

# CBT-LLM: A Chinese Large Language Model for Cognitive Behavioral Therapy-based Mental Health Question Answering

Hongbin Na

Australian Artificial Intelligence Institute,  
University of Technology Sydney, Sydney, Australia  
hongbin.na@student.uts.edu.au

## Abstract

The recent advancements in artificial intelligence highlight the potential of language models in psychological health support. While models trained on data from mental health service platform have achieved preliminary success, challenges persist in areas such as data scarcity, quality, and ensuring a solid foundation in psychological techniques. To address these challenges, this study introduces a novel approach to enhance the precision and efficacy of psychological support through large language models. Specifically, we design a specific prompt derived from principles of Cognitive Behavioral Therapy (CBT) and have generated the CBT QA dataset, specifically for Chinese psychological health Q&A based on CBT structured intervention strategies. Unlike previous methods, our dataset emphasizes professional and structured response. Utilizing this dataset, we fine-tuned the large language model, giving birth to CBT-LLM, the large-scale language model specifically designed for Cognitive Behavioral Therapy techniques. Empirical evaluations demonstrate that CBT-LLM excels in generating structured, professional, and highly relevant responses in psychological health support tasks, showcasing its practicality and quality. The model is available on Hugging Face: <https://huggingface.co/Hongbin37/CBT-LLM>.

**Keywords:** Large Language Model, Question Answering, Cognitive Behavioral Therapy, Mental Health Support.

## 1. Introduction

**Important:** Our research explores the potential of large language models to answer questions based on Cognitive Behavioral Therapy, but does **NOT** recommend their use as a substitute for psychological treatment without professional supervision.

The advancement of pre-trained language models (PLMs) has profoundly impacted various domains, such as finance (Shah et al., 2022), biology (Madani et al., 2023), heralding new frontiers in applications and research. Among these, the intersection of artificial intelligence and mental health has emerged as a particularly promising area. Here, the deployment of PLMs holds the potential to revolutionize mental health support, a notion underscored by the development of several preliminary systems aimed at leveraging these models for psychological assistance (Liu et al., 2021; Cheng et al., 2023; Lai et al., 2023). The core strengths of PLMs, namely their deep learning architectures and sophisticated attention mechanisms, enable them to parse and interpret complex emotional and cognitive information, positioning them as invaluable tools in mental health contexts.

Despite these advancements, the application of PLMs in mental health support is fraught with challenges, particularly regarding data scarcity and quality. The PsyQA dataset (Sun et al., 2021) attempts to mitigate the dearth of Chinese mental health data by collating structured question-answer pairs from online services. Similarly, the SMILE strategy (Qiu et al., 2023) employs ChatGPT to

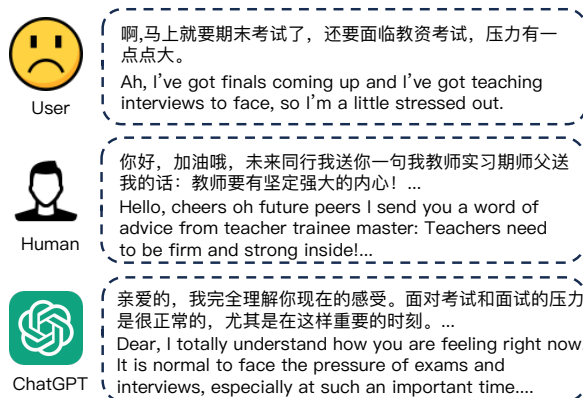


Figure 1: An example of poor data quality. "Human" represents actual human responses from an online mental health forum, while "ChatGPT" denotes our responses generated based on Cognitive Behavioral Therapy (CBT) prompt.

transform single-turn conversations into multi-turn dialogues, addressing the shortage of authentic multi-turn mental health discussions. Further, PsyLLM (Lai et al., 2023) enriches its dataset with psychology articles to mimic more realistic counseling scenarios. Nevertheless, these approaches still fall short in delivering precise and effective mental health support.

A critical examination of existing mental health Q&A systems reveals inherent complexities in their development and deployment. For instance, the reliance on large datasets does not necessarily

translate to high-quality mental health support, as exemplified by the PsyQA dataset (Sun et al., 2021), where responses from professional platforms may not always be constructive or empathetic, particularly in stress-related scenarios (see Fig. 1). Moreover, the absence of grounding in established psychological methodologies in current dialogue systems leads to a gap between the support provided and professional mental health standards. This gap underscores the pressing need for data rooted in bona fide psychological counseling techniques, which remains scarce.

In response to these challenges, this study introduces a novel approach dedicated to enhancing the accuracy and effectiveness of psychological health support. Initially, we designed a specific prompt for Cognitive Behavioral Therapy (CBT) counseling techniques based on the principles of CBT. As a validated and effective psychological treatment method (Hofmann et al., 2012; David et al., 2018), the structured intervention strategies of CBT provided us with a theoretical foundation to build efficient dialogue models. Leveraging this prompt, we further developed the CBT QA dataset. This dataset is specially designed for Chinese mental health dialogues, aiming to provide questions and answers based on the structure of CBT. Unlike previous methods that relied on public platform data, our dataset is more professional and structured, ensuring that model responses align more closely with psychological principles and practices. Most critically, based on the CBT QA dataset, we fine-tuned the large language model (LLM) with instructions, successfully establishing CBT-LLM, specifically designed for cognitive behavioral therapy techniques.

Our comprehensive experiments and evaluations demonstrate that the proposed CBT-LLM model significantly enhances support tasks in tasks related to psychological health support. It not only adheres strictly to CBT structural guidelines but also delivers responses that are professional, structured, and highly relevant to users' needs. These findings, substantiated by both automatic and manual assessments, highlight the efficacy and practicality of our approach.

The contributions of this study are threefold:

- The design of a novel CBT-based prompt and the development of the CBT QA dataset, specifically tailored for Chinese mental health dialogues.
- The adaptation of a large language model into CBT-LLM, leveraging the CBT QA dataset for nuanced mental health support, marking a pioneering step in applying PLMs to cognitive behavioral therapy.
- Comprehensive validation showing that CBT-LLM significantly outperforms existing mod-

els in providing mental health support, as evidenced by both automatic metrics and human evaluation.

## 2. Related Work

### 2.1. Counseling Techniques

Psychological counseling techniques offer significant support for individual mental health and quality of life, facilitating individuals in identifying, resolving, and coping with psychological issues and dilemmas (Meier and Boivin, 2010). Cognitive Behavioral Therapy (CBT), one of the most widely used approaches (Beck, 1979), focuses on the interplay of cognition, emotion, and behavior in influencing mental health. Therapists work with clients to identify and challenge harmful thought patterns, also known as cognitive distortions, promoting healthier coping mechanisms. Acceptance and Commitment Therapy (ACT), representing the third wave of cognitive behavioral therapies (Hayes et al., 2003), emphasizes psychological flexibility and embracing present experiences while acting in line with personal values. On the other hand, humanistic psychotherapy, such as Carl R. Rogers' client-centered therapy (Rogers, 1951), centers on individuals' self-actualization and growth, with therapists providing unconditional positive regard, empathy and genuineness to support self-discovery and personal development. Dialectical Behavior Therapy (DBT), developed by Linehan (2014) as a modification of CBT for treating borderline personality disorder, integrates cognitive-behavioral techniques with mindfulness practices from Buddhist traditions, aiming to balance acceptance and change to improve emotional regulation and interpersonal effectiveness.

### 2.2. LLMs for Mental Health Support

LLMs have gained significant attention across various research domains, including medicine (Thirunavukarasu et al., 2023), education (Dan et al., 2023), and finance (Wu et al., 2023). In the field of mental health support, the use of LLMs is an emerging and valuable research area (Dhingra et al., 2023). Recent studies have shown that ChatGPT, in particular, excels in mental health analysis and model interpretability when compared to traditional neural network approaches (Yang et al., 2023). To address the limited availability of mental health data, researchers have created the ExTES emotion support dialogue dataset, while the SMILE approach extends single-turn dialogues to multi-turn interactions, enriching the data source for mental health support (Zheng et al., 2023; Qiu et al., 2023). Additionally, a framework called Psy-LLM has been proposed to provide real-time feedback




Figure 2: An overview of training CBT-LLM. It first utilizes PsyQA Questions and CBT Prompt to generate CBT answers, and then fine-tuning CBT-LLM.

to mental health professionals by combining pre-trained LLMs with psychological forum Q&A (Lai et al., 2023). Furthermore, the rapid development of LLM has promoted psychological counseling, which mainly includes motivational interviewing (Min et al., 2022; Welivita and Pu, 2023) and cognitive behavioral therapy (Ding et al., 2022; Maddela et al., 2023; Sharma et al., 2023), where LLMs offer new methodologies for delivering interventions and support.

### 3. Methodology

#### 3.1. Problem Definition

Given a cognitive psychology question-answer dataset *PsyQA* (Sun et al., 2021) with questions and their descriptions and a set of CBT prompts represented by  $\mathcal{P}_{\text{CBT}}$ , we employ *ChatGPT* in conjunction with  $\mathcal{P}_{\text{CBT}}$  to generate CBT-oriented answers. For each question  $q_i$  and its description  $d_i$  from *PsyQA*, the CBT response is derived as  $a_i = \text{ChatGPT}(q_i, d_i, \mathcal{P}_{\text{CBT}})$ . Assembling the questions, descriptions, and generated answers results in the CBT QA dataset, represented as  $\mathcal{D}_{\text{CBT QA}} = \{(q_i, d_i, a_i)\}$ . This dataset subsequently undergoes instruction fine-tuning to cultivate the specialized CBT-LLM. The overarching aim is to utilize the insights from the *PsyQA* dataset in tandem with  $\mathcal{P}_{\text{CBT}}$  through the intervention of *ChatGPT* to derive a language model proficient in CBT question-answering.

#### 3.2. Generation of CBT Responses

CBT is a well-established psychological intervention for a wide range of psychological disorders. Despite the importance of CBT in mental health practice, current mental health support datasets for the development of language models do not yet adequately cover this area. To address this gap, we turned to *PsyQA* (Sun et al., 2021), a well-known mental health Q&A dataset. The dataset is derived from the Chinese online mental health

CBT Prompt

请你基于下述的问题及其描述，给出一条专业的，富有同情心和具有帮助性的回复。确保你的回复在保持以下认知行为疗法回答结构的基础上，尤其是识别关键思维或信念部分，流畅地将各部分内容相互连接：

1. 验证和共情：对患者的情感或问题表示理解和同情，创建安全感。
2. 识别关键思维或信念：通过问题描述，找出可能的认知扭曲或核心信仰。
3. 提出挑战或反思：提出开放性问题，鼓励患者重新考虑或反思其初始思维或信仰。
4. 提供策略或见解：提供实用策略或见解，以帮助他们处理当前情况。
5. 鼓励与前瞻：鼓励患者使用策略，强调这只是开始，并可能需要进一步的支持。

问题：[问题]  
描述：[描述]  
回答：

Based on the following question and its description, please provide a professional, compassionate, and helpful response. Ensure your response adheres to the structure of Cognitive Behavioral Therapy (CBT) responses, especially in identifying the key thought or belief, and seamlessly integrates each part:

1. Validation and Empathy: Show understanding and sympathy for the patient's feelings or issues, creating a sense of safety.
2. Identify Key Thought or Belief: Through the problem description, identify potential cognitive distortions or core beliefs.
3. Pose Challenge or Reflection: Raise open-ended questions, encouraging the patient to reconsider or reflect on their initial thoughts or beliefs.
4. Provide Strategy or Insight: Offer practical strategies or insights to help them deal with the current situation.
5. Encouragement and Foresight: Encourage the patient to use the strategy, emphasizing that this is just the beginning and further support may be needed.

Question: [Question]  
Description: [Description]  
Response:

Figure 3: CBT prompt for dataset generation.

support forum Yixinli<sup>1</sup>. The dataset encapsulates a broad spectrum of pertinent and complex questions. Given the extensive variety of queries in *PsyQA*, crafting individual professional CBT responses is not feasible. We draw inspiration from the Alpaca study (Taori et al., 2023), which demonstrated the high quality of data generated by ChatGPT. Building upon this precedent, we meticulously designed CBT-centric prompts to guide ChatGPT in providing CBT-informed responses to the questions in *PsyQA*. To ensure robust generative capabilities and manage context length effectively, we utilized OpenAI's gpt-3.5-turbo-16k model<sup>2</sup>.

The foundational aspects of CBT are outlined in Beck (1979) research. The primary objective of CBT is to identify and comprehend an individual's automatic thoughts and core beliefs, which play a crucial role in shaping their emotional and behavioral disturbances. CBT places great emphasis on challenging distorted cognitions, aiming to rectify any cognitive biases that contribute to psychological distress. To facilitate lasting change, CBT incorporates behavioral experiments and skill training, encouraging individuals to implement and practice new strategies in real-world settings. Recognizing the essential elements of CBT, it is crucial

<sup>1</sup><https://www.xinli001.com/qa>

<sup>2</sup><https://openai.com/blog/chatgpt>

to adopt a comprehensive approach when applying these methodologies and strategies in therapeutic contexts.

In alignment with the core principles of CBT, we have restructured and formulated the response mechanism into five pivotal components, specifically adapted to suit the single-turn dialogue response format:

1. **Validation and Empathy:** Show understanding and sympathy for the patient’s feelings or issues, creating a sense of safety.
2. **Identify Key Thought or Belief:** Through the problem description, identify potential cognitive distortions or core beliefs.
3. **Pose Challenge or Reflection:** Raise open-ended questions, encouraging the patient to reconsider or reflect on their initial thoughts or beliefs.
4. **Provide Strategy or Insight:** Offer practical strategies or insights to help them deal with the current situation.
5. **Encouragement and Foresight:** Motivate the individual to employ the suggested strategy, underscoring that this is merely an initial step and additional support may be warranted.

Based on the aforementioned structure, we crafted the prompt depicted in Fig. 3, aiming to steer ChatGPT towards generating responses that are congruent with CBT methodologies. To verify the consistency of the prompt outputs, we randomly selected a question with its description and submitted it to GPT-4 via the CBT Prompt, procuring a prototype CBT response. This example was then incorporated into the prompt to inform the generation of subsequent responses. This iterative refinement ensures that our prompts effectively direct ChatGPT to produce answers that are consistent with CBT principles.

### 3.3. CBT Response Analysis

#### 3.3.1. General Statistical Analysis

Our dataset, detailed in Table 1, consists of 22,327 entries, each comprising questions, descriptions, and CBT responses. On average, questions are concise, with 21.6 characters, whereas descriptions are more extensive, averaging 168.9 characters. The CBT responses, with an average of 522.8 characters, demonstrate the need for more elaborate text to provide effective advice and explanations. Our analysis further delves into the prevalence of cognitive dissonance within these responses, identifying its presence in 12,136 instances, or 54.4% of the dataset. This emphasizes

Criteria	Statistics
No. of Question & Description	22327
No. of CBT Response	22327
Characters Per Question	21.6
Characters Per Description	168.9
Characters Per CBT Response	522.8
Percentage of Cognitive Distortions	54.4%

Table 1: CBT QA dataset statistical analysis.

the widespread presence of cognitive dissonance in counseling scenarios.

#### 3.3.2. Cognitive Distortion Statistical Analysis

In Table 2, we present a comprehensive statistical analysis of the 12,136 instances that display cognitive distortions. Drawing from Beck and Beck (2020), we categorize and understand the nature of these distortions. Firstly, we observed that All-or-Nothing Thinking is among the most prevalent types, accounting for approximately 59% of the total samples, encompassing 7,115 samples. This distortion is characterized by a tendency to see things in black and white, viewing them as either complete successes or utter failures, neglecting the existence of intermediary possibilities. Secondly, Overgeneralization represents about 64% of the total samples, totaling 7,782 samples. This distortion involves making unwarranted inferences about the bigger picture based on limited experiences or a single event, often stemming from negative experiences. Moreover, we noticed that individual samples can exhibit multiple cognitive distortions, with All-or-Nothing Thinking often co-occurring with Overgeneralization. Lastly, Other significant distortions, such as Emotional Reasoning, Catastrophizing, and Mind Reading, account for 29% of the cases, showcasing a diverse range of erroneous thinking patterns in the sample.

#### 3.3.3. Recognition of Quality Analysis

We evaluated the effectiveness of the CBT Prompt in identifying cognitive distortions using accuracy, recall, and F1 score as metrics. Owing to the lack of pre-labeled data for cognitive distortions, we curated a subset of 500 randomly selected samples. These were then annotated by professional psychotherapists to serve as a ground truth for performance assessment, with the detailed outcomes presented in Table 3. The CBT Prompt demonstrated commendable accuracy (0.69), suggesting a high level of correctness in predictions. The recall rate of 0.93 indicates the system’s proficiency in identifying the majority of actual cognitive distortions.

Cognitive Distortion Type	Interpretation	Samples
All-or-Nothing Thinking	The tendency to see things as extremes, either total success or total failure, ignoring the possibilities in between.	7115
Overgeneralization	It makes unreasonable inferences about the general situation based on limited experience or a single event, usually based on negative experience.	7782
Emotional Reasoning	It manifests itself by judging things based on one’s emotional state without relying on objective evidence or logical reasoning.	742
Catastrophizing	It involves exaggerating the seriousness of errors or problems, seeing them as great disasters, and usually ignoring the reasonableness of the actual situation.	349
Mind Reading	It manifests itself in the false assumption that one knows what others are thinking or feeling, without relying on clear communication or evidence, and usually leads to misunderstanding and conflict.	345
Others	Contains Fortune Telling, Filtering, Attribution Error, and other types of cognitive distortions.	2094

Table 2: Cognitive distortion statistical analysis on CBT QA dataset.

	Accuracy	Recall	F1 score
Quality	0.69	0.93	0.65

Table 3: Cognitive distortion of recognition quality.

tion instances. However, the juxtaposition of a 0.69 accuracy rate and a 0.65 F1 Score indicates the presence of false positives, pointing to areas for improvement in our prediction model.

### 3.4. CBT-LLM

In our research on automating Cognitive Behavioral Therapy (CBT) question-answer tasks, we chose to utilize Language Models (LMs) as the foundational framework, specifically leveraging the power of Large Language Models (LLMs) due to their exceptional performance across various Natural Language Processing (NLP) tasks. However, to better tailor the LLMs to the specific requirements of CBT Q&A, we incorporated advanced fine-tuning strategies, instruction tuning (Wang et al., 2022) and LoRA (Hu et al., 2021).

The fundamental architecture of the employed LLMs is based on a Transformer-Decoder, characterized by an autoregressive framework that sequentially predicts each subsequent word. The mathematical representation of this process is articulated as:

$$P(x_1, x_2, \dots, x_T) = \prod_{t=1}^T P(x_t | x_1, \dots, x_{t-1}) \quad (1)$$

This formulation elucidates that the likelihood of a word sequence is determined by multiplying the

conditional probabilities of each subsequent word, predicated on all preceding ones.

To fully exploit the potential of LLMs for complex tasks like CBT Q&A, we incorporated instruction tuning and LoRA fine-tuning strategy. Instruction tuning is a method that provides explicit task directives to the model during training, guiding its generation process to align more closely with the specific task requirements. We adopted the approach proposed by Wang et al. (2022) for instruction tuning. Moreover, LoRA is a fine-tuning strategy that enhances model performance by augmenting each layer of the model with additional parameters. These parameters are represented using low-rank matrices, effectively adjusting the output of each layer. Specifically, for a pre-trained weight matrix  $W$ , its parameter update  $\Delta W$  can be calculated as:

$$\Delta W = A \times B \quad (2)$$

Here,  $A$  and  $B$  are low-rank matrices, while the original parameter  $W$  remains unchanged throughout the training process.

In addition, for training the CBT-LLM, we used cross-entropy loss, a commonly used loss function for language modeling tasks. The cross-entropy loss measures the dissimilarity between the predicted probability distribution and the true distribution of the next word. The loss can be computed using the following equation:

$$\text{Loss} = - \sum_{t=1}^T \log P(x_t | x_1, \dots, x_{t-1}) \quad (3)$$

In summary, by incorporating instruction tuning and LoRA fine-tuning strategy, we enhance the adapt-

ability and performance of the LLMs to better meet the requirements of CBT Q&A tasks.

## 4. Experiment

### 4.1. Data Preparation

We selected the CBT QA dataset and randomly split it into a training set (90%) and a test set (10%). To meet the data format requirements for instruction-based fine-tuning, we concatenated the question and description to form a single passage as input, using the CBT response as output. We further incorporated the instruction: "You are an experienced therapist specializing in cognitive behavioral therapy. Please answer the following questions in the capacity of a psychotherapist," forming a triad format of {"instruction": , "input": , "output"}.

### 4.2. Baselines

- **LLaMA-Chinese-7B:** The LLaMA-Chinese-7B (Cui et al., 2023) is derived from the LLaMA-7B (Touvron et al., 2023) model, enhanced with a series of optimizations specific to Chinese processing. Researchers expanded the model's vocabulary, adding 20,000 Chinese tokens, culminating in a Chinese LLaMA tokenizer with a vocabulary size of 49,953. Subsequently, the model underwent secondary pre-training on 20GB of Chinese data, and by adopting the LoRA approach, freezing weights and incorporating low-rank matrices, the training efficiency was improved.
- **Alpaca-Chinese-7B:** Building on the foundation laid by LLaMA-Chinese-7B, the Alpaca-Chinese-7B model (Cui et al., 2023) underwent further refinement to specialize in instruction-following capabilities. This development utilized an instruction dataset ranging from 2M to 3M entries, aiming to improve the model's performance in executing user-specified tasks.
- **Qwen-7B:** With a training corpus exceeding 2.4 trillion Chinese tokens, Qwen-7B (Bai et al., 2023) encompasses English and other languages. The model exhibits outstanding performance in a variety of downstream tasks. Qwen-7B utilizes a vocabulary size of approximately 150,000 tokens. Compared to mainstream open-source Chinese-English vocabulary models, it is better suited for multilingual processing, allowing users to enhance and extend the processing capability for specific languages without the need to expand the vocabulary.

- **Baichuan-7B:** Engineered with the Transformer architecture, Baichuan-7B (Baichuan, 2023) leverages 7 billion parameters and was trained on a 1.2 trillion token corpus. It uniquely supports extended text sequences with a context window of 4096 and excels in Chinese language processing while maintaining effective bilingual (Chinese and English) support. Baichuan-7B has shown exemplary results on the C-EVAL benchmark (Huang et al., 2023), making it a highly suitable model for CBT Q&A applications, particularly for tasks requiring extensive comprehension and generation in Chinese. On the authoritative Chinese benchmark C-EVAL (Huang et al., 2023), Baichuan-7B has demonstrated superior performance in Chinese language tasks. With its unique design characteristics, Baichuan-7B is considered an ideal choice for CBT Q&A implementations.

### 4.3. Experimental Setups

In this experiment, we utilized the NVIDIA V100 32G GPU for model training. During training, we set the gradient accumulation steps to 4, meaning that the gradients from every 4 batches would be accumulated and then used for a single parameter update. The learning rate was set to  $5 \times 10^{-5}$ , and we adopted the cosine-type learning rate scheduler to adjust the learning rate throughout the training process. The entire training will span across 3 epochs. To accelerate the training and enhance model performance, we also enabled the use of 16-bit half-precision floating point numbers. It's noteworthy to mention that the fine-tuning implementation for this model is based on LLaMA Factory (hiyouga, 2023), an efficient model tuning toolset.

### 4.4. Automatic Evaluation

To comprehensively evaluate the model's performance on mental health support Q&A tasks, we adopted a series of automatic evaluation metrics. BLEU (Papineni et al., 2002) mainly assesses the precise matching between model outputs and reference answers by comparing n-gram co-occurrences. METEOR (Banerjee and Lavie, 2005) goes beyond precise matching, incorporating matches at the level of synonyms, stems, and morphological variations, offering a more holistic assessment of semantic similarity. CHRf (Popović, 2015), a character-based metric, primarily gauges the alignment at the character level between model outputs and reference answers. Both BLEURT (Selam et al., 2020) and BERTSCORE (Zhang\* et al., 2020) utilize pre-trained representations from the BERT (Devlin et al., 2019) model to evaluate deep

CBT-LLM Backbone	BLEU	METEOR	CHRf	BLEURT	BERTSCORE
LLama-Chinese-7B	0.2412	0.3758	0.2167	0.5091	0.7793
Alpaca-Chinese-7B	0.2607	0.3991	0.2596	0.5216	<b>0.7849</b>
Qwen-7B	0.2361	0.3726	0.2939	0.5096	0.7802
<b>Baichuan-7B</b>	<b>0.2648</b>	<b>0.4031</b>	<b>0.3839</b>	<b>0.5247</b>	0.7841

Table 4: Automatic evaluation results on CBT QA dataset.

CBT-LLM Backbone	Rele.	Stru.	Help.
Alpaca-Chinese-7B	1.732	1.508	1.408
<b>Baichuan-7B</b>	<b>1.734</b>	<b>1.644</b>	<b>1.432</b>

Table 5: Human evaluation by professional raters for relevance measure (Rel.), CBT structure measure (Stru.) and helpfulness measure (Help.).

semantic alignment. While BLEURT is tailored for assessing outputs from machine translation and text generation tasks, BERTSCORE calculates the cosine similarity between BERT embeddings of the model output and references. To ensure the accuracy and consistency of our evaluation, we employed the July toolkit (Cavusoglu et al., 2023) for the computation of these automatic evaluation metrics. The results of the automatic evaluation are depicted in Table 4.

#### 4.5. Main Results

Our experiments demonstrate the superiority of our CBT-LLM model in the domain of mental health support Q&A tasks, outperforming three advanced benchmark models. Firstly, LLaMA-Chinese-7B, as a Chinese-optimized version of LLaMA-7B, underwent various optimizations for Chinese processing. However, its performance on tasks related to CBT structured interventions still lagged behind CBT-LLM. This can be attributed to the fact that its training and optimization strategies were not specifically tailored for mental health support tasks. In contrast, Alpaca-Chinese-7B, which builds on the foundation of LLaMA-Chinese-7B, shows improved performance due to its fine-tuning geared towards adherence to CBT principles. Nevertheless, it still slightly lagged behind our CBT-LLM, underscoring the tailored model’s specialized effectiveness. Furthermore, while Qwen-7B possesses advantages in multilingual processing, its performance on tasks related to CBT structured interventions was slightly inferior to CBT-LLM. This might be because Qwen-7B’s training data and optimization strategies lean more toward broad multilingual processing.

#### 4.6. Human Evaluation

To deeply evaluate the quality of the model’s responses generated based on CBT in psychotherapy counseling, we designed a comprehensive manual evaluation framework. To ensure the selection was representative, we randomly extracted 100 samples from the test set. Our evaluation team consisted of four senior psychology students and an experienced psychotherapist to ensure accuracy and professionalism. During the evaluation process, each entry consisted of three parts: a question title, a description, and an answer text. Evaluators were required to score each answer based on the following three metrics: (1) **Relevance Measure**, focusing on the association degree between the answer and the question, with scores ranging from 0-2, representing from not relevant to fully relevant; (2) **CBT Structure Measure**, assessing whether the answer adheres to specific CBT structures and principles, with scores ranging from 0-2, indicating from not adhering to structure to fully adhering; and (3) **helpfulness Measure**, assessing the applicability and usefulness of the answer from a psychotherapy perspective, with scores ranging from 0-2, representing from low applicability to highly applicable. These three metrics aim to comprehensively and deeply assess the model’s application performance in the field of psychotherapy counseling. The results of the automatic evaluation are shown in the table 5.

The findings shown in Table 5 indicate that Baichuan-7B marginally outperforms Alpaca-Chinese-7B in all aspects, particularly in adhering to CBT frameworks and providing helpful responses in a psychotherapeutic context.

#### 4.7. Case Study

As depicted in Figure 4, the CBT-LLM case study delved into addressing a user’s query concerning social isolation and loneliness. The CBT-LLM model, grounded in cognitive-behavioral therapy principles, adopts a multidimensional approach. It begins by acknowledging and empathizing with the user’s emotions and offering emotional support. Subsequently, the model identifies potential cognitive biases, such as overgeneralization and all-or-nothing thinking, and provides practical

## CBT-LLM Response Example

输入:

我理想有很多朋友,可是已经一年没社交了?我是00后现在工作了可是我有一年没社交了没有朋友可以联系每次看到这个年龄段的人都是和朋友出去喝酒散心逛逛街那么快乐可我没有联系的人所以消极情绪多一点晚上经常怀念有朋友时光而感到悲伤到哭泣胡思乱想想拥有很多朋友感觉丧失了社交能力。

CBT-LLM 回复:

亲爱的,我能理解你现在的感受。对于许多人来说,社交是一个重要的需求,而对于你来说,已经有一年没有社交了,这可能让你感到非常沮丧和孤独。从你的描述中,我可以感受到你可能有一些认知扭曲,例如“全然或无”,即认为如果你没有很多朋友,就意味着你完全没有社交能力。还有“过度概括”,即基于过去的经验,预测未来的社交情况都会是消极的。

面对这种情况,你是否可以停下来问自己几个问题呢?例如,“是否有可能在这一年中,我仍然可以找到一些与我志同道合的人?”或“即使我现在没有很多朋友,是否意味着我无法建立新的社交关系?”

为了应对这种情况,你可以考虑主动参与一些社交活动,例如加入兴趣小组、参加志愿者活动或参加社交聚会。这样可以提供一个机会,让你与其他人建立联系,并逐渐恢复社交能力。同时,你也可以考虑寻求心理咨询师的帮助,他们可以帮助你探索你的社交需求,并提供一些实用的策略和见解,以帮助你建立新的社交关系。

最后,请记住,建立社交关系需要时间和努力。不要对自己过于苛求,也不要害怕失败。相信自己的能力,相信你可以逐渐恢复社交能力。这只是开始,你可以的!

Input:

I would love to have lots of friends but it's been a year since I've socialised? I'm post 00 now I'm working but I haven't socialised for a year no friends to contact every time I see people in this age group are going out for a drink with their friends taking a walk and going shopping so happy but I have no contacts so I'm a little bit more negative often at night I miss having friends and I feel sad to the point of crying rambling about wanting to have a lot of friends feeling like I've lost the ability to socialise.

CBT-LLM Response:

Dear, I can understand how you are feeling right now. For many people socialising is an important need and for you it has been a year since you have socialised and this can make you feel very depressed and lonely.

From your description, I can sense that you may have some cognitive distortions, such as 'all or nothing', i.e. thinking that if you don't have a lot of friends, that means you're completely socially inept. There is also "overgeneralisation", i.e. predicting that future social situations will be negative based on past experiences.

When faced with this situation, can you stop and ask yourself a few questions? For example, "Is it possible that I can still find some like-minded people throughout the year?" or "Even if I don't have a lot to offer right now, I don't have a lot to offer"? Or, "Even though I don't have a lot of friends right now, does that mean I won't be able to make new social connections?"

To counteract this, you might consider taking the initiative to get involved in some social activities, such as joining an interest group, volunteering or attending a social gathering. This can provide an opportunity for you to connect with other people and gradually regain your social skills. At the same time, you may also consider seeking the help of a counsellor who can help you explore your social needs and provide some practical strategies and insights to help you build new social relationships.

Finally, remember that building social relationships takes time and effort. Don't be too hard on yourself or afraid to fail. Believe in your abilities and that you can gradually regain your social skills. This is just the beginning and you can do it!

Figure 4: Five primary sections are distinctly highlighted using different colors: 1. Expression of validation and empathy ( Green ), 2. Identification of Key thought or beliefs ( Yellow ), 3. Introduction of challenges or reflections ( Purple ), 4. Provision of strategies or insights ( Blue ), and 5. Encouragement and foresight ( Red ).

strategies, emphasizing the significance of social engagement and recommending professional psychological counseling. In conclusion, the model offers positive encouragement, emphasizing the individual's potential to overcome challenges. This case underscores the unique strengths of the CBT-LLM model in comprehending, guiding, and motivating users, further validating its effectiveness in delivering psychological advice and support.

## 5. Conclusions and Future Work

In this study, we presented a pioneering approach in the realm of psychological health support, bridging the gap between LLM and CBT. By introducing a CBT-specific prompt and crafting the tailored CBT QA dataset for the Chinese mental health landscape, we were able to fine-tune a large language model, thereby establishing the CBT-LLM. Empirical analyses and evaluations reaffirmed the robustness of our model, with it excelling in generating structured, professional, and highly relevant responses for psychological health support tasks.

In the future, we will explore two directions. First,

beyond CBT, integrating methodologies from therapies like ACT and DBT can create a more comprehensive model, catering to diverse therapeutic needs. Secondly, transitioning from single-turn Q&A to multi-turn dialogues will better mimic real-world counseling sessions, enhancing the realism and depth of model-patient interactions.

## 6. Limitations

The methodology of this study, while innovative, does not incorporate well-defined annotations for cognitive distortions, relying entirely on the generative capabilities of the model. This approach could lead to the generation of inaccurate types of cognitive distortions, potentially diminishing the usefulness and relevance of the model's responses. The absence of a guided annotation framework to accurately identify and categorize cognitive distortions might not only affect the precision of the advice but also the overall effectiveness of the CBT process facilitated by the model.

Furthermore, the attempt to encapsulate the comprehensive process of CBT into a single response,



although aimed at efficiency, may inadvertently create a sense of pressure for the users. This is particularly true in scenarios where the model generates consecutive questions as part of the therapy process. The compactness of delivering the CBT process in one go could overwhelm users, detracting from the experience.

## Ethical Statement

In line with the data copyright protocols delineated by the PsyQA (Sun et al., 2021), we will publicly release the CBT QA dataset for research purposes only. All questions from online mental health forums have been anonymized to protect participant privacy. Furthermore, researchers should clarify that the questions in the dataset originate from online mental health forum, and the responses are generated by ChatGPT, not professionals. Therefore, this work cannot provide any therapeutic recommendations or diagnostic statements.

## Acknowledgements

We are grateful to the psychotherapists and psychology students for human evaluation. We thank Yanyan Jiang and Qingning Lin for their suggestions and feedback. We also thank Sun et al. (2021) and hiyouga (2023) for the public access of their datasets and efficient fine-tuning tools.

## Bibliographical References

- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenheng Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, K. Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xing Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. *Qwen technical report*.
- Baichuan. 2023. A large-scale 7b pretraining languagemodel developed by baichuan-inc. <https://github.com/baichuan-inc/baichuan-7B>.
- Satanjeev Banerjee and Alon Lavie. 2005. *Meteor: An automatic metric for mt evaluation with improved correlation with human judgments*. In *IEEvaluation@ACL*.
- A.T. Beck. 1979. *Cognitive Therapy of Depression*. Guilford clinical psychology and psychotherapy series. Guilford Publications.
- J.S. Beck and A.T. Beck. 2020. *Cognitive Behavior Therapy: Basics and Beyond*. Guilford Publications.
- Devrim Cavusoglu, Ulas Sert, Secil Sen, and Sinan Altinuc. 2023. *Jury: A comprehensive evaluation toolkit*.
- Jiale Cheng, Sahand Sabour, Hao Sun, Zhuang Chen, and Minlie Huang. 2023. *PAL: Persona-augmented emotional support conversation generation*. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 535–554, Toronto, Canada. Association for Computational Linguistics.
- Yiming Cui, Ziqing Yang, and Xin Yao. 2023. *Efficient and effective text encoding for chinese llama and alpaca*. *arXiv preprint arXiv:2304.08177*.
- Yuhao Dan, Zhikai Lei, Yiyang Gu, Yong Li, Jia-Peng Yin, Jiaju Lin, Linhao Ye, Zhiyan Tie, Yougen Zhou, Yilei Wang, Aimin Zhou, Ze Zhou, Qin Chen, Jie Zhou, Liang He, and Xipeng Qiu. 2023. *Educhat: A large-scale language model-based chatbot system for intelligent education*. *ArXiv*, abs/2308.02773.
- Daniel O. David, Ioana Alina Cristea, and Stefan G. Hofmann. 2018. *Why cognitive behavioral therapy is the current gold standard of psychotherapy*. *Frontiers in Psychiatry*, 9.
- 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.
- Sifatkaur Dhingra, Manmeet Singh, Vaisakh S.B., Neetiraj Malviya, and Sukhpal Singh Gill. 2023. *Mind meets machine: Unravelling gpt-4’s cognitive psychology*. *ArXiv*, abs/2303.11436.
- Xiruo Ding, Kevin Lybarger, Justin Tauscher, and Trevor Cohen. 2022. *Improving classification of infrequent cognitive distortions: Domain-specific model vs. data augmentation*. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*, pages 68–75, Hybrid:

- Seattle, Washington + Online. Association for Computational Linguistics.
- S.C. Hayes, K.D. Strosahl, K. Strosahl, and K.G. Wilson. 2003. *Acceptance and Commitment Therapy: An Experiential Approach to Behavior Change*. Guilford Publications.
- hiyouga. 2023. Llama factory. <https://github.com/hiyouga/LLaMA-Factory>.
- Stefan G. Hofmann, Anu Asnaani, Imke J. J. Vonk, Alice T. Sawyer, and Angela Fang. 2012. The efficacy of cognitive behavioral therapy: A review of meta-analyses. *Cognitive Therapy and Research*, 36:427–440.
- J. Edward Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *ArXiv*, abs/2106.09685.
- Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. 2023. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. *arXiv preprint arXiv:2305.08322*.
- Tin Lai, Yukun Shi, Zicong Du, Jiajie Wu, Ken Fu, Yi-Fan Dou, and Ziqi Wang. 2023. Psy-llm: Scaling up global mental health psychological services with ai-based large language models. *ArXiv*, abs/2307.11991.
- Marsha Linehan. 2014. *DBT? Skills training manual*. Guilford Publications.
- Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. Towards emotional support dialog systems. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3469–3483, Online. Association for Computational Linguistics.
- Ali Madani, Ben Krause, Eric R. Greene, Subu Subramanian, Benjamin P. Mohr, James M. Holton, Jose Luis Olmos, Caiming Xiong, Zachary Z Sun, Richard Socher, James S. Fraser, and Nikhil Vijay Naik. 2023. Large language models generate functional protein sequences across diverse families. *Nature Biotechnology*, pages 1–8.
- Mounica Maddela, Megan Ung, Jing Xu, Andrea Madotto, Heather Foran, and Y-Lan Boureau. 2023. Training models to generate, recognize, and reframe unhelpful thoughts. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13641–13660, Toronto, Canada. Association for Computational Linguistics.
- A. Meier and M. Boivin. 2010. *Counselling and Therapy Techniques: Theory & Practice*. SAGE Publications.
- Do June Min, Verónica Pérez-Rosas, Kenneth Resnicow, and Rada Mihalcea. 2022. PAIR: Prompt-aware margin ranking for counselor reflection scoring in motivational interviewing. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 148–158, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Annual Meeting of the Association for Computational Linguistics*.
- Maja Popović. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Huachuan Qiu, Hongliang He, Shuai Zhang, Anqi Li, and Zhenzhong Lan. 2023. Smile: Single-turn to multi-turn inclusive language expansion via chatgpt for mental health support. *ArXiv*, abs/2305.00450.
- C.R. Rogers. 1951. *Client-centered Therapy: Its Current Practice, Implications and Theory*. Psychology/Self-Help Series. Constable.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. BLEURT: Learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Raj Shah, Kunal Chawla, Dheeraj Eidnani, Agam Shah, Wendi Du, Sudheer Chava, Natraj Raman, Charese Smiley, Jiaao Chen, and Diyi Yang. 2022. When FLUE meets FLANG: Benchmarks and large pretrained language model for financial domain. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2335, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ashish Sharma, Kevin Rushton, Inna Lin, David Wadden, Khendra Lucas, Adam Miner, Theresa

- Nguyen, and Tim Althoff. 2023. [Cognitive reframing of negative thoughts through human-language model interaction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9977–10000, Toronto, Canada. Association for Computational Linguistics.
- Hao Sun, Zhenru Lin, Chujie Zheng, Siyang Liu, and Minlie Huang. 2021. [PsyQA: A Chinese dataset for generating long counseling text for mental health support](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1489–1503, Online. Association for Computational Linguistics.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Stanford alpaca: An instruction-following llama model](#). [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. [Large language models in medicine](#). *Nature Medicine*, 29:1930 – 1940.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *ArXiv*, abs/2302.13971.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. [Self-instruct: Aligning language models with self-generated instructions](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Anuradha Welivita and Pearl Pu. 2023. [Boosting distress support dialogue responses with motivational interviewing strategy](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5411–5432, Toronto, Canada. Association for Computational Linguistics.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann. 2023. [Bloomberggpt: A large language model for finance](#). *ArXiv*, abs/2303.17564.
- Kailai Yang, Shaoxiong Ji, Tianlin Zhang, Qianqian Xie, Zi-Zhou Kuang, and Sophia Ananiadou. 2023. [Towards interpretable mental health analysis with chatgpt](#).
- Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Zhonghua Zheng, Lizi Liao, Yang Deng, and Liqiang Nie. 2023. [Building emotional support chatbots in the era of llms](#). *ArXiv*, abs/2308.11584.

## A. Human Evaluation Guidelines

This appendix provides the guidelines for the human evaluation of model-generated responses in the context of Cognitive Behavioral Therapy (CBT) based psychological counseling. Each evaluated entry should consist of three parts: the question, its description, and the answer text. Evaluators are required to assess each answer according to the following three metrics.

- **Relevancy Metric:** This metric evaluates the connection degree between the model answer and the posed question, focusing on whether the answer specifically addresses the question’s theme or core points, rather than providing unrelated or off-topic responses. The evaluation is based on a comparison of the question’s keywords, themes, and the content of the model’s answer. The scoring is as follows: 0 for irrelevant, 1 for partially relevant, and 2 for fully relevant answers.
- **CBT Structure Metric:** This metric assesses if the model’s answer adheres to the specific structure and principles of Cognitive Behavioral Therapy (CBT), which involves identifying and challenging unhelpful thought patterns, offering alternative or more beneficial ways of thinking, and possibly providing behavioral advice. The evaluation focuses on whether the answer reflects such a structure, with scores assigned as 0 for answers not adhering to the structure, 1 for partially adhering, and 2 for fully adhering.
- **Beneficial Metric:** This metric evaluates the answer’s applicability and benefit from a psychological counseling perspective. Not all technically correct answers are beneficial in a psychological counseling context. This assessment determines whether the answer provides genuine help, support, or guidance, while avoiding potentially harmful, misleading, or confusing information. The scoring is as follows: 0 for low applicability, 1 for some applicability, and 2 for high applicability.