

Socratic Reasoning Improves Positive Text Rewriting

Anmol Goel^a, Nico Daheim^a, Christian Montag^b, Iryna Gurevych^a

^aUbiquitous Knowledge Processing Lab (UKP Lab)

Department of Computer Science and Hessian Center for AI (hessian.AI)

Technical University of Darmstadt

^bCentre for Cognitive and Brain Sciences,

Institute of Collaborative Innovation,

University of Macau, Macau, China

www.ukp.tu-darmstadt.de

Abstract

Reframing a negative into a positive thought is at the crux of several cognitive approaches to mental health and psychotherapy that could be made more accessible by large language model-based solutions. Such reframing is typically non-trivial and requires multiple rationalization steps to uncover the underlying issue of a negative thought and transform it to be more positive. However, this rationalization process is currently neglected by both datasets and models which reframe thoughts in one step. In this work, we address this gap by augmenting open-source datasets for positive text rewriting with synthetically-generated Socratic rationales using a novel framework called SOCRATICREFRAME. SOCRATICREFRAME uses a sequence of question-answer pairs to rationalize the thought rewriting process. We show that such Socratic rationales significantly improve positive text rewriting for different open-source LLMs according to both automatic and human evaluations guided by criteria from psychotherapy research. We validate our framework and the synthetic rationalizations with expert judgements from domain experts and psychology students in an IRB-approved annotation study. Our findings highlight the potential of utilizing the synergy between LLM reasoning and established psychotherapy techniques to build assistive solutions for reframing negative thoughts.

1 Introduction

Negative thoughts can have a profound effect on human judgement and well-being. Oftentimes negative thoughts overshadow positive thoughts (Vaish et al., 2008), because they can be emotionally deep-rooted and triggering when brought to surface (Beck, 1979). *Cognitive Reframing* is a highly-validated Cognitive Behavioral Therapy (CBT) intervention technique that aims to address this by identifying and reframing negative thoughts into positive ones (Clark, 2013) and has proven use-

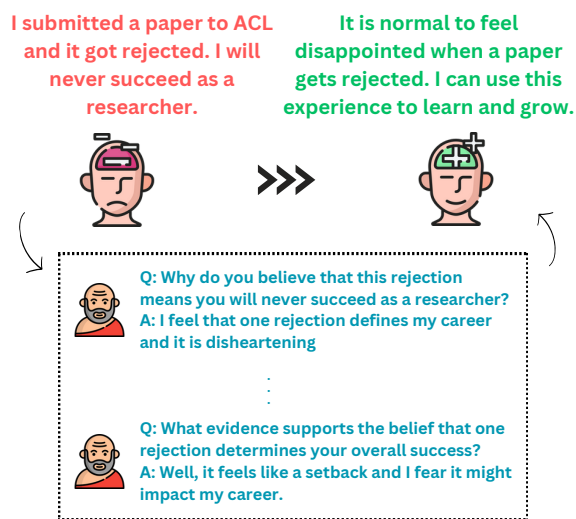


Figure 1: We use Socratic rationales consisting of question-answer pairs to improve positive text rewriting by verbalising the rewriting process. Here, we illustrate this: given a **negative thought**, we generate **Socratic rationales** grounded in Cognitive Behavioral Therapy to reframe the original thought into a **positive thought**.

ful in both clinical therapeutic and self-help settings (Williams, 2001). Therefore, reframing holds great potential as a strategy for clinicians to provide help in overcoming negative thoughts. Beyond this, teaching people to reframe their own thoughts as a self-guided mental health intervention could provide them with a coping strategy that does not rely on clinical help (Jorm et al., 2006). However, coming up with effective reframes can be challenging and training people, both for clinical and self-help settings, to effectively reframe negative thoughts is a time-consuming and laborious process. Altogether, this poses a risk of mental health improvements being slowed by a lack of availability of knowledge and clinicians.

Large Language Models (LLMs) hold great promise to overcome this lack of clinician support

and self-help resources by providing assistive solutions and could also be used as a training resource. Notably, it has been shown that they can effectively guide text generation towards desired attributes like non-toxicity (Liu et al., 2021; Zheng et al., 2023), style (Reif et al., 2022; Tikhonov et al., 2019), or persona (Song et al., 2020), and could therefore provide a diverse set of reframes. Recent work has also proposed LLM-based style transfer (Krishna et al., 2020; Reif et al., 2022) and empathetic response generation (Cao et al., 2025) methods in downstream applications. This motivates this work, where we aim to improve the reframing performance of LLMs.

While LLMs hold great promise for mental health applications due to their step-by-step reasoning and rationalization capabilities (Kojima et al., 2022, *inter alia*), they require careful investigation in tandem with domain experts. Several factors need to be taken into consideration before deploying LLM-based applications for client or therapist-facing applications. LLM-based solutions should not only be consistent with the client’s mental state, interpretable and avoid hallucinations, but importantly should also follow CBT guidelines.

Therefore, we aim to tackle this gap in literature by making the reasoning of LLMs explicit and include it as a supervision signal during training. We train LLMs for the positive text rewriting task by first generating a rationalization using the Socratic method¹, which uses questioning strategies for uncovering a person’s beliefs and motivations, and has shown to be a useful tool for CBT-based cognitive reframing. Socratic reasoning typically uses probing questions to elicit alternative perspectives. Table 1 lists different types of Socratic questions that are common in literature and practice. We call our method SOCRATICREFRAME.

We benchmark SOCRATICREFRAME with various state-of-the-art LLMs for the task of positive text rewriting using both automatic and human evaluations with clinical experts. Our results show that explicitly using Socratic reasoning as a training signal improves the text rewriting performance of LLMs while staying interpretable and faithful to CBT guidelines. We conduct extensive analysis on the synthetically-generated Socratic rationales and find that the socratic rationales are both informa-

¹The terminology for the Socratic method is not consistent within the literature (Carey and Mullan, 2004). In the context of this work, we use the term Socratic reasoning to refer to question-answer sequences that follow the Socratic method.

tive to the model and adhere to clinical standards according to a known Socratic-questioning evaluation scheme. Here, the rationales are rated as highly helpful and relevant for the reframing process.

We use the generated rationales to augment three existing cognitive reframing datasets from related works. We release all data and code publicly to enable subsequent research in the mental health domain, for example, for clinician training.

2 Background & Related Work

2.1 NLP and Mental Health

Research on NLP for mental health has primarily focused on utilizing linguistic features, as well as neural representations, to identify and analyze mental health conditions such as depression (Rinaldi et al., 2020; Yates et al., 2017), anxiety (Juhng et al., 2023; Wei et al., 2021), dyslexia (Björnsdóttir et al., 2023; Gala and Ziegler, 2016), autism (Cho et al., 2022; Goodkind et al., 2018), and schizophrenia (Mitchell et al., 2015; Sarioglu Kayi et al., 2017), among others. Typically, these studies rely on crowdsourced data or annotated social media posts to address the ethical and privacy concerns associated with medical data (Harrigian et al., 2020; Moßburger et al., 2020; Turcan and McKeown, 2019). Recently, efforts have been made to utilize datasets that are more representative of the interactions between a therapist and client in real world settings (Pérez-Rosas et al., 2017; Shapira et al., 2022; Howes et al., 2014; Lee et al., 2019; Cao et al., 2019; Tanana et al., 2015; Shreevastava and Foltz, 2021). Additionally, recent works have utilized synthetic data to enhance the performance of models in clinical (Kazi and Kahanda, 2019; Hiebel et al., 2023; Shim et al., 2021; Lindsay et al., 2022) and mental health (Wu et al., 2023b) settings. Chen et al. (2023b) propose Diagnosis-of-Thought prompting focusing only on diagnosing the type of cognitive distortion given a patient’s speech. In contrast, we develop a novel training framework for improving LLM performance on cognitive reframing. Our work aims to contribute to the growing body of literature leveraging positive psychology (Sheng et al., 2023) and LLMs to improve performance on mental health tasks.

2.2 Socratic Questioning

The Socratic method has found wide use in pedagogy (Bautista, 2014) and psychotherapy (Braun et al., 2015), because it can improve understand-

Question Type	Description	Exemplars
Clarification	Questions to go deeper into a thought	“Why do you say that?”
Probing assumptions	Questions to make someone think about unquestioned beliefs	“What could we assume instead?”
Probing reasons and evidence	Questions digging into the reasoning behind a thought	“How did you know that...?”
Probing implications	Questions probing the consequences of a thought	“What is likely to happen if ...?”
Probing alternative viewpoints	Questions about other, equally valid viewpoints	“What is another way to look at it?”
Question about the question	Meta-questions about the question itself	“Why do you think I asked this?”

Table 1: The six types of Socratic questions with representative exemplars from (Paul and Elder, 2019). We synthetically generate Socratic questions spanning all types to improve cognitive reframing with (Section 3).

ing and enable alternative perspectives without being explicit or direct. Instead, it uses questioning strategies to leave room for exploration which is also helpful for other applications, such as tutoring (Macina et al., 2023b). Recent works have used the step-by-step nature of Socratic questioning to improve NLP methods. Ang et al. (2023) collect a large-scale dataset from Reddit annotated with a question type to train Socratic question generation models. Wu et al. (2023a) develop a benchmark of Socratic-inspired deductive reasoning patterns to train state-of-the-art textual entailment and question answering models. Pagnoni et al. (2023) propose Socratic pretraining to enable a question-driven approach for summarizing documents. The scaffolding of Socratic questioning also enables LLMs to solve complex problems by decomposing them into smaller sub-problems. Qi et al. (2023); Shridhar et al. (2022, 2023) use Socratic methods to improve LLM’s performance on a variety of reasoning tasks, including math word problem solving and logical reasoning. Our work uses the Socratic method as the core principle to guide a language model while generating a reframed thought.

3 SOCRATICREFRAME

Cognitive reframing aims to reframe a negative thought into a positive thought. We approach this with LLMs that generate a positive thought conditioned on the negative thought and metadata. However, directly training an LLM to map positive to negative thought does not make the process behind the reframing explicit, which is important both for interpretability and model performance. Our method SOCRATICREFRAME overcomes this by training the model to verbalize a Socratic rationale before generating the reframing.

Formally, cognitive reframing transforms a negative thought $\mathbf{t} \in \mathcal{V}^*$ into a positive thought $\mathbf{r} \in \mathcal{V}^*$. Both are given as strings constructed from a (model) vocabulary \mathcal{V} . In addition to the negative

thought, metadata $\psi \in \mathcal{V}^*$ is often available which, for example, describes the person experiencing distress and the situation that caused it. One method for automatic reframing is using an autoregressive LLM with parameters θ to model the distribution

$$p_{\theta}(\mathbf{r} \mid \mathbf{t}, \psi) = \prod_{n=1}^{|\mathbf{r}|} p_{\theta}(r_n \mid \mathbf{t}, \psi, \mathbf{r}_{<n}) \quad (1)$$

paired with a decoding strategy, such as, sampling or greedy decoding. The model can be both finetuned using paired data $\mathcal{D}_{\text{train}} = \{(\mathbf{t}_i, \mathbf{r}_i, \psi_i)\}_{i=1}^N$ or used zero-shot with prompting. While intuitive, this does not allow the model to explicitly reason about and verbalize a rationale of the reframing. This is both less interpretable and might lead to overly-simplistic reframings.

Our proposed method SOCRATICREFRAME aims to overcome this by making the model rationalize the thought process behind a specific reframing by means of a sequence of Socratic question-answer pairs. This means that we introduce an additional variable $\mathbf{s} \in \mathcal{V}^*$ that is just a string of this question-answer sequence. The model is then tasked to first generate \mathbf{s} and only then is allowed to generate \mathbf{r} . This is outlined in Figure 2. Hence, the model now becomes

$$p_{\theta}(\mathbf{s} \circ \mathbf{r} \mid \mathbf{t}, \psi) = \prod_{n=1}^{|\mathbf{s} \circ \mathbf{r}|} p_{\theta}((s \circ r)_n \mid \mathbf{t}, \psi, (\mathbf{s} \circ \mathbf{r})_{<n}), \quad (2)$$

where $\mathbf{s} \circ \mathbf{r}$ means that \mathbf{s} is prepended to \mathbf{r} . For models that we finetune, this means that each training example is also augmented with a Socratic rationale \mathbf{s} such that $\mathcal{D}_{\text{train}} = \{(\mathbf{t}_i, \mathbf{r}_i, \psi_i, \mathbf{s}_i)\}_{i=1}^N$. During inference, \mathbf{s} is not known and the model generates it before generating the positive thought. This forces the model to explicitly reason about why a specific positive thought is generated to reframe a corresponding negative thought which we expect to

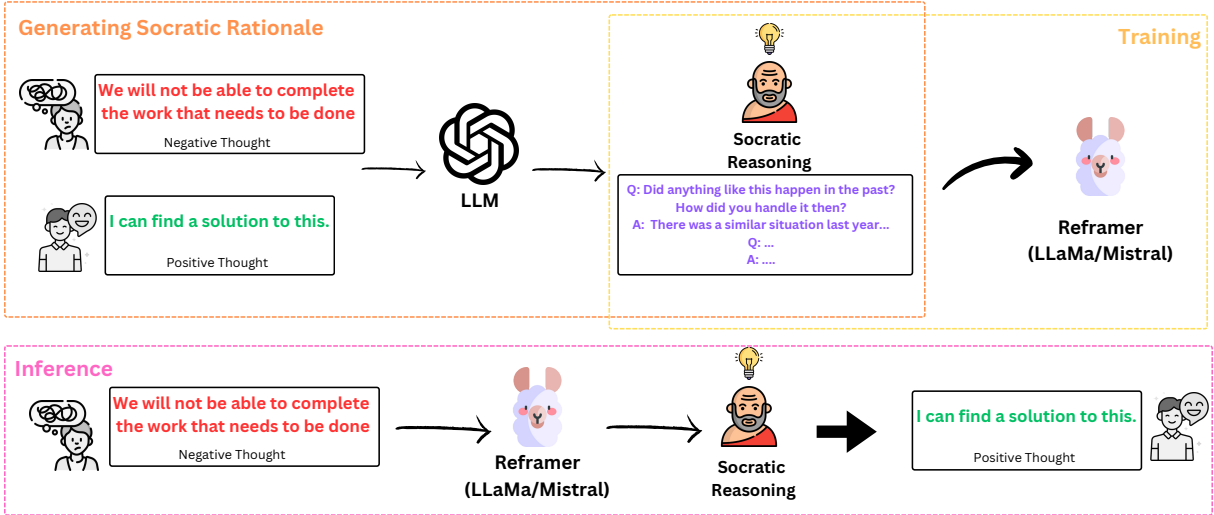


Figure 2: Detailed illustration of our framework SOCRATICREFRAME. First, we generate Socratic rationales using GPT-4 with a few-shot prompt. Then, we use the generated Socratic rationale to train models for cognitive reframing. During inference, the model generates the Socratic rationale before reframing the negative thought.

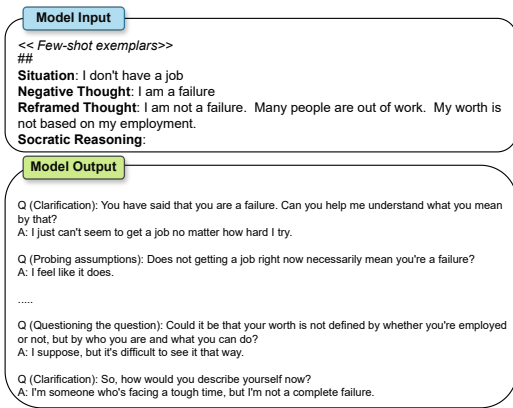


Figure 3: A sample Socratic rationale for an instance from COGREF (Sharma et al., 2023) generated with GPT-4 in a few-shot setting. We use clinical vignettes from psychology literature as few-shot exemplars.

improve both performance and interpretability. In the following section, we discuss how such Socratic rationales can reliably be synthetically generated for existing datasets by leveraging LLMs.

3.1 Generating Socratic Rationales

Our goal is to improve cognitive reframing with Socratic rationales. Such rationales can be time-consuming to write. Therefore, we use synthetically generated data which has recently been successfully used in text classification (Li et al., 2023), dialogue generation (Bao et al., 2023) and question-answering (Riabi et al., 2021), among others.

In particular, we few-shot prompt GPT-

4 (Achiam et al., 2023) to generate Socratic question-answer pairs based on negative thought t , positive thought r , metadata ψ , and in-context exemplars of Socratic rationales from clinical literature (Padesky, 1993) to ensure high quality and groundedness. The generated Socratic rationales then gives explicit reasoning steps for the particular reframing and can be used to augment the training data $\mathcal{D}_{\text{train}}$. This can be seen as a form of data augmentation and knowledge distillation into the model that is trained for cognitive reframing. As we will show in Section 6, the Socratic rationales are rated favourably by human annotators. An example rationale is shown in Figure 3. Details and examples are in Appendix A and Appendix D. In addition to generating rationales, we prompt GPT-4 to classify the specific type of each generated Socratic question according to the six types described in Table 1, namely “Clarification”, “Probing Assumptions”, “Probing reasons or evidence”, “Probing implications”, “Probing alternative viewpoints”, and “Question”. These can be used for further insights into the reframing process.

4 Empirical Analysis

We show the effectiveness of our method on multiple datasets for positive text rewriting that are outlined in Section 4.1. More details on the exact experimental set-up, such as the models used, are found in Section 4.2 and Appendix A and the used metrics are introduced in Section 4.3.

	POSREF	PATREF	COGREF
Train	6,679	5,249	400
Test	835	18,635	200
ψ	-	Persona	Situation

Table 2: Statistics of the used datasets.

4.1 Datasets

Table 2 describes the datasets used in this work which we further outline in the following.

We use three recently-released open-source datasets. First, Positive Psychology Frames (POSREF) (Ziems et al., 2022) which contains tweets where a hashtag has indicated stress. Each tweet is mapped to a corresponding reframed text which is grounded in a set of reframing strategies, for example, Growth Mindset, Optimism, Self-affirmation. POSREF does not contain any metadata. Then, we use Pattern Reframe (PATREF) (Maddela et al., 2023), a crowdsourced dataset of negative thoughts that are conditioned on personas. These personas describe the person experiencing the respective negative thought. Here, ψ contains the persona-specific information associated with each example. For instance, a persona from the dataset includes “My mother was a teacher. My favorite food is a salad. I enjoy nature. I teach a yoga class. I am single.” Finally, we use Cognitive Reframe (COGREF) (Sharma et al., 2023) which is an expert-annotated dataset of situations, thoughts and reframes. Mental health practitioners were prompted with a situation, which can be used as metadata, to generate a negative and a reframed thought. For example, one situation from the dataset is “I participated in a hackathon and I lost”.

4.2 Experimental Set-up

We use different LLMs that we either prompt or finetune. Namely, we prompt ChatGPT, LLaMa-2 7B (Touvron et al., 2023) and Mistral 7B (Jiang et al., 2023) and finetune LLaMa 7B and Mistral 7B for cognitive reframing. All models are finetuned using LoRA (Hu et al., 2021) due to computational constraints. For generating Socratic rationales, we use GPT-4. We use the transformers (Wolf et al., 2020) library for all our experiments.

We test the models on different setups. *Few Shot (FS)*: Few shot exemplars of negative and positive thought are given to the model which are handcrafted from literature on cognitive reframing. *Chain of Thought (CoT)*: The model is asked

to think step-by-step and rationalize the reframed thought. ² *Finetune (FT)*: The model is finetuned using LoRA to reframe a negative thought without a Socratic rationale. *Socratic CoT (SoC_{CoT})*: The typical CoT prompt is modified by replacing “think step-by-step” with an instruction to generate Socratic rationale before generating the reframe. *Finetune with Socratic Rationales (SoC)*: The model is finetuned using LoRA to reframe a negative thought by first rationalising it using Socratic questioning and then generating the positive reframing. The prompt template used for finetuning is shown in Figure 9.

4.3 Evaluating Cognitive Reframers

We aim to reframe a negative into a positive thought while preserving the original meaning and semantics. Hence, following Hu et al. (2022); Qu et al. (2023), we evaluate them along the criteria of Transfer Strength, which describes how well the thought was transferred to a positive one, and Content Preservation, which evaluates how well the original meaning of the thought was preserved. We describe how we measure both of them in the following.

Transfer strength defines how well negative sentiment is turned positive. To measure it, we first use a finetuned RoBERTa model to evaluate the sentiment scores of the original and reframed thoughts. Then, we use the pairwise difference of sentiment scores between the original and reframed thoughts and report the average, denoted by (Δ Pos). We report the number of samples with an increase in positivity with respect to the original thought as accuracy. A decrease in the score from the original thought is considered a failed case (Acc). Sharma et al. (2023) show that people tend to prefer more empathetic reframes over overly positive reframes. Hence, we also report Δ Emp as the difference between empathy scores of original and reframed thoughts. The scores are computed using a pre-trained RoBERTa empathy classifier. Our aim is to reframe a thought but also preserve the original meaning and not drastically change it. We rely on the following automatic metrics which have been widely used in the text reframing literature: BLEU (Papineni et al., 2002) and BLEURT (Sellam et al., 2020). Following all three datasets we consider

²We observed similar performance for zero-shot and few-shot CoT and therefore only report zero-shot CoT. We believe the similar performance in both setups stems from the nature of the reframing task where there is extra contextual information in the form of ψ and the LLMs are able to generalize well even without few-shot exemplars.

Dataset	Model	Content Preservation (\uparrow)						Transfer Strength (\uparrow)								
		BLEU			BLEURT			Δ Pos			Acc. (%)			Δ Emp		
		LLaMa	Mistral	ChatGPT	LLaMa	Mistral	ChatGPT	LLaMa	Mistral	ChatGPT	LLaMa	Mistral	ChatGPT	LLaMa	Mistral	ChatGPT
POSREF	FS	12.6	13.1	15.1	0.53	0.55	0.57	0.54	0.49	0.54	89.15	90.10	92.00	0.58	0.66	0.70
	CoT	12.1	13.3	15.6	0.51	0.55	0.59	0.55	0.48	0.57	90.02	91.16	92.30	0.61	0.69	0.73
	FT	14.3	15.1	-	0.56	0.58	-	0.55	0.51	-	91.46	91.10	-	0.63	0.72	-
	SoC _{CoT}	13.9	13.6	14.5	0.49	0.51	0.55	0.51	0.49	0.51	89.82	90.42	92.12	0.55	0.59	0.71
	SoC	15.7	15.9	-	0.58	0.61	-	0.58	0.52	-	92.45	92.09	-	0.69	0.76	-
PATREF	FS	71.9	73.0	73.8	0.64	0.65	0.63	0.49	0.78	0.81	97.02	96.51	97.61	0.81	0.80	0.87
	CoT	72.0	73.5	73.9	0.64	0.65	0.65	0.48	0.77	0.83	97.23	96.01	97.60	0.83	0.81	0.89
	FT	72.3	74.1	-	0.66	0.67	-	0.51	0.80	-	97.67	96.79	-	0.89	0.85	-
	SoC _{CoT}	70.2	73.5	71.2	0.59	0.61	0.62	0.48	0.62	0.78	95.02	95.64	96.10	0.80	0.82	0.86
	SoC	75.2	75.1	-	0.69	0.68	-	0.52	0.82	-	97.90	97.00	-	0.90	0.88	-
COGREF	FS	27.1	27.9	30.0	0.59	0.61	0.63	0.67	0.68	0.69	90.12	90.94	91.90	0.79	0.81	0.88
	CoT	27.0	27.5	31.6	0.59	0.60	0.63	0.67	0.69	0.71	90.0	91.67	92.55	0.82	0.84	0.90
	FT	29.0	30.0	-	0.61	0.62	-	0.69	0.70	-	91.48	91.88	-	0.89	0.91	-
	SoC _{CoT}	26.4	27.1	30.6	0.55	0.59	0.60	0.65	0.68	0.70	90.44	91.54	92.0	0.80	0.85	0.89
	SoC	30.1	31.7	-	0.62	0.63	-	0.70	0.72	-	92.30	92.56	-	0.91	0.94	-

Table 3: Automatic evaluation results. FS=Few shot. CoT = Chain of Thought. FT=Finetune with LoRA. SoC=finetune with Socratic rationale. Items in **bold** represent the best performance. Note that we do not finetune ChatGPT. We observe that using Socratic rationales for fine-tuning models significantly improves the text rewriting performance across all datasets. Mean values over three runs are reported.

(Ziems et al., 2022; Maddela et al., 2023; Sharma et al., 2023), we report the BLEU scores for each pair of original and reframed thoughts. Since it has been shown that BLEU scores do not always correlate well with human judgements on semantic similarity, we also use a BERT-based BLEU variation, BLEURT-20 which is trained on synthetic samples to get accurate semantic similarity scores. We report the average over a dataset. Each score is a value between 0 and 1, ranging from no to complete semantic similarity.

4.4 Evaluating Socratic rationales

For Socratic rationales to be useful for a model, they should intuitively contain new salient information that is not contained in the negative thought. We use the recently-proposed information-theoretic metric REV (Chen et al., 2023a) to measure this. REV uses conditional \mathcal{V} -information to compute the usable information that can be extracted from a variable (s) by a model to predict another variable (r), conditioned on a third variable (t, ψ). In our setting, this measures the extra information provided by the Socratic rationales beyond what is contained in the original negative thought. More formally, we compute

$$\text{REV}(t, r, \psi, s) = -\log p_{\theta'}(r | t, \psi) + \log p_{\theta}(r | t, \psi, s), \quad (3)$$

where θ' and θ are the parameters of two models trained to minimize cross-entropy, respectively. The REV metric can then be computed for an entire corpus \mathcal{D} by averaging the pointwise (per-example) scores. While Chen et al. (2023a) use fine-tuned models on the specific datasets for computing the

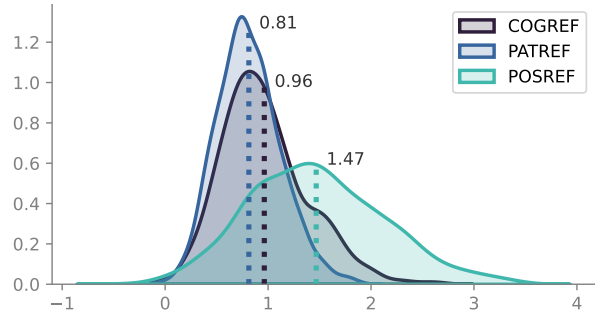


Figure 4: The Socratic rationales generated for all three datasets have positive REV scores, meaning that they indeed provide useful information for text rewriting. Mean of the example-wise REV values are plotted for all three datasets.

score, Lu et al. (2023) show that it can even be calculated directly from pretrained LLMs with the same effectiveness as using fine-tuned models (Lu et al., 2023).

5 Results

In this section, we detail our results obtained with different strategies. First, our main results in Section 5.1 show that the Socratic rationales improve text rewriting. Then, we show that the rationales are indeed informative for the model in Section 4.4.

5.1 Main Results

Table 3 shows that using Socratic rationales consistently outperforms both vanilla finetuning and prompting strategies on all datasets. We show an example in Table 4 that highlights how the positive rewritings appear much more thoughtful, because the Socratic rationale forces the model to explicitly reason about the rewriting process. We empirically

Method	Reframed Thought
Mistral-FS	I will improve.
Mistral-CoT	I will improve and learn to be better.
Mistral-FT	I will learn from the feedback and grow.
Mistral-SoC	It is normal to feel disappointed. I can use this experience to learn and grow.

Table 4: Sample reframed thoughts generated by Mistral variants for the original thought: “I submitted a paper to ACL and it got rejected. I will never succeed as a researcher.” The additional socratic rationale leads to a more detailed and well-founded reframing.

check whether the generated reframes are close in length to the ground truth reframes for all datasets. We observe that the average length of generated reframes with our best performing model on all datasets is only 5 tokens longer than the ground truth, on average. It is noteworthy that Socratic reasoning improves both the content preservation as well as transfer strength capabilities of smaller models beyond the performance of ChatGPT, despite the model being multiple times larger than LLaMa 7B and Mistral 7B. Across all datasets, we observe that Mistral performs better in content preservation while LLaMA performs better in the sentiment of the reframed thought. This could be due to the different pretraining datasets used for training the models. We also observe that our method consistently generates more empathetic reframes which is generally preferred by people (Sharma et al., 2023). We note that both positivity and empathy are lower for POSREF than other datasets, possibly because POSREF contains tweets which are known to be less positive in nature (Sokolova et al., 2017). We observe that using Socratic reasoning improves BLEU scores by almost 2 points on average when compared to finetuning without Socratic rationales. Distilling Socratic reasoning (from GPT-4) into smaller models helps to preserve the original meaning and improve the positive sentiment transfer. This is desirable for users and practitioners alike, because more positive sentiment will likely benefit users more and content preservation ensures that the reframing stays relevant for them.

5.2 Socratic Rationales are Informative

Similar to previous work for evaluating free-text rationales (Chen et al., 2023a), we use the pretrained GPT-Neo 2.7B (Gao et al., 2020) for computing REV. We report the average REV metric over each sample for all datasets. A REV value > 0 suggests

Socratic rationales support reframing by providing additional information, while < 0 indicates otherwise. In Figure 4, all datasets have $REV > 0$, highlighting Socratic rationale’s utility in enhancing reframing tasks. Table 5 shows the generated rationales for each of the dataset we consider. In particular, POSREF has the highest informativeness for the reframed thoughts. We attribute this to two main factors: First, POSREF lacks additional context like situations or personas, causing token dispersion where more information can lead to low values of REV. Second, POSREF samples from Reddit align better with GPT-Neo’s training data than the crowdsourced data in the other datasets.

5.3 Socratic Rationales progressively lead to a better reframe

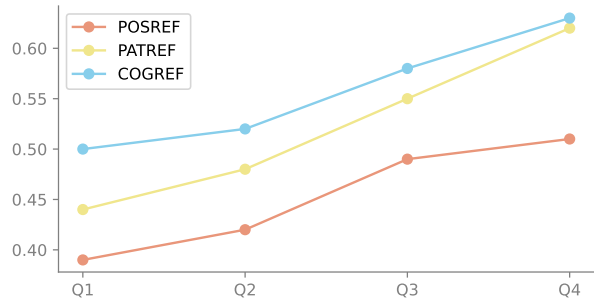


Figure 5: Sentiment scores for the generated Socratic question-answer pairs show that the synthetic rationales progressively get better to lead to an improved reframe.

We partition Socratic rationales into four quarters sequentially, each representing 25% of the data, and calculate average sentiment scores for each quarter. Figure 5 reveals a consistent trend across all datasets: intermediate answers lead to increasingly positive questions, suggesting that they aid in iterative reframing.

6 Human Evaluation

6.1 Evaluating Reframed Thoughts

Prior cognitive reframing studies (Maddela et al., 2023; Sharma et al., 2023) use Likert scales for assessing aspects like fluency and readability. Yet, our focus lies on helping users overcome unwanted thoughts. Li et al. (2024) show that pairwise human preferences are well-defined. In a similar setup as Sharma et al. (2023), we compare 100 randomly selected SOCRATICREFRAME generated reframes against ChatGPT or a baseline model without Socratic rationales. After consenting to participate, two computer science graduate students were re-

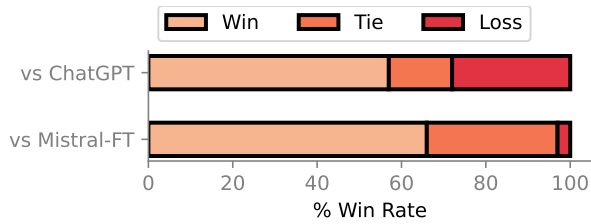


Figure 6: Human evaluation results for Mistral-SoC compared to Mistral-FT and ChatGPT. With a win rate of over 50% human annotators prefer reframes generated by our method to ChatGPT and simple fine-tuning.

cruited to conduct the comparisons. The win rate indicates a preference for Socratic reframes over the baseline by the raters. The participants were asked to select the reframed thought they find most relatable, helpful and memorable - a criteria for good reframes defined in Sharma et al. (2023).

As shown in Figure 6, Mistral-SoC outperforms Mistral-FT by a significant margin with a win rate of over 65% and tie rate of around 30%. We fit a Bradley-Terry Model (Bradley and Terry, 1952) on the preference data and obtain a strength of 1.3 for Mistral-SoC as compared to strengths -1.8 and 0.5 for Mistral-FT and ChatGPT, respectively. Intuitively, this means that Mistral-SoC will have a win rate of 95% over Mistral-FT and 67% over ChatGPT, confirming the effectiveness of our method.

6.2 Evaluating Socratic Rationales

To evaluate the synthetically generated rationales, we consider two key attributes: whether they represent the use of Socratic questioning in a clinically meaningful way and the general helpfulness of the questioning to overcome negative thoughts. Following Braun et al. (2015), we use the Socratic Questioning Scale (SQS) to evaluate the quality of generated rationales. SQS contains five Likert-scale questions evaluating the use of the Socratic method in a snippet. The final score is the sum of these ratings. Originally developed for therapy session transcripts, we include only three relevant questions for text evaluation: Open-endedness, Context, and Relevance.³ The full questions are reported in Appendix C.1. Additionally, we consider *helpfulness* as a metric to assess whether or not the generated rationale is generally helpful in

³We note that in our evaluation criteria, “relevant” questions aim to tackle extrinsic hallucinations (generated text cannot be verified given the source) and “contextual” questions aim to tackle intrinsic hallucinations (generated text contradicts the source text), following the definitions from Ji et al. (2023).

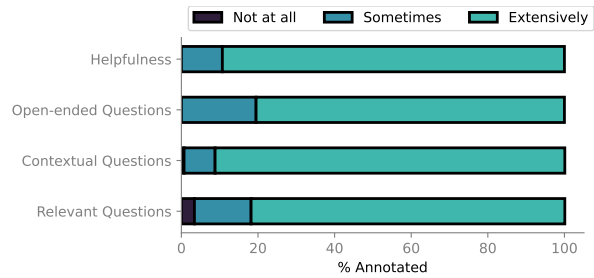


Figure 7: Human evaluation results for the Socratic rationales show that annotators judge the generated rationales as highly helpful and relevant for reframing negative thoughts.

overcoming negative thoughts. Annotators rate the generated rationales on a 3-point Likert scale.

Expert Evaluation We recruit 14 mental health experts and clinical psychology graduate students for evaluating the quality of the Socratic rationales. We verify the quality of randomly sampled Socratic rationales with a clinical expert before conducting the study with psychology students. The annotation task was deemed exempt from an ethics review by an institutional review board. Amidei et al. (2019) suggest the use of Cronbach’s α for checking the evaluation reliability in natural language generation tasks. We observe acceptably high values for helpfulness (0.67), diversity (0.69), reflectiveness (0.75) and relevance (0.65) of the questions in the rationale. We observe that raters validate the synthetically generated Socratic rationales as helpful to reframe the original thought. More than 50% of the raters also responded affirmatively on whether they are willing to use LLMs for cognitive reframing. More details are found in Appendix C.2.

Human Evaluations Three computer science graduate students annotated 100 randomly sampled Socratic rationales from the datasets. Figure 7 shows favorable ratings across all criteria. We obtained an Intra Class Correlation (ICC) of 0.806 for helpfulness and 0.601 for the SQS scale, indicating that while the SQS is challenging to evaluate, helpfulness is generally agreed upon by non-experts. These findings align with similar agreement scores in psychology (Braun et al., 2015). Additionally, a Pearson correlation coefficient of 0.72 ($p < 0.05$) between helpfulness and SQS ratings suggests that better Socratic questioning correlates with greater helpfulness in addressing negative thoughts.

7 Clinical Implications

Our findings suggest that large language models (LLMs) augmented with Socratic rationales can serve as valuable tools for mental health professionals in cognitive restructuring and psychotherapy. By breaking down the process of thought reframing into a structured sequence of rationalization steps, SOCRATICREFRAME aligns with established therapeutic approaches such as cognitive behavioral therapy (CBT) and Socratic questioning. This structured reasoning can assist clinicians by:

- **Enhancing Psychoeducation:** Mental health professionals can use LLM-assisted reframing to demonstrate thought restructuring techniques, making abstract cognitive principles more tangible for patients.
- **Supporting Self-Guided Interventions:** LLMs equipped with Socratic reasoning could supplement therapy by providing guided, interactive thought reframing outside of clinical sessions, helping patients develop cognitive flexibility in real time.
- **Reducing Cognitive Load for Clinicians:** – By automating initial rationalization steps, our approach may help therapists focus on more nuanced aspects of patient care, such as emotional processing and behavioral interventions, without sacrificing the rigor of cognitive restructuring.
- **Facilitating Training and Supervision:** Psychology students and early-career therapists could use LLM-generated Socratic rationales as a training aid, learning how to systematically guide patients through cognitive reframing with structured question-answer pairs.

Given these potential applications, our work highlights the importance of integrating AI-driven reasoning techniques with clinical expertise to create assistive tools that complement human-guided therapy rather than replace it. Future research should explore real-world deployment, patient engagement, and the ethical considerations of using LLM-based interventions in mental health contexts.

8 Conclusion & Future Work

In this work, we show that text rewriting models for cognitive reframing can be improved by using Socratic rationales to verbalize the reframing process.

By releasing our code and data we hope to enable future research, for example, on incorporating the different question types that are annotated, or using the data in therapist training.

Limitations

While synthetic data generated by LLMs like GPT-4 offers a promising solution to data scarcity and privacy concerns, a drawback could be a potential lack of diversity and complexity. Our generated Socratic rationales may not fully capture all intricate nuances or patterns in authentic therapist dialogues while performing cognitive change with the Socratic method. A significant limitation when working with mental health and psychotherapy-based datasets is the use of heavily curated or crowd-sourced data. This can lead to an imbalance in the demographics and language of our datasets. Since the datasets we use are sourced from social media or crowdsourcing platforms, the quality of the negative thoughts might also be limited.

In addition to the reframed thoughts, the datasets considered in this work also contain the cognitive distortions associated with each negative thought. Cognitive distortions are thought patterns that lead to negative feelings. We believe information like cognitive distortions and reframing strategies can be utilized to further improve the quality of the Socratic rationales and subsequently enhance the text reframing performance.

Finally, our focus in this work was to improve text reframing which is an effective short-term in-the-moment strategy. However, we emphasize that assessing long-term outcomes is imperative and is a future research direction.

Ethical Statement

While this work focused on generating Socratic rationales to improve positive text rewriting, open-ended LLMs still carry the risk of generating harmful outputs. Therefore, we advise careful consideration before they are applied in practice. Unsupervised use of the Socratic questioning data as is, without consulting trained professionals, could be harmful and requires careful design and implementation before being applied in real-world settings.

Datasets Our work augments open-source positive text rewriting datasets, some of which include social media-derived data. Social media users are not representative of clinical populations, and self-reported distress lacks clinical verification. We

mitigate this by incorporating expert evaluations to validate our approach. Public posts may contain sensitive disclosures. Our study adheres to IRB guidelines and we only use anonymized, public datasets, ensuring ethical data use. Social media data may reinforce biases or oversimplify mental health struggles. Our framework grounds reframing in psychotherapy principles, minimizing these risks. While social media data is valuable for scalable mental health interventions, expert oversight remains essential to ensure clinically sound applications.

IRB Approval We obtained approval from the Ulm University’s Institutional Review Board. In the board’s opinion, research using annotations and ratings on the anonymised data does not require consultation within the meaning of Article 23 of the Declaration of Helsinki, meaning that our study was exempt from further inquiry.

Acknowledgements

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. This work was supported by the DYNAMIC center, which is funded by the LOEWE program of the Hessian Ministry of Science and Arts (Grant Number: LOEWE/1/16/519/03/09.001(0009)/98).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Jacopo Amidei, Paul Piwek, and Alistair Willis. 2019. [Agreement is overrated: A plea for correlation to assess human evaluation reliability](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 344–354, Tokyo, Japan. Association for Computational Linguistics.
- Beng Heng Ang, Sujatha Das Gollapalli, and See-Kiong Ng. 2023. [Socratic question generation: A novel dataset, models, and evaluation](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 147–165, Dubrovnik, Croatia. Association for Computational Linguistics.
- Jianzhu Bao, Rui Wang, Yasheng Wang, Aixin Sun, Yitong Li, Fei Mi, and Ruifeng Xu. 2023. [A synthetic data generation framework for grounded dialogues](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10866–10882, Toronto, Canada. Association for Computational Linguistics.
- Lowell Bautista. 2014. [The socratic method as a pedagogical method in legal education](#). *University of Wollongong Faculty of Law, Humanities and the Arts - Papers*.
- Aaron T Beck. 1979. *Cognitive therapy and the emotional disorders*. Penguin.
- Marina Björnsdóttir, Nora Hollenstein, and Maria Barrett. 2023. [Dyslexia prediction from natural reading of Danish texts](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 60–70, Tórshavn, Faroe Islands. University of Tartu Library.
- Ralph Allan Bradley and Milton E. Terry. 1952. [Rank analysis of incomplete block designs: I. the method of paired comparisons](#). *Biometrika*, 39:324.
- Justin D Braun, Daniel R Strunk, Katherine E Sasso, and Andrew A Cooper. 2015. [Therapist use of socratic questioning predicts session-to-session symptom change in cognitive therapy for depression](#). *Behaviour research and therapy*, 70:32–37.
- Huiying Cao, Yiqun Zhang, Shi Feng, Xiaocui Yang, Daling Wang, and Yifei Zhang. 2025. [TOOL-ED: Enhancing empathetic response generation with the tool calling capability of LLM](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 5305–5320, Abu Dhabi, UAE. Association for Computational Linguistics.
- Jie Cao, Michael Tanana, Zac Imel, Eric Poitras, David Atkins, and Vivek Srikumar. 2019. [Observing dialogue in therapy: Categorizing and forecasting behavioral codes](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5599–5611, Florence, Italy. Association for Computational Linguistics.
- Timothy A Carey and Richard J Mullan. 2004. [What is socratic questioning?](#) *Psychotherapy: theory, research, practice, training*, 41(3):217–226.
- Hanjie Chen, Faeze Brahman, Xiang Ren, Yangfeng Ji, Yejin Choi, and Swabha Swayamdipta. 2023a. [REV: Information-theoretic evaluation of free-text rationales](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2007–2030, Toronto, Canada. Association for Computational Linguistics.
- Zhiyu Chen, Yujie Lu, and William Wang. 2023b. [Empowering psychotherapy with large language models: Cognitive distortion detection through diagnosis of thought prompting](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*,

- pages 4295–4304, Singapore. Association for Computational Linguistics.
- Sunghye Cho, Riccardo Fusaroli, Maggie Rose Pelella, Kimberly Tena, Azia Knox, Aili Hauptmann, Maxine Covello, Alison Russell, Judith Miller, Alison Hulink, Jennifer Uzokwe, Kevin Walker, James Fiumara, Juhi Pandey, Christopher Chatham, Christopher Cieri, Robert Schultz, Mark Liberman, and Julia Parish-morris. 2022. [Identifying stable speech-language markers of autism in children: Preliminary evidence from a longitudinal telephony-based study](#). In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 40–46, Seattle, USA. Association for Computational Linguistics.
- David A Clark. 2013. [Cognitive restructuring](#). *The Wiley handbook of cognitive behavioral therapy*, pages 1–22.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: efficient finetuning of quantized llms. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA.
- Núria Gala and Johannes Ziegler. 2016. [Reducing lexical complexity as a tool to increase text accessibility for children with dyslexia](#). In *Proceedings of the Workshop on Computational Linguistics for Linguistic Complexity (CLALC)*, pages 59–66, Osaka, Japan. The COLING 2016 Organizing Committee.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. [The pile: An 800gb dataset of diverse text for language modeling](#). *arXiv preprint arXiv:2101.00027*.
- Adam Goodkind, Michelle Lee, Gary E. Martin, Molly Losh, and Klinton Bicknell. 2018. [Detecting language impairments in autism: A computational analysis of semi-structured conversations with vector semantics](#). In *Proceedings of the Society for Computation in Linguistics (SCiL) 2018*, pages 12–22.
- Keith Harrigan, Carlos Aguirre, and Mark Dredze. 2020. [Do models of mental health based on social media data generalize?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3774–3788, Online. Association for Computational Linguistics.
- Nicolas Hiebel, Olivier Ferret, Karen Fort, and Aurélie Névoul. 2023. [Can synthetic text help clinical named entity recognition? a study of electronic health records in French](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2320–2338, Dubrovnik, Croatia. Association for Computational Linguistics.
- Christine Howes, Matthew Purver, and Rose McCabe. 2014. [Linguistic indicators of severity and progress in online text-based therapy for depression](#). In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 7–16, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. [Lora: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Zhiqiang Hu, Roy Ka-Wei Lee, Charu C. Aggarwal, and Aston Zhang. 2022. [Text style transfer: A review and experimental evaluation](#). 24(1):14–45.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Anthony F Jorm, Nicholas B Allen, Colin P O'Donnell, Ruth A Parslow, Rosemary Purcell, and Amy J Morgan. 2006. Effectiveness of complementary and self-help treatments for depression in children and adolescents. *Medical journal of Australia*, 185(7):368–372.
- Swanie Juhng, Matthew Matero, Vasudha Varadarajan, Johannes Eichstaedt, Adithya V Ganesan, and H. Andrew Schwartz. 2023. [Discourse-level representations can improve prediction of degree of anxiety](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1500–1511, Toronto, Canada. Association for Computational Linguistics.
- Nazmul Kazi and Indika Kahanda. 2019. [Automatically generating psychiatric case notes from digital transcripts of doctor-patient conversations](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 140–148, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). *Advances in neural information processing systems*, 35:22199–22213.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Fei-Tzin Lee, Derrick Hull, Jacob Levine, Bonnie Ray, and Kathy McKeown. 2019. [Identifying therapist](#)

- conversational actions across diverse psychotherapeutic approaches. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pages 12–23, Minneapolis, Minnesota. Association for Computational Linguistics.
- Junlong Li, Fan Zhou, Shichao Sun, Yikai Zhang, Hai Zhao, and Pengfei Liu. 2024. Dissecting human and llm preferences. *arXiv preprint arXiv:2402.11296*.
- Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. Synthetic data generation with large language models for text classification: Potential and limitations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10443–10461, Singapore. Association for Computational Linguistics.
- Hali Lindsay, Johannes Tröger, Mario Magued Mina, Philipp Müller, Nicklas Linz, Jan Alexandersson, and Inez Ramakers. 2022. Generating synthetic clinical speech data through simulated ASR deletion error. In *Proceedings of the RaPID Workshop - Resources and Processing of linguistic, para-linguistic and extra-linguistic Data from people with various forms of cognitive/psychiatric/developmental impairments - within the 13th Language Resources and Evaluation Conference*, pages 9–16, Marseille, France. European Language Resources Association.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. 2021. DExperts: Decoding-time controlled text generation with experts and anti-experts. 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 6691–6706, Online. Association for Computational Linguistics.
- Sheng Lu, Shan Chen, Yingya Li, Danielle Bitterman, Guergana Savova, and Iryna Gurevych. 2023. Measuring pointwise \mathcal{V} -usable information in-context-ly. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15739–15756, Singapore. Association for Computational Linguistics.
- Jakub Macina, Nico Daheim, Sankalan Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023a. MathDial: A dialogue tutoring dataset with rich pedagogical properties grounded in math reasoning problems. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5602–5621, Singapore. Association for Computational Linguistics.
- Jakub Macina, Nico Daheim, Lingzhi Wang, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023b. Opportunities and challenges in neural dialog tutoring. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2357–2372, Dubrovnik, Croatia. Association for Computational Linguistics.
- 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.
- Margaret Mitchell, Kristy Hollingshead, and Glen Coppersmith. 2015. Quantifying the language of schizophrenia in social media. In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 11–20, Denver, Colorado. Association for Computational Linguistics.
- Luis Moßburger, Felix Wende, Kay Brinkmann, and Thomas Schmidt. 2020. Exploring online depression forums via text mining: A comparison of Reddit and a curated online forum. In *Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task*, pages 70–81, Barcelona, Spain (Online). Association for Computational Linguistics.
- Christine A Padesky. 1993. Socratic questioning: Changing minds or guiding discovery. In *A keynote address delivered at the European Congress of Behavioural and Cognitive Therapies, London*, volume 24.
- Artidoro Pagnoni, Alex Fabbri, Wojciech Kryscinski, and Chien-Sheng Wu. 2023. Socratic pretraining: Question-driven pretraining for controllable summarization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12737–12755, Toronto, Canada. 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 *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- R. Paul and L. Elder. 2019. *The Thinker’s Guide to Socratic Questioning*. Thinker’s Guide Library. Foundation for Critical Thinking.
- Verónica Pérez-Rosas, Rada Mihalcea, Kenneth Resnicow, Satinder Singh, and Lawrence An. 2017. Understanding and predicting empathic behavior in counseling therapy. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1435, Vancouver, Canada. Association for Computational Linguistics.
- Jingyuan Qi, Zhiyang Xu, Ying Shen, Minqian Liu, Di Jin, Qifan Wang, and Lifu Huang. 2023. The art of SOCRATIC QUESTIONING: Recursive thinking with large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural*

- Language Processing*, pages 4177–4199, Singapore. Association for Computational Linguistics.
- Renyi Qu, Lyle Ungar, and João Sedoc. 2023. [Conditioning on dialog acts improves empathy style transfer](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13254–13271, Singapore. Association for Computational Linguistics.
- Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. [A recipe for arbitrary text style transfer with large language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848, Dublin, Ireland. Association for Computational Linguistics.
- Arij Riabi, Thomas Scialom, Rachel Keraron, Benoît Sagot, Djamé Seddah, and Jacopo Staiano. 2021. [Synthetic data augmentation for zero-shot cross-lingual question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7016–7030, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alex Rinaldi, Jean Fox Tree, and Snigdha Chaturvedi. 2020. [Predicting depression in screening interviews from latent categorization of interview prompts](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7–18, Online. Association for Computational Linguistics.
- Efsun Sarioglu Kayi, Mona Diab, Luca Pauselli, Michael Compton, and Glen Coppersmith. 2017. [Predictive linguistic features of schizophrenia](#). In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 241–250, Vancouver, Canada. Association for Computational Linguistics.
- Thibault Sellam, Amy Pu, Hyung Won Chung, Sebastian Gehrmann, Qijun Tan, Markus Freitag, Dipanjan Das, and Ankur Parikh. 2020. [Learning to evaluate translation beyond English: BLEURT submissions to the WMT metrics 2020 shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 921–927, Online. Association for Computational Linguistics.
- Natalie Shapira, Dana Atzil-Slonim, Rivka Tuvai Mashiach, and Ori Shapira. 2022. [Measuring linguistic synchrony in psychotherapy](#). In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 158–176, Seattle, USA. 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.
- Xu Sheng, Fumiyo Fukumoto, Jiye Li, Go Kentaro, and Yoshimi Suzuki. 2023. [Learning disentangled meaning and style representations for positive text reframing](#). In *Proceedings of the 16th International Natural Language Generation Conference*, pages 424–430, Prague, Czechia. Association for Computational Linguistics.
- Heereen Shim, Dietwig Lowet, Stijn Luca, and Bart Vanrumste. 2021. [Synthetic data generation and multi-task learning for extracting temporal information from health-related narrative text](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 260–273, Online. Association for Computational Linguistics.
- Sagarika Shreevastava and Peter Foltz. 2021. [Detecting cognitive distortions from patient-therapist interactions](#). In *Proceedings of the Seventh Workshop on Computational Linguistics and Clinical Psychology: Improving Access*, pages 151–158, Online. Association for Computational Linguistics.
- Kumar Shridhar, Jakub Macina, Mennatallah El-Assady, Tanmay Sinha, Manu Kapur, and Mrinmaya Sachan. 2022. [Automatic generation of socratic subquestions for teaching math word problems](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4136–4149, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. 2023. [Distilling reasoning capabilities into smaller language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7059–7073, Toronto, Canada. Association for Computational Linguistics.
- Marina Sokolova, Vera Sazonova, Kanyi Huang, Rudraneel Chakraborty, and Stan Matwin. 2017. [Studying positive speech on twitter](#). *arXiv preprint arXiv:1702.08866*.
- Haoyu Song, Yan Wang, Wei-Nan Zhang, Xiaojiang Liu, and Ting Liu. 2020. [Generate, delete and rewrite: A three-stage framework for improving persona consistency of dialogue generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5821–5831, Online. Association for Computational Linguistics.
- Michael Tanana, Kevin Hallgren, Zac Imel, David Atkins, Padhraic Smyth, and Vivek Srikumar. 2015. [Recursive neural networks for coding therapist and patient behavior in motivational interviewing](#). In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 71–79, Denver, Colorado. Association for Computational Linguistics.
- Alexey Tikhonov, Viacheslav Shibaev, Aleksander Nagaev, Aigul Nugmanova, and Ivan P. Yamshchikov. 2019. [Style transfer for texts: Retrain, report errors, compare with rewrites](#). In *Proceedings of the*

- 2019 *Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3936–3945, Hong Kong, China. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Elsbeth Turcan and Kathy McKeown. 2019. [Dreaddit: A Reddit dataset for stress analysis in social media](#). In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, pages 97–107, Hong Kong. Association for Computational Linguistics.
- Amrisha Vaish, Tobias Grossmann, and Amanda Woodward. 2008. [Not all emotions are created equal: the negativity bias in social-emotional development](#). *Psychological bulletin*, 134(3):383–403.
- Jason Wei, Kelly Finn, Emma Templeton, Thalia Wheatley, and Soroush Vosoughi. 2021. [Linguistic complexity loss in text-based therapy](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4450–4459, Online. Association for Computational Linguistics.
- Chris Williams. 2001. [Use of written cognitive-behavioural therapy self-help materials to treat depression](#). *Advances in Psychiatric Treatment*, 7(3):233–240.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yongkang Wu, Meng Han, Yutao Zhu, Lei Li, Xinyu Zhang, Ruofei Lai, Xiaoguang Li, Yuanhang Ren, Zhicheng Dou, and Zhao Cao. 2023a. [Hence, socrates is mortal: A benchmark for natural language syllogistic reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2347–2367, Toronto, Canada. Association for Computational Linguistics.
- Zixiu Wu, Simone Balloccu, Ehud Reiter, Rim Helaoui, Diego Reforgiato Recupero, and Daniele Riboni. 2023b. [Are experts needed? on human evaluation of counselling reflection generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6906–6930, Toronto, Canada. Association for Computational Linguistics.
- Andrew Yates, Arman Cohan, and Nazli Goharian. 2017. [Depression and self-harm risk assessment in online forums](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2968–2978, Copenhagen, Denmark. Association for Computational Linguistics.
- Chujie Zheng, Pei Ke, Zheng Zhang, and Minlie Huang. 2023. [Click: Controllable text generation with sequence likelihood contrastive learning](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1022–1040, Toronto, Canada. Association for Computational Linguistics.
- Caleb Ziems, Minzhi Li, Anthony Zhang, and Diyi Yang. 2022. [Inducing positive perspectives with text reframing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3682–3700, Dublin, Ireland. Association for Computational Linguistics.

A Experimental details

We few-shot prompt GPT-4 with the default parameters from the OpenAI Python package.⁴ We use three few-shot exemplars. Following similar works on synthetic data generation with LLMs (Macina et al., 2023a), we use temperature sampling with $T = 0.4$ and no top-k truncation for generating the reframes. For the fine-tuning experiments, following Dettmers et al. (2023), we apply LoRA to all linear layers; training for 5 epochs with ADAMW, a batch size of 8 and set the learning rate to $5e-4$. All other hyperparameters take the default value in the configuration. We use NVIDIA A100 80GB for the training.

B Prompt Design

B.1 Generating Socratic Rationales

System Prompt for GPT-4

You are an expert in Cognitive Behavioral Therapy and Cognitive Restructuring, focusing on guided discovery. The term cognitive restructuring refers to the process of challenging, and changing, irrational thoughts. Socratic questioning is one technique to encourage this process. Therapists use Socratic questioning verbally by asking probing questions about their clients' irrational thoughts. As clients improve their awareness of irrational thoughts, they can begin to consciously question their own thoughts. The six types of Socratic Questions are: 1) Clarification, 2) Probing assumptions, 3) Probing reasons and evidence, 4) Questioning perspectives, 5) Probing implications and 6) Questioning the question. They can be used in any order and not all of them might be needed for a given client.

Figure 8: This system prompt is used to prime GPT-4 with a personality grounded in the Socratic method.

An example of a clinical vignette taken from Padesky (1993) used as a few-shot exemplar is:

Negative Thought: "I'm a complete failure in every way."
Positive Thought: "I am not a failure."
Socratic Questioning:
Q (Clarification): You look defeated when you say that. Do you feel defeated?
A: Yes. I'm no good.

⁴<https://github.com/openai/openai-python>

Q (Probing assumptions): You say you are no good. Is it true that you haven't done anything at all good?
A: Nothing of importance.
Q (Probing reasons and evidence): How about for your children this week -- did you care for them at all?
A: Of course, I helped my wife put them to bed and took them to soccer practice.
Q (Questioning perspectives): Do you think that was important to them?
A: I suppose so.
Q (Probing implications): And did you do anything to make your wife happy this week?
A: She liked the fact that I came home from work on time.
Q (Probing implications): Would a "complete failure" be able to respond to his wife's request in such a successful way?
A: I guess not.
Q (Probing implications): So is it really accurate to say you are a complete failure in every way?
A: I suppose not.
Q (Clarification): So how do you feel now?
A: I guess a little better.

The prompt template for GPT-4 is shown in Figure 3 and the system prompt is shown in Figure 8.

B.2 Finetuning

Prompt template used during finetuning

Input: Given a situation: «SITUATION» and the associated negative thought: «NEGATIVE THOUGHT», generate the Socratic rationale for guided discovery and reframing the negative thought to a positive thought.

Output:
«SOCRATIC RATIONALE»
«POSITIVE THOUGHT»

Figure 9: This template is used to finetune LLMs using the Socratic rationales.

The prompt template used for finetuning is shown in Figure 9.

C Human Evaluation Protocol

C.1 Evaluating Socratic Rationales

Socratic Questioning Scale (SQS) The following questions were rated by annotators on a 3-point Likert scale ranging from 1 (not at all) to 3 (extensively) :

- **How frequently were questions asked that help develop alternative perspectives?** - *This question aims to assess the frequency*

of inquiries that encourage the exploration of diverse viewpoints or opinions. Annotators should focus on identifying questions that prompt respondents to consider different angles, challenge assumptions, or think beyond the conventional narrative.

- **Was the question answering focused on the emotions and situation of the person?** - *This question assesses whether the majority of answers provided focus on the emotional state and circumstances of the individual involved. Annotators should consider whether responses primarily address the feelings, experiences, or immediate context of the person in question.*
- **Were the questions open-ended and require thoughtful reflection?** - *This question is designed to evaluate whether the questions posed to respondents are open-ended and demand deep contemplation rather than eliciting simple, direct answers. Annotators should look for questions that encourage respondents to think critically and provide nuanced, reflective responses.*

Helpfulness Annotators rate the helpfulness of the question-answer pairs on a 3-point Likert scale ranging from 1 (not helpful at all) to 3 (very helpful). The following question was asked to the annotators:

How helpful was the questioning in general?

This question seeks to gauge the overall effectiveness and utility of the questions posed. Annotators should consider the extent to which the questions contributed to a meaningful discussion, elicited insightful responses, or facilitated a deeper exploration of the topic. Assessments should encompass the clarity, relevance, and engagement level of the questions.

We plan to release the generated rationales to complement the original datasets under the Apache license.⁵

C.2 Validation with experts

We recruit psychology students for our study. The students were given credits in return of their annotation efforts. The average age of participants was 24.2 with 80% participants being females. The participants were shown three diverse samples from

across the datasets and asked to rate the helpfulness and SQS, in the same setting as described in Appendix C.1. Additionally, participants were asked to answer the following question with either "Yes", "No" or "Not sure":

Would you use an LLM to help you restructure negative thoughts? *This question assesses the willingness and openness of psychology students towards incorporating Large Language Models (LLMs) into their therapeutic practices. Responses provide insights into whether LLMs are perceived as viable tools for cognitive restructuring, indicating their potential acceptance within psychological interventions.*

We observe that only 7% of the participants did not want to use LLMs for restructuring negative thoughts with 55% responding affirmatively and the rest being unsure. We plan to conduct a more large-scale annotation study with more experts and non-experts in the future.

C.3 Evaluating Reframes

For a given context and negative thought, annotators are shown two sample reframes and asked the following question:

Given the context and original negative thought, which reframed thought do you find more relatable, helpful and memorable (A vs B)?

D Qualitative Examples

Table 5 shows qualitative examples of generated socratic rationales for each of the three datasets used in our study.

⁵<https://www.apache.org/licenses/LICENSE-2.0>

Dataset	Input	Generated Rationale
COGREF	<p>Situation: Having problems coming up with words to write essay.</p> <p>Negative Thought: I will get a bad grade.</p> <p>Reframed Thought: I'm feeling anxious about this assignment. I can use skills to calm myself, try my best on the essay, ask for any appropriate help, and whatever grade I get I will know that I tried my hardest.</p>	<p>Q: You're worried about doing badly on this assignment. How have you done on past essays?</p> <p>A: I've gotten good grades, sometimes. But this one feels different.</p> <p>Q: Do you think getting a bad grade on this specific essay means you're generally bad at writing or at this subject overall?</p> <p>A: Well, I don't think so, but I'm unsure about this topic.</p> <p>Q: What evidence is leading you to believe that you will get a bad grade?</p> <p>A: I'm having trouble coming up with points to write about.</p> <p>Q: If a friend was in your position, what advice would you give them?</p> <p>A: I would probably say to not panic and just give it their best shot.</p> <p>Q: What would happen if you got a bad grade on this essay? How would that affect your overall academic achievement?</p> <p>A: One bad grade isn't going to ruin my overall performance I guess.</p> <p>Q: Can using some support like talking to your teacher or a study group help you in getting through this?</p> <p>A: That might be a good idea.</p> <p>Q: So given these factors, do you still think it's a definite that you will get a bad grade?</p> <p>A: No, I suppose it's not definite.</p> <p>Q: So, what will you do about your essay now?</p> <p>A: I think I'll start by focusing on writing a draft based on what I know and then ask for some help to see how I can improve it.</p>
PATREF	<p>Persona: my favorite band is rise against. my favorite team is the raptors. i usually go to bed late. i like sports. i am studying to be a dentist.</p> <p>Negative Thought: I am about to graduate and become a Dentist, but that means nothing. Anyone can do that.</p> <p>Reframed Thought: I am about to graduate and become a Dentist. Although it is possible for anyone to become a dentist, it is very difficult, and I am proud of myself.</p>	<p>Q: It sounds like you're doing something remarkable, yet you downplay its value. Do you genuinely believe becoming a dentist means nothing?</p> <p>A: It doesn't seem special when everyone else is doing similar things.</p> <p>Q: Do you think everyone could become a dentist?</p> <p>A: Maybe not everyone, but a lot of people could.</p> <p>Q: It required you to put in several years of hard study and practice, didn't it?</p> <p>A: Yes, I did work hard.</p> <p>Q: How many people do you think are willing or able to do what you did?</p> <p>A: I'm not sure, probably not that many.</p> <p>Q: So wouldn't that mean that your achievement is indeed quite significant and not something anyone could do?</p> <p>A: I guess so.</p> <p>Q: Isn't it possible then that your achievement of becoming a dentist is both valuable and hard-earned?</p> <p>A: Yes, it's starting to look that way.</p> <p>Q: How do you feel about your upcoming graduation now?</p> <p>A: I'm starting to recognize its importance and feel a bit proud.</p>
POSREF	<p>Negative Thought: This physics homework is the death of me every week. It's so hard. And then I have another physics exam tomorrow night.</p> <p>Reframed Thought: Physics is challenging, but is worth it. I have an exam tomorrow, which I hope to do well on.</p>	<p>Q: It sounds quite challenging, is Physics always this difficult for you or only certain aspects?</p> <p>A: Certain parts, like quantum mechanics.</p> <p>Q: So, is it accurate to state that the entire subject is difficult, or just the quantum mechanics part?</p> <p>A: I suppose just the quantum mechanics part.</p> <p>Q: Can you give me an example of a physics problem or concept you've been able to understand quite well?</p> <p>A: Yeah, I've been doing well in classical mechanics.</p> <p>Q: Do you think that possessing skills in one area of physics could potentially aid in understanding other areas, including quantum mechanics?</p> <p>A: Maybe. But quantum mechanics seems very different.</p> <p>Q: If classical mechanics and quantum mechanics are two aspects of the same discipline, isn't there a possibility that understanding one could help with the other?</p> <p>A: It's a possibility, I guess.</p> <p>Q: Do you think saying that the physics homework is the death of you is a fair and accurate statement?</p> <p>A: Probably not, it's just really frustrating sometimes.</p> <p>Q: What if you thought about the physics homework as a challenge rather than a death sentence? Could that change how you approach it?</p> <p>A: Maybe. I could try.</p> <p>Q: Now, how do you feel about the physics exam tomorrow?</p> <p>A: I'm nervous, but maybe if I look at it as a challenge, it might feel less overwhelming.</p>

Table 5: Socratic rationales generated with GPT-4 for different types of inputs from the three datasets we consider.