

Facilitating Self-Guided Mental Health Interventions Through Human-Language Model Interaction: A Case Study of Cognitive Restructuring

Ashish Sharma
ashshar@cs.washington.edu
University of Washington
Seattle, WA, USA

Kevin Rushton
krushton@mhanational.org
Mental Health America
Alexandria, VA, USA

Inna Wanyin Lin
ilin@cs.washington.edu
University of Washington
Seattle, WA, USA

Theresa Nguyen
tnguyen@mhanational.org
Mental Health America
Alexandria, VA, USA

Tim Althoff
althoff@cs.washington.edu
University of Washington
Seattle, WA, USA

ABSTRACT

Self-guided mental health interventions, such as “do-it-yourself” tools to learn and practice coping strategies, show great promise to improve access to mental health care. However, these interventions are often cognitively demanding and emotionally triggering, creating accessibility barriers that limit their wide-scale implementation and adoption. In this paper, we study how human-language model interaction can support self-guided mental health interventions. We take *cognitive restructuring*, an evidence-based therapeutic technique to overcome negative thinking, as a case study. In an IRB-approved randomized field study on a large mental health website with 15,531 participants, we design and evaluate a system that uses language models to support people through various steps of cognitive restructuring. Our findings reveal that our system positively impacts emotional intensity for 67% of participants and helps 65% overcome negative thoughts. Although adolescents report relatively worse outcomes, we find that tailored interventions that simplify language model generations improve overall effectiveness and equity.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**; • **Computing methodologies** → **Natural language processing**.

KEYWORDS

mental health, language models, human-AI collaboration, cognitive restructuring, field study, randomized trial

ACM Reference Format:

Ashish Sharma, Kevin Rushton, Inna Wanyin Lin, Theresa Nguyen, and Tim Althoff. 2024. Facilitating Self-Guided Mental Health Interventions Through

Human-Language Model Interaction: A Case Study of Cognitive Restructuring. In *CHI'24: ACM CHI conference on Human Factors in Computing Systems, May 11–16, 2024, Hawaii, USA*. ACM, New York, NY, USA, 29 pages. <https://doi.org/10.1145/3613904.3642761>

1 INTRODUCTION

As mental health conditions surge worldwide, healthcare systems are struggling to provide accessible mental health care for all [62, 63, 83]. Self-guided mental health interventions, such as tools to journal and reflect on negative thoughts, offer great promise to expand modes of care and help people learn coping strategies [64, 71, 72, 82]. While not a replacement for formal psychotherapy, these interventions provide immediate on-demand access to resources that can help develop techniques for mental well-being, especially for those who lack access to a trained professional, are on waiting lists, or seek to supplement therapy with other forms of care [1].

However, developing interventions that individuals can effectively use without the assistance of a professional therapist is challenging [34]. Currently, most self-guided interventions simply transform traditional manual therapeutic worksheets into digital online formats with limited instructions and support [82]. Using these worksheets without professional support often leads to cognitively demanding and emotionally triggering experiences, limiting engagement and usage [7, 32, 34, 94]. For example, a popular self-guided intervention involves independently practicing *Cognitive Restructuring of Negative Thoughts*, an evidence-based, well-established process that helps people notice and change their negative thinking patterns [9, 21]. However, the practice includes complex steps like *identifying thinking traps* (faulty or distorted patterns of thinking like “*all-or-nothing thinking*”), which pose a significant challenge for many [9, 21]. Most individuals lack the necessary knowledge or experience to successfully use such interventions independently without explicit training and support. Moreover, analyzing one’s own thoughts, emotions, and behavioral patterns can be emotionally triggering, especially for those actively experiencing distress. Such accessibility barriers inhibit the widespread adoption of self-guided mental health interventions.

Language models may be able to assist individuals in engaging with self-guided mental health interventions to improve intervention accessibility and effectiveness. Specifically, language models

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CHI'24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05.
<https://doi.org/10.1145/3613904.3642761>

can help individuals learn and practice cognitively demanding tasks (e.g., through automatic suggestions on potential thinking traps in a thought). Moreover, the language model support could help users manage emotionally triggering thoughts and potentially enable a reduction in the emotional intensity of negative thoughts.

Previous work exploring self-guided mental health interventions based on human-language model interaction has predominantly been limited to small-scale [54] and wizard-of-oz-style research [46, 49, 60, 85]. These studies were conducted in controlled lab settings and evaluated on online crowdworker platforms, such as MTurk, which may not accurately represent the people who actively seek mental health care or use such an intervention [59]. Much less is known about intervention effectiveness in ecologically valid settings with individuals experiencing mental health challenges and seeking care. This limits our understanding of end-user preferences within these emerging forms of intervention [15, 17, 59, 67]. Furthermore, language models can exhibit biases resulting in highly varied performance across people from diverse demographics and populations [16, 23, 53]. There is a need to investigate and improve the equity of language modeling-based interventions.

In this paper, we study supporting self-guided mental health interventions with human-language model interaction. Specifically, we take *Cognitive Restructuring*, an evidence-based self-guided intervention, as a case study. Cognitive Restructuring helps individuals to recognize when they are getting stuck in distorted patterns of thinking (*identifying thinking traps*) and to come up with new ways to think about their situation (*writing reframed thoughts*).

We design and evaluate a human-language model interaction based tool for cognitive restructuring. We conduct an ecologically valid and large-scale randomized field study on Mental Health America (screening.mhanational.org; MHA; a popular website that hosts mental health tools and resources) with over 15,000 participants (Section 3). We examine the design of our tool, investigate its impact on people seeking mental health care, and evaluate and improve its equity across key subpopulations. Specifically, we address the following research questions:

RQ1 – Design. How can we design a self-guided cognitive restructuring intervention that is supported through human-language model interaction?

RQ2a – Overall Effectiveness. To what extent does human-language model interaction based cognitive restructuring help individuals in alleviating negative emotions and overcoming negative thoughts?

RQ2b – Impact of Design Hypotheses. What is the impact of individual design hypotheses on the intervention effectiveness?

RQ3 – Equity. How equitable is the intervention and what strategies may improve equity?

To address these research questions, we first formulate the design hypotheses for this intervention through qualitative feedback from participants of early prototypes of the system and brainstorming with mental health professionals (Section 4.1). We hypothesize that assisting users in cognitively and emotionally challenging processes, contextualizing thought reframes by reflecting on situations and emotions, integrating psychoeducation, facilitating interactive refinement of reframes, and ensuring safety will result in a

more effective human-language model interaction system for cognitive restructuring. Taking these hypotheses into consideration, we design a new system that uses a language model to support people through various steps of cognitive restructuring, including the *identification* of thinking traps in thoughts and the *reframing* of negative thoughts. Our language model suggests possible thinking traps a given thought may have, as well as suggests possible ways of reframing negative thought (Section 4.2).

After systematic ethical and safety considerations, active collaboration with mental health professionals, advocates, and clinical psychologists, and IRB review and approval, we deploy this system on MHA. We utilize a mixed-methods study design, combining quantitative and qualitative feedback, to assess the outcomes of this system on the platform visitors. We find that 67.64% of participants experience a positive shift (i.e., reduction) in their emotion intensity and 65.65% report helpfulness in overcoming their negative thoughts through the use of our system. Moreover, participants indicate that the system assists them in alleviating cognitive barriers by simplifying task complexity and emotional barriers by providing a less triggering experience (Section 5). Also, enabling participants to iteratively improve reframes by seeking additional reframing suggestions leads to a 23.73% greater reduction in the intensity of negative emotions. Moreover, participants who choose to make their reframes actionable report superior outcomes compared to those who make them empathic or personalized (Section 6).

To address the needs of individuals from diverse demographics and subpopulations, it is critical to develop equitable solutions. Here, we evaluate the performance of our system across people of different demographics and subpopulations (Section 7). The intervention is found to be more effective for individuals identifying as females, older adults, individuals with higher education levels, and those struggling with issues, such as parenting and work. However, it is found to be less effective for individuals identifying as males, adolescents, those with lower education levels, and those struggling with issues such as hopelessness or loneliness. We investigate the potential benefits of customizing the intervention for adolescents, who we found to have one of the largest disparities in our intervention outcomes. We find that making the language model-generated suggestions simpler and more casual leads to a 14.44% increase in reframe helpfulness for adolescents in the age group of 13 to 14.

We discuss the implications of our study for the use of human-language interaction in the development of self-guided mental health interventions, emphasizing the need for effective personalization, appropriate levels of interactivity with the language model, equity across subpopulations, and its safety (Section 8).

To summarize, our contributions are as follows:

- (1) Design of a novel health and wellness technology for human-language model interaction based self-guided mental health intervention.
 - We **design and evaluate a new system** that uses language models to help people through cognitive restructuring of negative thoughts, a core, well-established intervention in Cognitive Behavioral Therapy (Section 4).
 - We **formulate design hypotheses/decisions** that researchers developing self-guided interventions must consider to increase efficacy and ensure safety (Section 4).

- We **conduct randomized controlled trials** to assess the impact of different design hypotheses/decisions (Sections 5 and 6), how equitable the technology is, and what strategies may improve its equity (Section 7).
- (2) Large-scale, randomized, empirical studies in an ecologically informed setting to understand how people with lived experience of mental health interact with this technology and key takeaways for researchers developing self-guided interventions that leverage human-language model interaction.
- **There are opportunities to assist users in cognitively challenging and emotionally triggering psychological processes:** Effective self-guided mental health interventions could rapidly increase access to care. However, there are many steps in such interventions that pose cognitive and emotional challenges, leading to well-documented, high rates of participant dropout. Our work shows how the use of human-language model interaction can help address these challenges and help people use these interventions more effectively (Section 5).
 - **Certain types of suggestions are most effective:** Our study suggests that users are more likely to seek more actionable suggestions and are more likely to find actionable suggestions more helpful (Section 6).
 - **Human-language model interaction systems are not necessarily equitable, but it is possible to adapt them to individual subgroups:** Our work shows that human-language model interaction systems are likely to have heterogeneous effects across key subpopulations that require further adapting the intervention (Section 7).

2 RELATED WORK

Our work builds upon previous research on digital mental health interventions (Section 2.1), AI for mental health (Section 2.2), and the design of human-language model interaction systems (Section 2.3).

2.1 Digital Mental Health Interventions

The critical gap between the overwhelming need for and limited access to mental health care has prompted clinicians, technologists, and advocates to develop digital interventions that provide accessible care for all. Several efforts have concentrated on facilitating digital, text-based supportive conversations, through peer-to-peer support networks, such as TalkLife (talklife.com) and Supportiv (supportiv.com), as well as through on-demand talk therapy platforms like Talkspace (talkspace.com), BetterHelp (betterhelp.com), and SanVello (sanvello.com). Researchers have conducted randomized controlled trials to study the efficacy of these interventions compared to traditional methods of care, such as in-person counseling and worksheet-based skill practices [42, 58, 88].

Another key focus in this pursuit has been the development of self-guided mental health interventions [64, 71, 72, 82]. These interventions are designed in various forms, such as “*Do-It-Yourself*” apps to improve mental health “in-the-moment” of crisis, self-help tools for learning and practicing therapeutic skills, and more. Popular examples include self-guided meditation such as Headspace (headspace.com) or Calm (calm.com). Researchers have also explored the design of apps to track mood changes [75], emotion

regulation [87] and to combat loneliness [18]. Krizan et al. [48] studied the process of online self-screening of mental illnesses and its role in help-seeking. Howe et al. [41] designed and evaluated a workplace stress-reduction intervention system and found that high-effort interventions reduced the most stress. Also, Cognitive Behavioral Therapy (CBT) is an evidence-based, well-established psychological treatment [9]. Researchers have designed digital self-guided interventions that streamline elements of CBT like cognitive restructuring of negative thoughts by transforming traditional manual worksheets into digital online formats [69, 82]. Other work has focused on digital mental health interventions based on Dialectical Behavioral Therapy (DBT) [73, 74].

Our work extends this literature by investigating how to design and evaluate self-guided digital mental health interventions that leverage human-language model interaction. We take CBT-based cognitive restructuring as a case study and design and evaluate a system that assists individuals in restructuring of negative thoughts through language models in a large-scale randomized field study.

2.2 AI for Mental Health

Our work is related to the growing field of research in Artificial Intelligence (AI) and Natural Language Processing (NLP) for mental health and wellbeing.

Prior work has proposed machine learning-based methods to measure key mental health constructs including adaptability and efficacy of counselors [66], personalized vs templated counseling language [2], psychological perspective change [2, 95], staying on topic [95], therapeutic actions [50], empathy [80], moments of change [68], counseling strategies [65, 76], and conversational engagement [77]. Research has also been conducted on building virtual assistants and chatbots for behavioral health [28, 61, 84] and counseling [89]. Further, researchers have designed AI-based systems to assist mental health support providers. Tanana et al. [92] proposed a machine learning system for training counselors. Sharma et al. [78, 79] developed a GPT-2 and reinforcement learning-based system for providing empathy-focused feedback to untrained online peer supporters. Our work informs how to support self-guided mental health interventions using AI-based systems.

Prior computational work on cognitive restructuring has often relied on small-scale, wizard-of-oz style studies in controlled lab settings [60, 85]. Further, Sharma et al. [81] and Ziems et al. [97] have developed language models for automating cognitive reframing and positive reframing respectively. Sharma et al. [81] also investigate what constitutes a “high-quality” reframe and find that people prefer highly empathic or specific reframes, as opposed to reframes that are overly positive. Our work expands on these studies by investigating the design of a system that supports people through various steps of cognitive restructuring of negative thoughts. We evaluate this intervention in an ecologically valid setting with individuals experiencing mental health challenges on a large mental health platform. In addition, through randomized controlled trials, we assess the impact of key design hypotheses associated with this intervention including personalizing the intervention to the participant, facilitating iterative interactivity with the language model, and pursuing equity. Building on Sharma et al. [81], we also develop mechanisms that allow people to actively seek specific

types of reframes that are likely to be helpful (e.g., actionable) and assess its effects on different restructuring outcomes.

In their human-centered study, Kornfeld et al. [45] sought to understand the adoption of automated text messaging tools. The study revealed that the participants were interested in making the tools more personalized, favored varying levels of engagement, and wanted to explore a broad range of concepts and experiences. The design of our self-guided intervention builds on these findings aiming to facilitate intervention personalization through situations and emotions of participants, iterative engagement with the language model, and to improve the equity of the intervention across different participant issues and demographics.

Prior work has also studied mental health bias in language models. Lin et al. [53] investigated gendered mental health stigma present in masked language models and showed that models captured social stereotypes, such as the perception that men are less likely to seek treatment for mental illnesses. We study the disparities of our intervention's effectiveness among individuals from diverse demographics and facing different issues. We also propose a way of improving the interventions for a key subpopulation.

2.3 Design of Human-Language Model Interaction Systems

Broadly, our work relates to the design of human-language model interaction systems that facilitate an interactive setting in which humans can effectively engage with language models to accomplish real-world tasks [3, 51]. Examples include systems for creative writing [25], programming (e.g., CoPilot (github.com/features/copilot)), and brainstorming ideas (e.g., Jasper (jasper.ai)). Our work studies how such human-language model interaction can support self-guided mental health interventions.

3 STUDY OVERVIEW

Our study was conducted over a nine-month process involving iterative ideation, prototyping, and evaluation in collaboration with mental health experts (some of whom are co-authors).

3.1 A Case Study of Cognitive Restructuring

Cognitive Restructuring is a well-established therapeutic technique that fosters awareness of and methods for changing negative thinking patterns [9, 10]. Cognitive Restructuring has been proven to be an effective treatment strategy for psychological disorders, especially anxiety and depression [24]. It is a process that is central to Cognitive Behavioral Therapy [9], a modality of treatment which has been demonstrated to be as effective as, or more effective than, other forms of psychological therapy or psychiatric medications [22, 39].

A participant initiates this process by writing the negative thought they are struggling with. Next, they try to *identify any potential thinking traps* (biased or irrational patterns of thinking) in their thought. Thinking Traps, alternatively known as *cognitive distortions*, refer to biased or irrational patterns of thinking that lead individuals to perceive reality inaccurately [9, 31]. These typically manifest as exaggerated thoughts, such as making assumptions about what others think ("*Mind reading*"), thinking in extremes ("*All-or-nothing thinking*"), jumping to conclusions based on one

experience ("*Overgeneralizing*"), etc. Finally, participants *write a reframed thought* which involves appropriately addressing their underlying thinking traps and coming up with a more balanced and helpful perspective on the situation.

As an example, consider a PhD student who has struggled with their research project and starts to worry, "*I'll never complete my PhD*". Some possible thinking traps that this thought is falling into include *Catastrophizing* (thinking of the worst-case scenario) and *Fortune telling* (trying to predict the future). Addressing these thinking traps, a possible way in which the student can reframe the thought could be to say, "*I'm imagining the worst-case scenario. This project did not work out, but I was able to formulate meaningful hypotheses for my next research project.*" This process, along with our proposed system (Section 4.2), is also illustrated in Figure 1. Here, we study how cognitive restructuring can be supported through language models.

3.2 Study Design

Platform. We used Mental Health America (MHA) for our study, a large mental health website that provides mental health resources and tools to millions of users. We deployed our system on this platform, where it was hosted with appropriate ethical considerations and participant consent process (see Privacy, Ethics, and Safety), along with other existing self-guided mental health tools.

Participants. Many MHA visitors are interested in mental health resources including self-guided systems. Our study participants were visitors to MHA, who chose to use our system and provided informed consent to participate in the research study. Our study comprised a total of 15,531 participants.¹ We carried out several experiments in parallel, including randomized trials with independent participants. The number of participants involved in each experiment is represented by "N" values throughout the paper. Participants were 13 years and older (also, see Privacy, Ethics, and Safety below for a discussion on our reasons for including minors).

Tasks and Procedures. We started our study by formulating design hypotheses for developing self-guided cognitive restructuring through human-language model interaction. This was done through feedback from users of early prototypes of the system, brainstorming with mental health experts, and leveraging pertinent findings from prior work (Section 4.1). Based on these design hypotheses, we developed our final system for human-language model interaction powered cognitive restructuring (Section 4.2).

We evaluated the effectiveness of our system on several psychotherapy based metrics through a field study on the MHA platform (Section 5). In addition, we evaluated the importance of our design hypotheses by conducting randomized trials on the platform, explicitly ablating specific design features, such as removing psychoeducation from the system (Section 6). Finally, we analyzed and improved the equity of our system (Section 7).

Privacy, Ethics, and Safety. We designed and conducted our field study after carefully reviewing the potential benefits and risks to

¹Overall, we had 43,347 participants in total who consented initially, but 27,816 (64.17%) dropped out before completing the outcome survey. Analysis in the paper only include the 15,531 (35.83%) participants who completed the outcome survey. Note that dropout rates of online self-guided tools are commonly between 70% to 99.5% [32, 44].

participants in consultation and collaboration with mental health experts. Our study, including the participation of minors, was approved by our Institutional Review Board. We included minors (those aged between 13 and 17) in this study as they represent a large, key demographic on MHA and were already frequently using similar self-guided interventions outside of our study. Therefore, we concluded that disallowing minors access to some self-guided tools would not decrease risk but potentially make an effective intervention inaccessible to a large fraction of users. We obtained informed consent from adults and informed assent from minors, ensuring they were fully aware of the study's purpose, risks, and benefits (Appendix Figure 12 and Figure 13). Participants were informed that they would interact with an AI-based model that automatically generates reframed thoughts, without any human supervision. Further, appropriate steps were taken to avoid harmful content generation (Section 4.2), but participants were informed about the possibility that some of the generated content may be upsetting or disturbing. Also, participants were given access to a crisis hotline. Examples in this paper have been anonymized using best privacy practices [56]. Also, see Section 8.3 for broader discussion on the frameworks that guided our ethical and safety considerations.

For minors, we requested and received a waiver of parental permission from the IRB. This is because discussing mental health issues with parents may pose additional risks including discomfort in disclosing psychological distress [70, 86] and reduced autonomy [96]. Therefore, obtaining parental permission is generally avoided in similar studies [72]. Also, given our online setting, obtaining parental permission is impractical and logistically difficult as we are not directly interacting with the participants. We instead obtained the assent of minors as approved by the IRB.

4 RQ1: HOW CAN WE DESIGN A SELF-GUIDED COGNITIVE RESTRUCTURING INTERVENTION THAT IS SUPPORTED THROUGH HUMAN-LANGUAGE MODEL INTERACTION?

Here, we design a novel system for human-language model interaction based self-guided cognitive restructuring.

4.1 Design Hypotheses

The design of our system was based on hypotheses that were formulated by incorporating qualitative feedback from users of early prototypes of the system, brainstorming on design decisions with mental health experts, and leveraging relevant insights from previous mental health studies. Here, we briefly describe the design hypotheses that surfaced through this process. Our primary contributions include the design (Section 4.2) and evaluation (Section 5) of a system based on these hypotheses. In addition, we evaluate the impact of H₂, H₃, and H₄ through randomized trials (Section 6). Moreover, we evaluate H₁ and H₅ through qualitative analysis of participant feedback (Sections 5 and 6.4 respectively).

H₁: Assisting participants in processes that are cognitively and emotionally challenging may improve intervention effectiveness. Previous research suggests that restructuring negative thoughts can be *cognitively* and *emotionally* challenging [82]. The

limited availability of mental health professionals and resources often acts as barriers to accessibility, typically resulting in a lack of knowledge and exposure to therapeutic processes among people. Moreover, the entrenched nature of thoughts makes them difficult to overcome in the moment. In fact, many participants during our design exploration phase highlighted these challenges (e.g., a participant wrote, “*I struggle with formulating these thought → situation → thinking trap → reframe scenarios by myself*”; another participant wrote, “*I find it hard to think of these reframes myself in the moment*”; another participant wrote, “*I have a hard time figuring out with cognitive distortions I’m using*”). Therefore, we aim to guide participants through the identification of thinking traps and assist them in writing effective reframes. Here, we leverage language models to achieve this, as detailed in Section 4.2. Section 5 evaluates the effectiveness of this assistance.

H₂: Contextualizing thought reframes through situations and emotions may improve intervention effectiveness. Cognitive Behavioral Therapy posits that our thoughts are shaped by our situations, subsequently influencing our beliefs and emotions [9]. Therefore, the ability to reflect on situations and emotions and recognize their connection with negative thoughts could be beneficial in accurately assessing thinking traps and writing effective reframes. Our system enabled participants to contextualize their thoughts by answering additional questions related to their situations and emotions (Section 4.2). However, it is well established that introducing additional burdens in the form of questions may lead to increased dropout, a core challenge in digital mental health [7, 94]. Section 6.1 describes the randomized trial that assesses the impact of this contextualization on the tradeoff between intervention effectiveness and overall participant engagement.

H₃: Integrating psychoeducation may improve intervention effectiveness. Cognitive Restructuring is a skill people can learn. Learning this skill enables participants to identify their negative thoughts, assess the thinking traps they often fall into, and develop the ability to reframe them into something more hopeful in the moment [43, 91]. However, acquiring this skill is not straightforward and typically requires comprehensive psychoeducation coupled with practice (e.g., one participant wrote “*I don’t think that everyone knows about thinking traps and the types or kinds of thinking traps there is, so I think there should be a description or definition about each thinking traps*”). Here, we introduced participants to different thinking traps through definitions and examples, along with strategies to reframe their specific thinking traps (Section 4.2). Section 6.2 details the randomized trial that evaluates the impact of integrating psychoeducation on skill learnability.

H₄: Facilitating interactive refinement of reframes may improve intervention effectiveness. A key component of our design involves the suggestion of automatically generated reframes through a language model (H₁; Section 4.2). However, we found that participants of early prototypes of the system desired the ability to have varied AI suggestions. One participant said, “[*I want*] *more choices and variation in the reframing*”. Another participant said “*I wish it included actionable insights*.” Moreover, many participants desired the ability to interactively refine the reframes. One participant said, “*I need to write the information and revise it as many times*

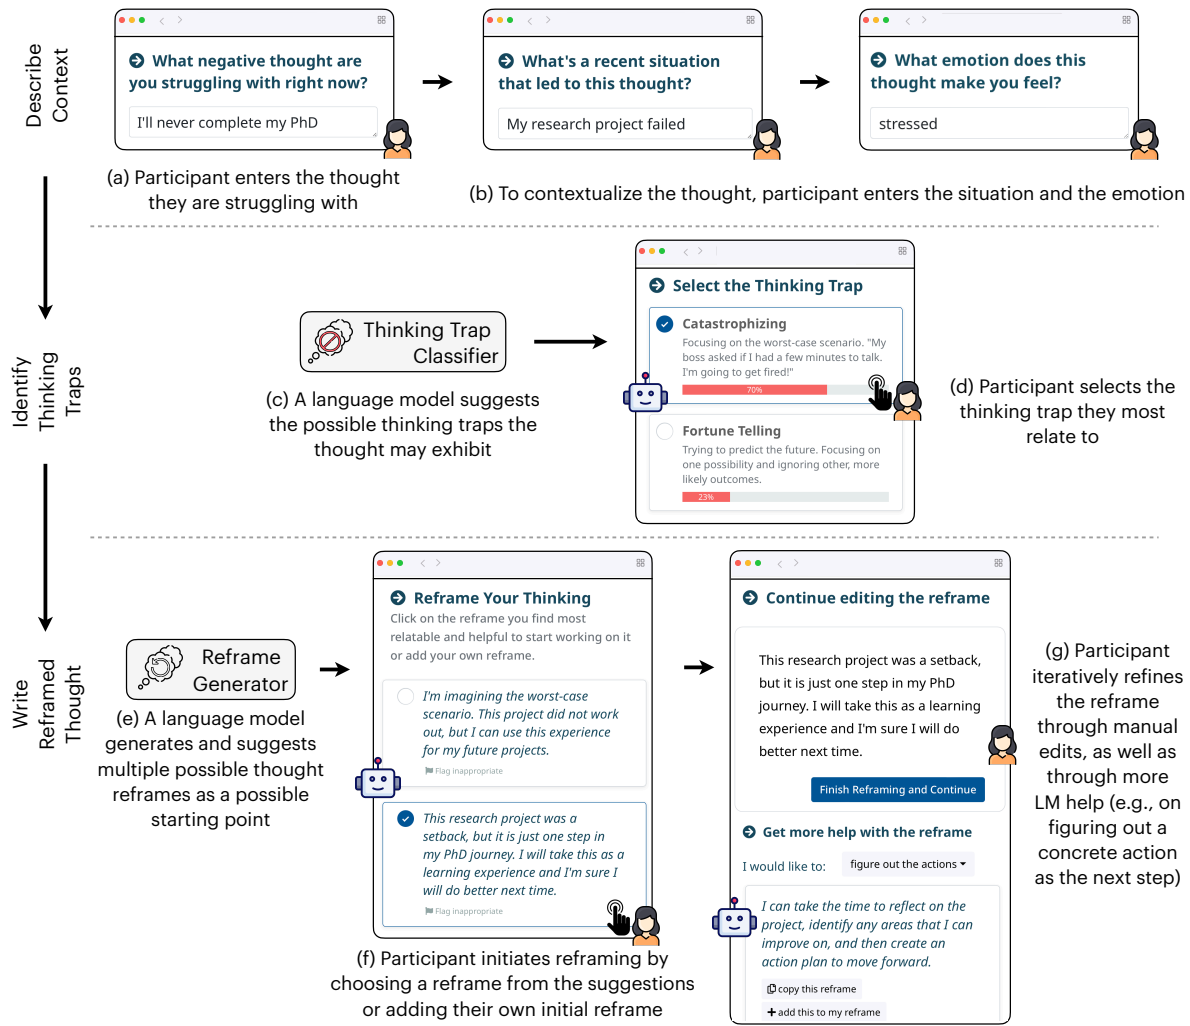


Figure 1: We design a human-language model interaction based system for self-guided cognitive restructuring of negative thoughts. The system involves (a-b) describing the context by participants, (c-d) LM-assisted identification of thinking traps, and (e-g) LM-assisted writing of reframed thoughts.

as possible.” Another participant said, “It could be more interactive, also go more in depth.”

Therefore, our design included the option for iterative editing and updating of reframes. This enabled participants to seek more specific suggestions from the language model, including “making it more relatable to their situation”, “figuring out the next steps and actions”, and “feeling more supported and validated” (Section 4.2). Section 6.3 evaluates the effects of this interactive interaction through a randomized trial.

H5: Mechanisms to avoid unsafe content generation and flag inappropriate content may make the intervention more safe. The need for safety considerations is crucial when intervening in high-stakes settings like mental health [52, 55]. There is a risk that AI might inadvertently harm, rather than help, individuals coping

with mental health challenges. Therefore, a key goal was to ensure safety and minimize risks. Our design included mechanisms to avoid unsafe content generation and flag inappropriate content, as described in Section 4.2. We discuss the content that was flagged by participants in Section 6.4.

4.2 System Design

In our system (Figure 1), we guided participants through a five-step process. This included (1) describing the thought, (2) detailing the situation, (3) reflecting on the emotion, (4) identifying the thinking traps, and (5) finally, reframing the thought.

Step I, II, and III: Participant describes the context. On selecting to use the system and after consenting to participate in our study (Appendix Figure 12 and Figure 13), a participant first

articulates the thought they are struggling with (e.g., “*I’ll never complete my PhD*”; Figure 1a). Next, to contextualize their thought (H₂; Figure 1b), the participant describes a recent situation that may have led to this thought (e.g., “*My research project failed*”). Additionally, the participant reflects on the emotion they are currently experiencing along with its intensity (e.g., “*stressed*”; 9 out of 10).

Step IV: LLM-assisted selection of thinking traps. The next step is to identify thinking traps. The typical process of identifying thinking traps involves participants navigating through a list to single out possible traps in their thoughts, which we identified as a cognitive and emotional barrier (H₁). Here, we use a language model to assist participants in the identification of thinking traps among 13 common thinking traps (see Appendix Table 6).

For this, we rank the thinking traps for the given thoughts using a language model and show those to the participants along with their individual estimated likelihoods (e.g., “*Catastrophizing – 70%*; *Fortune Telling – 23%*; *Overgeneralizing – 7%*”; Figure 1c). In order to incorporate psychoeducation, we provide definitions and examples of these thinking traps (H₃; Appendix Table 6). We use the GPT-3 model [19] finetuned over a dataset of thinking traps by Sharma et al. [81]. This model achieves a top-1 accuracy of 62.98% on the 13-class thinking trap classification problem.

The participant selects one or more thinking traps from this ranked list that they most closely identify with based on their thinking pattern (e.g., “*Catastrophizing*”; Figure 1d).

Step V: LLM-assisted writing of reframes. Finally, the participant writes a reframed thought addressing their thinking traps. Writing reframes to negative thoughts while maintaining composure in the moment is challenging (H₁), therefore, we use a language model to assist participants. Our language model is based on Sharma et al. [81], who propose a retrieval-enhanced in-context learning method using GPT-3 [19]. Given a new thought T_i and a situation S_i , this model retrieves k -similar examples from a dataset of $\{(situation, thought, reframe), \dots\}$ triples collected from mental health experts. Those k examples are used as an in-context prompt for the GPT-3 model to generate reframes for T_i and S_i ($k = 5$).

Using this language model, we generate multiple suggestions for possible thought reframes and show them to the participant as a potential starting point (Figure 1e).² The multiple suggestions are aimed at offering varying perspectives to the participant’s original thought (e.g., “*I’m imagining the worst-case scenario. This project did not work out, but I can use this experience for my future projects.*”, “*This research project was a setback, but it is just one step in my PhD journey. I will take this as a learning experience and I’m sure I will do better next time.*”, “*I am disappointed that my research project failed, but I can still complete my PhD if I keep working hard and don’t give up.*”³). In addition, to incorporate psychoeducation, we provide instructions on ways in which the participant can reframe the specific thinking traps selected by them in the previous step (e.g., “*Tips to overcome catastrophizing*”; H₃; Appendix Table 6).

The participant initiates reframing by choosing a reframe from the initial suggestion list or by writing a reframe on their own (Figure 1f). The participant then iteratively refines this reframe through manual edits, as well as through additional, optional help from the language model. For this, we provide participants with an option to seek more specific suggestions from the language model, including “*making it more relatable to their situation*”, “*figuring out the next steps and actions*”, and “*feeling more supported and validated*” (H₄; Figure 1g). For the option selected by the participant, we generate additional suggestions using the language model. Participants can either copy these additional suggestions, add them to their initial reframe, replace their initial reframe with it, or use it as an inspiration. Also, see Appendix Figure 14 for the detailed interface.

Safety considerations. To minimize harmful outputs generated by language models (H₅), we combined classification-based content filtering with rule-based content filtering. We use a classification-based content filtering system provided by Azure OpenAI (bit.ly/azure-content-filter) which identifies and filters out content related to “hate”, “sexual”, “violence”, and “self-harm” categories. In addition, we developed a rule-based method to filter out any generated content that contained words or phrases related to suicidal ideation or self-harm. To achieve this, we created a list of 50 regular expressions (e.g., to identify phrases like “*feeling suicidal*”, “*want to die*”, and “*harm myself*”) based on suicidal risk assessment lexicons such as the one by Gaur et al., 2019 [35]. A language model-generated reframe suggestion that matched any of the regular expressions was filtered out and not suggested to the participants. Also, participants were given the option to flag inappropriate reframing suggestions through a “*Flag inappropriate*” button (Figure 1f; Section 6.4).

We deployed this system on the MHA platform and studied its effectiveness with platform visitors (see Section 3.2). We make the code used to design the system publicly available at github.com/behavioral-data/Self-Guided-Cognitive-Restructuring.

5 RQ2A – TO WHAT EXTENT DOES OUR INTERVENTION HELP INDIVIDUALS IN ALLEVIATING NEGATIVE EMOTIONS AND OVERCOMING NEGATIVE THOUGHTS?

We used a mixed-method approach to evaluate the effectiveness of our system. Here, we first describe the different quantitative and qualitative measures used in our study, followed by the evaluation of our system on these metrics.

5.1 Quantitative Measures

Drawing from various metrics prevalent in the cognitive behavioral therapy literature, we conducted a comprehensive analysis of our system. We assessed the effects of our system on the participants’ emotions, the efficacy of the reframes they wrote, and learnability of the skill:

- (1) **Reduction in Emotion Intensity:** Intensity of the participant’s emotion before the system use – Intensity of the participant’s emotion after the system use. We collected the emotion associated with the participant’s negative thought before the system use (“*What emotion does this thought make*

²To generate multiple suggestions, we perform top-p sampling [40] multiple times. We show three suggestions by default, but participants are provided with an option to seek more suggestions if needed.

³The reframing suggestions in this example have been generated by our GPT-3 based model.




Figure 2: (a) Reduction in emotion intensity of participants before and after using the system (emotion scale: 1 to 10). We found that 67% of the participants reported having a positive reduction in (negative) emotions (N=1,922). (b) Participants with higher emotion intensity before using the system reported a higher reduction in emotion intensity post the system usage (N=1,922). Error bars represent 95% bootstrapped confidence intervals.

you feel?”) and collected its intensity both before and after the system use (“How strong is your emotion? (1 to 10)”).

- (2) **Reframe Relatability:** After the system use, we asked the participant: “How strongly do you agree or disagree with the following statement? – I **believe** in the reframe I came with” (1 to 5; 1: Strongly Disagree; 5: Strongly Agree).
- (3) **Reframe Helpfulness:** After the system use, we asked the participant: “How strongly do you agree or disagree with the following statement? – The reframe **helped** me deal with the thoughts I was struggling with” (1 to 5; 1: Strongly Disagree; 5: Strongly Agree).
- (4) **Reframe Memorability:** After the system use, we asked the participant: “How strongly do you agree or disagree with the following statement? – I will **remember** this reframe the next time I experience this thought” (1 to 5; 1: Strongly Disagree; 5: Strongly Agree).
- (5) **Skill Learnability:** After the system use, we asked the participant: “How strongly do you agree or disagree with the following statement? – By doing this activity, I **learned** how I can deal with future negative thoughts” (1 to 5; 1: Strongly Disagree; 5: Strongly Agree).

5.2 Qualitative Measures

We also collected subjective feedback from participants. At the end of the system usage, we asked an optional open-ended question “We would love to know your feedback. What did you like or dislike about the tool? What can we do to improve?”

5.3 Results: Quantitative

67.64% of the participants reported a positive change in emotions. We assessed the difference in the intensity of participants’ self-reported negative emotions before and after utilizing our system (N=1,922). Figure 2a shows the distribution. Our findings revealed a positive emotional shift in 67.64% (1,300) of the participants, while 24.56% (472) of the participants reported no change in their

Outcome Measure	Mean	Std
Reduction in Emotion Intensity (-10 to 10)	1.90	1.29
Reframe Relatability (1 to 5)	3.84	1.17
Reframe Helpfulness (1 to 5)	3.33	1.35
Reframe Memorability (1 to 5)	3.52	1.36
Skill Learnability (1 to 5)	3.39	1.39

Table 1: Mean and standard deviation of the five quantitative measures as reported by participants.

Outcome Measures	Initial Emotion Intensity	
	≤ 7	> 7
Reduction in Emotion Intensity (-10 to 10)	0.95	2.13
Reframe Relatability (1 to 5)	3.98	3.65
Reframe Helpfulness (1 to 5)	3.49	3.09
Reframe Memorability (1 to 5)	3.67	3.33
Skill Learnability (1 to 5)	3.57	3.14

Table 2: Participants with higher emotion intensity before using the system reported a higher reduction in emotion intensity post the system usage. Participants with higher initial emotion intensities reported worse reframing outcomes, suggesting that writing effective reframes and learning the cognitive restructuring skill was harder when individuals were emotionally agitated (N=1,922).

emotion intensity. A small 7.80% (150) of the participants reported a negative shift in their emotions, with the majority (72%; 108) of them experiencing a relatively minor negative shift of -1.

Participants with higher initial emotion intensity experienced a greater improvement in emotions. We checked the effects of our system on participants with different initial emotion intensities. We found that participants with more intense initial emotions (>7 out of 10) reported 124.21% more substantial positive shifts in their emotional state than those with less intense initial emotions (2.13 vs. 0.95; ≤ 7 out of 10; N=1,922; $p < 10^{-54}$; Fig 2b). Psychotherapy research suggests that a higher intensity of negative moods and depression is associated with stronger negative cognition and maladaptive thoughts [12]. Our findings indicate that participants with greater initial emotional intensity could have a greater benefit from a cognitive restructuring intervention like ours, potentially due to the positive effects it has on their cognition, which in turn may positively effect mood and emotion.

Majority of participants found the reframes believable, helpful, and memorable. We evaluated the effectiveness of the reframes that people are able to write using our system (N=1,922). Overall, we found that 80.49% of the participants found the reframes relatable to them, 65.65% of participants found the reframes helpful in overcoming negative thoughts, and 70.49% of participants found the reframes memorable or easy to remember. Further investigating people with different emotional intensities, we found that people with higher initial emotional intensities (> 7 out of 10) reported

⁴We use a Two-sided student’s t-test for all statistical tests in this paper.

8.29% lower reframe relatability (3.65 vs. 3.98; $p < 10^{-5}$), 11.46% lower reframe helpfulness (3.09 vs. 3.49; $p < 10^{-5}$), 9.26% lower reframe memorability (3.33 vs. 3.67; $p < 10^{-5}$), and 12.04% lower skill learnability (3.14 vs. 3.57; $p < 10^{-5}$) than those with lower initial emotional intensities (≤ 7 out of 10; $N=1,922$; Table 2). This suggests that when individuals are emotionally agitated, it is harder to come up with effective reframes and learn cognitive restructuring.

Most participants found the system helpful in learning cognitive restructuring. We assessed how effectively our system can be used to learn the skill of managing negative thoughts. We found that 67.38% of the participants reported that the system helped them in learning how to deal with negative thoughts.

We also report mean and standard deviations of the quantitative outcome measures in Table 1.

5.4 Results: Qualitative

On analyzing the qualitative feedback from participants, we observed that participants highlighted three key ways in which they found assistance from our system (H_1).

First, many participants indicated that the system helped them overcome cognitive barriers, especially when they “feel stuck”, and doing this exercise is “difficult”, “on their own” and “in the moment.” A participant wrote, “My own reframes are difficult, and AI gives multiple other perspectives to consider.” Also, some participants reported that it helped them find “the right words” or “ideas to start with.” A participant wrote, “Thank you for helping me to find the right words to clearly reframe a negative thought and how to apply the thought to my own thinking processes.” Another noted, “I appreciated that the option of having the AI tool walk you through the reframing process step by step (e.g., by choosing the negative thought you may be experiencing + giving possible reframing ideas to start with/add more details to).”

Second, participants expressed how the system enabled a less emotionally triggering experience. One participant wrote, “I felt in control and more comforted that I can handle difficult situations with confidence.” Another participant wrote, “This activity let me calm down...”. Another participant noted, “...this made the process much less daunting...”. This is perhaps consistent with the quantitative findings on reduced emotion intensity (Section 5.3).

Third, participants valued that the system allowed them to explore multiple viewpoints. One participant wrote, “...After reading several reframes and looking over them I realized that there are many options, many positive sides.” Another participant wrote, “I felt reassured to see multiple views, and reflect upon them...”

Overall, these results suggest that there are opportunities to assist participants in cognitively challenging and emotionally triggering psychological processes through human-language model interaction.

6 RQ2B – WHAT IS THE IMPACT OF INDIVIDUAL DESIGN HYPOTHESES ON THE EFFECTIVENESS OF THE INTERVENTION?

Here, we studied the impact of individual design hypotheses (Section 4.1) on the effectiveness of the intervention and overall participant engagement. To facilitate this, we deployed different design

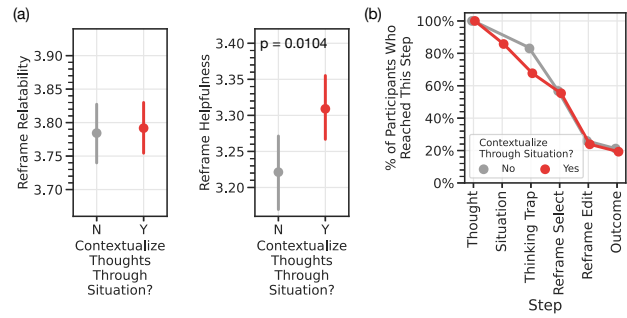


Figure 3: Randomized controlled trial to estimate the effects of contextualizing thoughts through situation ($N=1,636$). (a) Contextualizing participant thoughts through their situations led to 2.80% more helpful reframes (3.31 vs. 3.22; $p = 0.0104$) but did not lead to more relatable reframes. (b) Asking for additional context did not lead to a lower completion rate. Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$.

variations of our system by ablating specific design features (e.g., one variation that includes psychoeducation and another variation that removes it). For each design ablation, we conducted randomized controlled trials in which incoming participants were randomly assigned one of the two design variations (Appendix Table 5). To measure the impact, we evaluated the difference in outcomes between participants involved in the two design variations. In this section, we report the results from ablating contextualization (H_2), psychoeducation (H_3), and interactivity (H_4).

6.1 Contextualizing through Situations Improves Reframe Helpfulness

Reflecting on situations and emotions and understanding their connection with negative thoughts can help in writing more personalized and effective reframes (H_1). However, asking participants for additional information like descriptions of a relevant situation and emotion can potentially increase dropout which could prevent successful outcomes. Therefore, the reflection process comes with a tradeoff with higher participant burden and higher dropout rates. Here, we conducted two different randomized trials. One where we enabled contextualization through situations to half of the participants at random. Another where we enabled contextualization through emotions to half of the participants at random.

We found that contextualizing participant thoughts through their situations led to 2.80% more helpful reframes and similar levels of relatability (3.31 vs. 3.22; $N=1,636$; $p = 0.0192$; Figure 3a). This indicates the benefits of increased reflection in self-guided mental health interventions. Also, while typically an increased information request is correlated with a higher dropout rate, we found that a similar number of participants reached the end of the tool, regardless of the additional information requested (Figure 3b).

Surprisingly, participants who contextualized their thoughts through emotions reported 3.86% lower levels of relatability (3.87

vs. 3.72; $N=4,016$; $p < 0.001$; Appendix Figure 7). This may be because our language model does not incorporate emotions while identifying thinking traps or suggesting reframes, due to the lack of relevant cognitive restructuring dataset containing self-reported emotion annotations. Consequently, this could possibly lead participants to develop unwarranted expectations where they anticipate their emotional states to be addressed in the reframing suggestions, even when they are not. Note that this is different for descriptions of situations, which the language model does take into account and typically reflects in the generated reframes.

Qualitative feedback from participants indicated that they desired the inclusion of these steps as it helped them better process their thoughts (e.g., a participant wrote “*What made it especially helpful was being able to contextualize my feelings, which I feel allows for a more relatable reframe*”).

6.2 Integrating Psychoeducation has Limited Impact on Overall Effectiveness

Providing psychoeducation with the intervention may help people learn the cognitive restructuring skill more effectively. While users of early prototypes expressed interest in integrating psychoeducation (H₃; Section 4.1), our randomized trial indicated that it did not lead to significant quantitative improvement in outcomes including skill learnability as self-reported by participants (at $\alpha = 0.05$; $N=1,850$; Appendix Figure 7).

Nevertheless, qualitative feedback from participants indicated that they found the provided definitions, examples, and strategies helpful (e.g. one participant wrote, “*I like how the tool provided explanations*”; another participant wrote, “*I like the simple explanations and examples for each thought trap*”).

6.3 Increased Interactivity with the Language Model is Associated with Improved Outcomes

We provided participants with an option to seek more specific suggestions from the language model, including “*making the reframe more relatable to their situation*”, “*figuring out the next steps and actions*”, and “*feeling more supported and validated*” (H₄; Figure 1g; Appendix Figure 14).

In our randomized trial, where only half of the participants at random were given this option, we found that having this option available led to a 23.73% greater reduction in emotion intensity (2.19 vs. 1.77; $p = 0.0192$; $N=2,165$) and insignificant differences in other outcomes (Figure 4). Also, some participants appreciated having this option. One participant wrote, “*I’m glad there was an option to get supported and validated*”.

Further, among the study participants who were provided this option, we observed that 38.60% of them made use of it. Those who chose to use it to further interact with the language model to seek additional reframing suggestions of specific types (actionable, empathic, or personalized) reported 5.57% higher reframe helpfulness (3.41 vs. 3.23; $p < 10^{-5}$) and 4.86% higher skill learnability (3.45 vs. 3.29; $p < 0.001$) than participants who did not use it ($N=992$; Appendix Figure 9). Moreover, those who chose to make their reframes actionable during this step (by choosing the option, “*I want to figure*

out the next steps and actions”) reported significantly superior effectiveness across all five outcomes than those who did not (Figure 5a). Prior work has highlighted the importance of behavioral activation which involves engaging in behaviors or actions that may help in overcoming negative thoughts [20, 30]. Our work shows that participants explicitly seeking out actionable reframed thoughts are more likely to report better outcomes.

Moreover, those who chose to make their reframe empathic reported a 21.86% higher reduction in emotion intensity (2.23 vs. 1.83; $p = 0.0204$), 5.52% higher reframe helpfulness (3.44 vs. 3.26; $p < 0.001$) and 5.14% higher skill learnability (3.48 vs. 3.31; $p < 0.001$) and no significant differences based on reframe relatability and reframe memorability (at $\alpha = 0.05$; $N=992$; Figure 5b). We did not find significant differences based on whether a participant chose to make a reframe personalized or not (at $\alpha = 0.05$; $N=992$; Figure 5c).

6.4 Analysis of the Language Model Generated Content that was Flagged Inappropriate

We implemented a feature for participants to report any inappropriate content generated by the language model. This was achieved by including a distinct “Flag inappropriate” button for each generated reframe (Section 4.2). Overall, we found that 0.65% (301 out of 46,593) of the reframing suggestions shown were flagged. After conducting a qualitative review of these flagged reframes, we found a few dozen instances where the model’s suggestions repeated the negative factors/sentiment described by the participant in the original thought which may inadvertently reinforce negative beliefs about oneself (e.g., “*I may be a “failure”, but I’m still trying my best.*” in response to the thought “*I’m a failure*”). Note that in most cases repeating parts of the participant’s thought or situation helps to validate their experience and emotional reaction and personalize their reframe. Therefore, this highlights the importance of effectively differentiating which aspects of the participant’s thoughts to re-state in the reframing suggestions. Future work should look more closely at how to facilitate this differentiation.

Nevertheless, for many of the 301 flagged instances, because all participants were shown three reframes as starting points, participants were able to select a different reframing suggestion from the three options presented to them, eventually reporting favorable outcomes. We also checked user dropout between users who flagged content vs. those who did not and found no significant differences. Interestingly, the dropout rate was slightly lower among users who flagged content, at 38.2%, compared to 46.5% for those who did not. Notably, none of the flagged instances had references to suicidal ideation or self-harm, suggesting that the safety mechanisms designed to address these concerns were likely effective (Section 4.2).

7 RQ3 – HOW EQUITABLE IS THE INTERVENTION AND WHAT STRATEGIES MAY IMPROVE ITS EQUITY?

Next, we assess how equitable our intervention is across the issues expressed by participants (Section 7.1) and across participant demographics (Section 7.2). Moreover, we work towards improving equity of our system by improving its effectiveness for a specific

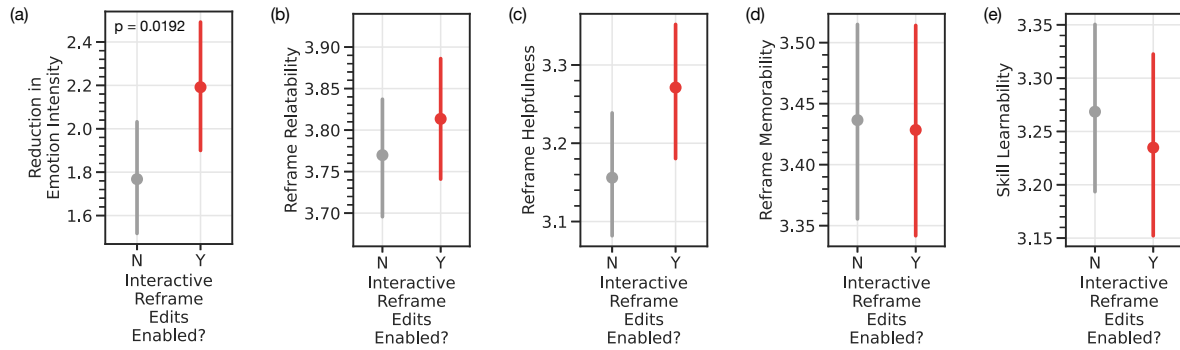


Figure 4: Randomized controlled trial to estimate the effects of enabling participants to iteratively edit reframes through increased interaction with the language model (N=2,165). 38% of participants chose to use this intervention. (a) Having the option of interactive reframe edits available to participants led to a 23.73% greater reduction in emotion intensity (2.19 vs. 1.77). (b-e) However, it did not lead to significant differences in other outcomes (at $\alpha = 0.05$). Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$.

subpopulation experiencing one of the worst outcome disparities, adolescents (Section 7.3).

7.1 Assessing Outcomes across Issues

To better determine the effectiveness of our intervention in various scenarios, we assessed the outcomes across different types of situations and thoughts that individuals might experience. We characterized participants' situations and thoughts based on the broader issues that they relate to. In collaboration with mental health experts (some of whom are co-authors), we manually labeled 500 thoughts and situations to identify the potential issues that they are associated with. The result of this iterative open-ended coding process was a set of 16 different issues expressed by participants. These include *Body Image*, *Dating & Marriage*, *Family*, *Fear*, *Friendship*, *Habits*, *Health*, *Hopelessness*, *Identity*, *Loneliness*, *Money*, *Parenting*, *School*, *Tasks & Achievement*, *Trauma*, and *Work*. See Appendix Table 7 for their definitions and examples.

We used this dataset to finetune a GPT-3 model (text-davinci), which achieved an accuracy of 73.00% on a held-out set (random performance 6.25%). We used this model to analyze the outcomes for people experiencing different issues and to identify the issues where our intervention performed better or worse (Table 3).

We found that participants expressing *Hopelessness* and *Loneliness* related thoughts reported worse outcomes relative to other issues. Participants with *Hopelessness* (e.g., “I will never be better”) reported 41.27% lower reduction in emotion intensity (1.11 vs. 1.89; $p = 0.0220$), 8.33% lower reframe relatability (3.42 vs. 3.72; $p = 0.0200$), 16.61% lower reframe helpfulness (2.66 vs. 3.19; $p < 0.001$), 10.53% lower reframe memorability (3.06 vs. 3.42; $p = 0.0169$), and 12.07% lower skill learnability (2.84 vs. 3.23; $p = 0.0119$) than the population means. Moreover, participants with *Loneliness* (e.g., “I feel like no one is with me”) reported 8.45% lower reframe relatability (3.43 vs. 3.72; $p = 0.0031$), 14.11% lower reframe helpfulness (2.74 vs. 3.19; $p < 0.001$), 11.40% lower reframe memorability (3.03 vs. 3.42; $p < 0.001$), and 14.24% lower skill learnability (2.77 vs. 3.23; $p < 0.001$) than the population means. These differences could suggest that thoughts related to some issues are

more challenging to overcome than others (as also suggested in psychology theory [11, 37, 38]) or that they represent a different subpopulation. However, we also found a lower reduction in emotion intensity (i.e., an outcome measured pre- and post-intervention), suggesting that our system might have had greater difficulty in assisting these issues. In fact, some participants commented that the reframing suggestions did not work well for issues that were too complex and nuanced. One participant wrote, “It might be too simple for more complicated problems.” Another participant wrote, “More complex problems need more precise results in my opinion.” Some participants thought that the suggestions to such complex problems were “superficial”, “artificial”, or “hard to relate to.” Future iterations of the system could benefit from designing more sophisticated language modeling solutions for complex issues.

We also observed *Tasks & Achievement* related thoughts (e.g., “I can’t finish my work”) to have 4.30% lower reframe relatability (3.56 vs. 3.72; $p = 0.0299$), 5.56% lower reframe memorability (3.23 vs. 3.42; $p = 0.0274$), and 7.42% lower skill learnability (2.99 vs. 3.23; $p = 0.0071$). Qualitative feedback from participants with such thoughts revealed that they often sought concrete actions beyond what the reframe suggestions could offer.

Moreover, those who used our system for *Parenting* and *Work* reported significantly better outcomes than the population means. Those with *Parenting* issues reported 12.63% higher reframe relatability (4.19 vs. 3.72; $p = 0.0102$), 15.67% higher reframe helpfulness (3.69 vs. 3.19; $p = 0.0145$), and 16.08% higher reframe memorability (3.97 vs. 3.42; $p = 0.0080$). And those with *Work* issues reported 22.22% higher reduction in emotion intensity (2.31 vs. 1.89; $p = 0.0197$), 4.57% higher reframe relatability (3.89 vs. 3.72; $p = 0.0198$), 10.97% higher reframe helpfulness (3.54 vs. 3.19; $p < 0.001$), 10.23% higher reframe memorability (3.77 vs. 3.42; $p < 0.001$), and 10.84% higher skill learnability (3.58 vs. 3.23; $p < 0.001$).

7.2 Assessing Outcomes across Demographics

Language modeling interventions are known to be biased toward people of specific demographics. For interventions targeting mental

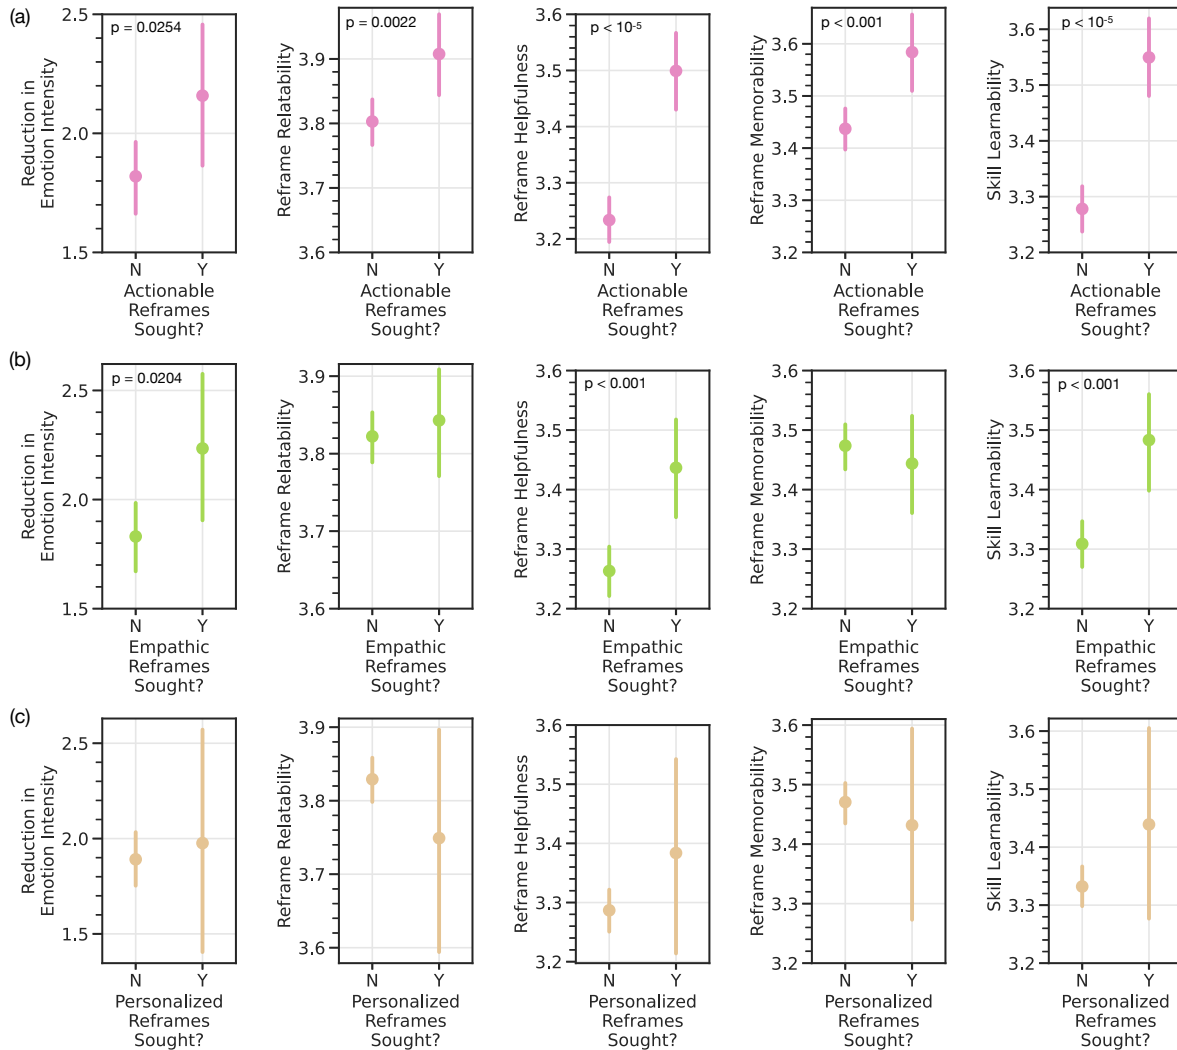


Figure 5: Participants were provided with an option to seek one or more of the following types of more specific suggestions from the language model – actionable, empathic, or personalized. Among those who used any of the options (N=992), (a) those who chose to make their reframes actionable experienced superior effectiveness across all five outcomes; (b) those who chose to make their reframes empathic reported 21.86% higher reduction in emotion intensity (2.23 vs. 1.83), 5.52% higher reframe helpfulness (3.44 vs. 3.26), and 5.14% higher skill learnability (3.48 vs. 3.31) and no significant differences based on reframe reliability and reframe memorability (at $\alpha = 0.05$); (c) those who chose to make their reframes personalized reported no significant differences in outcomes (at $\alpha = 0.05$). Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$.

health, previous research has found that language models are likely to perpetuate social stereotypes, for example, under-emphasizing men’s mental health [53]. Broadly, this corresponds to the principle of demographic parity in the fairness in machine learning literature [16, 23, 57].

Here, we studied the difference in outcomes of our intervention across participants of different demographics. We asked participants to optionally provide demographic information, including age (ranges between 13 to 65+), gender (Female, Male, or Non-Binary), race/ethnicity (American Indian or Alaska Native, Asian,

Black or African American, Hispanic or Latino, Middle Eastern or North African, Native Hawaiian and Pacific Islander, White, More than One, or Other), and education levels (Middle School, High School, Undergraduate, Graduate, or Doctorate).

Table 4 reports the outcome differences. We found that participants aged 17 or younger reported 4.84% lower reframe reliability (3.54 vs. 3.72; $p = 0.0091$), 10.03% lower reframe helpfulness (2.87 vs. 3.19; $p < 0.001$), 10.23% lower reframe memorability (3.07 vs. 3.42; $p < 0.001$), and 10.22% lower skill learnability (2.90 vs. 3.23; $p < 0.001$) compared to the population mean. On the other hand,

Issues	Reduction in Emotion Intensity	Reframe Relatability	Reframe Helpfulness	Reframe Memorability	Skill Learnability	N
Body Image	1.42	3.89	3.20	3.49	3.38	71
Dating & Marriage	2.05	3.79	3.20	3.47	3.33	328
Family	1.99	3.78	3.26	3.35	3.34	170
Fear	1.63	3.53	3.07	3.26	3.20	123
Friendship	1.91	3.65	3.20	3.48	3.20	159
Habits	1.72	3.98	3.50	3.52	3.57	42
Health	2.36	3.91	3.45	3.77	3.47	53
Hopelessness	1.11	3.41	2.66	3.06	2.84	70
Identity	2.54	4.00	3.55	3.64	3.09	11
Loneliness	1.56	3.43	2.74	3.03	2.77	146
Money	1.71	3.73	2.80	3.43	3.17	30
Parenting	2.06	4.19	3.69	3.97	3.61	36
School	1.94	3.79	3.20	3.34	3.13	181
Tasks & Achievement	1.65	3.56	3.04	3.23	2.99	232
Trauma	1.42	3.33	2.58	3.00	2.50	12
Work	2.31	3.89	3.54	3.77	3.58	258

Table 3: Effectiveness of our system across different issues expressed by participants. Numbers highlighted in green indicate outcomes that are significantly better than the population mean ($p < 0.05$). Numbers highlighted in red indicate outcomes that are significantly worse than the population mean ($p < 0.05$). We found that participants who expressed *Parenting* and *Work* related issues reported better outcomes than the population means. Moreover, participants who expressed *Hopelessness*, *Loneliness*, and *Tasks & Achievement* related issues reported worse outcomes.

those aged 25 or above reported significantly better outcomes overall. This suggests that our intervention is less effective for adolescents and more effective for adults. Section 7.3 explores improving the effectiveness of our intervention for adolescents.

Moreover, we found that male participants reported 7.52% lower reframe helpfulness (2.95 vs. 3.19; $p = 0.0162$), and 5.88% lower skill learnability (3.04 vs. 3.23; $p = 0.0364$) than the population mean. This is consistent with prior work that shows that language models are likely to be disparate toward men’s mental health [53]. Race or ethnicity of the participants was not consistently associated with better or worse outcomes. However, those who identified their race as “Other” reported a 58.73% lower reduction in emotion intensity than the population mean (0.78 vs. 1.89; $p = 0.0058$). Moreover, those who identified their race as “Middle Eastern or North African” reported 12.87% lower reframe memorability (2.98 vs. 3.42; $p = 0.0180$) and 13.31% lower skill learnability (2.80 vs. 3.23; $p = 0.0207$).

Finally, based on education levels, participants with a “Middle School” education reported 9.40% lower reframe helpfulness (2.89 vs. 3.19; $p = 0.0249$), 13.45% lower reframe memorability (2.96 vs. 3.42; $p < 0.001$), and 13.62% lower skill learnability (2.79 vs. 3.23; $p < 0.001$). On the other hand, those who identified as “Graduate” reported 6.99% higher reframe relatability (3.98 vs. 3.72; $p = 0.0020$), 8.78% higher reframe helpfulness (3.47 vs. 3.19; $p < 0.001$), 7.89% higher reframe memorability (3.69 vs. 3.42; $p = 0.0019$), and 8.98% higher skill learnability (3.52 vs. 3.23; $p = 0.0017$). Note that age and education are strongly correlated (pearson’s correlation = 0.62), especially for younger participants that did not yet have the time to

advance their education, suggesting that the relationship between education and outcomes may be at least partially explained by age.

7.3 Improving Intervention Equity by Improving the Experience of Adolescents

Because intervention effectiveness differs significantly across people’s issues and demographics, it is crucial to identify solutions that improve intervention equity. This may require adapting the intervention to different subpopulations.

Here, we performed one specific experiment to study how our language model-based intervention may be adapted to make it more equitable. We particularly focused on teenagers and adolescents whom we found to have one of the largest outcome discrepancy.⁵

Research suggests that current treatment methods are often structurally incompatible with the ways adolescents engage with, or wish to engage with, mental health care [47]. Although adolescents are more likely to use self-guided mental health interventions [71], our analysis suggests that our human-language modeling based intervention may be less effective for this demographic (Section 7.2). Given the escalating youth mental health crisis [5], it is important to develop solutions that bridge this gap. To achieve this, we tried and identified the challenges that uniquely affect adolescents.

We hypothesized that the linguistic complexity of our system may affect its performance among adolescents. Research in sociolinguistics has shown that language use varies with age [6]. On

⁵While we also observed a large outcome discrepancy across educational attainment, this can largely be explained through age (as a 15 year old almost certainly did not yet have a chance to complete a college education yet).

Demographics	Reduction in Emotion Intensity	Reframe Relatability	Reframe Helpfulness	Reframe Memorability	Skill Learnability	N
Age						
13–14	1.84	3.64	2.94	3.03	2.99	146
15–17	1.64	3.50	2.68	3.07	2.83	149
18–24	1.98	3.78	3.13	3.48	3.15	247
25–34	2.00	3.89	3.32	3.70	3.40	179
35–44	2.15	3.97	3.40	3.69	3.51	109
45–54	1.64	3.96	3.32	3.64	3.56	71
55–64	1.88	3.96	3.46	3.96	3.21	32
65+	1.20	4.00	3.38	3.88	4.13	8
Gender						
Female	1.92	3.78	3.21	3.44	3.26	646
Male	2.19	3.74	2.95	3.38	3.04	258
Non-Binary	1.94	3.76	3.30	3.46	3.22	54
Race/Ethnicity						
AIAN	2.17	2.50	2.67	3.17	2.67	6
Asian	1.91	3.79	3.08	3.43	3.12	216
Black / African Am.	2.43	3.85	3.30	3.62	3.49	47
Hispanic or Latino	1.97	3.91	3.43	3.59	3.47	76
MENA	1.90	3.78	2.90	2.98	2.80	50
NHPI	2.00	4.40	4.00	3.60	3.20	5
White	2.05	3.73	3.12	3.47	3.22	438
More than One	2.83	3.84	3.29	3.24	3.16	38
Other	0.78	3.75	3.06	3.29	3.02	48
Education						
Middle School	1.80	3.58	2.89	2.96	2.79	120
High School	1.80	3.65	2.97	3.31	3.09	313
Undergraduate	2.04	3.79	3.19	3.54	3.28	239
Graduate	2.30	3.98	3.47	3.69	3.52	211
Doctorate	1.46	4.21	2.96	3.93	3.07	28

Table 4: Effectiveness of our system across different demographic population. Numbers highlighted in green indicate outcomes that are significantly better than the population mean ($p < 0.05$). Numbers highlighted in red indicate outcomes that are significantly worse than the population mean ($p < 0.05$). AIAN: American Indian or Alaska Native; MENA: Middle Eastern or North African; NHPI: Native Hawaiian and Pacific Islander. We found that adolescents, males, and those with middle school education reported worse outcomes. Moreover, adults (age ≥ 25) and those with graduate and doctorate education reported better outcomes.

analyzing the reading complexity of the reframed thoughts authored by participants of different age groups, we found that those between the ages of 13 and 17 tend to write thoughts and reframes with the lowest levels of reading complexity (based on the Coleman–Liau Index [27]; Appendix Figure 10). Therefore, we tried to reduce the reading complexity of the reframing suggestions to adolescents. For this, given a reframing suggestion, we asked the GPT-3 language model [19] to make it easier to understand and more casual (using the prompt, “Revise the following text to make it easy to understand for a 5th grader. Also, make it more casual: {reframe}”), similar to other efforts targeting

people of different subpopulations (e.g., for scientific communication [4]). Also, see Appendix Table 8 for examples that illustrate this rewriting process.

Figure 6 reports the results of a randomized trial that only provides these easier to understand and more casual reframing suggestions to half of the participants at random. We found that adolescents in the age group 13 to 14 reported 8.60% higher reframe relatability (4.04 vs. 3.72; $p = 0.0161$) and 14.44% higher reframe helpfulness (3.17 vs. 2.77; $p = 0.0049$) when they were suggested reframes with lower reading complexity ($N=148$). Moreover, adolescents in the age group 15 to 17 reported 15.58% higher reframe

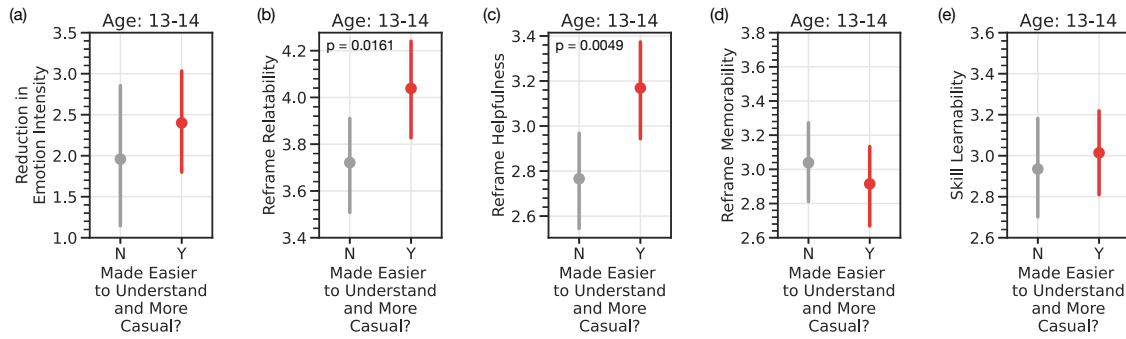


Figure 6: Randomized controlled trial to estimate the effects of making reframes easier to understand and more casual on adolescents in age group 13 to 14 (N=148). Adolescents reported 8.60% higher reframe relatability (4.04 vs. 3.72) and 14.44% higher reframe helpfulness (3.17 vs. 2.77) if they were suggested easier to understand and more casual reframes compared to instances where such reframes were not suggested. Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$.

helpfulness (3.19 vs. 2.76; $p = 0.0042$) when they were suggested such reframes (N=174). We did not find significant differences for adults (≥ 18) through this intervention (N=760; Appendix Figure 11). This suggests that a simpler and more casual language might be beneficial to many. However, based on qualitative feedback, certain adult participants expressed a preference for a less casual language. Future work could explore how to accommodate such individual preferences.

8 DISCUSSION

8.1 Supporting the Learning and Practice of Self-Guided Interventions

Our work demonstrates how language modeling interventions can support mental health. Approximately 20% of people worldwide are experiencing mental health problems, but less than half receive any treatment [62, 63]. Due to widespread clinician shortages, lengthy waiting lists, and lack of insurance coverage, many vulnerable individuals have limited access to therapy and counseling. In addition, mental health issues are heavily stigmatized, which frequently prevents individuals from seeking appropriate care [83].

Effective self-guided mental health interventions could rapidly increase access to care [64, 71, 72, 82]. However, despite their inherent promise, the wide-scale implementation of these interventions remains a challenge owing to the cognitive and emotional challenges that they pose [34, 82]. Most interventions that digitally facilitate self-guided interventions simply transform traditional manual therapeutic worksheets into digital online formats [82]. These provide limited instructions and support, which affects user engagement and usage [7, 32, 34, 94]. Other studies have used wizard-of-oz methods to assist users [46, 49, 54, 60, 85]. However, the controlled research setting of these studies limit their ecological validity, thereby limiting our understanding of user preferences when systems are deployed in real world [15, 17, 59, 67].

Here, we contribute the design of a novel system for human-language model interaction-based self-guided cognitive restructuring of negative thoughts. We conduct a large-scale, randomized,

empirical studies in an ecologically informed setting to understand how people with lived experience of mental health interact with it.

Our findings open up opportunities for improved learning and practicing of key mental health strategies and coping skills. Moreover, these interventions could complement traditional treatment options, e.g., by being accessible to users when they have difficulties finding a therapist, or in between sessions.

8.2 Implications on the Design of Self-Guided Mental Health Intervention

Several of our design hypotheses (Section 4.1) were observed to improve intervention outcomes. These include personalizing the intervention to the participant, facilitating iterative interactivity with the language model, and pursuing equity, all of which may generalize to support other self-guided mental health interventions.

Personalization. Effectively supporting humans through self-guided interventions necessitates personalization [45]. Our design incorporated personalization of reframes by not only seeking more information from the participants in the form of their situations, but also integrating it into the suggestions generated by the language model. We found that this form of personalization led to more helpful reframes than an intervention without it (Section 6.1). This was potentially beneficial in generating suggestions that were more realistic and made fewer assumptions about the participant. Moreover, it emphasized the benefits of increased self-reflection by participants, particularly when thinking about the situation associated with negative thoughts.

On the other hand, when we solicited *emotions* from the participants and failed to incorporate them into the generated suggestions, we observed a significant *decrease* in helpfulness (Section 6.1). This can likely be attributed to unrealistic expectations set up by our intervention, where participants might presume that their emotional states will be addressed in the reframing suggestions, even when that is not the case.

This shows the significance of personalizing language modeling suggestions when developing such self-guided interventions. Note

that we did not find significant improvements based on whether participants explicitly sought specific suggestions to make a reframe personalized (Section 6.3).

Interactivity. In our intervention, participants who interacted more actively with the language model (by seeking additional suggestions) achieved better outcomes (Section 6.3). This highlights the significance of designing interventions that enable a better interaction between the participant and the language model.

Moreover, qualitative feedback from participants revealed that many of them anticipated a continuous, back-and-forth interaction with the language model. Some participants desired a more in-depth exploration of thoughts. One participant wrote, “*I like it but it needs to go deeper with the thoughts.*” Our design facilitates the iterative refinement of reframes. However, a more in-depth exploration could potentially involve addressing new thoughts that arise during the writing of a reframe for the original thought, or even tackling multiple related thoughts at once. Future iterations of the system could work towards a design that is capable of processing multiple thoughts in parallel.

Some participants even expected a chat-like interaction, probably influenced by their experiences with popular systems like OpenAI’s ChatGPT (chat.openai.com). While such mechanisms may offer greater interactivity, their open-ended and uncontrollable nature creates challenges in making them conform to well-established therapeutic processes like cognitive restructuring [29, 52, 90, 93]. In evidence-based cognitive behavioral therapy, cognitive restructuring is typically exercised in a very structured manner (e.g., by asking a specific set of questions in a sequence). Our work shows the promise of extending such exercises with language models with similar digital interfaces.

While it is possible to replicate these steps within a pure chat-bot paradigm, the nature of chat-based interactions can become more complex and may differ from the current methods employed in these interventions. For example, reviewing multiple thinking traps or integration of expandable psychoeducation content is more complex in a purely chat-based interface. Due to the similarity of our intervention with well-established therapy exercises and worksheets, our interface is likely to be a lower burden than a new chat interface. Reducing participant burden has often been associated with improved engagement and completion outcomes [7, 32, 34, 94]. In particular, this has been observed many times in our studied population (MHA platform visitors with lived experience in mental health that are not driven by study compensation incentives). Therefore, appropriate care is required when designing interventions that offer the right trade-off between interactivity and principled adherence to theory.

Another key aspect related to interactivity is over-reliance on the language model assistance. Taking away any “productive struggle” and doing the restructuring *for* the user without fostering reflection and independent practice is likely unproductive. This relates to “desirable difficulty” that differentiates true, long-term “learning” of a skill from the short-term “performance” during skill acquisition [13, 14, 36]. A longer-term goal would be to modulate the difficulty of the self-guided intervention relative to the skill level of the user of the intervention such that the user can build their skills optimally, which forms an interesting direction of future research.

To assess if users are being over-reliant on a system like ours or being able to learn the skill through it, one could observe if users progressively apply the skill they are being taught in their daily lives. This would involve asking users whether they caught themselves thinking negatively, recognized the negative thinking patterns, and reframed the thought in-the-moment while they were having the thought. There exist standardized measures such as the “Competencies of Cognitive Therapy Scale” that operationalize this type of assessment [91]. While ours was a short single-use intervention, this kind of assessment requires a more longer term study, which may require different incentives, recruiting, and platforms and was therefore outside the scope of this paper.

Equity. We found that our intervention was less effective for adolescents, males, and individuals with lower levels of education (Section 7.2). This is consistent with prior research, which has revealed a bias in language models used in mental health contexts toward similar demographics [53]. These findings highlight the importance of adapting self-guided mental health interventions utilizing AI models to suit the needs of different demographics and key subpopulations.

Our work proposes a rewriting-based method to achieve this goal (Section 7.3). This could involve identifying the specific challenges associated with intervening in certain populations (e.g., reading complexity for adolescents) and then, designing appropriate solutions to address those challenges (e.g., lowering the reading complexity).

8.3 Ethics and Safety

The use of AI in mental health presents both opportunities and risks. The systematic ethical and safety considerations in our approach were based on a principle-based ethics framework, following Coghlan et al. [26], Floridi & Cowls [33], and Beauchamp & Childress [8].

Here, we discuss the five common principles derived from these frameworks, providing systematic guidance on how to responsibly navigate potential risks in our mental health setting.

(1) Non-maleficence. *Avoid causing physical, social, or mental harm to participants.* We co-designed the tool with mental health experts, patient advocates, and clinicians (several of whom are co-authors of this paper) to identify any potential risks early. We further studied the safety implications prior to deployment. For this, we developed a realistic sandbox environment to ensure thorough testing instead of immediate deployment. Our safety testing involved examining system inputs that could potentially produce harmful outputs. Our safety filtering mechanisms were updated accordingly to address the potentially harmful LM-generated content (e.g., by adding specific regular expressions and using a content moderation API). Analysis of LM-generated content that was flagged suggests that these efforts were largely successful (Section 6.4). Additionally, throughout the study, participants were given access to a crisis hotline and were able to quit at any point.

(2) Beneficence. *Ensure that interventions do good or provide real benefit to participants.* This involved co-designing with mental health experts, patient advocates, and clinicians to identify opportunities to benefit. This led us to focus on cognitive restructuring,

which is a well-established, evidence-based intervention that has been shown to positively impact people’s mental well-being. Our large-scale field study in an ecologically valid setting demonstrates that our system positively impacts emotional intensity for 67% of participants and helps 65% overcome negative thoughts (Section 5).

(3) Respect for Autonomy. *Respect participants’ values and choices.* We prioritized human agency and initiative by enabling participants to maintain control over what they want to write. Our language model initially offered suggestions on how to reframe thoughts. However, further interaction with the model was only possible when explicitly requested by the participant. Collection of demographic and outcome data was intentionally left optional.

(4) Justice. *Treat participants without unfair bias, discrimination or inequity.* Our study incorporated a broad set of participants with varied demographics, including underrepresented subpopulations. We explicitly evaluated the equity of positive outcomes and identified key subpopulations where the intervention was found to be less effective. We took measures to improve the effectiveness for adolescents for whom the outcome discrepancy was among the largest observed, by lowering the reading complexity (Section 7.3).

(5) Explicability. *Provide users with sufficient transparency about the nature and effects of the technology, and be accountable for its design and deployment.* We sought informed consent from participants and were transparent about the use of AI, risks, and data use (Appendix Figure 12 and Figure 13). We also continuously monitored quantitative and qualitative feedback for potential concerns, which also informed our design hypotheses. Finally, we make our code publicly available at github.com/behavioral-data/Self-Guided-Cognitive-Restructuring.

8.4 Limitations

Our evaluation was confined to a population from a single platform. Yet, our sample size was relatively large and diverse. However, only a few participants over the age of 65 and identifying as American Indian or Alaska Native and Native Hawaiian or Pacific Islander (Table 4). While outcomes varied especially across age rather than race/ethnicity, more work is needed to identify and evaluate opportunities for culturally responsive interventions. The focus of this study was limited to short-term outcomes, necessitating further research to evaluate the long-term effects on participants. Still, our study contributes multiple large-scale randomized trials to inform the design and efficacy of digital mental health interventions. Our outcomes rely significantly on the quality of language model-generated thinking traps and reframes and therefore could change as language models improve. However, the algorithms and models used for generation in this paper represent the current state-of-the-art [81].

9 CONCLUSION

In this paper, we designed a human-language model interaction system that leverages language models to support people through various steps of cognitive restructuring. Through a series of field studies and randomized trials on a large mental health website, we evaluated this system with 15,531 participants. Our findings demonstrated the effectiveness of this system in helping people reduce

the intensity of their negative emotions and effectively reframing negative thoughts. Also, we proposed and validated various design hypotheses including contextualizing people’s thoughts through their situations and facilitating iterative interaction with the language model. Moreover, we assessed the equity of our system across people with different issues and people of different demographics, and improved equity by demonstrably improving the experience of adolescents through lowering the reading complexity of the language model suggestions.

ACKNOWLEDGMENTS

We are grateful to MHA visitors for participating in our field study. We thank the UW Behavioral Data Science Group members for their suggestions and feedback. We also thank Justin Evans for their assistance in model deployment, and Adam Miner, David Wadden, Khendra Lucas, and Sebastin Santy for their input on the tool interface. T.A., A.S. and I.W.L. were supported in part by NIH grant R01MH125179, NSF CAREER IIS-2142794, NSF grant IIS-1901386, Bill & Melinda Gates Foundation (INV-004841), the Office of Naval Research (#N00014-21-1-2154), a Microsoft AI for Accessibility grant, a Garvey Institute Innovation grant, and UW Azure Cloud Computing Credits.

REFERENCES

- [1] Carol E Adair, Gerald M McDougall, Craig R Mitton, Anthony S Joyce, T Cameron Wild, Alan Gordon, Norman Costigan, Laura Kowalsky, Gloria Pasmeny, and Anora Beckie. 2005. Continuity of care and health outcomes among persons with severe mental illness. *Psychiatric Services* 56, 9 (2005), 1061–1069.
- [2] Tim Althoff, Kevin Clark, and Jure Leskovec. 2016. Large-scale analysis of counseling conversations: An application of natural language processing to mental health. *Transactions of the Association for Computational Linguistics* 4 (2016), 463–476.
- [3] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *CHI*. 1–13.
- [4] Tal August, Katharina Reinecke, and Noah A Smith. 2022. Generating scientific definitions with controllable complexity. In *ACL*. 8298–8317.
- [5] Shelli Avenevoli, Joel Swendsen, Jian-Ping He, Marcy Burstein, and Kathleen Ries Merikangas. 2015. Major depression in the national comorbidity survey-adolescent supplement: Prevalence, correlates, and treatment. *Journal of the American Academy of Child & Adolescent Psychiatry* 54, 1 (2015), 37–44.
- [6] Federica Barbieri. 2008. Patterns of age-based linguistic variation in American English. *Journal of sociolinguistics* 12, 1 (2008), 58–88.
- [7] Amit Baumel, Frederick Muench, Stav Edan, and John M Kane. 2019. Objective user engagement with mental health apps: systematic search and panel-based usage analysis. *JMIR* 21, 9 (2019), e14567.
- [8] Tom L Beauchamp and James F Childress. 2001. *Principles of biomedical ethics*. Oxford University Press, USA.
- [9] Aaron T Beck. 1976. *Cognitive therapy and the emotional disorders*. International Universities Press.
- [10] Aaron T Beck, Denise D Davis, and Arthur Freeman. 2015. *Cognitive therapy of personality disorders*. Guilford Publications.
- [11] Aaron T Beck, Maria Kovacs, and Arlene Weissman. 1975. Hopelessness and suicidal behavior: An overview. *JAMA* 234, 11 (1975), 1146–1149.
- [12] Christopher G Beevers, Tony T Wells, and Ivan W Miller. 2007. Predicting response to depression treatment: The role of negative cognition. *Journal of Consulting and Clinical Psychology* 75, 3 (2007), 422.
- [13] Robert A Bjork. 2017. Creating desirable difficulties to enhance learning. *Carmerthen: Crown House Publishing* (2017).
- [14] Robert A Bjork and Elizabeth L Bjork. 2020. Desirable difficulties in theory and practice. *Journal of Applied research in Memory and Cognition* 9, 4 (2020), 475.
- [15] Ann Blandford, Jo Gibbs, Nikki Newhouse, Olga Perski, Aneasha Singh, and Elizabeth Murray. 2018. Seven lessons for interdisciplinary research on interactive digital health interventions. *Digital health* 4 (2018), 2055207618770325.
- [16] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In *ACL*. 5454–5476.

- [17] Judith Borghouts, Elizabeth Eikey, Gloria Mark, Cynthia De Leon, Stephen M Schueller, Margaret Schneider, Nicole Stadnick, Kai Zheng, Dana Mukamel, and Dara H Sorkin. 2021. Barriers to and facilitators of user engagement with digital mental health interventions: systematic review. *JMIR* 23, 3 (2021), e24387.
- [18] Eliane M Boucher, Emily C McNaughton, Nicole Harake, Julia L Stafford, and Acacia C Parks. 2021. The impact of a digital intervention (Happify) on loneliness during COVID-19: qualitative focus group. *JMIR Mental Health* 8, 2 (2021), e26617.
- [19] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *NeurIPS* 33 (2020), 1877–1901.
- [20] Hannah A Burkhardt, George S Alexopoulos, Michael D Pullmann, Thomas D Hull, Patricia A Areán, and Trevor Cohen. 2021. Behavioral activation and depression symptomatology: longitudinal assessment of linguistic indicators in text-based therapy sessions. *JMIR* 23, 7 (2021), e28244.
- [21] David D Burns. 1980. *Feeling Good: The New Mood Therapy*. Avon New York.
- [22] Andrew C Butler, Jason E Chapman, Evan M Forman, and Aaron T Beck. 2006. The empirical status of cognitive-behavioral therapy: A review of meta-analyses. *Clinical psychology review* 26, 1 (2006), 17–31.
- [23] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [24] David A Clark. 2013. Cognitive restructuring. *The Wiley handbook of cognitive behavioral therapy* (2013), 1–22.
- [25] Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A Smith. 2018. Creative writing with a machine in the loop: Case studies on slogans and stories. In *IUI*. 329–340.
- [26] Simon Coghlan, Kobi Leins, Susie Sheldrick, Marc Cheong, Piers Gooding, and Simon D'Alfonso. 2023. To chat or bot to chat: Ethical issues with using chatbots in mental health. *DIGITAL HEALTH* 9 (2023), 20552076231183542.
- [27] Meri Coleman and Ta Lin Liao. 1975. A computer readability formula designed for machine scoring. *Journal of Applied Psychology* 60, 2 (1975), 283.
- [28] Alison Darcy, Jade Daniels, David Salinger, Paul Wicks, and Athena Robinson. 2021. Evidence of human-level bonds established with a digital conversational agent: cross-sectional, retrospective observational study. *JMIR Formative Research* 5, 5 (2021), e27868.
- [29] Munmun De Choudhury, Sachin R Pendse, and Neha Kumar. 2023. Benefits and harms of large language models in digital mental health. *arXiv preprint arXiv:2311.14693* (2023).
- [30] Sona Dimidjian, Manuel Barrera Jr, Christopher Martell, Ricardo F Muñoz, and Peter M Lewinsohn. 2011. The origins and current status of behavioral activation treatments for depression. *Annual review of clinical psychology* 7 (2011), 1–38.
- [31] Xiruo Ding, Kevin Lybarger, Justin Tauscher, and Trevor Cohen. 2022. Improving Classification of Infrequent Cognitive Distortions: Domain-Specific Model vs. Data Augmentation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*. 68–75.
- [32] Theresa Fleming, Lynda Bavin, Mathijs Lucassen, Karolina Stasiak, Sarah Hopkins, and Sally Merry. 2018. Beyond the trial: systematic review of real-world uptake and engagement with digital self-help interventions for depression, low mood, or anxiety. *Journal of medical Internet research* 20, 6 (2018), e199.
- [33] Luciano Floridi and Josh Cows. 2022. A Unified Framework of Five Principles for AI in Society. *Harvard Data Science Review* (2022), 535–545.
- [34] Sandra Garrido, Chris Millington, Daniel Cheers, Katherine Boydell, Emery Schubert, Tanya Meade, and Quang Vinh Nguyen. 2019. What works and what doesn't work? A systematic review of digital mental health interventions for depression and anxiety in young people. *Frontiers in psychiatry* 10 (2019), 759.
- [35] Manas Gaur, Amanuel Alambo, Joy Prakash Sain, Ugur Kursuncu, Krishnaprasad Thirunarayan, Ramakanth Kavuluru, Amit Sheth, Randy Welton, and Jyotishman Pathak. 2019. Knowledge-aware assessment of severity of suicide risk for early intervention. In *WWW*. 514–525.
- [36] Mark A Guadagnoli and Timothy D Lee. 2004. Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning. *Journal of motor behavior* 36, 2 (2004), 212–224.
- [37] Louise C Hawkey and John T Cacioppo. 2010. Loneliness matters: A theoretical and empirical review of consequences and mechanisms. *Annals of behavioral medicine* 40, 2 (2010), 218–227.
- [38] Liesl M Heinrich and Eleonora Gullone. 2006. The clinical significance of loneliness: A literature review. *Clinical psychology review* 26, 6 (2006), 695–718.
- [39] Stefan G Hofmann, Anu Asnaani, Imke JJ Vonk, Alice T Sawyer, and Angela Fang. 2012. The efficacy of cognitive behavioral therapy: A review of meta-analyses. *Cognitive therapy and research* 36 (2012), 427–440.
- [40] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. The curious case of neural text degeneration. In *ICLR*.
- [41] Esther Howe, Jina Suh, Mehrab Bin Morshed, Daniel McDuff, Kael Rowan, Javier Hernandez, Marah Ihab Abdin, Gonzalo Ramos, Tracy Tran, and Mary P Czerwinski. 2022. Design of digital workplace stress-reduction intervention systems: Effects of intervention type and timing. In *CHI*. 1–16.
- [42] Thomas D Hull and Kush Mahan. 2017. A study of asynchronous mobile-enabled SMS text psychotherapy. *Telemedicine and e-Health* 23, 3 (2017), 240–247.
- [43] Natalie E Hundt, Joseph Mignogna, Cathy Underhill, and Jeffrey A Cully. 2013. The relationship between use of CBT skills and depression treatment outcome: A theoretical and methodological review of the literature. *Behavior therapy* 44, 1 (2013), 12–26.
- [44] Eirini Karyotaki, Annet Kleiboer, Filip Smit, David Trevor Turner, A Mira Pastor, Gerhard Andersson, Thomas Berger, Cristina Botella, JM Breton, Per Carlbring, et al. 2015. Predictors of treatment dropout in self-guided web-based interventions for depression: an individual patient data meta-analysis. *Psychological medicine* 45, 13 (2015), 2717–2726.
- [45] Rachel Kornfield, Jonah Meyerhoff, Hannah Studd, Ananya Bhattacharjee, Joseph Jay Williams, Madhu Reddy, and David C Mohr. 2022. Meeting Users Where They Are: User-centered Design of an Automated Text Messaging Tool to Support the Mental Health of Young Adults. In *CHI*. 1–16.
- [46] Rachel Kornfield, Caitlin A Stamatis, Ananya Bhattacharjee, Bei Pang, Theresa Nguyen, Joseph J Williams, Harsh Kumar, Sarah Popowski, Miranda Beltzer, Christopher J Karr, et al. 2023. A text messaging intervention to support the mental health of young adults: User engagement and feedback from a field trial of an intervention prototype. *Internet Interventions* 34 (2023), 100667.
- [47] Kaylee P Kruzan, Ellen E Fitzsimmons-Craft, Mallory Dobias, Jessica L Schleider, and Abhishek Pratap. 2022. Developing, deploying, and evaluating digital mental health interventions in spaces of online help-and information-seeking. *Procedia computer science* 206 (2022), 6–22.
- [48] Kaylee Payne Kruzan, Jonah Meyerhoff, Theresa Nguyen, Madhu Reddy, David C Mohr, and Rachel Kornfield. 2022. "I Wanted to See How Bad it Was": Online Self-screening as a Critical Transition Point Among Young Adults with Common Mental Health Conditions. In *CHI*. 1–16.
- [49] Harsh Kumar, Yiyi Wang, Jiakai Shi, Ilya Musabirov, Norman AS Farb, and Joseph Jay Williams. 2023. Exploring the Use of Large Language Models for Improving the Awareness of Mindfulness. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–7.
- [50] Fei-Tzin Lee, Derrick Hull, Jacob Levine, Bonnie Ray, and Kathleen McKeown. 2019. Identifying therapist conversational actions across diverse psychotherapeutic approaches. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*. 12–23.
- [51] Mina Lee, Megha Srivastava, Amelia Hardy, John Thickett, Esin Durmus, Ashwin Paranjape, Ines Gerard-Ursin, Xiang Lisa Li, Faisal Ladhak, Frieda Rong, et al. 2022. Evaluating human-language model interaction. *arXiv preprint arXiv:2212.09746* (2022).
- [52] Ron C Li, Steven M Asch, and Nigam H Shah. 2020. Developing a delivery science for artificial intelligence in healthcare. *NPJ digital medicine* 3, 1 (2020), 107.
- [53] Inna Lin, Lucille Njoo, Anjalie Field, Ashish Sharma, Katharina Reinecke, Tim Althoff, and Yulia Tsvetkov. 2022. Gendered Mental Health Stigma in Masked Language Models. In *EMNLP*. 2152–2170.
- [54] Kien Hoa Ly, Ann-Marie Ly, and Gerhard Andersson. 2017. A fully automated conversational agent for promoting mental well-being: A pilot RCT using mixed methods. *Internet interventions* 10 (2017), 39–46.
- [55] Nicole Martinez-Martin, Karola Kreitmair, et al. 2018. Ethical issues for direct-to-consumer digital psychotherapy apps: addressing accountability, data protection, and consent. *JMIR mental health* 5, 2 (2018), e9423.
- [56] Tara Matthews, Kathleen O'Leary, Anna Turner, Manya Sleeper, Jill Palzkill Woelfer, Martin Shelton, Cori Manthorne, Elizabeth F Churchill, and Sunny Consolvo. 2017. Stories from survivors: Privacy & security practices when coping with intimate partner abuse. In *CHI*. 2189–2201.
- [57] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)* 54, 6 (2021), 1–35.
- [58] Christine Moberg, Andrea Niles, and Dale Beermann. 2019. Guided self-help works: randomized waitlist controlled trial of Pacifica, a mobile app integrating cognitive behavioral therapy and mindfulness for stress, anxiety, and depression. *JMIR* 21, 6 (2019), e12556.
- [59] David C Mohr, Aaron R Lyon, Emily G Lattie, Madhu Reddy, and Stephen M Schueller. 2017. Accelerating digital mental health research from early design and creation to successful implementation and sustainment. *JMIR* 19, 5 (2017), e7725.
- [60] Robert R Morris, Stephen M Schueller, and Rosalind W Picard. 2015. Efficacy of a web-based, crowdsourced peer-to-peer cognitive reappraisal platform for depression: randomized controlled trial. *JMIR* 17, 3 (2015), e72.
- [61] Ginger Nicol, Ruoyun Wang, Sharon Graham, Sherry Dodd, and Jane Garbutt. 2022. Chatbot-delivered cognitive behavioral therapy in adolescents with depression and anxiety during the COVID-19 pandemic: feasibility and acceptability study. *JMIR Formative Research* 6, 11 (2022), e40242.
- [62] Mark Olsson. 2016. Building the mental health workforce capacity needed to treat adults with serious mental illnesses. *Health Affairs* 35, 6 (2016), 983–990.
- [63] World Health Organization et al. 2022. *Mental health and COVID-19: early evidence of the pandemic's impact: scientific brief, 2 March 2022*. Technical Report. World Health Organization.

- [64] Shireen Patel, Athfah Akhtar, Sam Malins, Nicola Wright, Emma Rowley, Emma Young, Stephanie Sampson, and Richard Morriss. 2020. The acceptability and usability of digital health interventions for adults with depression, anxiety, and somatoform disorders: qualitative systematic review and meta-synthesis. *JMIR* 22, 7 (2020), e16228.
- [65] Verónica Pérez-Rosas, Kenneth Resnicow, Rada Mihalcea, et al. 2022. PAIR: Prompt-Aware margin Ranking for Counselor Reflection Scoring in Motivational Interviewing. In *EMNLP*. 148–158.
- [66] Verónica Pérez-Rosas, Xinyi Wu, Kenneth Resnicow, and Rada Mihalcea. 2019. What Makes a Good Counselor? Learning to Distinguish between High-quality and Low-quality Counseling Conversations. In *ACL*. 926–935.
- [67] Erika S Poole. 2013. HCI and mobile health interventions: how human-computer interaction can contribute to successful mobile health interventions. *Translational behavioral medicine* 3, 4 (2013), 402–405.
- [68] Yada Pruksachatkun, Sachin R Pendse, and Amit Sharma. 2019. Moments of Change: Analyzing Peer-Based Cognitive Support in Online Mental Health Forums. In *CHI*. 1–13.
- [69] Stefan Rennick-Egglestone, Sarah Knowles, Gill Toms, Penny Bee, Karina Lovell, and Peter Bower. 2016. Health Technologies' In the Wild' Experiences of Engagement with Computerised CBT. In *CHI*. 2124–2135.
- [70] Luzette A Samargia, Elizabeth M Saewyc, and Barbara A Elliott. 2006. Foregone mental health care and self-reported access barriers among adolescents. *The Journal of School Nursing* 22, 1 (2006), 17–24.
- [71] Jessica L Schleider, Mallory L Dobias, Jenna Y Sung, and Michael C Mullarkey. 2020. Future directions in single-session youth mental health interventions. *Journal of Clinical Child & Adolescent Psychology* 49, 2 (2020), 264–278.
- [72] Jessica L Schleider, Michael C Mullarkey, Kathryn R Fox, Mallory L Dobias, Akash Shroff, Erica A Hart, and Chantelle A Roulston. 2022. A randomized trial of online single-session interventions for adolescent depression during COVID-19. *Nature Human Behaviour* 6, 2 (2022), 258–268.
- [73] Jessica Schroeder, Jina Suh, Chelsey Wilks, Mary Czerwinski, Sean A Munson, James Fogarty, and Tim Althoff. 2020. Data-driven implications for translating evidence-based psychotherapies into technology-delivered interventions. In *Proceedings of the 14th EAI International Conference on Pervasive Computing Technologies for Healthcare*. 274–287.
- [74] Jessica Schroeder, Chelsey Wilkes, Kael Rowan, Arturo Toledo, Ann Paradiso, Mary Czerwinski, Gloria Mark, and Marsha M Linehan. 2018. Pocket Skills: A Conversational Mobile Web App To Support Dialectical Behavioral Therapy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [75] Stephen M Schueller, Martha Neary, Jocelyn Lai, and Daniel A Epstein. 2021. Understanding people's use of and perspectives on mood-tracking apps: interview study. *JMIR* 8, 8 (2021), e29368.
- [76] Raj Sanjay Shah, Faye Holt, Shirley Anugrah Hayati, Aastha Agarwal, Yi-Chia Wang, Robert E Kraut, and Diyi Yang. 2022. Modeling Motivational Interviewing Strategies On An Online Peer-to-Peer Counseling Platform. *CSCW* 6, 2 (2022), 1–24.
- [77] Ashish Sharma, Monojit Choudhury, Tim Althoff, and Amit Sharma. 2020. Engagement Patterns of Peer-to-Peer Interactions on Mental Health Platforms. In *ICWSM*, Vol. 14. 614–625.
- [78] 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*. 194–205.
- [79] 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* 5, 1 (2023), 46–57.
- [80] Ashish Sharma, Adam S Miner, David C Atkins, and Tim Althoff. 2020. A Computational Approach to Understanding Empathy Expressed in Text-Based Mental Health Support. In *EMNLP*. 5263–5276.
- [81] 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 *ACL*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 9977–10000.
- [82] Jane Shkel, Gavin Green, Stacey Le, Benjamin Kaveladze, Veronique Marcotte, Kevin Rushton, Theresa Nguyen, and Stephen M Schueller. 2023. Understanding Users' Experiences of a Novel Web-Based Cognitive Behavioral Therapy Platform for Depression and Anxiety: Qualitative Interviews From Pilot Trial Participants. *JMIR Formative Research* 7 (2023), e46062.
- [83] Amy E Sickel, Jason D Seacat, and Nina A Nabors. 2014. Mental health stigma update: A review of consequences. *Advances in Mental Health* 12, 3 (2014), 202–215.
- [84] Chaitali Sinha, Abby L Cheng, and Madhura Kadaba. 2022. Adherence and Engagement With a Cognitive Behavioral Therapy-Based Conversational Agent (Wysa for Chronic Pain) Among Adults With Chronic Pain: Survival Analysis. *JMIR* 6, 5 (2022), e37302.
- [85] C Estelle Smith, William Lane, Hannah Miller Hillberg, Daniel Kluver, Loren Terveen, and Svetlana Yarosh. 2021. Effective Strategies for Crowd-Powered Cognitive Reappraisal Systems: A Field Deployment of the Flip* Doubt Web Application for Mental Health. *CSCW* 5, CSCW2 (2021), 1–37.
- [86] Diana MY Smith, Sarah M Lipson, Shirley B Wang, and Kathryn R Fox. 2022. Online methods in adolescent self-injury research: Challenges and recommendations. *Journal of Clinical Child & Adolescent Psychology* 51, 2 (2022), 143–154.
- [87] Wally Smith, Greg Wadley, Sarah Webber, Benjamin Tag, Vassilis Kostakos, Peter Koval, and James J Gross. 2022. Digital emotion regulation in everyday life. In *CHI*. 1–15.
- [88] Jiyoung Song, Boris Litvin, Ryan Allred, Shiyu Chen, Thomas D Hull, and Patricia A Areán. 2023. Comparing Message-Based Psychotherapy to Once-Weekly, Video-Based Psychotherapy for Moderate Depression: Randomized Controlled Trial. *JMIR* 25 (2023), e46052.
- [89] Aseem Srivastava, Ishan Pandey, Md Shad Akhtar, and Tanmoy Chakraborty. 2023. Response-act guided reinforced dialogue generation for mental health counseling. In *The Web Conference*. 1118–1129.
- [90] Elizabeth Stade, Shannon Wiltsey Stirman, Lyle H Ungar, David Bryce Yaden, H Andrew Schwartz, João Sedoc, Robb Willer, Robert DeRubeis, et al. 2023. Artificial Intelligence Will Change the Future of Psychotherapy: A Proposal for Responsible, Psychologist-led Development. (2023).
- [91] Daniel R Strunk, Shannon N Hollars, Abby D Adler, Elizabeth A Goldstein, and Justin D Braun. 2014. Assessing patients' cognitive therapy skills: Initial evaluation of the competencies of cognitive therapy scale. *Cognitive therapy and research* 38 (2014), 559–569.
- [92] Michael J Tanana, Christina S Soma, Vivek Srikumar, David C Atkins, and Zac E Imel. 2019. Development and evaluation of ClientBot: Patient-like conversational agent to train basic counseling skills. *JMIR* 21, 7 (2019), e12529.
- [93] Steven Tate, Sajjad Fouladvand, Jonathan H Chen, and Chwen-Yuen Angie Chen. 2023. The ChatGPT therapist will see you now: Navigating generative artificial intelligence's potential in addiction medicine research and patient care. , 2249–2251 pages.
- [94] John Torous, Jessica Lipschitz, Michelle Ng, and Joseph Firth. 2020. Dropout rates in clinical trials of smartphone apps for depressive symptoms: a systematic review and meta-analysis. *Journal of affective disorders* 263 (2020), 413–419.
- [95] Dave Wadden, Tal August, Qisheng Li, and Tim Althoff. 2021. The Effect of Moderation on Online Mental Health Conversations. In *ICWSM*, Vol. 15. 751–763.
- [96] Coralie J Wilson and Frank P Deane. 2012. Brief report: Need for autonomy and other perceived barriers relating to adolescents' intentions to seek professional mental health care. *Journal of adolescence* 35, 1 (2012), 233–237.
- [97] Caleb Ziems, Minzhi Li, Anthony Zhang, and Diyi Yang. 2022. Inducing Positive Perspectives with Text Reframing. In *ACL*. 3682–3700.

A APPENDICES

Table 5: Randomized Controlled Trial Experiments

Randomized Controlled Trial	Section with Results
Enabling contextualization through situations vs. disabling contextualization through situations	Section 6.1
Enabling contextualization through emotions vs. disabling contextualization through emotions	Section 6.1
Adding psychoeducation vs. removing psychoeducation	Section 6.2
Enabling the option to seek more LM suggestions vs. disabling the option to seek more LM suggestions	Section 6.3

Figure 7: Randomized controlled trial to estimate the effects of contextualizing thoughts through emotions (N=4,016). (a) Contextualizing participant thoughts through their emotions led to 3.86% lower relatable reframes (3.87 vs. 3.72). Note that our language model does not necessarily incorporate emotions unless they are expressed in the thought or situation as well. (b) Asking for additional context did not lead to lower completion rate. Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$.




Figure 8: Randomized controlled trial to estimate the effects of integrating psychoeducation (N=1,850). We did not find significant quantitative improvement in outcomes on integrating psychoeducation. Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$.

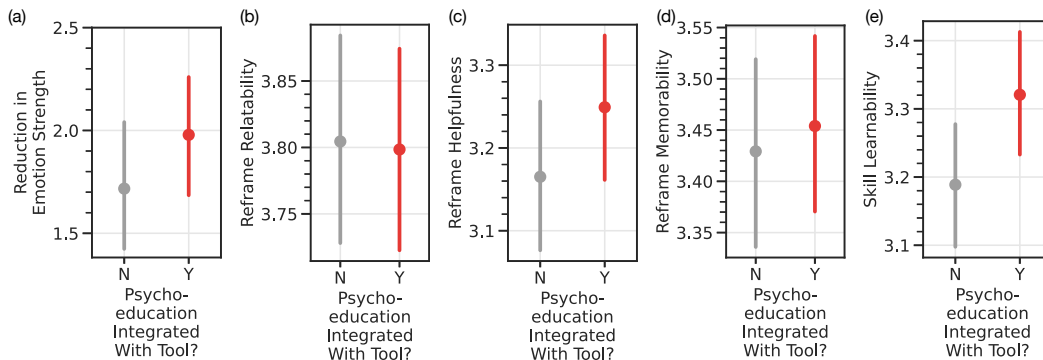


Figure 9: Participants who further interacted with the language model to seek additional reframing suggestions of specific types (actionable, empathic, or personalized) reported 5.57% higher reframe helpfulness (3.41 vs. 3.23) and 4.86% higher skill learnability (3.45 vs. 3.29) and no significant differences in reduction in emotion intensity, reframe relatability, and reframe memorability (at $\alpha = 0.05$; $N=992$). Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$. For a randomized trial assessing the effects of this intervention, see Figure 4.

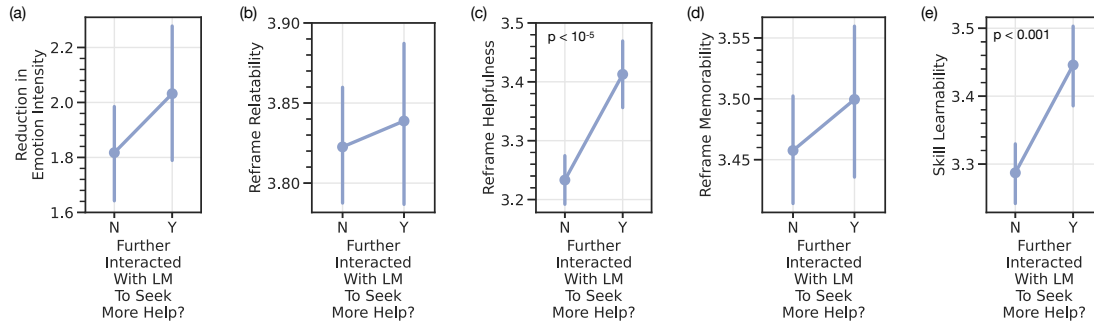


Figure 10: Reading Complexity (Coleman–Liau Index) of the thoughts written by participants based on their age. We find that adolescents (those below the age of 18) write thoughts with the least reading complexity. Error bars represent 95% bootstrapped confidence intervals.

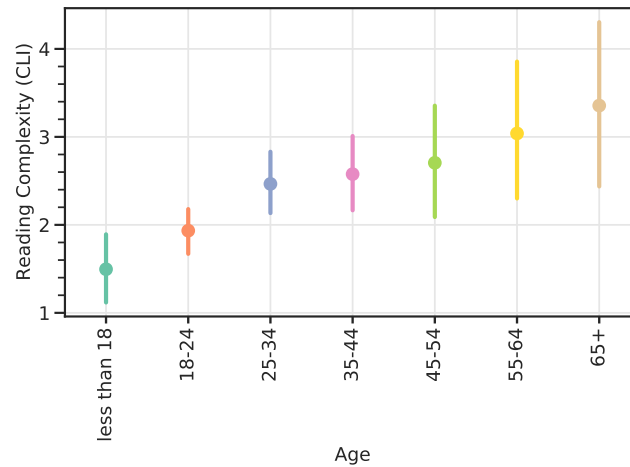


Figure 11: Randomized controlled trial to estimate the effects of making reframes easier to understand and more casual on age groups 15 to 17 and ≥ 18 . Participants in the age group 15 to 17 reported 15.58% higher reframe helpfulness (3.19 vs. 2.76) when they were suggested such reframes. We did not find significant differences for participants in age group ≥ 18 due to this intervention. Error bars represent 95% bootstrapped confidence intervals. Effects without p-values were not significant at $\alpha = 0.05$.

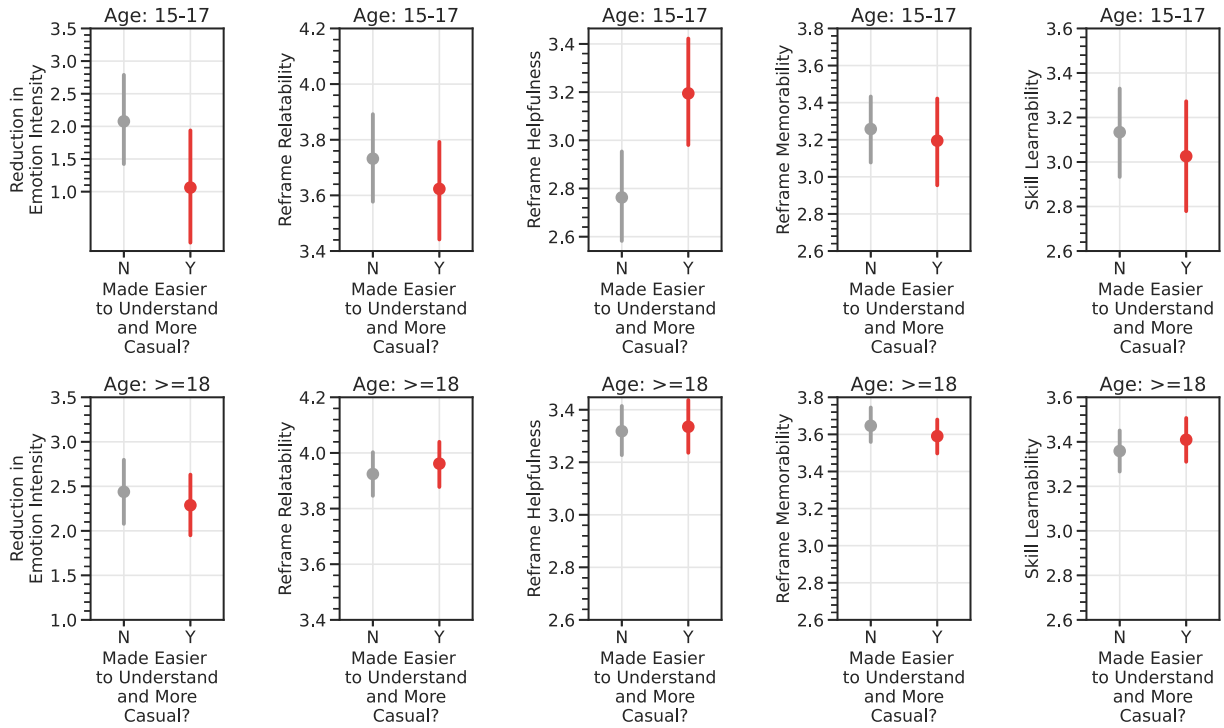


Figure 12: Consent form used in our study. The form continues on the next page (1/2).

Consent Form

This tool uses artificial intelligence to generate reframed thoughts and is part of a research study in collaboration with the University of Washington.

Purpose: The purpose of the study is to understand how digital tools can help people recognize thinking traps and practice reframing negative thoughts.

Procedure: You will be asked to describe a thought and a situation you are struggling with. You will then identify potential "thinking traps" (or cognitive distortions) in the thought and reframe it in a way that is more positive, realistic, or helpful. Finally, you will be asked to take an *optional* demographic survey, which can be skipped as preferred. The tool is expected to take ~5 minutes to complete.

Benefits: By using this tool, you may learn about thinking traps. You will practice identifying them and reframing negative thoughts and situations. However, there is no guarantee that the tool will help you reframe your thoughts.

Data Collection and Sharing: We will not ask you for your name or any identifiable personal information. Usage data will be made unidentifiable to the best of our extent, will be analyzed to improve the tool, and may be shared and used for future research.

Risks: Talking about situations and thoughts you are struggling with may be disturbing to you and may bring up negative emotional reactions. In addition, the tool uses artificial intelligence to generate reframed thoughts. Appropriate steps have been taken to avoid harmful reframes, but there is a possibility that the generated content might be upsetting to you. Also, the *optional* demographic survey asks for information that may be sensitive and could make you feel uncomfortable (e.g., "*What are the main things contributing to your mental health problems right now?*"). This tool is not being actively monitored by a human and should not be used as a "cry for

Figure 13: Consent form used in our study (2/2).

help" outlet. Should you have a strong negative reaction to some of the content, you can text MHA to 741741 or call or text 988.

Participation: Participation in this study is completely voluntary. You will not receive any payment for participation. You can refuse participation or stop participating at any time without penalty or loss of benefits to which you are otherwise entitled.

Contact Us: If you have questions or concerns about this research, or if you think you have been harmed from being in the study, please email us (Ashish Sharma at ashshar@cs.washington.edu, or Professor Tim Althoff at althoff@cs.washington.edu). If you have questions about your rights as a research participant, you can call the University of Washington Human Subjects Division at (206) 543-0098.

- By ticking this box, you are agreeing to use this tool. You are also confirming that you are at least 13 years old. Be sure that questions about the tool have been answered and that you understand what you are being asked to do. You may contact us if you think of a question later. You are free to stop using the tool at any time. To save a copy of this consent form, you can use [this link](#).

Table 6: Definitions, examples, and tips to overcome for each thinking trap. Definitions and examples borrowed from Sharma et al. [81]. We also use these as part of our psychoeducation integrated into the system.

Thinking Trap	Definition	Example	Tips to Overcome
All-or-Nothing Thinking	Thinking in extremes.	"If it isn't perfect, I failed. There's no such thing as 'good enough'."	Things in life are rarely black and white. Focus on what's positive or neutral about the situation.
Overgeneralizing	Jumping to conclusions based on one experience.	"They didn't text me back. Nobody ever texts me back."	Recall times when things went well for you. Imagine what it would be like for things to go well next time.
Labeling	Defining a person based on one action or characteristic.	"I said something embarrassing. I'm such a loser."	Consider all different aspects of a person.
Fortune Telling	Trying to predict the future. Focusing on one possibility and ignoring other, more likely outcomes.	"I'm late for the meeting. I'll make a fool of myself."	Be curious about what's going to happen next. Focus on what you can control and let go of what you can't.
Mind Reading	Assuming you know what someone else is thinking.	"She didn't say hello. She must be mad at me."	Try to imagine other, less negative possibilities. Try to ask the person what they're thinking, rather than just assuming.
Emotional Reasoning	Treating your feelings like facts.	"I woke up feeling anxious. I just know something bad is going to happen today."	Consider all the information you have.
Should Statements	Setting unrealistic expectations for yourself.	"I shouldn't need to ask for help. I should be independent."	Think about where your unrealistic expectations came from. Let your mistakes be an opportunity to learn and grow.
Personalizing	Taking things personally or making them about you.	"He's quiet today. I wonder what I did wrong."	Think about all the other things that could be affecting someone's behavior.
Disqualifying the Positive	When something good happens, you ignore it or think it doesn't count.	"I only won because I got lucky."	Go out of your way to notice the positive side.
Catastrophizing	Focusing on the worst-case scenario.	"My boss asked if I had a few minutes to talk. I'm going to get fired!"	Keep in mind that worst-case scenarios are very unlikely. Try to remind yourself of all the more likely, less severe things that could happen.
Comparing and Despairing	Comparing your worst to someone else's best.	"My niece's birthday party had twice the amount of people"	Remember that what you see on social media and in public is everyone showing off their best.
Blaming	Giving away your own power to other people.	"It's not my fault I yelled. You made me angry!"	Take responsibility for whatever you can—no more, no less.
Negative Feeling or Emotion	I'm having a negative feeling or emotion which isn't a thinking trap.	"I am feeling lonely."	Feeling negative emotions is a normal part of life. Think about what we can control, and what positive things we can be grateful for.

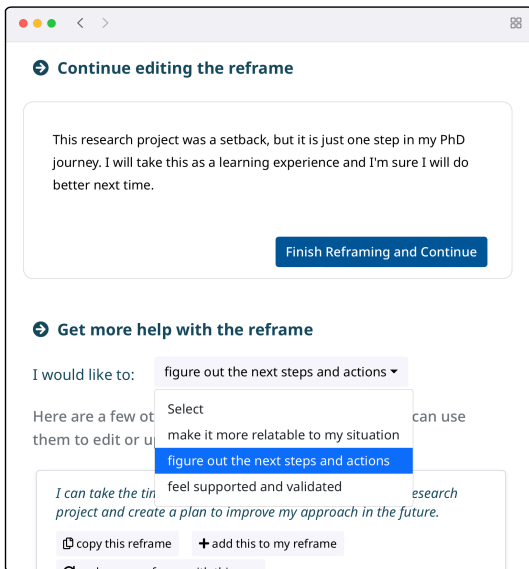
Table 7: Definitions and examples of the list of issues identified in our open coding process.

Issue	Definition	Example Thought
Body image	Feeling ugly, dieting, disordered eating	"I'm fat and ugly"
Dating & marriage	Insecurities around dating and sexuality ("No one will date me" or "I got rejected") or specific situations involving a significant other ("My husband argues with me" or "My wife left me"). Does not include abuse!	"I am scared my girlfriend is going to break up with me because I dont do much for her."
Family	Specific situations involving interactions with family. If the only family member involved is the person's child, that's usually classified as parenting.	"I am the worst daughter in the world."
Fear	Intrusive thoughts about bad things happening; worrying about the future; imagining worst-case scenarios	"Something will go wrong on my flight today"
Friendship	Specific situations involving interactions with specific people. Doesn't have to be hyper-specific, but different from "loneliness" which is more about just the general concept of feeling isolated or unloved.	"My friend doesn't like me anymore."
Habits	Drugs & alcohol, addiction, or just any habits the person is trying to break.	"I made a goal to quit smoking and failed. I keep failing at this"
Health	Real or imagined illness; access to healthcare. Does not include mental health.	"I might be having a serious illness"
Hopelessness	Feeling like things will never get better; feeling like there's no point in trying	"I have lost all hope"
Identity	Discrimination due to race, gender, sexual orientation, etc; or coming to terms with being LGBTQ+.	"I'm not who I want to be. I hate my appearance and my voice."
Loneliness	Not having friends; being isolated from loved ones; general social anxiety	"Why don't i have friends?"
Money	Financial troubles. Includes being jobless if the main worry is finances.	"My financial situation is going out of hand. I'm worried about my future"
Parenting	Feeling like a bad parent, or worried about becoming a parent	"I'm a bad mom"
School	Bad grades, getting into college, fears of graduating, etc.	"I will fail my exam"
Tasks & achievement	Worrying about not being good enough. Examples: hobbies, chores, executive function.	"I'm not good enough"
Trauma	Violence, rape, verbal or physical abuse, etc.	"I can't get through this"
Work	Situations in the workplace. Includes being jobless if the main worry is competency.	"I'm late for the meeting. This shows what a jerk I am"

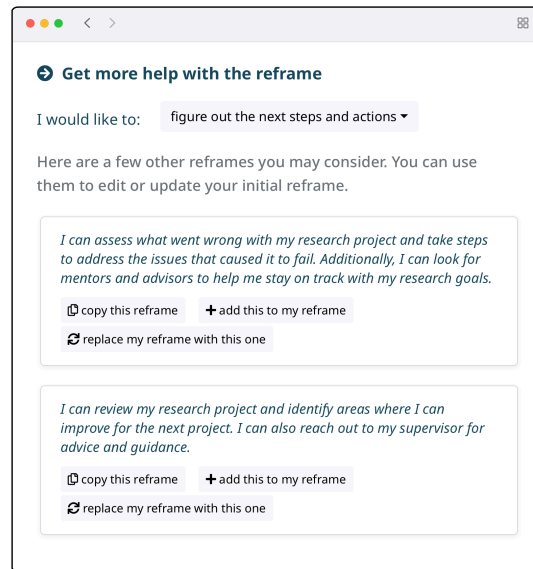
Table 8: Examples illustrating the rewriting of reframes to make them simpler and casual for adolescents. Simpler and casual reframes use more straightforward words and phrases (e.g., “valuable” → “awesome”, “things I have accomplished” → “stuff I’ve already done”, “potential I have” → “how much more I can do”, “improve my physical health” → “make my body healthier”) and use informal language (e.g., “But you know what?”, “Plus...”). This could potentially be more appealing especially to adolescents.

Thought	Original Reframe	Simpler and Casual Reframe
I’m worthless. I haven’t done anything valuable in my life.	I may not have achieved all of my goals yet, but I’m still valuable and capable of doing great things. I can focus on the things I have accomplished and the potential I have to achieve even more.	I haven’t reached all my goals yet, but that’s okay! I’m still awesome and can do awesome things. I can think about the stuff I’ve already done and how much more I can do in the future.
I will be alone my whole life. All of my friends are starting to be in relationships.	It’s normal to feel lonely and to worry about not having a partner. I can focus on developing meaningful relationships with people, and on loving and appreciating myself.	It’s totally normal to feel lonely sometimes and worry about not having a partner. But you know what? Instead of stressing about it, I can focus on building amazing friendships with people who really care about me. Plus, I can also learn to love and appreciate myself just the way I am.
I’m going to fail in school and my body will never change	I may not do well in school right now, but I can still make changes to my lifestyle and work hard to improve my grades. I can also take steps to improve my physical health.	Right now, school isn’t going so great for me. But that doesn’t mean I can’t do things to make it better. I can change the way I live and put in a lot of effort to make my grades improve. I can also do things to make my body healthier.

Figure 14: Detailed interface and process for iterative edits of reframes through further interaction with the language model.



(a) To iteratively refine the reframes, we provide three options to choose from — “make it more relatable to my situation”, “figure out the next steps and actions”, and “feel supported and validated”. Each option enables further interaction with the language model



(b) For the option selected by the participant, a language model generates additional reframing suggestions. Participants can either copy this reframe, add it to their current reframe, replace their original reframe with it, or use it as an inspiration