

How AI can help people be more empathetic about mental health

January 23 2023



Credit: Pixabay/CC0 Public Domain

Empathy is critical to having supportive conversations about mental health. But this skill can be tricky to learn, especially in the moment when a person is sharing something hard.

A team led by researchers at the University of Washington studied how artificial intelligence could help people on the platform [TalkLife](#), where

people give each other mental health support. The researchers developed an AI system that suggested changes to participants' responses to make them more empathetic. The system helped people communicate [empathy](#) more effectively than traditional training did. In fact, the best responses resulted from a collaboration between AI and people.

The researchers published these findings Jan. 23 in *Nature Machine Intelligence*.

UW News reached out to senior author Tim Althoff, UW assistant professor in the Paul G. Allen School of Computer Science & Engineering, for details about the study and the concept of AI and empathy.

Why did you choose the TalkLife platform to study?

Tim Althoff: Prior research suggests that peer-support platforms could have a significant positive impact on mental health care because they help address the massive challenge of access. Because of insurance issues, stigma or isolation, many people find that free online peer-support platforms are all they have access to. TalkLife is the biggest peer-support platform globally and it has a large number of motivated peer supporters.

Also, TalkLife leadership recognized the importance and potential impact of our research on how computing can empower [peer support](#). They kindly supported our research through collaboration, feedback, participant recruiting and data-sharing.

What inspired you to help people communicate with more empathy?

TA: It is well established that empathy is critical for helping people feel supported and for forming trusted relationships. But empathy is also complex and nuanced. It can be challenging for people to find the right words in the moment.

While counselors and therapists are trained in this skill, [our prior research](#) established that peer supporters currently miss many opportunities to respond to each other more empathetically. We also found that peer supporters do not learn how to express empathy more effectively over time, which suggests that they could benefit from empathy training and feedback.

On the surface it seems counterintuitive to have AI help with something like empathy. Can you talk about why this is a good problem for AI to solve?

TA: What the AI feedback can do is be very specific and be "contextual" and give suggestions about concretely responding to a message that's right in front of someone. It can give someone ideas in a "personalized" way rather than through generic training examples or with rules that may not apply to every single situation a person will face. It also only pops up if someone needs it—if their response is great, the system can give a light touch of positive feedback.

People might wonder "why use AI" for this aspect of human connection. In fact, we designed the system from the ground up not to take away from this meaningful person-person interaction. For example, we only show feedback when needed and we train the model to make the smallest possible changes to a response to communicate empathy more effectively.

How do you train an AI to 'know' empathy?

TA: We worked with two [clinical psychologists](#), Adam Miner at Stanford University and David Atkins in the UW School of Medicine, to understand the research behind empathy and adapt existing empathy scales to the asynchronous, text-based setting of online support on TalkLife. Then we had people annotate 10,000 TalkLife responses for various aspects of empathy to develop AI models that can measure the level of expressed empathy in text.

To teach the AI to give actionable feedback and concrete suggestions, we developed a reinforcement learning-based system. These systems need a lot of data to be trained, and while empathy isn't expressed as often as we would like on platforms such as TalkLife, we still found thousands of good examples. Our system learns from these to generate helpful empathy feedback.

In your evaluation of this system, did you see people becoming reliant on AI for empathy or did people learn how to be more empathetic over time?

TA: Our randomized trial demonstrated that peer supporters with access to feedback expressed between 20% and 40% more empathy than supporters in the control group that did not have access to such feedback.

Among our participants, 69% of peer supporters reported that they feel more confident at writing supportive responses after this study, indicating increased self-efficacy.

We further studied how participants made use of the feedback and found that peer supporters did not become overly reliant on the AI. For example, they would use the feedback indirectly as a broader inspiration rather than "blindly" following the recommendations. They also flagged

feedback in the few cases when it was not helpful or even inappropriate. I was excited that the collaboration between human peer supporters and AI systems led to better outcomes than either one on their own.

I also want to highlight the significant efforts we took to consider and address ethical and safety risks. Those include having the AI work with the peer supporter instead of the person currently in crisis, conducting the study in a TalkLife-like environment that is intentionally not integrated into the TalkLife platform, giving all participants access to a crisis hotline and allowing peer supporters to flag feedback for review.

What do these results mean in terms of the future of human-AI collaboration?

TA: One area of human-AI collaboration that I am particularly excited about is AI-supported communication. There are so many challenging communication tasks with critical outcomes—from helping someone feel better to challenging misinformation on [social media](#)—where we seem to expect people to do well, without any form of training or support. In most cases, all we are given is an empty chat box.

We can do better, and I believe that [natural language](#) processing technology can play a big role in helping people achieve their conversational goals. Specifically, our study shows that human-AI collaboration can be effective even for complex and open-ended tasks such as having empathetic conversations.

More information: Tim Althoff, Human–AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support, *Nature Machine Intelligence* (2023). [DOI: 10.1038/s42256-022-00593-2](https://doi.org/10.1038/s42256-022-00593-2).
www.nature.com/articles/s42256-022-00593-2

Provided by University of Washington

Citation: How AI can help people be more empathetic about mental health (2023, January 23) retrieved 12 May 2024 from <https://medicalxpress.com/news/2023-01-ai-people-empathetic-mental-health.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.