Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)
Sister Conferences Best Papers Track

Towards Facilitating Empathic Conversations in Online Mental Health Support: A Reinforcement Learning Approach (Extended Abstract)*

Ashish Sharma[1] , **Inna W. Lin**[1] , **Adam S. Miner**[2,3] , **David C. Atkins**[4] and **Tim Althoff**[1]
[1]Paul G. Allen School of Computer Science & Engineering, University of Washington
[2]Department of Psychiatry and Behavioral Sciences, Stanford University
[3]Center for Biomedical Informatics Research, Stanford University
[4]Department of Psychiatry and Behavioral Sciences, University of Washington
{ashshar, ilin, althoff}@cs.washington.edu

Abstract

Online peer-to-peer support platforms enable conversations between millions of people who seek and provide mental health support. If successful, web-based mental health conversations could improve access to treatment and reduce the global disease burden. Psychologists have repeatedly demonstrated that *empathy*, the ability to understand and feel the emotions and experiences of others, is a key component leading to positive outcomes in supportive conversations. However, recent studies have shown that highly empathic conversations are rare in online mental health platforms.

In this paper, we work towards improving empathy in online mental health support conversations. We introduce a new task of *empathic rewriting* which aims to transform low-empathy conversational posts to higher empathy. Learning such transformations is challenging and requires a deep understanding of empathy while maintaining conversation quality through text fluency and specificity to the conversational context. Here we propose PARTNER, a deep reinforcement learning (RL) agent that learns to make sentence-level edits to posts in order to increase the expressed level of empathy while maintaining conversation quality. Our RL agent leverages a policy network, based on a transformer language model adapted from GPT-2, which performs the dual task of generating candidate empathic sentences and adding those sentences at appropriate positions. Through a combination of automatic and human evaluation, we demonstrate that PARTNER successfully generates more empathic, specific, and diverse responses and outperforms NLP methods from related tasks such as style transfer and empathic dialogue generation. We share our code publicly at bdata.uw.edu/empathy.

*This is an extended abstract of the paper "Towards Facilitating Empathic Conversations in Online Mental Health Support: A Reinforcement Learning Approach" [Sharma *et al.*, 2021], that appeared and won the Best Paper Award at *The Web Conference, 2021*, for the IJCAI 2022 Sister Conference Best Paper Track.

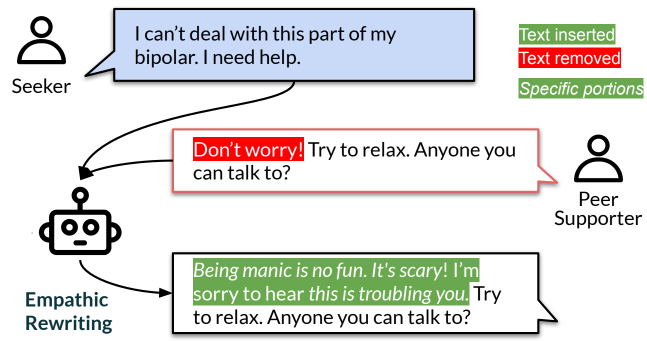

Figure 1: **Empathic Rewriting task.** Given a post from support seeker and a low-empathy response, the task is to rewrite the response to make it more empathic, through text insertions and deletions. This task requires inferring *specific* feelings and experiences from seeker's post and using them for making appropriate changes to the response through empathic mechanisms like emotional reactions, interpretations, and explorations [Sharma *et al.*, 2020].

1 Introduction

Online mental health support platforms such as TalkLife (talklife.com) are used by millions of users for expressing emotions, sharing stigmatized experiences, and receiving peer support. These platforms might help improve access to mental health support as mental health care remains a global challenge with widespread shortages of workforce [Olfson, 2016] and other barriers like stigma [White and Dorman, 2001]. A key component of providing successful support is *empathy*, the ability to understand or feel the emotions and experiences of others [Elliott *et al.*, 2011]. Quantitative evidence shows that empathic interactions have strong associations with symptom improvement in mental health support [Elliott *et al.*, 2018] and are instrumental in building therapeutic alliance and rapport [Bohart *et al.*, 2002; Robert *et al.*, 2011]. Yet, highly empathic conversations are rare on online support platforms [Sharma *et al.*, 2020].

Empowering peer supporters on online support platforms with feedback and training, for example through machine-in-the-loop writing systems [Clark *et al.*, 2018; Tanana *et al.*, 2019], has the potential to help supporters express higher levels of empathy and in turn improve the effectiveness of these platforms [Miner *et al.*, 2019; Sharma *et al.*, 2020]. Traditional methods for improving empathy (e.g., in-person training) do not scale to the millions of users of online support platforms. However, computational methods that can support peer-supporters by suggesting ways to modify existing conversation utterances to make them more empathic may help meet this need and indirectly benefit support seekers.

In this paper, we introduce **Empathic Rewriting**, a new task that aims to transform low-empathy conversations to higher empathy (Figure 1). For example, given a post from a support seeker "*I can't deal with this part of my bipolar. I need help.*" and a low-empathy response "*Don't worry! Try to relax. Anyone you can talk to?*", we want to increase empathy in the response by transforming it to "*Being manic is no fun. It's scary! I'm sorry to hear this is troubling you. Try to relax. Anyone you can talk to?*"; the rewritten response should communicate more empathy through an understanding of feelings and experiences ("*Being manic is no fun. It's scary*") and display of felt emotions ("*I'm sorry to hear this is troubling you*").

Learning such transformations is challenging and requires a deep understanding of empathy while maintaining conversation quality through text fluency and specificity to the conversational context. Here, we propose PARTNER,[1] a deep reinforcement learning (RL) model for empathic rewriting (Section 3). We design an RL agent which learns to add new empathic sentences to posts or replace existing sentences in posts with more empathic ones. The agent operates on a pair of seeker post and the original response post (which rarely is highly empathic [Sharma *et al.*, 2020]) and makes edits to the response at the level of a sentence by simultaneously (a) identifying positions in the original response post where changes are required, and (b) generating empathic sentences for insertion or replacement at the identified positions (Section 3.2). We model this agent using a policy network based on a transformer decoder model adapted from GPT-2 [Radford *et al.*, 2019]. We build upon existing large-scale pre-training of GPT-2 on conversations, as done in DialoGPT [Zhang *et al.*, 2020], and modify it to perform the two simultaneous actions of identifying positions and generating empathic sentences for empathic rewriting (Section 3.3). Through carefully constructed scoring functions, we reward transformations that increase empathy in posts while maintaining text fluency, context specificity, and diversity (Section 3.4).

Our experiments demonstrate that PARTNER can effectively increase empathy in posts in fluent, specific, and diverse ways and outperforms baselines used in related text generation tasks by $> 35\%$ in empathy improvement (Section 4). We view our approach and findings as a key step towards building AI systems for facilitating empathic conversations on online mental health support platforms, but these insights may generalize beyond mental health to other conversational settings on web-based platforms.

2 Problem Definition and Goals

2.1 Empathic Rewriting

We introduce *empathic rewriting*, a new task that aims to transform low-empathy conversational posts to higher empathy. In contrast with empathic dialogue generation [Rashkin *et al.*, 2019], where the objective is to generate empathic posts from scratch, this task requires making changes to existing posts to make them empathic. This is more consistent with realistic use-cases in difficult, high-stakes settings such as online support systems, which are likely to augment, rather than replace humans [Miner *et al.*, 2019]. Formally, let $\mathbf{S_i}$ be a seeker post and $\mathbf{R_i}$ be a corresponding response post. We aim to transform $\mathbf{R_i}$ into its more empathic counterpart $\hat{\mathbf{R}}_\mathbf{i}$.

2.2 Goals

For empathic rewriting to be useful in improving supportive conversations, the rewriting process should achieve specific goals related to empathy, conversation and natural language generation quality, and purposeful and precise feedback:

Theoretically-grounded empathy. Computational research typically defines empathy as reacting with emotions of warmth and compassion [Buechel *et al.*, 2018]. However, psychotherapy research emphasizes aspects of empathy related to communicating cognitive understanding of feelings and experiences of others [Selman, 1980]. For empathic rewriting to be useful and potentially adopted in online mental health support, we need to design methods grounded in psychology and psychotherapy research. Here, we adopt the theoretically-grounded framework of empathy designed by Sharma et al. [2020] as reward signals (Section 3.4)

Context specificity and response diversity. Consider a rewriting approach that transforms every response to a generic but empathic response (e.g., "*That must have been really hard for you*"). While this approach may seem to "solve" empathic rewriting, it suffers from the issues of low specificity and diversity, which are critical for obtaining purposeful transformations. In this work, we learn rewriting actions that simultaneously achieve the goals of context specificity and response diversity using reinforcement learning (Section 3.4).

Text fluency and sentence coherence. In addition, only generating empathic words or phrases may not be sufficient. Without appropriate measures, the rewriting process may lead to an ungrammatical, non-fluent final response with non-coherent sentences. In this paper, we avoid such responses through carefully constructed reward functions (Section 3.4).

Rewriting for feedback and training. An important way in which the task of empathic rewriting can be used is for providing feedback and training to people through machine-in-the-loop writing systems [Clark *et al.*, 2018]. For humans to adopt such feedback, however, the rewriting process should make changes that are precise and minimal. Here, we train a reinforcement learning agent which learns when to stop making changes through a special "stopping" action (Section 3.2).

3 PARTNER: Empathic Rewriting Using Reinforcement Learning

Here, we present PARTNER, a reinforcement learning model for the task of empathic rewriting (Figure 2). Conceptually, our agent leverages context from the seeker post which it uses for making specific empathic changes. Alongside, it operates on the response post, looks for areas where empathy could be

[1]em**PA**thic **R**ewri**T**ing in me**N**tal h**E**alth suppo**R**t

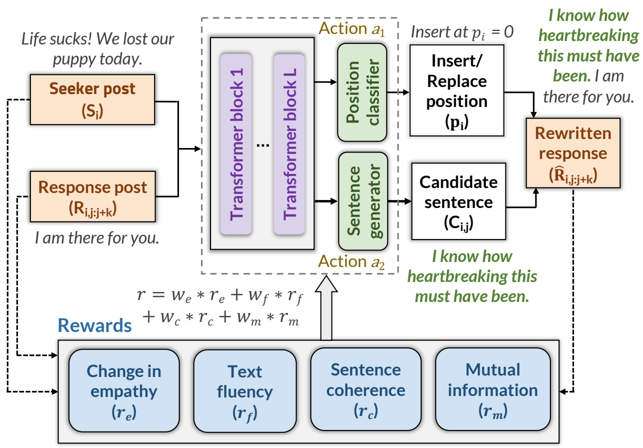

Figure 2: PARTNER uses a deep reinforcement learning approach for Empathic Rewriting. It leverages a transformer language model for performing the two actions of (1) selecting positions for insertion or replacement and (2) generating candidate empathic sentences. It uses four reward functions that promote increase in empathy, text fluency, sentence coherence, context specificity, and diversity.

improved, and works on those improvements in fluent, coherent, specific, and diverse ways.

3.1 State: Seeker Post & Fixed-Length Contiguous Spans of Response Post

Our agent simultaneously operates on seeker post and fixed-length contiguous spans of response post. Formally, let $\mathbf{R_i}$ contain n sentences $\mathbf{R_{i,1}}, ..., \mathbf{R_{i,n}}$. At each step, we focus on a contiguous window of k sentences starting from the jth sentence $\mathbf{R_{i,j:j+k}} = \mathbf{R_{i,j}}, ..., \mathbf{R_{i,j+k-1}}$. Then, our state $s \in \mathcal{S}$ is denoted by the pair $(\mathbf{S_i}, \mathbf{R_{i,j:j+k}})$.

3.2 Actions: Sentence-Level Edits

Our agent takes actions at the level of a sentence, i.e. it either *inserts* new sentences or *replaces* existing sentences with newer ones. A deletion operation is equivalent to replacing a sentence with an empty string. Our agent can make word-level changes by replacing the original sentence with a slightly different sentence containing only word-level edits. In a state $(\mathbf{S_i}, \mathbf{R_{i,j:j+k}})$, our agent simultaneously takes two actions – (a_1) select a position in $\mathbf{R_{i,j:j+k}}$ for insertion or replacement, (a_2) generate a candidate empathic sentence. The action space $\mathcal{A}_1$ of a_1 consists of 2k+2 actions – k+1 positions for insertions, k positions for replacements, and one *special* action to stop the agent from making any further changes.

3.3 Policy

At its core, our policy has a transformer language model consisting of a stack of masked multi-head self-attention layers, based on GPT-2 [Radford *et al.*, 2019]. It takes as input an encoded representation of our state $(\mathbf{S_i}, \mathbf{R_{i,j:j+k}})$ and generates the action $a = (a_1, a_2)$.

(a_1) **Selecting a position for insertion or replacement.** A k sentence window $\mathbf{R_{i,j:j+k}}$ has $k+1$ positions for insertions and k positions for replacement. Then, our task is to select one of these $2k+1$ positions. For selecting this position, we first encode the input string "$\mathbf{S_i}$ <SPLIT> $\mathbf{R_{i,j:j+k}}$" using the transformer block of GPT-2. We then pass this encoded representation through a linear layer to get the prediction $\hat{\mathbf{p}}_\mathbf{i}$ of the position for insertion or replacement.

(a_2) **Generating a candidate sentence.** We first encode our input string through GPT-2. We then compute a probability distribution over vocabulary tokens by transforming the encoded representation into a vocabulary-sized vector through a softmax layer. Finally, we use top-p sampling [Holtzman *et al.*, 2020] over this probability distribution to generate the desired candidate sentence $\mathbf{C_{i,j}}$.

3.4 Rewards

Our reward function, $r = w_e{*}r_e + w_f{*}r_f + w_c{*}r_c + w_m{*}r_m$, aims to increase empathy in posts and maintain text fluency, sentence coherence, context specificity, and diversity:

Change in empathy (r_e)**.** We reward actions that increase empathy of $\mathbf{R_i}$ and penalize actions that decrease empathy of $\mathbf{R_i}$. Let $f_e(\cdot)$ be a function that measures empathy of posts. Then, the change in empathy reward, r_e, is defined as $f_e(\hat{\mathbf{R}}_\mathbf{i}) - f_e(\mathbf{R_i})$, where $f_e(\cdot)$ is estimated using the empathy classification model of Sharma et al. [2020].

Text fluency (r_f)**.** We aim to prevent actions that lead to outputs that are highly empathic but lack text fluency. Therefore, we reward actions that lead to fluent outputs and penalize actions resulting in non-fluent outputs. Here, we operationalize *text fluency* as the inverse of perplexity of the generated $\hat{\mathbf{R}}_\mathbf{i}$s.

Sentence coherence (r_c)**.** While the candidate sentence that is added to the original response $\mathbf{R_i}$ might be highly empathic and fluent, it may not be well-suited for the response, leading to incoherent sentences in the transformed response $\hat{\mathbf{R}}_\mathbf{i}$. Here, we design a reward function, r_c that measures the average sentence coherence probability between the candidate sentence and existing sentences in the response.

Mutual information for specificity and diversity (r_m)**.** In the process of empathic rewriting, the final rewritten response may become generic (e.g., "*I understand how you feel*") thereby affecting the overall conversation quality. In order to ensure specificity to the seeker post and diversity of responses, we exploit the idea of maximizing mutual information between seeker post and the rewritten response post [Li *et al.*, 2016].

4 Experiments and Results

4.1 Warm-Start Using Supervised Learning

We use the pre-trained weights of DialoGPT [Zhang *et al.*, 2020] for initializing our transformer language model. In addition, we use a warm-start strategy using supervised learning on a parallel dataset of (low empathy, high empathy) pairs.

4.2 Dataset: TalkLife

We use a dataset of mental health conversations from the TalkLife (talklife.com) platform. The dataset contains 10.9M conversational threads and 26.9M pairs of (seeker post, response post) interactions from 642K users.

	Model	Change in empathy (↑)	Perplexity (↓)	Specificity (↑)	Diversity (↑) distinct-1	distinct-2	Sentence coherence (↑)	Edit rate (↓)
Dialogue Generation	DialoGPT [Zhang *et al.*, 2020]	0.4698	8.6500	0.8921	0.0382	0.1334	0.6683	1.3520
	MIME [Majumder *et al.*, 2020]	1.2069	9.0171	0.8837	0.0031	0.0198	0.3687	1.8193
Seq-to-Seq Generation	Latent Seq. [He *et al.*, 2019]	0.9745	8.7143	0.8512	0.0001	0.0002	**0.9252**	7.8853
	BART [Lewis *et al.*, 2020]	-0.0611	**7.2040**	0.8878	**0.0722**	**0.3945**	0.4560	**0.7496**
	PARTNER	**1.6410**	7.3641	**0.9052**	0.0659	0.3807	0.3030	0.9654

Table 1: Performance of PARTNER and comparisons with dialogue generation and other sequence-to-sequence generation baselines on the set of automatic metrics. PARTNER outperforms all baselines in empathy improvement and generates fluent, specific, and diverse outputs with lower edits. (↑) indicates higher is better, (↓) indicates lower is better.

(A) BLEU Score

Model	BLEU score
PARTNER	**0.1391**
DialoGPT	0.0722
MIME	0.0808
Latent Seq.	0.0254
BART	0.0956

(B) Human Evaluation

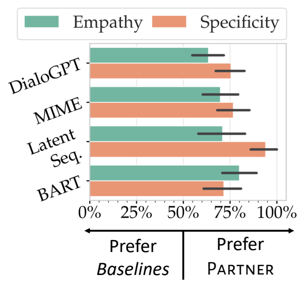

(C) Qualitative Example

Example	Model	Rewritten response post
Seeker post: I hope I didn't lose my only two best friends. They are the only ones I could relate to. **Response post:** What happened between you two?	PARTNER	That's awful. What happened between you two? What caused you to break?
	DialoGPT	I hope you didn't lose them. I don't want to lose them.
	MIME	That is a good thing to do.
	Latent Seq.	I'm sorry to hear that.
	BART	What happened? What happened between you two?

Figure 3: (A) BLEU scores against empathic rewritings from experts. (B) Human evaluation results. (C) An example of empathic rewriting using PARTNER and baselines. Check our full paper [Sharma *et al.*, 2021] for more examples.

4.3 Results

Automatic metrics. We find that empathic rewriting through PARTNER achieves the largest change in empathy (35% more than the next best baseline, MIME [Majumder *et al.*, 2020]) and is more specific than all baselines (Table 1). Also, MIME generates empathic outputs (+1.21 change in empathy) but the generations have low diversity (86% less than PARTNER) indicating similar responses for most seeker posts. BART [Lewis *et al.*, 2020] generates outputs with lowest perplexity, highest diversity, and lowest edit rate, which is consistent with substantial improvements to language models in recent years [Brown *et al.*, 2020]. However, to our surprise, the rewritten responses through BART receive an overall drop of 0.06 in empathy, indicating that the model is unable to perform the task of empathic rewriting well and typically generates non-empathic, fluent, diverse text.

Human evaluation. Since automatic evaluation in language generation is often not robust, we perform a human evaluation on our key metrics (empathy, fluency, and specificity) through A/B testing. We recruit six graduate students in clinical psychology with expertise in empathy and mental health support and ask them to compare outputs from PARTNER against other baseline models, ablations given the same input. We find that rewritten responses from PARTNER are preferred for empathic and specific responses over all baselines (Figure 3B). DialoGPT is judged more fluent but generates responses following similar templates (e.g., "*I'm sorry you.... I hope you....*").

BLEU scores. We create a small parallel dataset of 180 pairs of corresponding low and rewritten high empathy response posts with rewritings from people having substantial expertise in empathy, mental health, and therapy (six graduate students in clinical psychology; none are co-authors). We compare outputs of PARTNER and baselines based on this ground truth using the BLEU metric [Papineni *et al.*, 2002] (Figure 3A). We find that the outputs from PARTNER are closest to expert rewritings (86% better than the next best baseline, BART).

5 Discussion and Conclusion

The burden of mental illness globally is overwhelming, and common mental disorders are some of the most debilitating illnesses worldwide [Collins *et al.*, 2011]. Existing mental health resources and interventions are ill-suited to the size of the need. Online mental health support platforms that make use of peer supporters is one route to scaling up support, but the biggest challenge is to effectively train or scaffold the peer supporters. Our empathic rewriting approach represents a foundational proof-of-concept of how computational methods may help peer supporters online.

Rewriting human-generated responses may be an effective approach to balancing the benefits and risks of using artificial intelligence in mental health settings. By combining human knowledge of context and experience, our approach can both provide feedback to online peer-supporters with actionable, real-time examples, and provide support seekers with more empathic responses. Importantly, this machine-in-the-loop approach can help mitigate some of the risks related to toxicity and safety of AI systems in settings of suicidal ideation, self-harm, or insensitive comments related to race/ethnicity/gender [Li *et al.*, 2020; Luxton *et al.*, 2012].

Acknowledgements

We would like to thank TalkLife and Jamie Druitt for supporting this work and for providing us access to a TalkLife dataset. T.A., A.S., and I.W.L. were supported in part by NSF grant IIS-1901386, NSF grant CNS-2025022, NIH grant R01MH125179, Bill & Melinda Gates Foundation (INV-004841), the Office of Naval Research (#N00014-21-1-2154), a Microsoft AI for Accessibility grant, and a Garvey Institute Innovation grant. A.S.M. was supported by NIH NCATS grants KL2TR001083 and UL1TR001085 and the Stanford Human-Centered AI Institute.

References

[Bohart *et al.*, 2002] Arthur Bohart, Robert Elliott, Leslie Greenberg, and Jeanne Watson. Empathy. *J. C. Norcross (Ed.)*, 2002.

[Brown *et al.*, 2020] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

[Buechel *et al.*, 2018] Sven Buechel, Anneke Buffone, Barry Slaff, Lyle Ungar, and João Sedoc. Modeling empathy and distress in reaction to news stories. In *EMNLP*, 2018.

[Clark *et al.*, 2018] Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A Smith. Creative writing with a machine in the loop: Case studies on slogans and stories. In *IUI*, 2018.

[Collins *et al.*, 2011] Pamela Y Collins, Vikram Patel, Sarah S Joestl, Dana March, Thomas R Insel, Abdallah S Daar, Isabel A Bordin, E Jane Costello, Maureen Durkin, Christopher Fairburn, et al. Grand challenges in global mental health. *Nature*, 2011.

[Elliott *et al.*, 2011] Robert Elliott, Arthur Bohart, Jeanne Watson, and Leslie S Greenberg. Empathy. *Psychotherapy*, 2011.

[Elliott *et al.*, 2018] Robert Elliott, Arthur Bohart, Jeanne Watson, and David Murphy. Therapist empathy and client outcome: An updated meta-analysis. *Psychotherapy*, 2018.

[He *et al.*, 2019] Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. A probabilistic formulation of unsupervised text style transfer. In *ICLR*, 2019.

[Holtzman *et al.*, 2020] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *ICLR*, 2020.

[Lewis *et al.*, 2020] Mike Lewis, Yinhan Liu, and et al. Goyal, Naman. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL*, 2020.

[Li *et al.*, 2016] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A diversity-promoting objective function for neural conversation models. In *NAACL-HLT*, 2016.

[Li *et al.*, 2020] Ron C Li, Steven M Asch, and Nigam H Shah. Developing a delivery science for artificial intelligence in healthcare. *NPJ Digital Medicine*, 2020.

[Luxton *et al.*, 2012] David Luxton, Jennifer June, and Jonathan Fairall. Social media and suicide: a public health perspective. *American journal of public health*, 2012.

[Majumder *et al.*, 2020] Navonil Majumder, Pengfei Hong, Shanshan Peng, Jiankun Lu, Deepanway Ghosal, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. Mime: Mimicking emotions for empathetic response generation. In *EMNLP*, 2020.

[Miner *et al.*, 2019] Adam Miner, Nigam Shah, Kim Bullock, Bruce Arnow, Jeremy Bailenson, and Jeff Hancock. Key considerations for incorporating conversational ai in psychotherapy. *Frontiers in psychiatry*, 10, 2019.

[Olfson, 2016] Mark Olfson. Building the mental health workforce capacity needed to treat adults with serious mental illnesses. *Health Affairs*, 2016.

[Papineni *et al.*, 2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002.

[Radford *et al.*, 2019] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.

[Rashkin *et al.*, 2019] Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *ACL*, 2019.

[Robert *et al.*, 2011] Elliot Robert, Arthur Bohart, JC Watson, and LS Greenberg. Empathy. *Psychotherapy*, 2011.

[Selman, 1980] Robert L Selman. *Growth of interpersonal understanding*. Academic Press, 1980.

[Sharma *et al.*, 2020] Ashish Sharma, Adam Miner, Dave Atkins, and Tim Althoff. A computational approach to understanding empathy expressed in text-based mental health support. In *EMNLP*, 2020.

[Sharma *et al.*, 2021] Ashish Sharma, Inna Lin, Adam Miner, Dave Atkins, and Tim Althoff. Towards facilitating empathic conversations in online mental health support: A reinforcement learning approach. In *TheWebConf*, 2021.

[Tanana *et al.*, 2019] Michael Tanana, Christina Soma, Vivek Srikumar, David Atkins, and Zac Imel. Development and evaluation of clientbot: Patient-like conversational agent to train basic counseling skills. *JMIR*, 2019.

[White and Dorman, 2001] Marsha White and Steve M Dorman. Receiving social support online: implications for health education. *Health education research*, 2001.

[Zhang *et al.*, 2020] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation. In *ACL, system demonstration*, 2020.