

# Knowledge Planning in Large Language Models for Domain-Aligned Counseling Summarization

Aseem Srivastava<sup>1</sup>, Smriti Joshi<sup>2</sup>, Tanmoy Chakraborty<sup>3</sup>, Md. Shad Akhtar<sup>1</sup>

<sup>1</sup>IIT Delhi, <sup>2</sup>Wysa, <sup>3</sup>IIT Delhi, India

{aseems, shad.akhtar}@iiitd.ac.in, smriti@touchkin.com, tanchak@iitd.ac.in

## Abstract

In mental health counseling, condensing dialogues into concise and relevant summaries (*aka* counseling notes) holds pivotal significance. Large Language Models (LLMs) exhibit remarkable capabilities in various generative tasks; however, their adaptation to domain-specific intricacies remains challenging, especially within mental health contexts. Unlike standard LLMs, mental health experts first plan to apply domain knowledge in writing summaries. Our work enhances LLMs' ability by introducing a novel *planning engine* to orchestrate structuring knowledge alignment. To achieve high-order planning, we divide knowledge encapsulation into two major phases: (i) holding dialogue structure and (ii) incorporating domain-specific knowledge. We employ a planning engine on Llama-2, resulting in a novel framework, PIECE. Our proposed system employs knowledge filtering-cum-scaffolding to encapsulate domain knowledge. Additionally, PIECE leverages sheaf convolution learning to enhance its understanding of the dialogue's structural nuances. We compare PIECE with 14 baseline methods and observe a significant improvement across ROUGE and Bleu scores. Further, expert evaluation and analyses validate the generation quality to be effective, sometimes even surpassing the gold standard. We further benchmark PIECE with other LLMs and report improvement, including Llama-2 (+2.72%), Mistral (+2.04%) and Zephyr (+1.59%), to justify the generalizability of the planning engine.

## 1 Introduction

Mental health counseling serves as a crucial front-line defense against mental illness. In a typical counseling session, clients articulate their issues while therapists provide support. An essential component of these sessions involves building a strong therapeutic bond and documenting the entire dialogue, commonly known as counseling note,

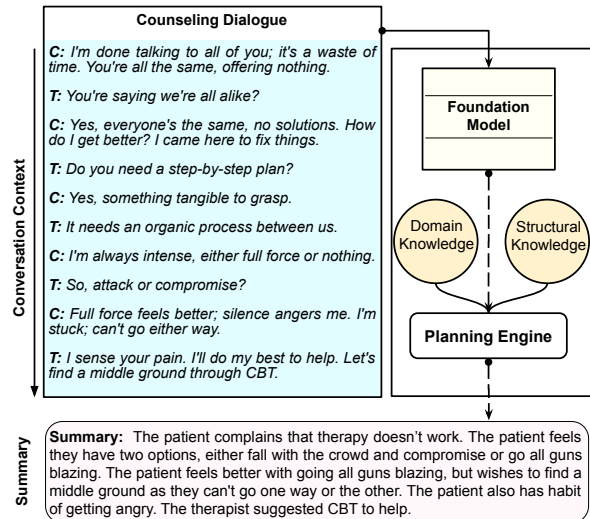


Figure 1: The proposed pipeline allows LLMs to first plan and then generate. In our approach, prioritizing planning before generation enriches summarization with conversational structure and domain knowledge.

record, or summary<sup>1</sup>. However, traditional methods of note-taking during sessions can present significant challenges, where therapists are required to divert their attention to note-taking, disrupting counseling interaction and support quality. This distraction not only hinders the therapist's focus but also deteriorates the required therapeutic bond, necessitating the need to automate this process.

While recent advancements in Large Language Models (LLMs) for mental health are enormous (Thirunavukarasu et al., 2023; Heerden et al., 2023), it is crucial to incorporate fundamental domain knowledge and understanding of the structural nuances of the dialogue, which current research lacks in ability (Casella et al., 2023). Consequently, we require a knowledge planner to adeptly capture domain intricacies and determine generation priorities. For instance, Figure 1 illustrates a sample counseling conversation between a therapist (T) and a client (C). Unlike conventional

<sup>1</sup>[www.apa.org/gradpsych/2007/01/track](http://www.apa.org/gradpsych/2007/01/track)

approaches that directly generate summaries using foundational models, our proposed pipeline introduces a planner. Through this planner, we infuse nuanced dialogue structure and domain-specific knowledge into LLMs.

Earlier efforts in counseling summarization (Srivastava et al., 2022) focused on utilizing domain knowledge by integrating patient health questionnaires (PHQ) and counseling components. However, it fell short compared to recent LLMs. The broader scope of such models has been explored for incorporating knowledge into generated text (Feng et al., 2023) and controlling text generation (Yu et al., 2022). While these efforts are focused on general-purpose downstream tasks, the complexity of counseling dialogue necessitates introducing a more (a) domain-centric and (b) knowledge-structured approach to cater to this problem.

Our study introduces a novel *planning engine*, designed specifically to guide the generation of LLMs. Focused on enhancing counseling summarization, we employ MentalLlama as the foundation to develop our framework – PIECE, **planning engine for mental counseling note generation**. Our model’s planning engine exploits knowledge scaffolding and sheaf learner by integrating domain-specific and structural knowledge into LLMs. We evaluate our PIECE against several LLMs, such as Mistral, Zephyr, and Llama. To assess the effectiveness of our approach, we compare PIECE’s performance against 14 baseline methods. Our quantitative evaluation consists of four automatic summary evaluation metrics – ROUGE: R-1, R-2, R-L, and Bluert, along with a domain-centric metric, Mental Health Information Capture (MHIC) proposed by (Srivastava et al., 2022). We observe a clear improvement of 3.42%, 10.11%, and 6.01% across R-1, R-2, and R-L metrics, respectively. Furthermore, to assess the generalizability of PIECE, we experiment on ACI-BENCH, a clinical note generation dataset (Yim et al., 2023). We observe that PIECE surpasses the state-of-the-art LLMs with planning-engine on ACI-BENCH as well.

Additionally, we perform an extensive expert evaluation through an established clinical relevance framework on a set of six dedicated metrics and three task-relevant survey questionnaires. The results affirm the superiority of PIECE against the baselines and demonstrate the adaptability of the planning engine across various LLMs. Our contributions are summarized below:

- We propose PIECE that integrates the planning

engine with MentalLlama to address the issue of unreliable generation by LLMs. This engine plans the LLM’s generation by filtering dialogue and injecting domain and structural knowledge.

- We extensively evaluate PIECE against 14 baseline methods. We present PIECE’s significant improvement evaluated across both automatic, human, and expert evaluation metrics.
- We demonstrate the adaptability of the planning engine as it seamlessly integrates with alternative LLMs like Mistral, Zephyr, and Llama. This may also expand the research in planning the generations of diverse LLMs.

PIECE is open sourced for research purpose at <https://github.com/LCS2-IIITD/PIECE>

## 2 Related Work

We present our literature review under three major segments to understand and build a planner for LLM-based counseling summarization.

**Generative AI in Mental Health Domain.** Recent advances in generative research in mental health have led to a diverse range of investigations. For instance, Sharma et al. (2023, 2021) utilized GPT-2 for reinforced feedback generation within peer counseling setups. Subsequent studies employed similar methods with language models (LMs) for empathetic generation (Sharma et al., 2021) and facilitated human-AI collaboration. Several studies have concentrated on controlled dialogue generation using LMs, leveraging reinforcement learning (Saha et al., 2022a,b; Priya et al., 2023; Srivastava et al., 2023; Mishra et al., 2023a,b,c). Despite this, there remains a scarcity of research aimed at enhancing LLMs in this domain (Liu et al., 2023). The concept of counseling summarization was previously explored by Srivastava et al. (2022), where they proposed a filtering mechanism based on annotation and domain knowledge to generate counseling summaries using LMs. As the focus has gradually shifted towards LLMs, recent findings by Yang et al. (2023a) have indicated that while LLMs exhibit robust capabilities, they still exhibit significant gaps compared to domain-specific methods.

**Knowledge Enhancement in LLMs.** There has been a significant effort to incorporate external and domain-specific knowledge into LLMs, as noted in various studies (Jiang et al., 2020, 2021; Choi et al., 2023). Tailoring LLMs to specific domains

has notably enhanced their ability to handle downstream tasks. Within the clinical domain, a series of specialized LLMs have emerged (Singhal et al., 2023). There exist a few dedicated LLMs specifically trained on mental health corpora, such as the variants of Llama-2 (Touvron et al., 2023), BART (Lewis et al., 2019), and T5 (Raffel et al., 2020) as MentalLlama, MentalBART and MentalT5, respectively (Yang et al., 2023b). While domain-specific LLMs facilitate more contextual relevance, they do not entirely mitigate risks associated with hallucinations or missing context. Planning LLM generations is one of the many solutions here (Valmeekam et al., 2023). Planning methodologies have been explored in various generative tasks, including reasoning (Wang et al., 2024), temporal generation (Zhang et al., 2022a), and code generation (Zhang et al., 2023). However, existing planning systems often lack domain-specific planning, notably within the mental health space.

**Text Summarization.** Knowledge-guided summarization has long been a research focus, and its variations include information-aware techniques (Sharma et al., 2019), perspective-based strategies applied in educational dialogues (Jain et al., 2023) and scientific document summarization (Kumar et al., 2023; Atri et al., 2023b), commonsense-driven clinical summarization (Tiwari et al., 2023), as well as approaches focusing on topic-awareness (Mukherjee et al., 2023) and attention mechanisms (Atri et al., 2023c). However, many of these methods largely rely on annotated corpora and the LMs to generate coherent text, neglecting deeper exploration into structural comprehension of input data. Limited research efforts, exemplified by Bodnar et al. (2022), discuss the use of sheaf for enhanced structural understanding, while Atri et al. (2023a) proposed sheaf to encapsulate structural information for summarization.

### 3 Dataset

We utilize MEMO, a common counseling summarization dataset with 191 counseling dialogues with 11,543 utterances (Srivastava et al., 2022). Each dyadic counseling dialogue contains a conversation between a therapist (5,722 utterances) and a client (5,814 utterances), along with an expert annotated counseling summary. In addition to this, MEMO further contains each utterance labeled with a counseling component – Symptom and History (SH), Patient Discovery (PD), Reflecting (RT), or Discus-

sion Filler (DF) (c.f. Appendix D for definitions). The dataset contains 2379 utterances tagged with SH, 5428 utterances tagged with PD, and 1242 utterances tagged with RT. These counseling components act as labels of relevance for utterances.

## 4 Proposed Methodology

Our proposed model, PIECE, operates on the MEMO dataset and aims to generate knowledge-enriched counseling summaries for a counseling dialogue containing  $n$  utterances,  $D = \langle u_1, u_2, u_3, \dots, u_n \rangle$ . There are two specific kinds of knowledge that PIECE focuses on – (a) (mental health) domain knowledge and (b) structural understanding of dialogue. The model achieves this knowledge integration through a novel *planning engine* that filters and fuses relevant information to guide the generation of knowledge-enriched summaries by the underlying LLM. Figure 2 presents a schematic diagram of PIECE.

### 4.1 Domain Knowledge Encapsulation

Knowledge encapsulation module involves *knowledge filtering* and *knowledge scaffolding*.

**Knowledge Filtering.** Knowledge filtering requires carefully identifying the most relevant utterances for crafting a knowledge-rich summary. This process initiates with classifying each utterance  $U_i$  into counseling components (SH, PD, RT, and DF). Among these, SH, PD, and RT are considered essential for generating comprehensive summaries, while DF utterances are regarded as non-essential and are therefore masked by leveraging the dedicated *Counseling Component Classifier*. This classifier is purposefully crafted to process the context of utterances,  $\langle u_1, u_2, u_3, \dots, u_i \rangle$ , through a GRU unit coupled with a self-attention block. Subsequently, a sequence of two dense layers learns these context-rich utterance representations to predict counseling components and filter non-essential discussion fillers.

The subsequent step involves examining the filtered utterances and further scrutinizing them using a mental health knowledge selection module, MH-Know. Similar to the approach opted by Srivastava et al. (2022) with the MEMO dataset, our module leverages the lexicon derived from the Patient Health Questionnaire (PHQ-9) to gauge the degree of mental health-specific content embedded within each utterance (Kroenke et al., 2001). The similarity scores ( $s_i$ ) are computed between

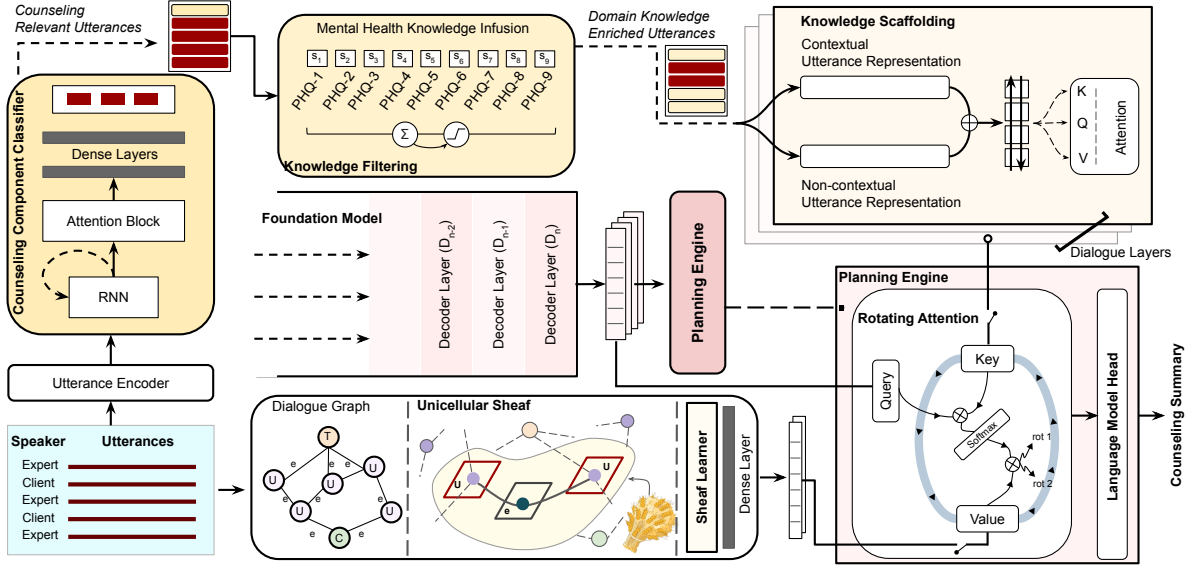


Figure 2: Architecture of PIECE. We propose a novel *planning engine* consisting of two primary sections: (a) integrating knowledge filtering-cum-scaffolding and (b) encapsulating structural understanding of dialogues. Filtration of relevant utterances utilizes component labeling counseling within the MEMO dataset to mask filler utterances, followed by knowledge scaffolding. Additionally, sheaf learners are employed for the structural understanding of counseling dialogue. The planning engine operates using a rotating attention mechanism using knowledge from both segments for better LLM generation.

individual utterances and the PHQ-9 lexicons. Subsequently, a predetermined threshold ( $T^\circ = 50\%$ , in our study) is applied to ascertain whether an utterance merits retention or should be masked. The resulting filtered dialogue, composed of essential and knowledge-rich utterances, then proceeds to the *knowledge scaffolding* stage to ensure optimal preservation of the embedded knowledge.

**Knowledge Scaffolding.** To ensure the cohesion and effective organization of the filtered knowledge segments, PIECE employs a knowledge scaffolding module. We adopt the scaffolding methodology proposed by Cohan et al. (2019), where the authors incorporated structural information of scientific papers into citations for citation-intent classification. We modify their knowledge scaffolding for dialogue settings by employing a mix of contextual and non-contextual representations for each relevant utterance, as shown in Equation 1.

$$R_k = \text{Attention}[\overleftarrow{\chi}(C(u_i) \oplus C^\circ(u_i))] \quad (1)$$

Here, the contextual embeddings  $C(x)$  are captured through BERT (Devlin et al., 2019), while the non-contextual embeddings  $C^\circ(x)$  are extracted using Glove (Pennington et al., 2014). These representations are then carefully concatenated and fed into a bidirectional LSTM layer  $\overleftarrow{\chi}$ , ensuring the

underlying semantic context. The resulting context-rich representations, enriched with domain-specific knowledge, subsequently serve as key, query, and value vectors within a self-attention block. This orchestrated interplay emerges in generating filtered scaffolding representations  $R_k$  enriched with the essence of extracted knowledge. Before understanding the role of these representations within the planning engine, we first understand the encapsulation of structural knowledge, another crucial component of PIECE’s knowledge-enriched summary generation pipeline.

## 4.2 Structural Knowledge Encapsulation

To ensure that the generated summaries accurately reflect the overall structure and flow of the original dialogue, PIECE incorporates a dedicated structural knowledge module. Sheaf has been long studied for its ability to hold structural knowledge in graphs via sheaf diffusion (Shepard, 1985; Curry, 2014; Bodnar et al., 2022). With a similar objective, we first construct a dialogue graph  $G = (V, E)$  that captures the intricate network of relationships between the utterances within the dialogue  $D$ . Each utterance  $u_i$  is represented as a vertex  $v \in V$  within the graph, while edges  $e \in E$  are directed toward the progression of dialogue. In the adjacency matrix  $A_d$  of the graph  $G$  for each dialogue instance

$d \in D$  with  $o$  utterances, we place each node’s feature  $f$  with its encoded utterance representations  $f \leftarrow bert(u_i)$ . As a result, we construct  $G$  to preserve the inherent dialogue structure while learning the graph representations via *sheaf learner*.

**Sheaf Learner.** Graphs inherit a notion of the neighborhood but not distance or direction. Sheaf theory provides a mathematical framework for imbuing graphs with geometric structure by associating vector spaces, known as stalks, with each node and edge, along with defining restriction maps (which are essentially linear transformations) between the stalks of incident nodes and edges. These restriction maps capture the local-to-global context, ensuring consistency of information across the graph. The resulting structure, comprising the graph along with the decorated stalks and restriction maps, is termed a cellular sheaf (Shepard, 1985). To facilitate the understanding of complex graph relationships, we construct a sheaf, which is a collection of interconnected cellular sheaves. Drawing inspiration from Bodnar et al. (2022), we employ a sheaf convolution network (*SCN*) (see Equation 2) to learn these sheaves.

$$R_{scn} = \eta((I - \Delta)(I \otimes W_1)A^{(o \times f)}W_2) \quad (2)$$

Here,  $A$  is the input to the network, whereas  $W_1$  and  $W_2$  are learnable weights.  $\Delta$  acts as a sheaf laplacian,  $I$  is an identity matrix, and  $\otimes$  represents direct matrix multiplication. We apply an activation function  $\eta$  on top as a *relu* function. For the resultant representation  $R_{scn}$ , we employ a dense layer to learn the graph geometry over topological space, inheriting better structural knowledge. The resultant structure-rich representations  $R_s$  act as an input to the planning engine.

### 4.3 Planning Engine

The planning engine is the core module that integrates the structural and domain-specific knowledge using a rotating attention mechanism. Acknowledging the equivalent significance of both structural and domain-specific insights, the rotating (cyclic) attention mechanism aims to retrieve information evenly from each segment. Here, domain-rich scaffolded representations  $R_k$  and structure-rich sheaf representations  $R_s$  act as both key ( $K$ ) and value ( $V$ ) once per cycle, yielding two rich dialogue representations. The query, the foundation model’s hidden representations, acts as query ( $Q$ )

for both cycles, as shown in Equation 3.

$$R_{rep} = sm\left(\frac{QK^T}{\sqrt{d_k}}\right)V \oplus sm\left(\frac{QV^T}{\sqrt{d_k}}\right)K \quad (3)$$

The language model head (LM-Head) operates on top of the fused representations of both dialogue representations to generate a counseling summary.

## 5 Experiments and Results

Here, we discuss the selection of baseline methods followed by the performance comparison with baselines, the ablation of PIECE, and analyses.

### 5.1 Baselines

We choose the following systems as our baselines – (i) **Segmented Modeling (SM)** leverages dialogue-acts on BiLSTM (Goo and Chen, 2018). (ii) **BART** is a standard transformer with a BERT encoder and GPT decoder (Lewis et al., 2019). (iii) **Pegasus** is pretrained with an objective of gap-sentence-generation (Zhang et al., 2020). (iv) **T5** uses a shared framework on transformer to pretrain on huge C4 corpus (Raffel et al., 2020). (v) **Pre-trained Language Model (PLM)** uses DialoGPT to segment topics and generate summaries with BART (Feng et al., 2021). (vi) **RankAE** uses an encoder to filter dialogue utterances into segments and generates summaries using denoising auto-encoder (Zou et al., 2021). (vii) **Summ<sup>N</sup>** is a language model pretrained for dialogue summarization (Zhang et al., 2022b). (viii) **DialogLM** is pretrained for long conversational summarization task (Zhong et al., 2022). (ix) **ConSum** is the reported state-of-the-art for counseling summarization task on the MEMO dataset, marking as the most relevant baseline method to compare with (Srivastava et al., 2022). (x) **Flan-T5** is a T5-based LLM, instruction-tuned on a mixture of tasks (Chung et al., 2022). (xi) **Mistral-7B** is an LLM instruction tuned with a sliding window attention mechanism for efficient and longer context (Jiang et al., 2023). (xii) **Zephyr-7B** is a fine-tuned Mistral LLM to surpass the large chat models (Tunstall et al., 2023). (xiii) **Llama-2-7B** is an optimized auto-regressive LLM currently state-of-the-art for many language generation tasks (Touvron et al., 2023). (xiv) **MentalLlama** is a standard Llama-2 model pretrained on a huge Reddit-based mental health corpus (Yang et al., 2023b). Our work is proposed on top of MetalLlama because of its domain-specific knowledge and performance.

Model	R-1	R-2	R-L	BS
SM (Goo and Chen, 2018)	20.46	3.80	18.87	-0.9454
BART (Lewis et al., 2019)	34.92	12.66	18.83	-0.7118
Pegasus (Zhang et al., 2020)	29.71	7.77	27.57	-0.6130
T5 (Raffel et al., 2020)	31.44	5.63	27.38	-0.5655
PLM (Feng et al., 2021)	34.24	11.19	33.35	-0.8678
RankAE Zou et al. (2021)	25.57	3.43	24.16	-1.0630
SUMM <sup>N</sup> (Zhang et al., 2022b)	34.06	11.32	20.99	-0.6088
DialogLM (Zhong et al., 2022)	28.14	9.21	17.57	-0.7377
ConSum (Srivastava et al., 2022)	45.36	15.71	24.75	0.3407
Flan-T5 (Chung et al., 2022)	41.30	16.00	29.05	0.2281
Mistral-7B (Chung et al., 2022)	46.74	14.98	32.48	0.3056
Zephyr-7B (Tunstall et al., 2023)	38.90	9.90	26.35	0.0094
Llama-7B (Touvron et al., 2023)	47.22	16.74	33.16	<b>0.4106</b>
MentalLlama (Yang et al., 2023b)	47.92	16.90	35.81	0.3953
<b>PIECE (<math>f_m</math>: MentalLlama)</b>	<b>49.62</b>	<b>18.61</b>	<b>38.10</b>	0.4102
$\Delta_{PIECE-BEST}(\%)$	$\uparrow 3.42$	$\uparrow 10.11$	$\uparrow 6.01$	$\downarrow 0.09$

Table 1: Results obtained on the MEMO counseling summarization dataset. We report Rouge-1 (R-1), Rouge-2 (R-2), Rouge-L (R-L), and Bleu Score (BS).

To evaluate the performance, we employ widely recognized metrics: ROUGE and Bleu Score.

## 5.2 Performance Comparison

Table 1 shows the performance of the baseline models, revealing the clear superiority of PIECE across three out of four metrics. Notably, MentalLlama turns out to be the best-performing baseline. PIECE excels in capturing both semantic and syntactic structures, as evidenced by the improvements of best-performing LLMs like Llama variants. Specifically, PIECE demonstrates improvements of +3.42%, +10.11%, and +6.01% in R-1, R-2, and R-L points, respectively. On the other hand, Llama yields the best scores for the Bleu metric; however, PIECE closely matches its performance, exhibiting a marginal drop of only  $-0.0004$  Bleu points. However, compared with ConSum, the benchmarked state-of-the-art model for counseling summarization, our findings reveal a significant boost in performance metrics. Notably, PIECE shows improvements of +9.39%, +18.45%, +53.93% and +20.39% in R-1, R-2, R-L and Bleu points, respectively.

**Generalizability.** MEMO is the only publicly available dataset that fits into the definition of counseling summarization, to the best of our knowledge. However, we further add experiments on another clinical note generation dataset, ACI-BENCH (Yim et al., 2023), as shown in Appendix (§Table 6).

## 5.3 Ablation Study

We perform an ablation study to assess the performance of underlying components contributing to PIECE. By systematically deconstructing the plan-

Ablations	R-1	R-2	R-L	BS
PIECE				
- StructKnow	48.25 ( $\downarrow 1.37$ )	16.22 ( $\downarrow 2.39$ )	37.40 ( $\downarrow 0.70$ )	0.4085 ( $\downarrow 0.0017$ )
- DomainKnow	48.36 ( $\downarrow 1.26$ )	15.78 ( $\downarrow 2.83$ )	36.96 ( $\downarrow 1.14$ )	0.3994 ( $\downarrow 0.0108$ )
PIECE ( $f_m = \alpha$ )				
$\alpha$ : +Mistral	46.81 ( $\uparrow 0.07$ )	13.28 ( $\downarrow 1.70$ )	34.52 ( $\uparrow 2.04$ )	0.3080 ( $\uparrow 0.0024$ )
$\alpha$ : +Zephyr	41.63 ( $\uparrow 2.73$ )	11.36 ( $\uparrow 1.46$ )	27.94 ( $\uparrow 1.59$ )	0.0106 ( $\uparrow 0.0012$ )
$\alpha$ : +Llama	48.25 ( $\uparrow 1.03$ )	14.02 ( $\downarrow 2.72$ )	35.88 ( $\uparrow 2.72$ )	0.4098 ( $\downarrow 0.0008$ )

Table 2: Ablation study of the proposed model, PIECE and generalizability of the planning engine on top of notable LLMs as foundation model ( $\alpha$ ), including Mistral, Zephyr, and Llama, illustrating a clear **improvement** ( $\uparrow$ ) in LLM generations by integrating planning engine.

ning engine and analyzing various elements of our model architecture, we present our findings in Table 2. The impact of the planning engine on the acting foundation models ( $f_m$ ) assesses the generalizability of the planning engine. Evidently, the role of the addition of knowledge is highlighted as we observe a clear decline in the performance of PIECE across all four metrics (c.f. Appendix A for comparative study). We observe the same trend for both structural and domain knowledge, with the performance declining by a significant margin of  $-2.83$  (R-2) points. At the same time, we perform a generalizability check for the planning engine to be inducted on top of multiple LLMs such as Mistral, Zephyr, and Llama, finding that each LLM, in general, benefits from the planning engine.

## 5.4 Qualitative Analysis

We assess PIECE’s text generation capabilities, apart from quantitative metrics, via detailed comparative analysis between PIECE and the two most relevant models: Llama and MentalLlama. Given that PIECE builds upon the MentalLlama, we carefully compare its performance against these two models and the gold standard. Table 3 presents an illustrative comparison. Our scrutiny centers on discerning the understanding of core counseling components, including discussions on *symptoms and reasons*, *reflecting client’s story*, and *critical routine details* of the dialogue along with the *additional domain understanding*.

In Table 3, the first row showcases the gold standard summary. Evidently, Llama excels in text coherence but falls short in producing domain-specific insights and barely touches upon the critical domain information, for example, *client’s referral from GP*, *self-emotional-assessment rating*, and *declaration of anxiety*. In contrast, MentalLlama demonstrates better mental health-specific nuances, highlighting counseling components such as *battling depression and loneliness*, *the burden*

Model	Summaries
Gold	Patient visited the therapist for feeling bad. Therapist knew that patient had such feelings in the past, but this time patient found it harder. Two months ago patient visited GP after that GP referred to therapist. Earlier patient had a breakup and felt lonely. Therapist asks how things have been since the patient went to GP. Patient feels nothing changed and unable to accept failure. Patient's mom was not happy with the patient's behavior. Therapist suggested CBT. Patients accept that the first year was tough. Therapist asks that it doesn't sound like you want to go back. Patient wants to but just doesn't feel like meeting the man. Patient used to sit in a friend's restaurant for money and stuff in the meantime. Patient used to like it because the patient used to speak to people. Therapist asks to scale the thoughts out of 100. Patient says 80. Patient feels empty when asked about negative thoughts. Patient doesn't want her mom to be too worried.
Llama	Patient adrift after a <b>brutal breakup</b> finds solace in a therapist's office. The weight of past struggles, a university hiatus, and crushing expectations press down, leaving her numb and weary. <b>Work, once at a restaurant, becomes a reminder of self-doubt.</b> But as secrets unfurl – <b>anxieties, fears, and the facade she wears for her mom</b> – a glimmer of hope flickers. Maybe, just maybe, this conversation, this courageous first step, can illuminate a path back to wholeness.
MentalLlama	Patient, <b>battling depression and loneliness</b> , seeks therapeutic guidance. Past <b>heartbreak</b> and struggles weigh heavily, <b>leaving her drained and empty.</b> <b>Work offers a temporary escape</b> , but anxieties about "not being good enough" resurface. The therapist recognizes the patient's inner strength and the <b>burden of hiding pain from her mother.</b> Through the conversation, a flicker of hope emerges, hinting at the possibility of a different path forward rooted in understanding and compassion.
PIECE	The patient visited the therapist feeling unable to recover from recent emotional difficulties, notably following a <b>breakup with a long-term boyfriend.</b> Expressing disappointment in seeking therapy and <b>reluctance to admit feelings of failure after being referred by the GP,</b> the <b>therapist highlighted the normalcy of anxiety.</b> <b>Concerns about returning to university after a break were discussed,</b> reflecting the patient's desire for success amidst <b>overwhelming pressure and self-doubt.</b> Additionally, feelings of inadequacy arose during a recent <b>work shift at a friend's restaurant, causing the patient to feel good but confused in talking to people.</b> The therapist further explored the <b>patient's sense of guilt for not meeting expectations, especially with mom,</b> revealing a deep emotional impact. The <b>patient admitted to feeling empty and drained, rating their emotional distress at 80 out of 100 and describing physical sensations resembling nervousness.</b> The therapist understood the <b>patient's struggle to appear unaffected for their mother's sake,</b> despite internal turmoil, ultimately feeling worse for doing so.

Table 3: A comparative analysis of three most relevant models – Llama, MentalLlama, and PIECE. The key emphasis on being analyzed here is understanding the core counseling components, including discussions on *symptoms and reasons*, *reflecting client's story*, and *critical routine details* along with the *additional intricate domain understanding*. While Llama generates general-purpose summaries, MentalLlama, being pretrained on mental health data, captures nuanced knowledge beyond Llama's scope. In contrast, PIECE is able to capture in-depth domain knowledge, surpassing other models in touching upon counseling nuances. Despite the grammatical proficiency of baselines, PIECE stands out for its structural understanding, focusing on intricate details, in some cases, better than gold.

of hiding pain, and past heartbreaks and integrating intricate domain information such as *struggles weigh heavy and the possibility of a different path forward*. Apparently, MentalLlama's performance lacks a grasp of the structural intricacies of dialogue. On the other hand, the summaries by PIECE are descriptive, emphasizing the complete dialogue structure and crucial counseling components inherited from the conversation such as *therapist highlighting normalcy, self-assessment, reluctance to admit feelings*, instances which MentalLlama and Llama skipped. Additionally, PIECE excels in enhancing domain understanding by incorporating spans from the LLM's vocabulary rather than the conversation's vocabulary.

**Error Analysis:** We extend the analyses for cases where our model falls short in capturing the intended details (c.f. Appendix C). We discuss two

Model	Relevance	Consistency	Fluency	Coherence
Llama	3.12	3.22	<b>3.76</b>	3.68
MentalLlama	3.57	3.31	3.75	3.71
PIECE	<b>3.73</b>	<b>3.39</b>	3.72	<b>3.75</b>

Table 4: Human evaluation on the summaries generated from PIECE model. The average interrater's agreement score ( $\kappa$ ) for PIECE is 0.82.

important cases: a) PIECE tends to include exaggerated information in shorter dialogues, leading to extra details, and b) PIECE summaries are usually longer than gold summaries and, on average, struggle to capture patient's behavior.

**Human Evaluation** We present human evaluation on four standard linguistic parameters, namely *relevance* (selection of relevant content), *consistency* (factual alignment between the summary and the source), *fluency* (linguistic quality of each sentence), and *coherence* (structure and organization of summary) as shown in Table 4. We employed 12 linguistics experts to rate each parameter on the Likert scale of 1 to 5. Out of twelve, seven were female, whereas five were male, all of them aged between 23 - 35. As shown in Table 4, PIECE surpasses across three out of four parameters. Notably, it achieved a score of 3.73 for relevance, indicating that PIECE's summaries capture the core knowledge. Additionally, PIECE scored 3.39 and 3.75 for consistency and coherence, respectively, showing the logical flow of the original conversation. Finally, PIECE competes with *Llama* on *fluency* metric. The superior performance underscores the model's linguistic quality and structural coherence. The average Cohen's kappa score ( $\kappa$ ) for PIECE is

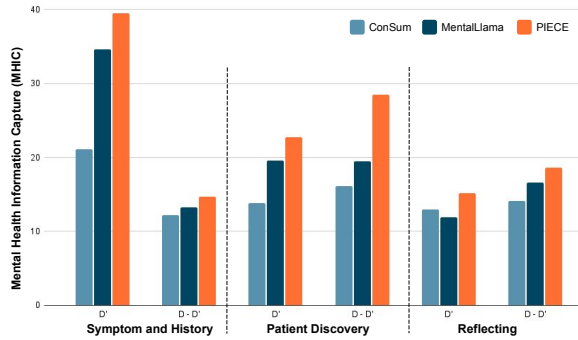


Figure 3: Domain-centric evaluation using Mental Health Info Capture (MHIC) metric. The proposed model, PIECE, distinctly excels in capturing domain knowledge compared to the two most relevant models.

0.82, which falls under the substantial category.

### 5.5 Mental Health Information Capture

The Mental Health Information Capture (MHIC) computes the intersection of utterances predicted by the counseling components classifier and the generated summary using the ROUGE-1 score to provide a qualitative evaluation of these summaries. Figure 3 presents a comparative metric study of the three top-performing models – ConSum, MentalLlama, and PIECE. Evidently, the earlier state-of-the-art benchmarked model, ConSum, on the MEMO dataset appears to be easily surpassed by MentalLlama without any additional effort. However, after integrating the *planning engine* to MentalLlama, i.e., PIECE, performance surpassed both variants, highlighting the efficacy of PIECE.

### 5.6 Mental Health Expert Validation

Our model, PIECE acts as an assistance for complex counseling processes, as experts can utilize the summaries by PIECE as counseling notes, reducing both time and cost. To validate the generation’s effectiveness, we collaborated with a clinical psychologist (having 10+ years of clinical experience) who assessed the summaries on two major fronts: clinical acceptability and LLM relevance.

**Clinical Acceptability.** Experts exploit a clinical acceptability framework (Sekhon et al., 2017) that evaluates parameters such as *affective attitude*, *burden*, *ethicality*, *coherence*, *opportunity costs*, and *perceived effectiveness* (c.f. Appendix E.1 for definitions). The expert rated summaries on a scale from 0 to 2, with higher scores indicating better acceptability (c.f. Table 5). The expert evaluation shows an average rating of 1.20 out of 2.00, falling well within the standard acceptability range of 0.70

	Af. Att.	Burden	Ethic.	Intv. Coh.	Cost	Percv. Eff.
$\mu$	1.23	0.36	1.03	1.20	0.85	1.27
$(\Delta)$	(0.21)	(0.34)	(0.23)	(0.33)	(0.19)	(0.27)

Table 5: Expert evaluation across six domain-centric metrics: *affective attitude* (Af. Att.), *burden*, *ethicality* (Ethic), *intervention coherence* (Intv. Coh.), *opportunity cost*, and *perceived effectiveness* (Percv. Eff.). We present mean ( $\mu$ ) and standard deviation ( $\Delta$ ) scores.

to 1.40. As a result, the expert concluded that the generated summaries demonstrate suitability for therapists, sometimes surpassing the gold standard.

**LLM Relevance.** To assess common LLM flaws, quality of generation, and relevance to experts, we ask experts to address the following aspects:

- Did the expert observe our model hallucinate?
- How does the model’s generation compare to that of the most competitive model (MentalLlama)?
- Is the generated summary relevant to the expert?

Evidently, 75% of instances having ‘negligible’ hallucinations in PIECE and 56.3% cases outperformed the most competitive model, MentalLlama. Consequently, 93.8% of cases were deemed either ‘totally relevant’ or ‘relevant to some extent’ for practical applications by experts. However, with 12.25% of instances marked *hallucinated*, there remains room for improvement in the future.

**Impact.** The end goal of this study is to prepare PIECE for a pilot study for experts via an assistive framework that can help professionals streamline note-taking. Subsequently, this will allow more focus on patient care and improve time efficiency.

## 6 Conclusion

In our work, we explored the research space of mental health counseling summarization. Our research focused on strategically orchestrating LLMs to plan before generating content, thereby incorporating relevant dimensions of conversational structure- and domain-centric knowledge into the summarization process. We proposed a novel *planning engine* for LLMs, integrated and presented in this research with MentalLlama as PIECE. We compared the performance of PIECE against a spectrum of 14 potential baseline methods, including state-of-the-art LLMs like Mistral, Llama-2, Flan-T5, Zephyr, and MentalLlama. Next, we presented an extensive evaluation spanning automatic, human, and expert evaluations. Expert analysis indicated that the summaries generated by PIECE exhibit better relevance and structure, often surpassing the closest



competitive LLMs and, in certain instances, even surpassing the gold standard. We conclude by discussing our research’s ethical considerations and applicability, emphasizing our approach as an assistive module tailored exclusively for mental health experts, which mitigates potential risks associated with direct client interactions, ensuring a safeguard against unintended harm. This strategic positioning underscores our commitment to safe utilization in mental health contexts.

## 7 Limitations

Research in the space of mental health counseling summarization, in general, poses several complex challenges. Firstly, there’s a substantial scarcity of diverse and high-quality datasets tailored for counseling summarization. Our research predominantly relies on the MEMO dataset. To the best of our knowledge, it is the only publicly available dataset. The limited dataset diversity hinders robustness, emphasizing the critical need for more datasets in this research area. Another challenge arises from the scale of the LLMs utilized in our research. Pretraining these super-large models on domain-specific corpora incurs significant costs and environmental implications. MentalLlama, chosen as the foundation model due to its pretraining corpus, underscores the necessity for more domain-centric models specifically tailored for mental health counseling. Additionally, evaluating LLM generations in a sensitive domain like mental health necessitates a nuanced understanding of the domain’s intricacies. LLMs, by nature, exhibit tendencies to generate information that might not accurately represent the domain context and could potentially "hallucinate". Hence, our research, on one end, reduces hallucination by planning the generation and, on the other end, is positioned as an assistive tool tailored for mental health experts rather than direct client-facing applications. This approach mitigates the risks associated with LLM-generated content and underscores the need for expert oversight and intervention in sensitive mental health contexts.

## 8 Ethical Considerations and Future Work

Our research is an augmentation of established state-of-the-art foundation models, focusing on enhancing the generational proficiency of Language Models (LLMs) via planning LLM generation rather than introducing LLMs from scratch. This

approach ensures continuous refinement, specifically targeting improvements in mental health counseling summarization within existing models. **Given the sensitive nature of research in the domain of mental health, we portray our research as an assistive module designed exclusively for mental health experts. Such a strategic approach mitigates risks associated with direct support-seeking client interactions, safeguarding against unintended impacts on client sentiments or mental states.** Mental health professionals retain autonomy, choosing whether to accept PIECE’s generated summary or further tailor it based on their expertise. However, our research scope acknowledges the evolving nature of complexities and diversity, emphasizing the need for not merely larger but planned LLMs. This leads to aligning LLMs with optimal parameter sizes and prioritizing the planned generational capabilities. Hence, paving the way for more responsible LLMs, upholding ethical standards in the research for sensitive domains like mental health counseling.

## Acknowledgement

The authors acknowledge the support of ihub-Anubhuti-iiitd Foundation and Infosys foundation through CAI at IIT-Delhi.

## References

- Yash Atri, Arun Iyer, Tanmoy Chakraborty, and Vikram Goyal. 2023a. [Promoting topic coherence and inter-document consorts in multi-document summarization via simplicial complex and sheaf graph](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2154–2166, Singapore. Association for Computational Linguistics.
- Yash Kumar Atri, Vikram Goyal, and Tanmoy Chakraborty. 2023b. [Fusing multimodal signals on hyper-complex space for extreme abstractive text summarization \(tl;dr\) of scientific contents](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD ’23*, page 3724–3736, New York, NY, USA. Association for Computing Machinery.
- Yash Kumar Atri, Vikram Goyal, and Tanmoy Chakraborty. 2023c. [Multi-document summarization using selective attention span and reinforcement learning](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:3457–3467.
- Cristian Bodnar, Francesco Di Giovanni, Benjamin Paul Chamberlain, Pietro Lio, and Michael M. Bronstein.

2022. [Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs](#). In *Advances in Neural Information Processing Systems*.
- M. Cascella, J. Montomoli, and V. et al. Bellini. 2023. [Evaluating the Feasibility of ChatGPT in Healthcare: An Analysis of Multiple Clinical and Research Scenarios](#). *J Med Syst*.
- Minje Choi, Jiaxin Pei, Sagar Kumar, Chang Shu, and David Jurgens. 2023. [Do LLMs understand social knowledge? evaluating the sociability of large language models with SockKET benchmark](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11370–11403, Singapore. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).
- Arman Cohan, Waleed Ammar, Madeleine van Zuylen, and Field Cady. 2019. [Structural scaffolds for citation intent classification in scientific publications](#). 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 3586–3596, Minneapolis, Minnesota. Association for Computational Linguistics.
- Justin Curry. 2014. [Sheaves, cosheaves and applications](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Xiachong Feng, Xiaocheng Feng, Libo Qin, Bing Qin, and Ting Liu. 2021. [Language model as an annotator: Exploring DialoGPT for dialogue summarization](#). In *Proceedings of the 59th ACL and the 11th IJCNLP*, Online.
- Zhangyin Feng, Weitao Ma, Weijiang Yu, Lei Huang, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting liu. 2023. [Trends in integration of knowledge and large language models: A survey and taxonomy of methods, benchmarks, and applications](#).
- Chih-Wen Goo and Yun-Nung Chen. 2018. [Abstractive dialogue summarization with sentence-gated modeling optimized by dialogue acts](#). *SLT*.
- Alastair C. Heerden, Julia R. Pozuelo, and Brandon A. Kohrt. 2023. [Global Mental Health Services and the Impact of Artificial Intelligence–Powered Large Language Models](#). *JAMA Psychiatry*, 80(7):662–664.
- Raghav Jain, Tulika Saha, Jhagrut Lalwani, and Sriparna Saha. 2023. [Can you summarize my learnings? towards perspective-based educational dialogue summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3158–3173, Singapore. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#).
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. [How can we know when language models know? on the calibration of language models for question answering](#). *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know? Transactions of the Association for Computational Linguistics](#), 8:423–438.
- Kurt Kroenke, Robert L Spitzer, and Janet B W Williams. 2001. [The PHQ-9](#). *Journal of General Internal Medicine*, 16(9).
- Sandeep Kumar, Guneet Singh Kohli, Tirthankar Ghosal, and Asif Ekbal. 2023. [Mup-scidocsum: Leveraging multi-perspective peer review summaries for scientific document summarization](#). In *Leveraging Generative Intelligence in Digital Libraries: Towards Human-Machine Collaboration*, pages 250–267, Singapore. Springer Nature Singapore.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#).
- Siyang Liu, Naihao Deng, Sahand Sabour, Yilin Jia, Minlie Huang, and Rada Mihalcea. 2023. [Task-adaptive tokenization: Enhancing long-form text generation efficacy in mental health and beyond](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15264–15281, Singapore. Association for Computational Linguistics.
- Kshitij Mishra, Priyanshu Priya, Manisha Burja, and Asif Ekbal. 2023a. [e-THERAPIST: I suggest you to cultivate a mindset of positivity and nurture uplifting thoughts](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13952–13967, Singapore. Association for Computational Linguistics.

- Kshitij Mishra, Priyanshu Priya, and Asif Ekbal. 2023b. [Help me heal: A reinforced polite and empathetic mental health and legal counseling dialogue system for crime victims](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(12):14408–14416.
- Kshitij Mishra, Priyanshu Priya, and Asif Ekbal. 2023c. [PAL to lend a helping hand: Towards building an emotion adaptive polite and empathetic counseling conversational agent](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12254–12271, Toronto, Canada. Association for Computational Linguistics.
- Sourajit Mukherjee, Adam Jatowt, Raghendra Kumar, Anubhav Jangra, and Sriparna Saha. 2023. [Can multimodal pointer generator transformers produce topically relevant summaries?](#) In *2023 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Priyanshu Priya, Kshitij Mishra, Palak Totala, and Asif Ekbal. 2023. [Partner: A persuasive mental health and legal counselling dialogue system for women and children crime victims](#). In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 6183–6191. International Joint Conferences on Artificial Intelligence Organization. AI for Good.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*.
- Tulika Saha, Vaibhav Gakhreja, Anindya Sundar Das, Souhitya Chakraborty, and Sriparna Saha. 2022a. [Towards motivational and empathetic response generation in online mental health support](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, page 2650–2656, New York, NY, USA. Association for Computing Machinery.
- Tulika Saha, Saichethan Reddy, Anindya Das, Sriparna Saha, and Pushpak Bhattacharyya. 2022b. [A shoulder to cry on: Towards a motivational virtual assistant for assuaging mental agony](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2436–2449, Seattle, United States. Association for Computational Linguistics.
- Mandeep Sekhon, Martin Cartwright, and Jill J. Francis. 2017. Acceptability of healthcare interventions: an overview of reviews and development of a theoretical framework. *BMC Health Services Research*, 17(1).
- Ashish Sharma, Inna W. Lin, Adam S. Miner, David C. Atkins, and Tim Althoff. 2021. [Towards facilitating empathic conversations in online mental health support: A reinforcement learning approach](#). In *Proceedings of the Web Conference 2021, WWW '21*, page 194–205, New York, NY, USA. Association for Computing Machinery.
- Ashish Sharma, Inna W. Lin, Adam S. Miner, David C. Atkins, and Tim Althoff. 2023. [Human-ai collaboration enables more empathic conversations in text-based peer-to-peer mental health support](#). *Nature Machine Intelligence*.
- Ashish Sharma, Koustav Rudra, and Niloy Ganguly. 2019. [Going beyond content richness: Verified information aware summarization of crisis-related microblogs](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 921–930, New York, NY, USA. Association for Computing Machinery.
- A. Shepard. 1985. [A cellular description of the derived category of a stratified space](#). In *Brown University*.
- K. Singhal, S. Azizi, and T. Tu. 2023. Large language models encode clinical knowledge. *Nature*.
- Aseem Srivastava, Ishan Pandey, Md Shad Akhtar, and Tanmoy Chakraborty. 2023. [Response-act guided reinforced dialogue generation for mental health counseling](#). In *Proceedings of the ACM Web Conference 2023, WWW '23*, page 1118–1129, New York, NY, USA. Association for Computing Machinery.
- Aseem Srivastava, Tharun Suresh, Sarah P. Lord, Md Shad Akhtar, and Tanmoy Chakraborty. 2022. [Counseling summarization using mental health knowledge guided utterance filtering](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*, page 3920–3930, New York, NY, USA. Association for Computing Machinery.
- A.J. Thirunavukarasu, D.S.J. Ting, and K. Elangovan. 2023. Large language models in medicine. *Nature Medicine*.
- Abhisek Tiwari, Anisha Saha, Sriparna Saha, Pushpak Bhattacharyya, and Minakshi Dhar. 2023. [Experience and evidence are the eyes of an excellent summarizer! towards knowledge infused multi-modal clinical conversation summarization](#). In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23*, page 2452–2461, New York, NY, USA. Association for Computing Machinery.
- Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude

- Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Cl  mentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct distillation of lm alignment](#).
- Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. 2023. [On the planning abilities of large language models - a critical investigation](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Xinyi Wang, Lucas Caccia, Oleksiy Ostapenko, Xingdi Yuan, William Yang Wang, and Alessandro Sordoni. 2024. [Guiding language model math reasoning with planning tokens](#).
- Kailai Yang, Shaoxiong Ji, Tianlin Zhang, Qianqian Xie, Ziyang Kuang, and Sophia Ananiadou. 2023a. [Towards interpretable mental health analysis with large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6056–6077, Singapore. Association for Computational Linguistics.
- Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, Sophia Ananiadou, and Jimin Huang. 2023b. [Mentallama: Interpretable mental health analysis on social media with large language models](#).
- Wen-wai Yim, Yajuan Fu, Asma Ben Abacha, Neal Snider, Thomas Lin, and Meliha Yetisgen. 2023. [Acibench: a novel ambient clinical intelligence dataset for benchmarking automatic visit note generation](#). *Nature Scientific Data*.
- Wenhao Yu, Chenguang Zhu, Zaitang Li, Zhiting Hu, Qingyun Wang, Heng Ji, and Meng Jiang. 2022. [A survey of knowledge-enhanced text generation](#). *ACM Computing Survey (CSUR)*.
- Haichao Zhang, Wei Xu, and Haonan Yu. 2022a. [Generative planning for temporally coordinated exploration in reinforcement learning](#). In *International Conference on Learning Representations*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020. [PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization](#). *arXiv:1912.08777*.
- Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B. Tenenbaum, and Chuang Gan. 2023. [Planning with large language models for code generation](#). In *The Eleventh International Conference on Learning Representations*.
- Yusen Zhang, Ansong Ni, Ziming Mao, Chen Henry Wu, Chenguang Zhu, Budhaditya Deb, Ahmed Awadallah, Dragomir Radev, and Rui Zhang. 2022b. [Summ<sup>n</sup>: A multi-stage summarization framework for long input dialogues and documents](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1592–1604, Dublin, Ireland. Association for Computational Linguistics.
- Ming Zhong, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2022. [Dialoglm: Pre-trained model for long dialogue understanding and summarization](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):11765–11773.
- Yicheng Zou, Jun Lin, Lujun Zhao, Yangyang Kang, Zhuoren Jiang, Changlong Sun, Qi Zhang, Xuanjing Huang, and Xiaozhong Liu. 2021. [Unsupervised summarization for chat logs with topic-oriented ranking and context-aware auto-encoders](#). *Proceedings of the AAAI Conference*, 35(16).

## A Discussion on Structural vs Domain Knowledge

In this study, we underscore two fundamental sources of knowledge: understanding conversational structure and integrating domain-specific knowledge. This arrangement helps cover the entire context precisely and, at the same time, adds domain knowledge to the LLM generations. In our proposed method, PIECE, we incorporate structural information through sheaf learning and infuse domain knowledge through knowledge filtering-cum-scaffolding.

**Sheaf Learners.** To capture the entire structure of the dialogue, we initiate the construction of a dialogue graph. This initial step required the conversion of dialogue interactions into a graph representation, facilitating the subsequent embedding process for LLM. Graph representations are natively considered to be generated via Graph Neural Networks (GNNs); however, their performance diminishes notably in heterophilic settings (similar neighbor nodes). We apply a vector space structure of conversational graphs called sheaves to introduce the notion of geometry in the structural understanding of conversation. Sheaf learners, as discussed in Section 4.2, hold structural context by learning each sheaf in an end-to-end fashion by learning the associated vector space with sheaf to understand geometry (*aka* conversational structure).

Yet, an intriguing question emerges – *Can sheaf learners also encapsulate domain knowledge?* The answer pivots on the breadth of knowledge required. Unlike conventional scenarios, our research explores both knowledge filtering and the integration of external domain-specific knowledge, necessitating the introduction of scaffolding.

**Knowledge Scaffolding.** We modified the classical scaffolding method introduced by Cohan et al. (2019) and refined this method to align with conversational understanding intricately. Our objective revolves around capturing filtered nuances and structuring them to provide domain-rich input to our LLM planner. This process ensures knowledge-rich delivery of information to the planner rather than fragmented knowledge post-filtering as done by Srivastava et al. (2022). This adaptation allows for the strategic incorporation of domain-specific knowledge, enhancing the LLM’s planning process and enabling a more holistic understanding of the

Model	R-1	R-2	R-L	BS
Llama	46.22	27.32	32.85	41.70
MentalLlama	48.90	28.16	30.44	41.00
PIECE ( <i>f<sub>m</sub></i> : MentalLlama)	48.05	28.91	33.27	42.76
$\Delta_{PIECE-BEST}(\%)$	↓ 1.73	↑ 2.59	↑ 1.27	↑ 2.54

Table 6: Results obtained on the ACI-BENCH clinical note generation dataset to assess the generalizability of PIECE. We report Rouge-1 (R-1), Rouge-2 (R-2), Rouge-L (R-L), and Bleurt Score (BS).

conversational context.

We utilized the pre-set split ratio of 70:20:10 for the train, validation, and test set.

## B Generalizability

It is worth noting that to the best of our knowledge, MEMO is the only publicly available dataset that directly fits into the definition of counseling summarization. However, in order to assess the generalizability of PIECE, we further add experiments on another clinical note generation dataset - the ACI-BENCH dataset. We present the results in Table 6, evaluated across the same set of metrics (R1/2/L + BS) as presented in Table 1. Our findings show that the proposed model surpasses the Llama model on the R1 metric by +3.95% but shows a minute decline (0.85 R1 pts) when compared with MentalLama. On the other hand, our model surpasses Llama by +5.81%, +1.27%, and +2.54% on R2, RL, and BS metrics, respectively. At the same time, our model surpasses MentalLama by +2.59%, +9.29%, and +4.29% on R2, RL, and BS metrics, respectively.

## C Error Analysis

We present our findings from a meticulous assessment of a sample of outputs using our proposed model, PIECE.

**Finding 1:** While the patient’s overwhelming self-doubt and anxiety are mentioned, the specific details about their job dissatisfaction, concerns about their children, and intervention with CBT techniques are not adequately captured. Evidently, our model tends to include exaggerated information from the domain in shorter dialogues. In the case of Case 1, where only 20 utterances were present, the model included additional information beyond what was necessary. This highlights a trade-off between incorporating external information and

	Gold Summaries	PIECE Summaries
1	The patient is fed up with the job and is feeling anxious and fearful about taking responsibility. The patient’s kids are doing drugs. The patient is conscious about kids. The patient prepared a report. Therapist provides a routine of CBT after reading report. The routine includes evaluating evidence of facts and carefully reviewing the facts, then resolving it. The patient finds no time for family and is anxious about kids. Therapist understood the negative thoughts of the patient. The therapist made the patient understand making a report is not going to work on kids, so they suggested making a good decision for welfare.	The patient expresses overwhelming self-doubt and anxiety in a recent professional situation. They question their ability to handle responsibility and fear the potential negative outcomes. Despite concerns, they also express reluctance to change jobs. The counselor employs CBT techniques to help the patient, highlighting the positive impact of their intervention. However, the patient continues to struggle with guilt. Ultimately, their anxiety stems from the power and responsibility of job. The counselor works to address these concerns, emphasizing professional responsibility.
2	The patient almost quit smoking. The patient got late for a gig because he forgot to look at his to-do list. The therapist told patient to stick it on his door, so whenever the patient opens the door, he can look at his to-do list. The therapist reminded patient about the application for school. The patient started writing on a calendar to remember.	The patient reports significant progress in quitting smoking, noting three weeks without a cigarette and expressing improvements in taste and breathing. However, they mention a setback in their job search, having visited a restaurant during dinnertime for a potential gig, which they acknowledge as a short-sighted decision. The counselor offers practical advice, suggesting strategies to improve organization and memory, such as using reminders and checklists. They also discuss progress in the patient’s research for schooling and encourage them to set application deadlines. The patient expresses gratitude for the support and suggests celebrating their success in quitting smoking with a treat for the counselor as well. Overall, the session focuses on acknowledging achievements, addressing setbacks, and providing tools for success in various aspects of the patient’s life.

Table 7: Error Analysis: Comparison of the gold summaries with the PIECE summaries, highlighting key shortcomings. The table highlights two specific cases where the model tends to include exaggerated information in shorter dialogues and struggles to capture specific details and the patient’s intrinsic traits, such as emotions and sentiments.

maintaining the relevance of generated summaries, marking it as a potential area for future study.

**Finding 2:** In another case (Case 2), we observed the disparity in length between the gold and PIECE summaries. This is generally not an issue but we further observe that PIECE summarizes the dialogue with a length that doubles the brevity of the former (gold) and still conveys the same essential information. Furthermore, while appropriateness is generally maintained across both gold and PIECE summaries, there persists a notable deficiency in encapsulating the intrinsic traits of the patient, including emotions, tone, and sentiments, which are typically discerned and incorporated by therapists.

The error analysis highlights the conditional shortcomings in the PIECE summarization framework. We touched upon a few of these topics in Section 5.6 and Section 5.7; however, we understand the importance of explicitly including it in the paper. Hence, we commit to include a separate section on ‘error analysis’ in the camera-ready.

## D MEMO Dataset

Drawing from the MEMO dataset source, we present the label definitions below:

- Symptom and History (SH): Captures utterances that provide the most insightful information for the therapist to assess the patient’s situation.
- Patient Discovery (PD): In counseling ses-

sions, patients often present with complex thoughts. The therapist builds a therapeutic relationship to help patients unravel these thoughts.

- Reflecting (RT): Therapist utterances are often concise to allow patients ample space to express themselves. Therapists may use imaginary scenarios to help understand the patient’s perspective.
- Discussion Filler (DF): These are peripheral utterances in the conversation, such as pleasantries (‘Good morning!’), non-lexical fillers (‘Ummm’), acknowledgments (‘Right’), and restatements of affirmations (‘Yeah. Yeah’). These utterances carry little to no relevance in summary generation.

## E Experimental Setup

### E.1 Expert Evaluation Metrics

Experts followed an established framework as discussed in Section 5.6. The evaluation criteria involve the following metrics taken directly from the framework:

- Affective Attitude: This dimension assesses how individuals feel about an intervention based on their clinical knowledge and perceptions. It examines subjective feelings toward the intervention.
- Burden: The burden dimension measures the perceived effort required to participate in or understand the intervention. It considers factors such

as spelling, grammar, and overall interpretation, reflecting the ease of use.

- **Ethicality:** This dimension evaluates the extent to which an intervention aligns with an organization’s value system and code of ethics. It explores whether there are any ethical concerns or conflicts with established ethical guidelines.
- **Intervention Coherence:** Here, the focus is on understanding the intervention itself. It measures how well the intervention is comprehended, highlighting the clarity and coherence of the intervention’s content or instructions.
- **Opportunity Cost:** Opportunity cost assesses the benefits and drawbacks of using a particular intervention in a given context. It considers the potential gains and losses associated with adopting the intervention.
- **Perceived Effectiveness:** This dimension evaluates how well the intervention is expected to perform in its intended setting. It assesses the perceived effectiveness of the intervention based on the expectations and experiences of individuals involved in its implementation.

## E.2 Quantitative Evaluation Metrics

We employ ROUGE and Bleu scores. The computation of ROUGE scores is carried out using the *py-rouge* library<sup>2</sup>, while Bleu scores are calculated using the *Hugging Face - Bleu* library<sup>3</sup>.

## E.3 Hardware and Parameters

To accelerate model training and enhance performance, we utilized A100 GPUs, which significantly expedited the training process. The maximum load we deployed on GPUs was standard 7-billion parameter LLM models *viz.* Llama-7B, Mistral-7B, etc. We thoroughly explored the hyperparameter space to identify the optimal configuration for our model and present it in our supplementary code. We used a full fine-tuning technique for our models. The learning rate was managed using a PyTorch scheduler, which gradually decayed the rate, starting from 1e-3 and decreasing by a factor of 0.1 with each epoch. The batch size for the training was set to 4, and the training was conducted over 10 epochs. Additionally, in all large language models (LLMs) used, only the last two layers were unfrozen during the fine-tuning process.

---

<sup>2</sup><https://pypi.org/project/py-rouge/>

<sup>3</sup><https://huggingface.co/spaces/evaluate-metric/bleu>