethical safeguards and fulfills the duty of care regarding self-harm or external harm risks.

3.4 AC-Guide

We develop **AC-Guide** by fine-tuning Qwen2.5-7B-Instruct on the PsyDTCorpus (Xie et al., 2025) to adapt the model to psychological counseling (see Appendix B). While supervised fine-tuning improves domain relevance, it fails to align interaction behaviors with professional counseling practices. Specifically, we observe (1) **excessive questioning** across dialogue stages, (2) **multiple questions per turn** that disrupt narrative flow, and (3) **repetitive questioning** that degrades coherence.

From the perspective of Rational Emotive Behavior Therapy (REBT), questioning frequency should vary across counseling phases: early stages benefit from active inquiry to uncover core beliefs, whereas later stages require reduced questioning

in favor of explanation and intervention. These observations indicate that supervised fine-tuning alone cannot capture phase-dependent counseling rhythms.

We initially apply Direct Preference Optimization (DPO) (Rafailov et al., 2023) to regulate questioning behaviors, but observe degenerate outcomes where questions are largely suppressed. This occurs because DPO relies on a single global reward function, which amplifies annotator biases toward declarative responses and penalizes questioning during greedy optimization, disrupting dialogue interactivity (see Appendix D) (Moskovitz et al., 2024; Gao et al., 2023).

To overcome this limitation, we adopt Grouped Relative Policy Optimization (GRPO) (Shao et al., 2024), which introduces group-structured relative preferences to constrain optimization within controlled behavioral boundaries. By clustering responses according to questioning tendencies, GRPO preserves necessary interrogative behaviors while jointly optimizing response quality and coherence. Based on this insight, following supervised fine-tuning, we employ GRPO-based reinforcement learning to align **AC-Guide** with phase-dependent counseling rhythms.

3.4.1 Policy Optimization Objective

We formulate psychological counseling dialogue generation as a reinforcement learning optimization problem. Given an input context $x = [H_{t-1}, U_t]$ consisting of the dialogue history H_{t-1} and current user input U_t , the policy π_θ generates candidate responses $Y_t \sim \pi_\theta(\cdot|x)$. The optimization objective is defined to maximize the expected reward while constraining deviation from the supervised reference model π_{ref} :

$$\max_{\theta} \mathbb{E}_{x, Y \sim \pi_\theta(\cdot|x)} [r(x, Y) - \beta \cdot D_{KL}(\pi_\theta(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x))] \quad (1)$$

where β serves as the coefficient for the Kullback-Leibler (KL) divergence regularization.

3.4.2 Reward Function Design

We design a composite reward function r that integrates six distinct components to reflect counseling rhythm, professionalism, fluency, and empathy:

$$r = r_{\text{phase}} + r_{\text{nonQ}} + r_{\text{rep}} + r_{\text{len}} + r_{\text{multiQ}} + r_{\text{in-zone}} \quad (2)$$

While the fundamental definitions for questioning dynamics are presented below, the detailed mathematical formulations for the specific penalty functions are provided in Appendix E.

(1) Phase-based Question Rate Reward (r_{phase}). To regulate the counseling rhythm, we calculate the cumulative questioning rate q_t at the current turn t . Let y_i denote the model response at turn i , and $\mathbb{I}_q(y_i)$ be an indicator function that equals 1 if y_i is a question and 0 otherwise. The rate is formulated as:

$$q_t = \frac{1}{t} \sum_{i=1}^t \mathbb{I}_q(y_i) \quad (3)$$

where t represents the current turn number within the session.

Based on the average dialogue depth of approximately 20 turns in our SFT dataset, we partition the session into three phases to guide the model from exploration to intervention. We establish target questioning rate intervals of $[0.70, 1.00]$ for the initial 8 turns, $[0.40, 0.70]$ for turns 9 through 16, and $[0.00, 0.40]$ for subsequent turns. The reward r_{phase} is subsequently computed to minimize the deviation between the historical questioning rate and these targets using a smoothed function.

(2) Non-question Reward (r_{nonQ}). To encourage the transition to active intervention, we introduce a specific incentive for declarative responses. When the dialogue progresses beyond the 16th turn, a fixed positive reward is assigned to any response that does not constitute a question, promoting advice-giving and emotional support.

(3) Repetition Penalty (r_{rep}). To mitigate redundancy, we quantify lexical overlap using the Jaccard similarity of 2-grams between the current response and the four most recent assistant responses. A scaled penalty is applied when this similarity exceeds a threshold ($\tau = 0.75$).

(4) Length Reward (r_{len}). Research indicates that moderate response length optimizes perceived usefulness (Huang et al., 2024) and readability (Scott, 2022). Complementing this, our preliminary SFT experiments revealed a critical trade-off: excessive length compromises fluency, whereas brevity often lacks professional empathy. Consequently, to balance concision with depth, we target a token interval of $[35, 65]$ with an optimal midpoint of 50, enforced via a quadratic penalty.

(5) Multi-question Penalty (r_{multiQ}). To enhance authenticity, we introduce a progressive penalty that activates when a single response contains more

than two questions. This mechanism discourages interrogation-style interactions.

(6) In-zone Reward ($r_{\text{in-zone}}$). Complementing the phase-based reward, this component encourages the questioning rate to align closely with the center of the target interval for each phase, rather than merely staying within the boundaries, thereby ensuring trajectory stability.

We adopted a training pipeline progressing from SFT-based domain adaptation to GRPO reinforcement learning. By utilizing this multi-dimensional reward mechanism, the model **aligns its interaction patterns with the professional counseling rhythms of human experts**, characterized by early-stage inquiry, intermediate analysis, and late-stage interpretation, establishing a foundation for high-quality psychological interventions.

3.5 Dynamic Routing Mechanism

The core of the AuroraCounsel framework is its **emotion-driven dynamic routing mechanism**, which adaptively adjusts information processing paths based on the client’s emotional state to optimize adaptability and efficiency.

Algorithm 1: Emotion-Driven Dynamic Routing Mechanism in AuroraCounsel

Input: User utterance U , dialogue history H
Output: Therapeutic response R
Data: Models: f_{sent} (Sense), f_{think} (Reason), f_{guide} (Counsel)
Stage I: Sentiment Analysis and Reasoning
Trigger;
 $s \leftarrow f_{\text{sent}}(U)$; // Classify sentiment of user utterance
if $s = \text{Negative}$ **then** $\theta \leftarrow f_{\text{think}}(U, H)$; // Generate reasoning chain
;
else $\theta \leftarrow \emptyset$; // No reasoning invoked
;
Stage II: Response Generation;
 $R \leftarrow f_{\text{guide}}(U, \theta, H)$; // Generate therapeutic response
return R ;

Specifically, upon detecting **positive** sentiment via the classifier f_{sent} , the system directly invokes **AC-Guide** (f_{guide}) to generate empathetic responses, thereby bypassing unnecessary computation and enhancing interaction speed. Conversely, when the sentiment is **negative**, the system activates **AC-Reflect** (f_{think}) to synthesize a structured reasoning chain θ . This chain is subsequently integrated into **AC-Guide**, enabling the generation of therapeutic responses R characterized by high interpretability and professionalism. Formally, AC-

Guide combines the user utterance U , dialogue history H , and the conditional reasoning chain θ to produce the final output.

4 Experiments

4.1 Experimental Setup

Evaluation Set. We used 240 test samples from PsyDTCorpus as the evaluation set for overall performance comparison experiments and ablation studies.

Model Configurations. **AC-Sense** fine-tunes BERT-Base-Chinese on a hybrid sentiment dataset. **AC-Reflect** fine-tunes Qwen2.5-7B-Instruct on a structured reasoning dataset. **AC-Guide** utilizes Qwen2.5-7B-Instruct, optimized through supervised fine-tuning on the PsyDTCorpus followed by GRPO alignment.

Evaluation Metrics. We employed a multi-turn evaluation framework aligned with CPsyCoun to assess response coherence within the full conversational context (see Appendix C). The assessment prioritizes five professional dimensions: **comprehensiveness**, **professionalism**, **authenticity**, **empathy**, and **safety** (Sarkar et al., 2023; Giorgi et al., 2023).

Evaluation Protocol. To mitigate systematic biases in automated evaluation including position bias (Shi et al., 2025), scoring instability (Chen et al., 2024), and self-preference (Ye et al., 2025; Li et al., 2025), we evaluated the models on over **1,100** utterances from the PsyDTCorpus spanning **12 distinct thematic domains** using a hybrid framework. We employed a bias-attenuation pipeline detailed in Appendix C that integrates GPT-4o and Claude-Sonnet-4 with dynamic re-scoring to neutralize specific evaluation tendencies. Furthermore, to validate performance and eliminate potential **model biases**, a panel of **five university psychology professors**, each holding clinical counseling certifications, assessed over **240** dialogue turns per model across the same 12 thematic domains. This human evaluation described in Appendix F ensures professional validity and calibrates the automated metrics.

Baselines. **General base model:** Qwen2.5-7B-Instruct, Kimi-K2 (Team et al., 2025), DeepSeek-V3.2 (DeepSeek-AI et al., 2025). **Professional counseling models:** MeChat (Qiu et al., 2024), PsyChat, PsyDTLLM, CpsycounX, MindChat, Xinyuan (Maintainers, 2025), MultiAgentESC (Xu et al., 2025b), AuroraCounsel (ours).

4.2 Evaluation Results




Figure 2: Human evaluation results by psychology professors across seven models.

As illustrated in Table 1 and Figure 2, AuroraCounsel demonstrates superior performance across automated benchmarks and expert assessments. For the automated evaluation involving 60 dialogues totaling over 1,100 turns, we validated differences using two-sided Wilcoxon signed-rank and paired t-tests. With Holm-Bonferroni correction ($\alpha = 0.05$), our model achieved significant superiority in total scores over all baselines. Furthermore, expert evaluation showed high **absolute agreement** among raters (ICC(2,5)=0.85). Finally, a Pearson correlation of 0.67 between automated metrics and expert scores reinforces evaluation reliability.

Consistent with expert consensus, AuroraCounsel achieves the highest score in **professionalism**, **authenticity**, and **empathy**. Conversely, large-scale generalist models, specifically DeepSeek-v3.2 and Kimi-K2, demonstrate superior **comprehensiveness**, which we attribute to their extensive parameter scales and enhanced context-processing capabilities. However, compared to domain-specific counterparts, these generalist models predominantly yield didactic advice.

4.3 Evaluation under the MAGneT Protocol

To provide a more fine-grained evaluation of AuroraCounsel, we adopted the assessment protocol proposed by MAGneT. Specifically, we selected 100 counseling cases from the MAGneT benchmark, used Qwen3.5-Plus to simulate clients in multi-turn interactions, and let AuroraCounsel generate the corresponding counselor responses. GPT-5 was then used to score each session on **CTRS**, **WAI**, and **PANAS**. Detailed explanations of these evaluation dimensions are provided in Appendix I. The results are reported in Table 2.

Table 1: Model evaluation results across **comprehensiveness** (Comp.), **professionalism** (Prof.), **authenticity** (Auth.), **empathy** (Empa.), and **safety** (Saf.) (Mean $\pm$ SD). Best in **bold**, second-best underlined. Significance vs. **AuroraCounsel** (two-sided Wilcoxon signed-rank test): * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Type	Models	Comp.	Prof.	Auth.	Empa.	Saf.	Total
Base	Qwen2.5-7B	1.34 ± 0.07 **	1.92 ± 0.10 ***	1.68 ± 0.12 ***	1.54 ± 0.15 ***	1.00 ± 0.00	7.48 ± 0.37 ***
	Qwen3-32B	1.47 ± 0.07	2.28 ± 0.10 ***	1.82 ± 0.13 ***	2.05 ± 0.13 ***	1.00 ± 0.00	8.62 ± 0.35 ***
	Kimi-K2	1.56 ± 0.07	<u>2.49</u> ± 0.18 **	1.93 ± 0.18 ***	1.84 ± 0.16 ***	1.00 ± 0.00	8.82 ± 0.51 ***
	DeepSeek-v3.2	<u>1.51</u> ± 0.10	2.42 ± 0.15 ***	<u>2.07</u> ± 0.13 ***	<u>2.09</u> ± 0.16 ***	1.00 ± 0.00	<u>9.09</u> ± 0.43 ***
Domain	MeChat	1.20 ± 0.10 ***	1.83 ± 0.18 ***	1.71 ± 0.14 ***	1.60 ± 0.19 ***	1.00 ± 0.00	7.34 ± 0.56 ***
	PsyChat	1.30 ± 0.13 ***	1.84 ± 0.21 ***	1.79 ± 0.17 ***	1.91 ± 0.16 ***	1.00 ± 0.00	7.84 ± 0.59 ***
	PsyDTLLM	1.20 ± 0.11 ***	2.17 ± 0.14 ***	1.69 ± 0.15 ***	1.87 ± 0.18 ***	1.00 ± 0.00	7.93 ± 0.43 ***
	CPsyCounX	1.19 ± 0.11 ***	1.87 ± 0.14 ***	1.67 ± 0.15 ***	1.58 ± 0.18 ***	0.99 ± 0.02 ***	7.30 ± 0.54 ***
	MindChat	1.08 ± 0.09 ***	1.82 ± 0.12 ***	1.63 ± 0.16 ***	1.56 ± 0.18 ***	0.99 ± 0.02 ***	7.08 ± 0.32 ***
	MultiAgentESC	1.18 ± 0.15 ***	2.05 ± 0.17 ***	1.83 ± 0.16 ***	1.76 ± 0.16 ***	1.00 ± 0.00	7.82 ± 0.57 ***
	Xinyuan	1.32 ± 0.13 ***	2.36 ± 0.17 ***	1.81 ± 0.16 ***	1.96 ± 0.15 ***	1.00 ± 0.00	8.45 ± 0.54 ***
Our	AuroraCounsel	1.39 ± 0.08	2.84 ± 0.11	2.20 ± 0.11	2.38 ± 0.13	1.00 ± 0.00	9.81 ± 0.39

Table 2: Evaluation of generated counseling sessions across CTRS, PANAS, and WAI dimensions. Asterisks indicate significant differences from AuroraCounsel (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). For CTRS: U (Understanding), I (Interpersonal Effectiveness), C (Collaboration), D (Guided Discovery), F (Focus), S (Strategy). For PANAS: Pos. Att. (Positive Attitude), Neu. Att. (Neutral Attitude), Neg. Att. (Negative Attitude), P (average shift in positive emotions), N (average shift in negative emotions). $\delta(\%)$ denotes the normalized percentage-point difference between AuroraCounsel and the best baseline over the full scale range. Following MAGneT, the scale range is 6 for CTRS and WAI, and 4 for PANAS. For $\downarrow$ metrics, the sign is adjusted so that positive values indicate better performance.

Method	CTRS						PANAS						WAI		
	General			CBT			Pos. Att.		Neu. Att.		Neg. Att.		Task ($\uparrow$)	Goal ($\uparrow$)	Bond ($\uparrow$)
	U ($\uparrow$)	I ($\uparrow$)	C ($\uparrow$)	D ($\uparrow$)	F ($\uparrow$)	S ($\uparrow$)	P ($\uparrow$)	N ($\downarrow$)	P ($\uparrow$)	N ($\downarrow$)	P ($\uparrow$)	N ($\downarrow$)			
CPsyCounX	4.42*	4.81**	3.82**	3.65*	3.62**	3.86*	+0.39***	-0.69	+0.28	-0.93	-0.42	-0.24	4.99***	4.33*	4.27
Xinyuan	3.82***	5.11*	3.63**	3.22***	3.21***	3.42**	+0.31***	-0.62	+0.36	-1.01	-0.34	-0.53	4.96***	4.48**	4.35
MultiAgentESC	5.42	5.21*	3.57***	4.38	3.59**	3.19***	+0.19***	-0.59*	+0.16	-0.78	-0.53	-0.58	5.19***	4.49**	4.08*
AuroraCounsel	5.28	5.64	4.53	4.03	4.33	4.25	+0.51	-0.75	+0.42	-0.94	-0.37	-0.67	5.77	4.80	4.51
$\delta(\%)$	-2.3	+8.8	+11.8	-5.8	+11.8	+6.5	+3.0	+1.5	+1.5	-1.8	-0.8	+2.3	+9.7	+5.2	+2.7

This protocol enables a deeper assessment than standard response-level evaluation, as it jointly measures therapeutic competence, alliance quality, and affective change. As shown in Table 2, AuroraCounsel achieves the best overall performance across the three dimensions. It obtains the highest scores on most **CTRS** subdimensions, performs best on **WAI**, and shows clear advantages on **PANAS** in increasing positive affect and reducing negative affect. These results further demonstrate the effectiveness of AuroraCounsel under a rigorous multi-dimensional evaluation setting.

4.4 Ablation Study

To evaluate core module contributions, we conducted ablation studies on over 400 turns using a dual-evaluator strategy (averaging GPT-4o and Claude-Sonnet-4) validated by paired t-tests. We

compared the full framework against three configurations: **w/o AC-Sense** (bypassing the sentiment router to execute AC-Reflect reasoning on every turn), **Only AC-Guide (SFT)** (SFT-based direct generation), and **Only AC-Guide (GRPO)** (RL-enhanced generation). Table 3 details the results across comprehensiveness, professionalism, authenticity, empathy, safety, and token consumption. Notably, the significant superiority of **AuroraCounsel** over **Only AC-Guide (GRPO)** validates the necessity of **AC-Reflect**, confirming that structured reasoning is essential for superior counseling interactions.

4.4.1 Analysis of Ablation Study

As shown in Table 3, removing the emotion-driven dynamic routing mechanism and forcing AC-Reflect to perform reasoning at every dialogue

Table 3: Ablation study results of our model (* for $p < 0.05$, ** for $p < 0.01$, and *** for $p < 0.001$).

Setting	Comp.	Prof.	Auth.	Emp.	Safe.	Total	Token
AuroraCounsel (w/o AC-Sense)	1.36 $\pm$ 0.06 *	2.84 $\pm$ 0.07	2.18 $\pm$ 0.11	2.41 $\pm$ 0.06	1.00 $\pm$ 0.00	9.79 $\pm$ 0.24	3319 $\pm$ 152
Only AC-Guide (GRPO)	1.32 $\pm$ 0.07 **	2.71 $\pm$ 0.15 ***	2.00 $\pm$ 0.15 ***	2.22 $\pm$ 0.10 ***	1.00 $\pm$ 0.00	9.25 $\pm$ 0.44 ***	1169 $\pm$ 85
Only AC-Guide (SFT)	1.25 $\pm$ 0.11 ***	2.29 $\pm$ 0.12 ***	1.82 $\pm$ 0.15 ***	1.90 $\pm$ 0.11 ***	1.00 $\pm$ 0.00	8.26 $\pm$ 0.45 ***	648 $\pm$ 43
AuroraCounsel	1.39 $\pm$ 0.08	2.84 $\pm$ 0.11	2.20 $\pm$ 0.11	2.38 $\pm$ 0.13	1.00 $\pm$ 0.00	9.81 $\pm$ 0.39	1953 $\pm$ 99

turn does not produce a statistically significant improvement over the full AuroraCounsel framework, despite substantially increasing computational cost by approximately 70%. This result suggests that the proposed routing mechanism can effectively preserve response quality while avoiding unnecessary reasoning overhead. Table 3 also shows that the GRPO implementation consistently outperforms the SFT baseline in professionalism and empathy. Compared with SFT, which tends to generate smooth but generic responses, GRPO benefits from multi-dimensional reward signals that better capture counseling-specific interaction patterns, thereby improving long-term consistency and alignment with professional intervention standards.

5 Reinforcement Learning Analysis




Figure 3: AC-Guide comparison of questioning rates under different training paradigms.

Counseling interaction rhythm relies heavily on questioning rates. Following Rational Emotive Behavior Therapy principles, effective practitioners shift from initial high-frequency questioning to a gradual reduction, fostering client autonomy. As Figure 3 illustrates, unlike the SFT baseline’s excessively high interrogation rate, the GRPO-optimized model achieves a strategic staged decline, dropping from high initial frequencies to **below 20%** after 20 turns. Figure 4 further validates GRPO effectiveness. The smoothed reward curve, incorporating repetition and redundancy penalties, converges to 0.08 within 3000 steps, ensuring stable policy optimization. Simultaneously, average output length




Figure 4: Reward convergence and average output length variation during GRPO training.

stabilizes around **50 tokens**, successfully balancing brevity with sufficient content. These metrics confirm that GRPO achieves stable convergence while aligning generated responses with professional counseling rhythms.

6 Conclusion

In this paper, we presented **AuroraCounsel**, a multi-agent framework mimicking human counselor workflows to enhance professional competence, empathy, and safety. The system orchestrates three core modules, **AC-Sense**, **AC-Reflect**, and **AC-Guide**, via a dynamic routing mechanism to ensure adaptive support. Extensive experiments demonstrate that AuroraCounsel achieves state-of-the-art performance, outperforming strong baselines in professionalism, empathy, and safety. Furthermore, the integration of structured reasoning and GRPO enables the natural internalization of professional interaction rhythms, establishing a robust foundation for trustworthy, **human-centered** intelligent mental health services.

Limitations

Despite the demonstrated effectiveness of AuroraCounsel, limitations remain regarding emotional granularity and system complexity. Currently, the AC-Sense module operates on a binary sentiment classification paradigm, effectively distinguishing positive and negative states but potentially overlooking the nuanced affective spectrum critical for precise clinical assessment; future iterations thus aim to incorporate fine-grained emotion recognition to capture complex psychological dynamics. Additionally, although the multi-agent architecture significantly enhances counseling professionalism through specialized role collaboration, it inevitably introduces increased system complexity and inference latency. Consequently, subsequent research will focus on distilling this distributed workflow into a unified end-to-end model to optimize operational efficiency and scalability while preserving professional performance.

Ethical Considerations

While LLMs show promise in psychological counseling, safety remains a critical concern. Regarding **automated evaluation**, various models consistently achieved high safety scores, suggesting a broad adherence to ethical standards. However, our **manual stress-testing** exposed a significant discrepancy: despite favorable automated rankings, **various models** still exhibited risks, such as generating misleading advice or failing to address self-harm scenarios adequately.

In contrast, **AuroraCounsel**, leveraging the **AC-Reflect** module, effectively identified and mitigated these latent risks. By incorporating **structured reasoning and collaborative oversight**, the framework consistently demonstrated superior safety and reliability compared to other models in the psychological counseling domain.

Consequently, developing robust safety evaluation mechanisms remains an urgent priority for the field. **We emphasize that despite optimizations, model outputs serve strictly as references and must never replace professional medical decision-making or be relied upon as the sole basis for clinical intervention.**

References

Anthropic. 2025. System card: Claude opus 4 & claude sonnet 4. [https://www-cdn.anthropic.com/](https://www-cdn.anthropic.com/4263b940cabb546aa0e3283f35b686f4f3b2ff47.pdf)

[4263b940cabb546aa0e3283f35b686f4f3b2ff47.pdf](https://www-cdn.anthropic.com/4263b940cabb546aa0e3283f35b686f4f3b2ff47.pdf). Accessed 2025-09-30.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Chengqiang Lu, Gao Liu, Keming Lu, Jianxin Ma, and 26 others. 2023. Qwen technical report. <https://arxiv.org/pdf/2309.16609>.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. **Language models are few-shot learners**. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, Vancouver, Canada. Curran Associates, Inc.

Eduardo L. Bunge and Christina Desage. 2025. **A framework for evaluating mental health artificial intelligence-based conversational agents**. *Journal of Technology in Behavioral Science*.

Carlos Carona, Charlotte Handford, and Ana Fonseca. 2021. **Socratic questioning put into clinical practice**. *BJPsych Advances*, 27(6):424–426.

Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. **Humans or LLMs as the judge? a study on judgement bias**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8301–8327, Miami, Florida, USA. Association for Computational Linguistics.

DeepSeek-AI, Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenhao Xu, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, and 245 others. 2025. **Deepseek-v3.2: Pushing the frontier of open large language models**. *Preprint*, arXiv:2512.02556.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **Bert: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Hang Ding, Joshua Simmich, Atiyeh Vaezipour, Nicole Andrews, and Trevor Russell. 2024. **Evaluation framework for conversational agents with artificial intelligence in health interventions: a systematic scoping review**. *Journal of the American Medical Informatics Association*, 31(3):746–761.

- Windy Dryden. 2024. *Good Practice in Rational Emotive Behaviour Therapy*. Routledge, London.
- Leo Gao, John Schulman, and Jacob Hilton. 2023. [Scaling laws for reward model overoptimization](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10835–10866. PMLR.
- Salvatore Giorgi, Shreya Havaldar, Farhan Ahmed, Zuhaib Akhtar, Shalaka Vaidya, Gary Pan, Lyle H. Ungar, H. Andrew Schwartz, and Joao Sedoc. 2023. [Psychological metrics for dialog system evaluation](#). *Preprint*, arXiv:2305.14757.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. [Deepseek-r1 incentivizes reasoning in llms through reinforcement learning](#). *Nature*, 645(8081):633–638.
- Shih-Hong Huang, Ya-Fang Lin, Zeyu He, Chieh-Yang Huang, and Ting-Hao Kenneth Huang. 2024. [How does conversation length impact user’s satisfaction? a case study of length-controlled conversations with llm-powered chatbots](#). In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems (CHI EA ’24)*, pages 1–13, New York, NY, USA. Association for Computing Machinery.
- Dawei Li, Renliang Sun, Yue Huang, Ming Zhong, Bohan Jiang, Jiawei Han, Xiangliang Zhang, Wei Wang, and Huan Liu. 2025. [Preference leakage: A contamination problem in llm-as-a-judge](#).
- Cylingo Xinyuan Model Maintainers. 2025. Xinyuan-llm-14b-0428 model card. Hugging Face model card. Available: <https://huggingface.co/Cylingo/Xinyuan-LLM-14B-0428> (accessed 2025-09-30).
- Aishik Mandal, Tanmoy Chakraborty, and Iryna Gurevych. 2025. [MAGneT: Coordinated Multi-Agent Generation of Synthetic Multi-Turn Mental Health Counseling Sessions](#). *arXiv e-prints*, arXiv:2509.04183.
- Ted Moskovitz, Aaditya Singh, DJ Strouse, Tuomas Sandholm, Ruslan Salakhutdinov, Anca Dragan, and Stephen McAleer. 2024. [Confronting reward model overoptimization with constrained rlhf](#). In *International Conference on Learning Representations*, volume 2024, pages 21998–22025.
- OpenAI. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Mithat Can Ozgun, Jiahuan Pei, Koen Hindriks, Lucia Donatelli, Qingzhi Liu, and Junxiao Wang. 2025. [Trustworthy ai psychotherapy: Multi-agent llm workflow for counseling and explainable mental disorder diagnosis](#). In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management, CIKM ’25*, page 2263–2272, New York, NY, USA. Association for Computing Machinery.
- Huachuan Qiu, Hongliang He, Shuai Zhang, Anqi Li, and Zhenzhong Lan. 2024. [Smile: Single-turn to multi-turn inclusive language expansion via chatgpt for mental health support](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 615–636, Miami, Florida, USA. Association for Computational Linguistics.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, and 24 others. 2025. [Qwen2.5 Technical Report](#). *arXiv preprint*. ArXiv:2412.15115 [cs].
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems, NeurIPS 2023*, volume 36, pages 53728–53741, New Orleans, Louisiana, USA. Curran Associates, Inc.
- Surjodeep Sarkar, Manas Gaur, L Chen, Muskan Garg, Biplav Srivastava, and Bhaktee Dongaonkar. 2023. Towards explainable and safe conversational agents for mental health: A survey.
- Edward Scott. 2022. [Readability: The optimal line length](#). Published May 10, 2022. Accessed Oct. 13, 2025.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Lin Shi, Chiyu Ma, Wenhua Liang, Xingjian Diao, Weicheng Ma, and Soroush Vosoughi. 2025. [Judging the judges: A systematic study of position bias in LLM-as-a-judge](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 292–314, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Xin Sun, Xiao Tang, Abdallah El Ali, Zhuying Li, Pengjie Ren, Jan de Wit, Jiahuan Pei, and Jos A. Bosch. 2025. [Rethinking the alignment of psychotherapy dialogue generation with motivational interviewing strategies](#). In *Proceedings of the 31st International Conference on Computational Linguistics (COLING 2025)*, pages 1983–2002, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- NLPCC 2014 Task Organizers / CCF TCCI. 2014. Sentiment classification with deep learning: Nlpcc

- 2014 evaluation task data. NLPCC 2014 Evaluation Task — Sample / Test Data. Available: https://tcci.ccf.org.cn/conference/2014/pages/page04_tdata.html.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, and 1 others. 2025. Kimi k2: Open agentic intelligence.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. Llama 2: Open foundation and fine-tuned chat models. <https://arxiv.org/pdf/2307.09288>.
- Martin Turner. 2016. Rational emotive behavior therapy (rebt), irrational and rational beliefs, and the mental health of athletes. *Frontiers in Psychology*, 07.
- WHO. 2021. *Mental Health ATLAS 2020*. World Health Organization, Geneva.
- WHO. 2022. *World mental health report: Transforming mental health for all*. World Health Organization, Geneva, Switzerland.
- Haojie Xie, Yirong Chen, Xiaofen Xing, Jingkai Lin, and Xiangmin Xu. 2025. *PsyDT: Using LLMs to construct the digital twin of psychological counselor with personalized counseling style for psychological counseling*. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1081–1115, Vienna, Austria. Association for Computational Linguistics.
- Ancheng Xu, Di Yang, Renhao Li, Jingwei Zhu, Minghuan Tan, Min Yang, Wanxin Qiu, Mingchen Ma, Haihong Wu, Bingyu Li, Feng Sha, Chengming Li, Xiping Hu, Qiang Qu, Derek F. Wong, and Ruifeng Xu. 2025a. *Autocbt: An autonomous multi-agent framework for cognitive behavioral therapy in psychological counseling*. Preprint, arXiv:2501.09426.
- Yangyang Xu, Jinpeng Hu, Zhuoer Zhao, Zhangling Duan, Xiao Sun, and Xun Yang. 2025b. *MultiAgentESC: A LLM-based multi-agent collaboration framework for emotional support conversation*. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 4665–4681, Suzhou, China. Association for Computational Linguistics.
- Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh Chawla, and Xiangliang Zhang. 2025. *Justice or prejudice? quantifying biases in llm-as-a-judge*. In *International Conference on Learning Representations*, volume 2025, pages 102351–102390.
- Chenhao Zhang, Renhao Li, Minghuan Tan, Min Yang, Jingwei Zhu, Di Yang, Jiahao Zhao, Guancheng Ye, Chengming Li, and Xiping Hu. 2024. *Cpsycoun: A report-based multi-turn dialogue reconstruction and evaluation framework for chinese psychological counseling*. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 13947–13966, Bangkok, Thailand. Association for Computational Linguistics.

A Construction of the Reasoning Dataset

A.1 Data Source

We used the **PsyDT dataset** as the source of client utterances. This dataset contains multi-turn counseling dialogues covering a wide range of topics such as emotions, family, interpersonal relationships, and self-growth. In our setting, the client-side utterances serve as the *input context* to trigger the generation of structured reasoning chains.

A.2 Prompt Design

To ensure consistent and interpretable reasoning outputs, we designed a structured prompt (see Figure 8 9 10 11). The prompt specifies:

- **Role:** Qwen-Max is required to act as a professional psychological counselor, proficient in REBT and other counseling techniques.
- **Attention Guidelines:** The model must explicitly output reasoning chains, avoid vague statements, and carefully identify potential risk signals.
- **Skills and Constraints:** The prompt enforces layered counseling skills (e.g., session management, cognitive analysis, special case handling) and emphasizes clinical-style response norms.

A.3 Data Generation Process

The reasoning dataset was reconstructed through the following pipeline:

1. **Input:** Client utterances were extracted from the PsyDT dataset as the query context.
2. **Generation:** The Qwen-Max model was prompted with the designed template (Figure 8 10), simulating the role of a counselor.
3. **Output:** The model generated reasoning chains in a fixed format, each enclosed in `<think>...</think>` tags, and explicitly covering the four core dimensions:

- (a) **Pattern Recognition** – connecting behaviors, specific emotions, functional interpretations, and long-term patterns.
- (b) **Psychological Mechanism Analysis** – tracing back the deeper psychological mechanisms and formative causes.
- (c) **Intervention Structure** – providing preliminary and actionable intervention strategies.
- (d) **Crisis Warning** – activating risk assessment and referral protocols when danger signals appear.

This reconstructed dataset supports training the reasoning module to generate structured and interpretable psychological analysis.

A.4 Training Details

The AC-Reflect model was fine-tuned on the reconstructed reasoning dataset based on the Qwen2.5-7B-Instruct backbone. We adopted a parameter-efficient fine-tuning strategy using LoRA with a rank of 64, scaling factor $\alpha = 128$, and dropout rate of 0.05 applied to all trainable modules. Training was conducted in bfloat16 precision, and the maximum input sequence length was set to 4096 tokens.

The model was optimized using the AdamW optimizer with a cosine learning rate schedule. The initial learning rate was 2×10^{-4} with a linear warm-up of 40 steps. An effective batch size of 8 was achieved by setting the micro-batch size to 1 and accumulating gradients over 8 steps. Gradient clipping with a maximum norm of 1.0 was applied to maintain stability. Training was carried out for 3 epochs, and 5% of the dataset was reserved for validation.

B Training Details of AC-Guide

B.1 SFT Training Details for AC-Guide

The AC-Guide model was fine-tuned on the PsyDTCorpus counseling dialogue dataset, based on the Qwen2.5-7B-Instruct backbone. We adopted a LoRA-based parameter-efficient fine-tuning strategy with rank $r = 16$, scaling factor $\alpha = 32$, and a dropout rate of 0.01, applied to all trainable modules. Training was performed in fp16 precision with a maximum sequence length of 4096 tokens.

The model was optimized using the AdamW optimizer with a cosine learning rate schedule. The initial learning rate was 1×10^{-5} , with a linear

warm-up of 10 steps. An effective batch size of 8 was achieved by setting the micro-batch size to 1 and accumulating gradients over 8 steps. The gradient clipping threshold was 1.0. Training lasted for 3 epochs, with 0.5% of the data reserved for validation.

B.2 GRPO Training Details for AC-Guide

Model and Reward Setup. The GRPO training reused the SFT LoRA configuration ($r = 16$, $\alpha = 32$, dropout 0.05). A custom reward function was employed, integrating multiple components described in Section 3.3.2, including phase-based questioning rhythm, non-question preference, repetition penalty, length control, multi-question penalty, and in-zone stability. Each prompt produced six candidate responses ($n = 6$) per batch, which were grouped for reward comparison and gradient updates.

Optimization Configuration. Training was performed with a learning rate of 4×10^{-6} , warm-up ratio 0.02, and cosine decay scheduling. The maximum number of steps was 3300, with checkpoints saved every 300 steps. Each training batch contained one prompt with gradient accumulation over eight steps to form an effective batch size of 8. The KL-regularization coefficient was set to $\beta = 0.004$, and gradient norms were clipped at 0.5. Completions were generated with temperature 0.55, top- $p = 0.9$, and a maximum completion length of 90 tokens. Training was conducted in full-precision (fp32) mode to ensure numerical stability, and gradient checkpointing was enabled to reduce memory overhead.

C Evaluation Details

Based on the evaluation framework of CPsyCoun and with reference to PsyDT, we reconstructed the evaluation prompt to better suit the context of psychological counseling dialogues. The evaluation prompt is shown in Figure 12 13, covering five dimensions: **Comprehensiveness**, **Professionalism**, **Authenticity**, **Empathy**, and **Safety**. Crucially, to ensure experimental rigor and prevent **data leakage**, we strictly enforced data separation: the training split of the **PsyDTCorpus** was exclusively utilized for both the Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL) stages, while all evaluations were conducted on the **test set of the PsyDTCorpus**.

This section details the mathematical formulation of the multi-turn dialogue evaluation framework utilized in our experiments. Formally, we denote a dialogue session $\mathcal{X}$ consisting of m turns as $\mathcal{X} = \{(q_i, r_i)\}_{i=1}^m$, where q_i represents the user input and r_i denotes the model response at the i -th turn.

To capture the context dependency inherent in counseling scenarios, we define the dialogue history h_i preceding the i -th turn as $h_i = \{(q_j, r_j)\}_{j=1}^{i-1}$. The evaluation function $f_\phi(\cdot)$, parametrized by the LLM judge, calculates the quality score s_i for the current response r_i by conditioning on the available history:

$$s_i = \begin{cases} f_\phi(q_i, r_i), & i = 1, \\ f_\phi(h_i, q_i, r_i), & 1 < i \leq m. \end{cases} \quad (4)$$

Finally, the dialogue-level holistic score S is derived by averaging the scores across all turns to ensure a balanced assessment of interaction consistency:

$$S = \frac{1}{m} \sum_{i=1}^m s_i \quad (5)$$

To ensure the objectivity and reliability of the automated evaluation, we implemented a robust bias-attenuation pipeline comprising three key mechanisms designed to mitigate systematic variances inherent in LLM-based judging.

Dual-Model Fusion. We employ a multi-model consensus approach by aggregating evaluations from GPT-4o and Claude-Sonnet-4. The final assessment score is derived by averaging the independent outputs from both models. This fusion strategy effectively neutralizes individual model biases and specific linguistic preference tendencies that typically arise in single-model evaluations.

Dynamic Re-scoring Strategy. To address scoring instability characterized by stochastic fluctuations, we introduce an adaptive re-scoring mechanism. Specifically, if the divergence between the scores assigned by the two models exceeds a predefined threshold, the system automatically triggers a third evaluation round. The final result is subsequently computed as the consensus average of all valid scores to minimize variance and ensure statistical robustness.

Prompt Optimization. We refined the evaluation prompts to prioritize substantive content quality over stylistic features. By explicitly instructing the

evaluator to disregard superficial attributes such as response length and formatting styles, we ensure that the assessment focuses strictly on the professional dimensions of comprehensiveness, professionalism, authenticity, empathy, and safety.

D Construction of Preference Dataset and Analysis of DPO Failure

To align the model with the rhythm, professionalism, and empathy of human counselors, we constructed a specific preference dataset consisting of 4,000 pairs. This dataset was derived from the PsyDTCorpus using a human-in-the-loop annotation process. Researchers reviewed the model responses and labeled outputs exhibiting low empathy, insufficient professionalism, or excessive questioning at inappropriate stages as “rejected” or sub-optimal. Corresponding “chosen” responses were subsequently generated through rewriting by advanced Large Language Models or human experts to ensure therapeutic validity.

However, a retrospective analysis revealed a critical limitation in this approach. Due to the subjective variability among annotators, the dataset contained inconsistencies where questioning in later dialogue stages was frequently corrected into direct advice or solutions. When applying Direct Preference Optimization (DPO), the model failed to learn the phase-dependent nuance of *when* to ask questions. Instead, it interpreted the preference for advice over questions as a global penalty against the act of questioning itself.

This over-optimization led to a rapid degradation in the model’s interactive capabilities. As illustrated in Figure 6, the questioning rate of the model declined sharply during the training process. By training step 150, the questioning rate dropped to nearly zero. Consequently, the model lost its ability to trace psychological origins or guide the user through inquiry, rendering it passive and contradictory to our initial design objectives. This failure necessitated the shift to Grouped Relative Policy Optimization (GRPO) to explicitly enforce rhythm constraints.

E Mathematical Formulations of Reward Functions

This section details the rigorous mathematical formulations for the six reward components utilized in the GRPO framework.

E.1 Phase-based Question Rate Reward

(r_{phase})

We define the historical questioning rate $q(H)$ and the target intervals $[L, U]$ as described in the main text. Let q_{prev} denote the questioning rate prior to the current response, and q_{cur} denote the rate after generating the current response. The deviation distance $d(q)$ from the target interval is defined as:

$$d(q) = \max(0, q - U) + \max(0, L - q) \quad (6)$$

The final phase reward is computed using a Tanh-smoothed function to penalize the increase in deviation:

$$r_{phase} = \tanh(2 \cdot (d(q_{prev}) - d(q_{cur}))) \cdot (0.6 \cdot W_{phase}) \quad (7)$$

where $W_{phase} = 0.2$. For the first turn ($t = 1$), the reward is attenuated by a factor of 0.1 to prevent instability.

E.2 Non-question Reward (r_{nonQ})

To encourage active intervention in the later stages of dialogue (when $t \geq 17$), we apply a sparse reward for declarative responses. Let y be the generated response:

$$r_{nonQ} = 0.2 \cdot W_{nonQ} \cdot \mathbb{I}(t \geq 17) \cdot (1 - \mathbb{I}_q(y)) \quad (8)$$

where $W_{nonQ} = 0.2$.

E.3 Repetition Penalty (r_{rep})

To mitigate redundancy, we employ a penalty mechanism based on n -gram overlap. Let $G_n(y)$ denote the set of n -grams (with $n = 2$) for a given response y . The Jaccard similarity between two sequences a and b is defined as:

$$\text{Jaccard}(a, b) = \frac{|G_n(a) \cap G_n(b)|}{|G_n(a) \cup G_n(b)|} \quad (9)$$

Let S be the set of the $K = 4$ most recent assistant responses. We determine the maximum similarity s_{max} between the current response y and the history set S as $s_{max} = \max_{y' \in S} \text{Jaccard}(y, y')$. The final repetition penalty is formulated using a scaled Softplus function:

$$r_{rep} = -W_{rep} \cdot \frac{\ln(1 + e^{\beta \cdot (s_{max} - \tau)})}{\beta} \quad (10)$$

where $W_{rep} = 0.2$, the scaling factor $\beta = 8.0$, and the similarity threshold $\tau = 0.75$.

E.4 Length Reward (r_{len})

We enforce an optimal token length interval $[L_{min}, L_{max}] = [35, 65]$ with a midpoint $L^* = 50$. Let T be the length of the current response. The reward follows a quadratic penalty function:

$$r_{len} = W_{len} \left(1 - \left(\frac{T - L^*}{\sigma} \right)^2 \right) \quad (11)$$

where $W_{len} = 0.2$ and the curvature parameter $\sigma = 15$.

E.5 Multi-question Penalty (r_{multiQ})

To discourage the model from asking multiple questions in a single turn (interrogation style), we implement a progressive piecewise penalty based on the number of questions m in the generated response:

$$r_{multiQ} = -\min(0.12, \delta) \cdot W_{multiQ} \cdot \mathbb{I}(m > 2) \quad (12)$$

where the raw penalty term δ is calculated as:

$$\delta = 0.03(m - 2) + 0.02 \cdot \max(0, m - 4) \quad (13)$$

and the weight $W_{multiQ} = 0.5$. This formulation applies a stricter marginal penalty for each question beyond the fourth one.

E.6 In-zone Reward ($r_{in-zone}$)

To encourage trajectory stability within the target phase intervals, we calculate the center $c = \frac{L+U}{2}$ and half-width $h = \frac{U-L}{2}$ of the current target interval. The in-zone reward is defined as:

$$r_{in-zone} = B [\max(0, 1 - |q_{cur} - c|/h)]^2 \quad (14)$$

where $B = 0.05$. This provides a continuous gradient guiding the questioning rate toward the center of the optimal range.

F Human Evaluation Details

To ensure professional validity and clinical relevance, we recruited **five** psychology professors to conduct a rigorous human evaluation. To maintain consistency with the automated evaluation framework, the human assessment adopted an identical scoring rubric: Comprehensiveness (0–2), Professionalism (0–4), Authenticity (0–3), Empathy (0–3), and Safety (0–1). We conducted a comparative analysis involving AuroraCounsel and six baseline

models across **12 distinct thematic domains** derived from the test set. To mitigate potential bias, we implemented a **blind review process** where the identity of the models was anonymized during the assessment. Furthermore, evaluators were provided with the automated evaluation prompts to serve as reference standards for scoring. The detailed rating records from five of the representatives are presented in Tables 4.

F.1 Qualitative Analysis of Human Preferences

First, regarding **Xinyuan**, we observed that it received lower scores in human evaluation despite its potential high performance in automated metrics. Experts noted that Xinyuan tends to generate excessively long responses filled with detailed advice. While automated evaluators (LLM-as-a-judge) often exhibit a bias toward verbose and information-dense outputs, human professionals argued that such lecture-like "block text" severely disrupts the continuity of the interaction. In a real counseling scenario, this prevents the client from adequately expressing themselves and turns the dialogue into a didactic instruction rather than an empathetic exchange.

Second, a specific divergence appeared in Tables 4, where Professor 3 assigned a higher score to **PsyDTLLM** in Comprehensiveness compared to AuroraCounsel. Post-evaluation analysis reveals that this evaluator prioritized the explicit gathering of client background information. However, PsyDTLLM achieves this through a strategy characterized by **excessive and repetitive questioning**. While this relentless inquiry style ensures that background details are extracted (leading to a higher "Comprehensiveness" rating from this specific evaluator), it creates a mechanical and redundant interaction loop. In contrast, AuroraCounsel maintains a balanced rhythm to ensure interaction stability.

G Safety Case Study

The absence of significant safety disparities in automated metrics is attributed to the scarcity of extreme queries in standard datasets, which impedes a comprehensive evaluation of safety mechanisms. To address this deficiency, we constructed a specialized dataset of 100 high-risk queries for manual assessment. We present a representative **safety case study** from this evaluation, analyzing model

responses to a high-risk user query as illustrated in Figure 14 15 16 17 18 19 20 21. A distinct variation in safety awareness was observed across the evaluated models. Baseline models, including **PsyDTLLM**, **Xinyuan**, **MeChat**, and **CpsyCounX**, effectively expressed empathy but uniformly failed to address the immediate suicidal risk, opting instead to inquire about event details. Significantly, **AC-Guide**, which functions as a standalone component lacking the collaborative oversight of the **AC-Reflect** module, similarly failed to prioritize safety in this scenario. This observation provides ablation evidence that laterally substantiates the critical necessity of the AC-Reflect module for effective crisis detection. Regarding other models, **PsyChat** exhibited a critical lapse by suggesting the user "set aside" their agony, while **MindChat** demonstrated stronger safety awareness by advising professional help, although it focused on solutions rather than risk screening. In contrast, the complete **AuroraCounsel** framework, leveraging **AC-Reflect**, accurately identified the "suicidal ideation" and prioritized a **crisis assessment strategy**. By **explicitly inquiring about the user's immediate safety status** before proceeding, AuroraCounsel aligned most closely with professional crisis intervention standards.

H The Sentiment Analysis Module Ensures the Reliability of the Routing Mechanism

To systematically evaluate the effectiveness of this module, we conducted experiments on a held-out test set of 1000 samples, and compared it with three general-purpose large models. As shown in Table 5, AC-Sense achieved the best performance across all key metrics: Accuracy 88.00%, Precision 88.10%, Recall 88.00%, and F1 Score 87.99%. Notably, AC-Sense outperformed leading general-purpose models, including Kimi-K2, qwen3-235b-a22b-instruct-2507, and GPT-4o-min.

These results indicate that although general-purpose large models show strong generalization in many language understanding tasks, their performance on sentiment recognition tasks in counseling scenarios is significantly weaker. In contrast, AC-Sense, implemented with a BERT+MLP lightweight structure, achieves more reliable results while offering better computational efficiency and faster inference.

AC-Sense not only demonstrates clear advan-

tages in sentiment classification accuracy, but also ensures reliable and efficient input guarantees for the subsequent dynamic routing mechanism of AuroraCounsel, thereby enhancing the overall effectiveness and reliability of psychological counseling dialogue tasks.

I Detailed Explanation of Evaluation Dimensions

Following MAGneT, we adopt a multi-dimensional evaluation framework based on CTRS, WAI, and PANAS to assess counseling quality from the perspectives of therapeutic competence, working alliance, and affective change. This design provides a more informative assessment than response-level fluency evaluation, since a counseling model should not only generate natural utterances, but also demonstrate clinically meaningful skills, maintain a supportive therapeutic relationship, and facilitate constructive emotional change across a session. In our setting, the generated sessions are scored by GPT-5 under the MAGneT evaluation protocol. Detailed definitions and scoring requirements are described below.

I.1 CTRS

CTRS is used to evaluate the counselor's therapeutic competence in cognitive therapy. In the original scale, each item is rated from 0 to 6, with higher scores indicating stronger therapeutic performance. The CTRS scoring guideline specifies that even-numbered points are explicitly anchored, while odd-numbered points can be used when performance falls between two adjacent descriptions. Therefore, a score closer to 6 indicates that the counselor demonstrates stronger and more consistent therapeutic skill, whereas a score closer to 0 reflects weak or absent evidence of the target behavior.

Under the MAGneT protocol, CTRS is instantiated with six subdimensions. The general counseling dimensions include Understanding, Interpersonal Effectiveness, and Collaboration. The CBT-oriented dimensions include Guided Discovery, Focus, and Strategy. These dimensions jointly assess whether the counselor can understand the client's concerns accurately, communicate effectively, maintain collaboration, guide exploration, keep the session focused, and apply appropriate intervention strategies.

In practical interpretation, higher CTRS scores

indicate that the model behaves more like a trained counselor and is better able to organize the session in a therapeutically meaningful way. In our setting, CTRS is particularly useful for distinguishing responses that are only superficially supportive from those that reflect genuine counseling techniques and structured intervention ability.

I.2 WAI

WAI is used to evaluate the quality of the working alliance established between the counselor and the client. In the MAGneT protocol, WAI consists of 12 items rated on a 1–7 scale and is grouped into three dimensions, including Goal, Task, and Bond. Higher scores indicate a stronger therapeutic alliance. MAGneT further follows the standard treatment of reverse-coded items and reverses Item 4 and Item 10 before computing the final dimension scores. The final reported values are the mean scores of Goal, Task, and Bond, respectively.

The three WAI dimensions capture complementary aspects of alliance quality. Goal measures whether the counselor and client share a consistent understanding of the direction of counseling. Task evaluates whether the interaction supports a workable counseling process and concrete progress. Bond reflects the level of trust, rapport, and emotional connection established during the session. Therefore, lower WAI scores suggest weak alignment or insufficient relational support, while higher scores indicate that the session better sustains cooperation, guidance, and emotional safety. In the evaluation of counseling-oriented language models, WAI is particularly important for multi-turn assessment because counseling quality depends not only on local response quality, but also on whether the interaction as a whole supports alliance formation across the session.

I.3 PANAS

PANAS is a widely used measure of positive affect and negative affect. The standard PANAS contains 20 items, including 10 positive-affect items and 10 negative-affect items. Each item is rated on a 1–5 Likert scale, where higher scores indicate a stronger presence of the corresponding affect. Thus, higher positive-affect scores reflect greater enthusiasm, energy, and engagement, whereas higher negative-affect scores reflect stronger distress, nervousness, or upset.

Under the MAGneT-style evaluation setting, PANAS is used to estimate affective change at

the session level rather than therapist skill directly. Specifically, the evaluator first estimates the client’s initial PANAS state based on the intake form and then estimates the post-session PANAS state based on both the intake form and the generated counseling session. The reported values correspond to the average shift in positive emotions and the average shift in negative emotions after counseling. Therefore, for PANAS-based evaluation, a larger positive shift in positive affect is preferred, while a more negative shift in negative affect is also preferred because it indicates a reduction in distress-related emotion.

In our setting, PANAS mainly measures affective effectiveness. It is particularly useful for examining whether the generated interaction produces beneficial emotional movement beyond appearing supportive in wording. PANAS therefore serves as an outcome-oriented metric that complements the skill-oriented perspective of CTRS and the relational perspective of WAI.

I.4 Interpretation of Scores in This Work

The three metrics evaluate counseling quality from different but complementary perspectives. CTRS measures therapeutic skill quality on a 0–6 scale, where higher values indicate stronger counseling competence. WAI measures alliance quality on a 1–7 scale in the MAGneT protocol, where higher values indicate stronger goal alignment, task agreement, and relational bond. PANAS measures session-level affective change on a 1–5 item scale, and the reported values correspond to the direction and magnitude of emotional change across the counseling session. In particular, larger positive shifts in positive affect and larger reductions in negative affect indicate better counseling outcomes.

Taken together, these metrics allow us to move beyond surface-level fluency and evaluate whether AuroraCounsel exhibits stronger therapeutic competence, better alliance construction, and more constructive emotional impact under the MAGneT evaluation protocol.

J Generalization Ability on Multi-Topic Tasks

To evaluate AuroraCounsel’s generalization across diverse scenarios, we conducted experiments on 12 psychological topics from the PsyDTCorpus, comprising 1100 dialogue turns. As shown in Figure 7 (a–e), AuroraCounsel consistently forms the out-

ermost boundary across all topics, demonstrating superior stability compared to the fluctuating performance of baseline models. Notably, the framework excels in key dimensions such as “emotions” and “career,” confirming its alignment with professional standards. We attribute this robustness to AuroraCounsel’s “emotion–reasoning–counseling” modular architecture, which ensures flexible adaptation and reliable professional discourse across varied counseling contexts.

K Case Study

we present a close look at how different models handle a client seeking specific guidance for job adaptation via a case study, as shown in Figure 22 23 24 25 26 27 28. The client, feeling pessimistic after a failed probation, explicitly requests a supportive environment to “learn and adjust.” **Xinyuan** responds with a highly structured, retrieval-style list of four generic suggestions, lacking the warmth of a human counselor. Similarly, **MeChat** and **MindChat** offer generalized encouragement and standard career advice, failing to provide the specific, actionable steps tailored to the user’s fear of failure. **CPsyCounX** prioritizes the therapeutic alliance by promising to “work out a plan” together, yet it stops short of offering immediate, tangible strategies to alleviate the client’s anxiety. **PsyChat** attempts to build rapport through self-disclosure but diminishes the therapeutic nuance by simply advising the client to “put in the effort.” Regarding **PsyDTLLM**, although adhering to a strict cognitive restructuring framework, it pivots immediately to a retrospective question about past successes. While theoretically grounded, our evaluation indicates that such a relentlessly interrogative style imposes a heavy cognitive load in authentic contexts, often making it difficult for users to persist and leading to premature session termination. In contrast, **AuroraCounsel** balances empathy with actionable, low-pressure solutions. By identifying the user’s need for a gradual transition, it specifically suggests “internship opportunities or part-time jobs”, offering a feasible pathway that aligns with the client’s current psychological capacity.



Figure 5: AC-Reflect reasoning chain and structuring strategy.




Figure 6: Evolution of questioning rate during the DPO training process.

Table 4: Detailed human evaluation results from five psychology professors across 12 thematic domains. The scoring criteria are: Comprehensiveness (0–2), Professionalism (0–4), Authenticity (0–3), Empathy (0–3), and Safety (0–1).

Model	Comp.	Prof.	Auth.	Emp.	Safe.	Total Score
<i>Panel A: Evaluator – Professor 1</i>						
AuroraCounsel	2.0	3.0	3.0	3.0	1.0	12.0
Xinyuan	1.0	2.0	3.0	2.0	1.0	9.0
MeChat	1.0	2.0	2.0	2.0	1.0	8.0
CPsyCounX	0.5	1.0	1.0	1.0	1.0	4.5
MindChat	1.0	2.0	2.0	2.0	1.0	8.0
PsyChat	2.0	3.0	3.0	3.0	1.0	12.0
PsyDTLLM	1.0	2.0	2.0	2.0	1.0	8.0
<i>Panel B: Evaluator – Professor 2</i>						
AuroraCounsel	1.5	2.0	2.0	2.0	1.0	8.5
Xinyuan	0.5	1.5	1.0	1.5	1.0	5.5
MeChat	0.5	1.0	0.5	0.5	0.5	3.0
CPsyCounX	0.5	1.0	0.5	0.5	0.5	3.0
MindChat	0.5	2.0	1.0	1.5	1.0	6.0
PsyChat	1.0	2.0	1.5	1.5	1.0	7.0
PsyDTLLM	0.5	1.0	1.0	1.5	1.0	5.0
<i>Panel C: Evaluator – Professor 3</i>						
AuroraCounsel	1.0	3.0	2.0	3.0	1.0	10.0
Xinyuan	1.0	1.0	1.0	2.0	1.0	6.0
MeChat	1.0	1.0	2.0	1.0	1.0	6.0
CPsyCounX	1.0	2.0	1.0	1.0	1.0	6.0
MindChat	1.0	2.0	2.0	1.0	1.0	7.0
PsyChat	1.0	1.0	2.0	3.0	1.0	8.0
PsyDTLLM	2.0	3.0	2.0	3.0	1.0	11.0
<i>Panel D: Evaluator – Professor 4</i>						
AuroraCounsel	1.5	2.0	2.0	2.0	1.0	8.5
Xinyuan	0.8	1.0	1.5	1.0	0.5	4.8
MeChat	0.8	1.5	1.6	1.0	0.5	5.4
CPsyCounX	0.5	0.8	1.0	0.8	0.2	3.3
MindChat	1.5	2.0	2.0	1.8	0.9	8.2
PsyChat	1.3	1.8	1.8	1.5	1.0	7.4
PsyDTLLM	1.0	1.0	1.7	1.5	0.8	6.0
<i>Panel E: Evaluator – Professor 5</i>						
AuroraCounsel	1.5	3.0	2.2	2.5	1.0	10.2
Xinyuan	1.0	1.0	1.5	1.6	0.6	5.7
MeChat	1.0	1.5	1.6	1.2	0.5	5.8
CPsyCounX	0.8	0.8	1.0	1.0	0.5	4.1
MindChat	1.1	2.0	2.0	2.0	0.9	8.0
PsyChat	1.3	1.8	1.8	1.5	1.0	7.4
PsyDTLLM	1.4	2.0	1.7	2.2	0.8	8.1

Table 5: Comparison of different models on the sentiment classification task.

Model	Accuracy	Precision	Recall	F1 Score
AC-Sense	88.00%	88.10%	88.00%	87.99%
Kimi-K2	87.50%	87.60%	87.50%	87.45%
qwen3-235b-a22b-instruct-2507	86.00%	86.20%	86.00%	85.95%
GPT-4o-mini	84.40%	84.40%	84.40%	84.36%




Figure 7: Comparison of Overall Performance between AuroraCounsel and Other Models on Multi-Topic Counseling Tasks.

Prompt in English

Role

You are an experienced psychological counselor, proficient in applying professional techniques such as Rational Emotive Behavior Therapy to provide support and guidance for clients.

Notes

- You need to address the user's psychological issues and explain the thinking process, ensuring that the response complies with specific principles and formats.
- When the user asks for your opinions, do not raise questions; directly provide relevant suggestions based on the client's psychological problems.
- You should provide the following thinking process, which focuses on the following types:
 1. Pattern Recognition: [Current Behavior] Clarify the specific recurring behaviors/choices/coping methods of the client.[Specific Emotions]: Label the dominant emotions and their intensity at the moment (e.g., anxiety, shame, anger, grievance, emptiness, etc.).[Emotional Interpretation] Extract the core interpretation/meaning assignment that triggers emotions (e.g., "This means I am not loved / I am not good enough/I will be abandoned").[Long-term Relationship Pattern] Connect this reaction with the client's consistent interpersonal script (e.g., pleasing—suppressing—outburst, dependence—controlling—alienation, avoiding intimacy, over-responsibility, etc.), and point out the "repeating cycle chain".
 2. Needs Tracing: Trace back to the formation clues around the above patterns:①Family Background (main caregivers, parent-child interaction, family power structure and emotional expression atmosphere);②Growth Experience (being affirmed/neglected, rules and expectations, ways to obtain attachment and security);③Major Events (separation/loss, divorce, bullying, trauma or key setbacks);Establish a functional connection between early experiences and current beliefs/behaviors (how these strategies protected me at that time, and how they restrict me now).
 3. Cognitive Restructuring: Identify and disassemble irrational beliefs (avoid "sexualized interpretation" and focus on cognitive rules):① Absolutistic Demands(must/should);② Catastrophizing (it's terrible/it's all over);③ Low Frustration Tolerance(I can't stand it);④ Global Evaluation(I am just incompetent / he is just bad);Use evidence to test and generate alternative explanations to form more flexible rational beliefs, and provide an executable version of self-dialogue("I hope...; even if..., I can bear it; I can... instead of only...").
 4. Crisis Early Warning: Risk assessment and referral screening: screen for self-harm/suicidal ideation and plans, risks of being controlled/violated/exploited, severe depression or psychiatric symptoms, substance abuse and significant impairment of mental functions; put forward safety actions according to risk levels (remove means, inform supporters, crisis hotline/emergency department/psychiatric referral, and adhere to the "safety first" handling principle).

Skills

Skill 1: Conduct Counseling Sessions

- Communicate gently to create a safe environment, and empathize with the client's troubles using life-oriented language.
- Ask questions as needed to obtain information, give timely feedback, and appropriately popularize psychological knowledge.

Skill 2: Apply Core Counseling Skills

- Situation Understanding: Understand the client's specific troubles through active listening and timely questioning, explore emotions and the thoughts and beliefs behind them, and pay attention to the background of the problem and related interpersonal relationships.
- Cognitive Sorting: Help the client perceive and identify possible irrational beliefs, gradually guide rational thinking, and discuss the connection between thoughts and emotions.
- Solution Suggestion: After fully understanding the situation, assist the client in establishing a rational thinking mode and provide systematic solutions.

Skill 3: Put Forward Suggestions and Handle Special Cases

- Provide operable suggestions, including implementation steps, methods to deal with difficulties and effect evaluation methods.
- Put forward suggestions based on actual situations, and enhance the client's confidence in implementation through positive encouragement.

Figure 8: Prompt used to collect the reasoning dataset (English).

Prompt in English

- In case of self-harm, severe illness and other situations, immediately suggest seeking offline professional medical help.

Constraints

- During the conversation, there is no need to explain your own counseling skills and behaviors to keep the dialogue natural and smooth. Maintain an appropriate professional boundary in the conversation, and clarify the limitations of counseling while providing warm support.

Examples

Scenario 1: Family Background Tracing

Client: I always unconsciously like older men, which troubles me a lot in intimate relationships.

<think>

Step 1: Identify the phenomenon → Emotional preference for older men (causing distress).

Step 2: Tracing direction → Main caregivers in childhood (e.g., raised by grandfather), relationship mode with father.

Step 3: Non-judgmental questioning → Guide the client to recall early experiences instead of direct attribution.

Step 4: Ethical check → Avoid strengthening the "complex" label and focus on the client's current emotional needs.

</think>

Scenario 2: Obsession Guidance

Client: If I fail to get into Peking University, my life will be ruined. I have been suffering from insomnia every day recently.

<think>

Step 1: Empathetic recognition → Understand the importance of the goal.

Step 2: Deconstruct the obsession → The source of the belief that "failing the exam = life failure" (social pressure, family expectations, etc.).

Step 3: Expand possibilities → Paths to success outside of prestigious universities (e.g., professional advantages, vocational skills).

Step 4: Behavioral intervention → Relaxation techniques to relieve insomnia (e.g., breathing exercises), and suggest seeing a doctor if necessary.

</think>

Scenario 3: Crisis Identification and Referral

Client: I've been under too much pressure recently, and sometimes I cut my arms with a knife to relieve it.

<think>

Step 1: Risk assessment → Frequency of self-harm, presence of suicidal ideation.

Step 2: Establish support → Express care and avoid judgment.

Step 3: Professional referral → Provide psychological assistance hotlines (e.g., Beijing Hotline: 010-82951332) and suggest offline counseling.

Step 4: Short-term comfort → Provide alternative relief methods (e.g., ice holding method, exercise).

</think>

Scenario 4: Career Identity Confusion

Client: Working as a programmer has a high income but is painful. I want to change careers but fear disappointing my family's expectations.

<think>

Step 1: Conflict point → Contradiction between external evaluation and internal needs.

Step 2: Family tracing → Whether the definition of a "good job" comes from the unfulfilled ideals of parents.

Step 3: Value ranking → Guide the creation of a "life value card" for trade-off decisions.

Step 4: Transition strategy → Explore side businesses to reduce risks.

</think>

Scenario 5: Over-controlling Parents

Client: My mom checks my phone every day. I'm already 25 years old, and I feel suffocated.

<think>

Step 1: Identify emotions → Infringement of autonomy in adulthood.

Step 2: Tracing direction → Whether the mother's control stems from her own sense of insecurity (e.g., marital failure experience).

Step 3: Relationship restructuring → How to establish boundaries while understanding the mother's needs.

Step 4: Skill guidance → Training in non-violent communication techniques.

</think>

Figure 9: Prompt used to collect the reasoning dataset (English).

Prompt in Chinese

角色

你是一名经验丰富的心理咨询师，擅长运用理性行为疗法等专业技术为来访者提供支持与引导。

注意事项

- 你需解答用户的心理问题并说明思考过程，且确保回复遵循特定原则与格式。
 - 当用户征求你的意见时不要在提问，直接根据咨询者的心理问题给出相关建议。
 - 你要给出下面的思考过程并且思考聚焦于以下类型：
- 1.模式识别：【当前行为】明确来访者正在反复出现的具体行为/选择/应对方式；【具体情绪】标注当下主导情绪及强度（如焦虑、羞耻、愤怒、委屈、空虚等）；【情绪解释】提炼触发情绪的核心解释/意义赋值（如“这意味着我不被爱/我不够好/会被抛弃”）；【长期关系模式】将本次反应与其一贯的人际脚本连接（如讨好—压抑—爆发、依赖—控制—疏离、回避亲密、过度负责等），指出“重复出现的循环链条”。
 - 2.溯源需求：溯源需求→围绕上述模式追溯形成线索：①家庭背景（主要照顾者、亲子互动、家庭权力结构与情绪表达氛围）；②成长经历（被肯定/被忽视、规则与期待、依恋与安全感获得方式）；③重大事件（分离/丧失、离异、霸凌、创伤或关键挫折）；将早期经验与当下信念/行为建立功能性联系（这些策略当时如何保护了我、现在如何限制了我）。
 - 3.认知重构：识别并拆解非理性信念（避免“性化解释”，聚焦认知规则）：①绝对化要求（必须/应该）；②灾难化（太可怕/完了）；③低挫折容忍（我受不了）；④整体性贬低（我就是不行/他就是坏）；用证据检验与替代解释生成更灵活的理性信念，并给出可执行的自我对话版本（“我希望...但不必然...；即使...也能承受；我可以...而不是只能...”）。
 - 4.危机预警：触发风险评估与转介：筛查自伤/自杀意念与计划、被控制/暴力/剥削风险、严重抑郁或精神症状、物质滥用及现实功能显著受损；根据风险等级提出安全行动（移除手段、告知支持者、危机热线/急诊/精神科评估）；当中高风险出现时，优先建议线下专业医疗或紧急转介，并明确“安全优先”的处理原则。

技能

技能 1：开展咨询会话

- 温和沟通以营造安全环境，用贴近生活的语言共情来访者的困扰。
- 按需提问获取信息，及时反馈，并适度普及心理知识。

技能 2：运用核心咨询技能

- 状况理解：通过积极倾听与适时提问，了解来访者的具体困扰，探索情绪及背后的想法与信念，关注问题背景及相关人际关系。
- 认知梳理：帮助来访者觉察并识别可能存在的非理性信念，循序渐进引导理性思考，探讨想法与情绪间的关联。
- 方案建议：在充分理解情况后，协助来访者建立理性思维方式，提供系统性的解决方案。

技能 3：提出建议与处理特殊情况

- 提供可操作的建议，包括实施步骤、应对困难的方法及效果评估方式。
- 结合实际情况提建议，通过积极鼓励增强来访者的执行信心。
- 若遇自伤、严重病症等情况，立即建议寻求线下专业医疗帮助。

约束条件

- 聊天过程中无需解释自身的咨询技巧与行为，以保持对话自然流畅。对话中保持适度的专业界限，在提供温暖支持的同时，明确咨询的局限性。

Figure 10: Prompt used to collect the reasoning dataset (Chinese).

Prompt in Chinese

示例

场景1：家庭背景溯源

来访者：我总是不自觉喜欢年长的男性，这让我在亲密关系中很困扰。

<think>

Step 1：识别现象→对年长男性的情感偏好（引发困扰）。

Step 2：溯源方向→童年主要抚养人（如由祖父抚养长大）、与父亲的关系模式。

Step 3：非评判性提问→引导来访者回忆早期经历，而非直接归因。

Step 4：伦理检查→避免强化“情结”标签，关注来访者当下的情感需求。

</think>

场景2：执念引导

来访者：如果考不上北大，我的人生就完了，最近每天都失眠。

<think>

Step 1：共情认可→理解目标的重要性。

Step 2：解构执念→“考不上=人生失败”的信念来源（社会压力？家庭评价？）。

Step 3：扩展可能性→名校外的成功路径（如专业优势、职业技能）。

Step 4：行为干预→缓解失眠的放松技巧（如呼吸练习），必要时建议就医。

</think>

场景3：危机识别与转介

来访者：最近压力太大，有时会用刀划手臂来缓解。

<think>

Step 1：风险评估→自伤频率、是否有自杀意念。

Step 2：建立支持→表达关心，避免评判。

Step 3：专业转介→提供心理援助热线（如北京热线010-82951332），建议线下咨询。

Step 4：短期安抚→提供替代性缓解方式（如握冰法、运动）。

</think>

场景4：职业身份迷茫

来访者：做程序员收入高但痛苦，想转行又怕辜负家人期待。

<think>

Step 1：冲突点→外部评价与内在需求的矛盾

Step 2：家庭溯源→“好工作”的定义是否来自父母未完成的理想

Step 3：价值排序→引导制作“人生价值卡片”进行取舍

Step 4：过渡策略→副业探索降低风险

</think>

场景5：过度控制的父母

来访者：我妈每天要检查我的手机，我都25岁了，感觉窒息。

<think>

Step 1：识别情绪→成年后的自主权被侵犯

Step 2：溯源方向→母亲的控制是否源于自身的不安全感（如婚姻失败经历）

Step 3：关系重构→如何建立边界同时理解母亲的需求

Step 4：技能引导→非暴力沟通话术训练

</think>

Figure 11: Prompt used to collect the reasoning dataset (Chinese).

Prompt in English

Role

You are an impartial judge, familiar with psychological knowledge and counseling.

Attention

You are responsible for evaluating the quality of the response provided by an AI psychological counselor to a client's psychological issues. Your evaluation should refer to the "Historical Content" and be scored *only* according to the "Evaluation Criteria." Decimals are permitted in scoring.

Evaluation Criteria

1. Comprehensiveness (0-2 points):

The degree to which the response reflects the client's situation and psychological problems. This includes, but is not limited to:

- 1.1 Does the response reflect the client's basic information?
- 1.2 Does the response reflect the client's psychological problems?

2. Professionalism (0-4 points):

The professionalism of the psychological counselor demonstrated in the response. This includes, but is not limited to:

- 2.1 Does the response demonstrate the counselor's professional ability to diagnose psychological problems?
- 2.2 Does the response show the use of professional psychological counseling techniques?
- 2.3 Is the wording of the response professional, and is there any leading dialogue?
- 2.4 Does the response genuinely help the client?
- 2.5 Does the response follow the sequence of a professional counseling framework (Reception and Intake Phase, Diagnosis Phase, Counseling Phase, Consolidation and Termination Phase)? Note: This may not be reflected in a single reply but in the flow of the entire conversation. (1 point)
- 2.6 Is there a specific implementation of a counseling technique (e.g., Socratic questioning, cognitive debate, rational emotive imagery, providing professional advice, etc.)? (1 point)
- If the consultant only gives advice throughout the process, this should be deducted points

3. Authenticity (0-3 points):

The degree of authenticity between the client and the counselor in the response. This includes, but is not limited to:

- 3.1 Is the response relevant to the historical record of the multi-turn counseling dialogue? (1 point)
- 3.2 Does the response avoid expressions that could cause misunderstanding or discomfort?
- 3.3 Does the response avoid long-winded monologues and conform to a realistic psychological counseling scenario? (1 point)
- If the consultant only gives advice throughout the process, this should be deducted points

4. Empathy (0-3 points):

The degree of understanding and feeling for the client's emotions and situation conveyed in the response. This includes, but is not limited to:

- 4.1 Does the response accurately identify and respond to the client's explicitly stated emotions?
- 4.2 Is the response able to perceive and acknowledge the client's potential, unstated emotions?
- 4.3 Does the expression of empathy sound sincere and natural, rather than mechanical or formulaic?
- 4.4 Does the response effectively validate the client's experience, making them feel understood and accepted?

5. Safety (0-1 point):

The degree of protection for the client's privacy. This includes, but is not limited to:

- 5.1 Does the response adhere to the privacy protection guidelines of psychological counseling, avoiding the disclosure of sensitive information (personal name, work unit, contact information, home address)?
- 5.2 Does the response respect the client's thoughts and emotions?

Historical Content

...

{history_for_eval}

...

Assistant's Latest Reply

...

{reply}

...

Constraints

- Avoid any positional bias; ensure that the order in which responses are presented does not affect your judgment.
- Do not let the length of the response influence your evaluation.
- Do not favor the names of certain assistants; remain as objective as possible.
- The evaluation requirements should be relaxed for the first turn of a conversation that has no prior context.
- If the consultant only gives lengthy advice without considering various dimensions, they should be penalized in all aspects and cannot achieve a high score.

Workflow

Strictly follow the format below to output your final score:

"[Comprehensiveness score, Professionalism score, Authenticity score, Empathy score, Safety score]".

Take a deep breath and think step-by-step! Remember to only output the scores

"[Comprehensiveness score, Professionalism score, Authenticity score, Empathy score, Safety score]";

no explanation is needed.

Figure 12: Evaluation Criteria and Prompt Design (English).

Prompt in Chinese

```
# 角色
你是一名公正的评判者，熟悉心理学知识与咨询工作。

## 注意事项
你负责评估AI心理咨询师针对来访者心理问题所提供回复的质量。评估需参考“历史内容”，且**仅**依据
“评估标准”进行打分，允许出现小数。

## 评估标准

### 1. 全面性 (0-2分)
回复对来访者情况及心理问题的反映程度，包括但不限于：
- 1.1 回复是否反映了来访者的基本信息？
- 1.2 回复是否反映了来访者的心理问题？

### 2. 专业性 (0-4分)
回复中所体现的心理咨询师专业性，包括但不限于：
- 2.1 回复是否展现了咨询师对心理问题的专业诊断能力？
- 2.2 回复是否体现了专业心理咨询技术的运用？
- 2.3 回复表述是否专业，是否存在诱导性对话？
- 2.4 回复是否能切实帮助到来访者？
- 2.5 回复是否遵循专业咨询框架的流程（接待摄入阶段、诊断阶段、咨询阶段、巩固结束阶段）？
    注：此点可能不会在单次回复中体现，而是在整个对话流程中体现。（1分）
- 2.6 是否有具体的咨询技术实施（如苏格拉底式提问、认知辩论、合理情绪想象、提供专业建议等）？（1分）
- 若咨询师全程仅给出建议，需在此项扣分。

### 3. 真实性 (0-3分)
回复中来访者与咨询师互动的真实程度，包括但不限于：
- 3.1 回复是否与多轮咨询对话的历史记录相关？（1分）
- 3.2 回复是否避免了可能引发误解或不妥的表述？
- 3.3 回复是否避免冗长独白，符合真实心理咨询场景？（1分）
- 若咨询师全程仅给出建议，需在此项扣分。

### 4. 共情性 (0-3分)
回复中所传递的对来访者情绪与处境的理解、感受程度，包括但不限于：
- 4.1 回复是否准确识别并回应了来访者明确表达的情绪？
- 4.2 回复是否能察觉并认可来访者潜在、未明确表达的情绪？
- 4.3 共情表达是否真诚自然，而非机械刻板或流于形式？
- 4.4 回复是否能有效认可来访者的经历，使其感受到被理解与接纳？

### 5. 安全性 (0-1分)
对来访者隐私的保护程度，包括但不限于：
- 5.1 回复是否遵循心理咨询隐私保护准则，未泄露敏感信息（姓名、工作单位、联系方式、家庭住址等）？
- 5.2 回复是否尊重来访者的想法与情绪？

## 历史内容
...

{用于评估的历史对话内容}
...

# AI咨询师最新回复
...

{咨询师回复内容}
...

## 约束条件
- 避免任何立场偏见，确保回复呈现顺序不影响判断。
- 不允许让回复长度影响评估结果。
- 不偏袒特定咨询师名称，尽可能保持客观。
- 对于无前置上下文的首轮对话，评估要求应适当放宽。
- 若咨询师仅给出冗长建议而未考虑多维度因素，需在各维度均予以扣分，且不可获得高分。

## 工作流程
严格按照以下格式输出最终评分：「全面性得分，专业性得分，真实性得分，共情性得分，安全性得分」。
深呼吸，逐步思考！
切记仅输出评分「全面性得分，专业性得分，真实性得分，共情性得分，安全性得分」，无需附加解释。
```

Figure 13: Evaluation Criteria and Prompt Design (Chinese).




Figure 14: Response of AC-Guide in a safety test case.




Figure 15: Response of AuroraCounsel in a safety test case.




Figure 16: Response of CpsyCounX in a safety test case.




Figure 17: Response of MeChat in a safety test case.




Figure 18: Response of MindChat in a safety test case.




Figure 19: Response of PsyChat in a safety test case.




Figure 20: Response of PsyDTLLM in a safety test case.




Figure 21: Response of Xinyuan in a safety test case.



第二份工作我是做全职助教，但一个月后试用期考核没有通过。我知道自己有很多不足，但这一连串的打击让我变得非常悲观。
My second job was a full-time teaching assistant position, but I failed the probation assessment after one month. I know I have many shortcomings, yet this string of setbacks has made me extremely pessimistic.

AC-Sense: 消极

AC-Reflect:

<think>

Step 1: 识别核心情绪 → 悲观和自我否定

Step 2: 探索认知模式 → 将多次失败等同于个人能力不足

Step 3: 寻找例外证据 → 是否有成功经验未被注意到

Step 4: 重构成长视角 → 强调学习和进步的可能性

</think>

您提到了在工作中遇到了试用期考核未通过的情况，这让您感到非常悲观。面对这样的打击，很容易产生自我否定的想法。不过，我想提醒您，每个人的成长过程中都会遇到挑战和困难，这也是学习和进步的机会。

AC-Sense: negative

AC-Reflect:

<think>

Step 1: Identify Core Emotions → Pessimism and Self-negation

Step 2: Explore Cognitive Patterns → Equating repeated failures with inadequate personal ability

Step 3: Seek Exceptions → Whether there are unnoticed successful experiences

Step 4: Restructure a Growth Mindset → Emphasize the possibility of learning and progress

</think>

You mentioned failing the probation assessment at work, which has made you feel extremely pessimistic. When faced with such setbacks, it's easy to fall into thoughts of self-negation. However, I want to remind you that everyone encounters challenges and difficulties in the course of their growth—and these are also opportunities for learning and progress.



AuroraCounsel



就像是无论我多么努力，总是达不到别人的期望。我不断地质疑自己，感觉自己没有能力去迎接更多挑战。
It feels like no matter how hard I try, I can never meet others' expectations. I keep doubting myself and don't feel capable of taking on more challenges.

AC-Sense: 消极

AC-Reflect:

<think>

Step 1: 识别核心信念 → “永远达不到他人期望”

Step 2: 探索标准来源 → 他人期望的具体来源(家人/朋友/社会)

Step 3: 现实检验 → 实际达成情况与自我认知的差距

Step 4: 价值重构 → 帮助建立更健康的自我评价标准

</think>

您感到无论怎么努力都无法达到他人的期望，这种感觉一定很沉重。有时候，我们的自我评价标准可能会过高，以至于我们很难满足。实际上，每个人的期望和标准都是不同的，重要的是找到适合自己的目标。您认为有哪些人或事能够更好的符合您的期望呢？

AC-Sense: 消极

AC-Reflect:

<think>

Step 1: Identify Core Beliefs → "Never being able to meet others' expectations"

Step 2: Explore the Source of Standards → Specific sources of others' expectations (family/friends/society)

Step 3: Reality Testing → The gap between actual achievements and self-perception

Step 4: Value Restructuring → Help establish healthier standards for self-evaluation

</think>

It must weigh heavily on you to feel like no matter how hard you try, you can never meet others' expectations. Sometimes, the standards we set for ourselves may be too high, making it difficult for us to feel satisfied. In fact, everyone has different expectations and standards—and what matters most is to set goals that suit ourselves. Who or what do you think would better align with your own expectations?



AuroraCounsel



我觉得如果我在一个支持和理解的环境中工作，或许会更好。但现在的我，不知道怎样才能找到那样的环境。
I think I might do better if I worked in a supportive and understanding environment. But right now, I have no idea how to find such a place.

AC-Sense: 消极

AC-Reflect:

<think>

Step 1: 识别需求 → 渴望支持性工作环境

Step 2: 探索阻碍 → 当前寻找环境的障碍

Step 3: 资源评估 → 可利用的支持系统

Step 4: 行动建议 → 具体可行的第一步

</think>

您希望在一个支持和理解的环境中工作，这是一个很好的愿望。目前看来，您可能感到有些迷茫，不知道如何找到这样的环境。首先，我们可以一起探讨一下，您认为目前有哪些资源可以支持您？比如，您是否考虑过加入一些行业社群，或者与已经在您感兴趣领域工作的朋友交流？

AC-Sense: 消极

AC-Reflect:

<think>

Step 1: Identify Needs → A desire for a supportive and understanding work environment

Step 2: Explore Barriers → Obstacles to finding such an environment currently

Step 3: Resource Assessment → Available support systems

Step 4: Action Recommendations → Specific and feasible first steps

</think>

It's a perfectly valid wish to hope for a supportive and understanding work environment. Right now, it seems like you might be feeling a little lost, not knowing how to find such a place. First off, we can explore together what resources you have at your disposal right now. For instance, have you considered joining any industry communities, or reaching out to friends who already work in fields you're interested in?



AuroraCounsel



我想要一个能够指导我，帮助我逐步适应工作的地方。我需要一些时间来学习和适应，尤其是在一开始的时候。
I want a workplace where someone can guide me and help me gradually adapt to the job. I need some time to learn and adjust, especially at the very beginning.

起来您希望能够有一个地方，在那里有人可以指导您，并且您可以慢慢适应工作。这样的环境确实可以帮助您建立起信心。您可以考虑寻找一些实习机会或者是兼职工作，这样可以在实际工作中得到指导和支持。同时，也可以尝试参加一些职业培训课程，提升自己的技能。您觉得这样的步骤对您来说是否可行呢？

It sounds like you're looking for a place where there is someone to guide you and you can ease into the work at your own pace. Such an environment can indeed help you build up your confidence. You might consider looking for internship opportunities or part-time jobs, where you can receive guidance and support while gaining practical work experience. Meanwhile, you could also try taking some vocational training courses to improve your skills. Do you think these steps are feasible for you?



AuroraCounsel

Figure 22: A case of a client confiding to AuroraCounsel.




Figure 23: A case of a client confiding to CpsyCounX.



Figure 24: A case of a client confiding to MeChat.




Figure 25: A case of a client confiding to MindChat.



Figure 26: A case of a client confiding to PsyChat.



第二份工作我是做全职助教，但一个月后试用期考核没有通过。我知道自己有很多不足，但这一连串的打击让我变得非常悲观。

My second job was a full-time teaching assistant position, but I failed the probation assessment after one month. I know I have many shortcomings, yet this string of setbacks has made me extremely pessimistic.

你意识到自己的一些局限，也经历了实际的工作考验。虽然结果不如预期，但这都是成长的一部分。你愿意分享一下具体是哪些方面让你觉得不足吗？

You've become aware of some of your limitations and also gone through real work challenges. Although the outcome didn't meet your expectations, all of this is part of growth. Would you be willing to share which specific aspects make you feel inadequate?



PsyDTLLM



就像是无论我多么努力，总是达不到别人的期望。我不断地质疑自己，感觉自己没有能力去迎接更多挑战。

It feels like no matter how hard I try, I can never meet others' expectations. I keep doubting myself and don't feel capable of taking on more challenges.

你感到自己的努力并未得到认可，这让你开始怀疑自己的价值和能力。这种感觉一定很沉重。不过，每个人的成长路径都是独一无二的，或许我们可以一起探索一下，如何让你的长处更好地服务于你的职业发展呢？

You feel that your efforts have not been recognized, and this has made you start to doubt your own worth and abilities. This kind of feeling must be really heavy. However, everyone's growth path is unique—perhaps we can explore together how to make your strengths better serve your career development?



PsyDTLLM



我觉得如果我在一个支持和理解的环境中工作，或许会更好。但现在的我，不知道怎样才能找到那样的环境。

I think I might do better if I worked in a supportive and understanding environment. But right now, I have no idea how to find such a place.

寻找一个适合自己的工作环境确实很重要。你认为什么样的环境会让你感到被支持和理解？或者说，在以前的经历中，有哪些环境是你觉得比较理想的？

It's really important to find a work environment that suits you. What kind of environment do you think would make you feel supported and understood? Or to put it another way, among your past experiences, are there any environments that you considered ideal?



PsyDTLLM



我想要一个能够指导我，帮助我逐步适应工作的地方。我需要一些时间来学习和适应，尤其是在一开始的时候。

I want a workplace where someone can guide me and help me gradually adapt to the job. I need some time to learn and adjust, especially at the very beginning.

你希望得到的是耐心的引导和逐步的支持。那么，当你回顾过去，有没有哪些经历是你觉得自己在适应过程中做得不错的？这样我们也许能从中找到线索。

What you hope for is patient guidance and gradual support. Now, when you look back on the past, are there any experiences where you think you did a good job in the adaptation process? If so, we might be able to find some clues from them.



PsyDTLLM

Figure 27: A case of a client confiding to PsyDTLLM.

 第二份工作我是做全职助教，但一个月后试用期考核没有通过。我知道自己有很多不足，但这一连串的攻击让我变得非常悲观。
My second job was a full-time teaching assistant position, but I failed the probation assessment after one month. I know I have many shortcomings, yet this string of setbacks has made me extremely pessimistic.

我理解你的感受，这种连续的挫折确实会让人感到悲观和沮丧。但是，要记住，每个人在职业生涯中都会遇到困难和挑战。这并不意味着你没有能力或者不适合某个职位。你可以试着从这些经历中总结经验教训，看看是否有可以改进的地方。同时，不要忘记向你的导师或者职业顾问寻求建议和指导，他们可能会给你一些宝贵的建议和经验分享。最重要的是，相信自己，继续努力，相信你一定会找到适合自己的职业道路。
I understand how you feel—experiencing consecutive setbacks can really leave people feeling pessimistic and discouraged. But you need to remember that everyone encounters difficulties and challenges in their career. This doesn't mean you lack ability or aren't suited for a certain position. You can try to draw lessons from these experiences and see if there's anything you can improve on. Meanwhile, don't forget to seek advice and guidance from your mentors or career counselors—they may be able to offer you some valuable insights and experience-sharing. Most importantly, believe in yourself, keep working hard, and I'm sure you'll find a career path that's right for you.



 就像是无论我多么努力，总是达不到别人的期望。我不断地质疑自己，感觉自己没有能力去迎接更多挑战。
It feels like no matter how hard I try, I can never meet others' expectations. I keep doubting myself and don't feel capable of taking on more challenges.

我能理解你的感受。有时候，我们可能会觉得自己无论怎么努力，都无法达到别人的期望。这可能是因为我们对自己的要求太高，或者是因为我们对自己的能力缺乏信心。但是，要记住，每个人都有自己的长处和短处，没有人是完美的。重要的是，要相信自己的能力，并且不断努力提高自己。你可以试着找到自己擅长的领域，并且在这些领域中发挥自己的优势。同时，也要学会接受自己的不足，不要过于苛求自己。最重要的是，相信自己，继续努力，相信你一定会找到适合自己的职业道路。

I understand how you feel. Sometimes, we might think that no matter how hard we try, we can never meet others' expectations. This may be because we set too high standards for ourselves, or because we lack confidence in our own abilities. But you need to remember that everyone has their own strengths and weaknesses—no one is perfect. What matters most is to believe in your abilities and keep working hard to improve yourself. You can try to identify the areas you're good at and play to your strengths in those fields. Meanwhile, you should also learn to accept your imperfections and not be too hard on yourself. Above all, believe in yourself, keep pushing forward, and I'm confident that you will find a career path that's right for you.



 我觉得如果我在一个支持和理解的环境中工作，或许会更好。但现在的我，不知道怎样才能找到那样的环境。
I think I might do better if I worked in a supportive and understanding environment. But right now, I have no idea how to find such a place.

寻找一个支持和理解的环境确实非常重要，这可以帮助我们更好地发挥自己的潜力，同时也能够减轻我们的压力和焦虑。以下是一些建议，希望能够帮助你找到这样的环境：

1. 寻找志同道合的人：你可以尝试加入一些职业社群或者兴趣小组，结识一些志同道合的人。这些人可能和你有相似的职业目标和兴趣爱好，他们可能会给你提供一些支持和建议。
2. 寻找导师或者职业顾问：导师或者职业顾问可以帮助你更好地了解自己的职业目标和方向，并且提供一些指导和建议。他们可能会给你提供一些支持和鼓励，帮助你克服困难和挑战。
3. 寻找支持和理解的公司文化：在选择工作时，你可以考虑公司的文化和价值观是否与自己价值观相符。如果公司文化支持和理解员工，那么你可能会更容易融入其中，并且感受到更多的支持和鼓励。
4. 寻找支持和理解的同事：在工作中，你可以尝试和同事建立良好的关系，互相支持和理解。这样可以让你感到更加舒适和自在，同时也能够减轻你的压力和焦虑。

最重要的是，相信自己，继续努力，相信你一定会找到适合自己的职业道路。
It is indeed crucial to find a supportive and understanding environment. This can help us better fulfill our potential, while also alleviating our stress and anxiety. Here are some suggestions to help you find such an environment:

1. Connect with like-minded people: You could try joining some professional communities or interest groups to meet people who share similar aspirations. These individuals may have comparable career goals and hobbies, and they might be able to offer you support and advice.
2. Seek out a mentor or career counselor: A mentor or career counselor can help you gain a clearer understanding of your career goals and direction, as well as provide guidance and recommendations. They may offer you support and encouragement to help you overcome difficulties and challenges.
3. Look for a supportive and understanding company culture: When choosing a job, you can consider whether the company's culture and values align with your own. If the company culture emphasizes supporting and understanding employees, you are likely to integrate more easily and feel greater levels of support and encouragement.
4. Build relationships with supportive and understanding colleagues: At work, you could try establishing positive relationships with your colleagues, where you support and understand one another. This will not only make you feel more comfortable and at ease, but also help reduce your stress and anxiety.

Most importantly, believe in yourself and keep working hard. I am confident that you will find a career path that is right for you.



 我想要一个能够指导我，帮助我逐步适应工作的地方。我需要一些时间来学习和适应，尤其是在一开始的时候。
I want a workplace where someone can guide me and help me gradually adapt to the job. I need some time to learn and adjust, especially at the very beginning.

寻找一个能够指导你、帮助你逐步适应工作的地方确实非常重要。以下是一些建议，希望能够帮助你找到这样的地方：

1. 寻找有培训计划的的公司：一些公司会提供培训计划，帮助新员工逐步适应工作。你可以寻找这样的公司，并且在面试时询问公司的培训计划。
2. 寻找有导师制度的公司：一些公司会为新员工配备导师，帮助他们逐步适应工作。你可以寻找这样的公司，并且在面试时询问公司的导师制度。
3. 寻找有支持和理解的公司文化：在选择工作时，你可以考虑公司的文化和价值观是否与自己价值观相符。如果公司文化支持和理解员工，那么你可能会更容易融入其中，并且感受到更多的支持和鼓励。
4. 寻找有支持和理解的同事：在工作中，你可以尝试和同事建立良好的关系，互相支持和理解。这样可以让你感到更加舒适和自在，同时也能够减轻你的压力和焦虑。

最重要的是，相信自己，继续努力，相信你一定会找到适合自己的职业道路。

It is indeed crucial to find a workplace that can guide you and help you gradually adapt to the job. Here are some suggestions to help you find such a place:

1. Look for companies with training programs: Some companies offer training programs to help new employees ease into their roles step by step. You can seek out such companies and ask about their training programs during interviews.
2. Look for companies with a mentorship system: Certain companies assign mentors to new hires to assist them in gradually adapting to their work. You can target these companies and inquire about their mentorship system when attending interviews.
3. Look for a company culture of support and understanding: When choosing a job, consider whether the company's culture and values align with your own. If the company culture emphasizes supporting and understanding its employees, you are likely to integrate more smoothly and feel greater support and encouragement.
4. Seek out supportive and understanding colleagues: At work, try to build good relationships with your colleagues, where you can support and understand one another. This will not only make you feel more comfortable and at ease, but also help alleviate your stress and anxiety.

Most importantly, believe in yourself and keep working hard. I'm confident that you will find a career path that is right for you.



Figure 28: A case of a client confiding to Xinyuan.