



Kent Academic Repository

She, Wan-Jou, Tabbaa, Luma, Covaci, Alexandra, Siriaraya, Panote, Yip, Benjamin, Perusquia-Hernandez, Monica, Nakajima, Shinsuke, Yamashita, Naomi, Shibuya, Yu and Ang, Chee Siang (2026) *Designing for GenAI-mediated mental health peer support: Lessons from human peer support interactions in a wellbeing app*. International Journal of Human-Computer Interaction . pp. 1-30. ISSN 1044-7318.

Downloaded from

<https://kar.kent.ac.uk/115954/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.1080/10447318.2026.2708444>

This document version

Publisher pdf

DOI for this version

Licence for this version

CC BY-NC-ND (Attribution-NonCommercial-NoDerivatives)

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal** , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).



Designing for GenAI-Mediated Mental Health Peer Support: Lessons from Human Peer Support Interactions in a Wellbeing App

Wan-Jou She, Luma Tabbaa, Alexandra Covaci, Panote Siriaraya, Benjamin Yip, Monica Perusquia-Hernandez, Shinsuke Nakajima, Naomi Yamashita, Yu Shibuya & Chee Siang Ang

To cite this article: Wan-Jou She, Luma Tabbaa, Alexandra Covaci, Panote Siriaraya, Benjamin Yip, Monica Perusquia-Hernandez, Shinsuke Nakajima, Naomi Yamashita, Yu Shibuya & Chee Siang Ang (03 Aug 2026): Designing for GenAI-Mediated Mental Health Peer Support: Lessons from Human Peer Support Interactions in a Wellbeing App, International Journal of Human-Computer Interaction, DOI: [10.1080/10447318.2026.2708444](https://doi.org/10.1080/10447318.2026.2708444)

To link to this article: <https://doi.org/10.1080/10447318.2026.2708444>



© 2026 The Author(s). Published with license by Taylor & Francis Group, LLC



Published online: 03 Aug 2026.



Submit your article to this journal [↗](#)



Article views: 260



View related articles [↗](#)



View Crossmark data [↗](#)

Designing for GenAI-Mediated Mental Health Peer Support: Lessons from Human Peer Support Interactions in a Wellbeing App

Wan-Jou She^a , Luma Tabbaa^b , Alexandra Covaci^b , Panote Siriaraya^a , Benjamin Yip^c ,
Monica Perusquia-Hernandez^d , Shinsuke Nakajima^e , Naomi Yamashita^f , Yu Shibuya^a 
and Chee Siang Ang^g 

^aDepartment of Information and Human Sciences, Kyoto Institute of Technology, Kyoto, Japan; ^bSchool of Engineering, Mathematics & Physics, University of Kent, Canterbury, UK; ^cDepartment of Geriatrics, Kent and Medway Medical School, Canterbury, UK; ^dDivision of Information Science, Nara Institute of Science and Technology, Nara, Japan; ^eDepartment of Information Science and Engineering, Kyoto Sangyo University, Kyoto, Japan; ^fDepartment of Informatics, Kyoto University, Kyoto, Japan; ^gSchool of Computing, University of Kent, Canterbury, UK

ABSTRACT

Mental health services remain difficult to access worldwide, and many young people turn instead to friends or online communities for support, particularly where mental illness is highly stigmatized. We study how such peer support is practised, and what this implies for the design of Generative AI (GenAI). We deployed a peer-support app with 28 Japanese university students over two weeks and qualitatively analyzed their peer-to-peer chat conversations. We identify seven recurring relational patterns, including tentative openings, routes into depth, layered accounts of lived experience, reciprocity, and careful navigation of medical disclosures. Building on these findings, we distinguish two roles for GenAI in this ecosystem: GenAI as support for human peers (backstage tools that scaffold and strengthen human-to-human support) and GenAI as a peer (a conversational companion in its own right). For each, we outline opportunities, limitations, and ethical boundaries, advancing HCI discourse beyond therapist-oriented GenAI, toward more reciprocal, peer-based mental health support.

KEYWORDS

Peer support; mental health; social support dynamics; conversational agents

1. Introduction

Mental health disorders are among the most pressing global health challenges; as of 2019, 1 in 8 people worldwide live with a mental health disorder, most prominently anxiety and depression disorders, a 48% increase from 1990 (GBD 2019 Mental Disorders Collaborators, 2022; World Health Organisation (WHO), 2022). Suicide is now the second leading cause of death among adolescents in the USA (National Alliance on Mental Illness (NAMI), 2025), and Japan reports the highest suicide rate among G7 countries (OECD, 2020). Despite the increased need, access to care remains limited. Only 23% of individuals with depression in high-income countries and just 3% in low- and middle-income countries receive adequate treatment (Moitra et al., 2022). In the USA, many individuals refrain from seeking formal care due to cost, limited availability, and geographic barriers (NAMI, 2025). In the UK, patients may wait 18 months or more for care due to a shortage of professionals (Punton et al., 2022), a scarcity that is echoed globally (WHO, 2022).

Disturbances in mental health could leave individuals experiencing isolation and loneliness, amplifying feelings of disconnection and stigma (Ando et al., 2013; Prizeman et al., 2023). Moreover, the scarcity of social support can further complicate these struggles (Wang et al., 2018). Hence, peer support has gradually emerged as a form of support provided by individuals with shared lived experiences. Through reciprocal emotional support and mutual understanding (Chung et al., 2024; Singer &

CONTACT Wan-Jou She  wjs2004@kit.ac.jp  Department of Information and Human Sciences, Kyoto Institute of Technology, Kyoto, 606-8585, Japan

© 2026 The Author(s). Published with license by Taylor & Francis Group, LLC

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Klimecki, 2014), peers can offer unique forms of empathy (Chung et al., 2024; Singer & Klimecki, 2014), emotional validation and normalization (Repper & Carter, 2011). Peer support can also reduce stigma by creating nonprofessional spaces of solidarity (Pfeiffer et al., 2011), while fostering empowerment and personal agency through reciprocal relationships (Solomon, 2004). Consequently, peer support is increasingly recognized as a vital component of contemporary mental health ecosystems (Gillard et al., 2015).

In the modern era, digital platforms are being widely used to amplify these benefits, enabling people with shared struggles to connect online. In the Human–Computer Interaction (HCI) community, research has identified unique values in digital peer support spaces, including extended inclusivity, enabling scalable and flexible support across geographical and demographic boundaries, and allowing anonymity (Andalibi et al., 2017; Drysdale et al., 2022; Kingod et al., 2017; Saha & De Choudhury, 2017). Despite these efforts, current accessible technology is not without its limitations. Finding a peer with whom one truly connects can be difficult, especially for those in low-prevalence or stigmatized groups (Kingod et al., 2017; Naslund et al., 2016). Even when a match is made, the support quality could be uneven, with reciprocity not guaranteed, and engagement often declines as conversations fade or lapse (Wang et al., 2023; Yang et al., 2017; Yao et al., 2021; Yuan et al., 2025). Studies of web-based and app-based peer-support platforms (e.g., Buddy Project (Andalibi & Flood, 2021), Togetherall¹) echo these struggles around matching, safety, and uneven responsibility between peers (Andalibi & Flood, 2021; O’Leary et al., 2017; Thomson et al., 2023). Furthermore, even online support services that provide a “live chat” with peers cannot guarantee timely availability, because suitable peers may simply not be present at the precise moment of need (She et al., 2023), let alone individuals with low-prevalence conditions or those experiencing stigmatized circumstances.

Recent advances in Generative Artificial Intelligence (GenAI), and in particular Large Language Models (LLMs), has emerged as a technology that could scaffold, augment, or even emulate these interactions to deliver mental health support. For example, existing work has investigated Conversational Agents (CAs) framed as GenAI “therapists,” reporting success in delivering structured therapy such as Cognitive Behavioral Therapy (CBT) and psychoeducation (Fitzpatrick et al., 2017; Karkosz et al., 2024; Nicol et al., 2022). Other studies have examined GenAI-powered prompts to scaffold wellbeing support conversations for trained counselors (Hsu et al., 2025; Syed et al., 2024). However, the majority of these approaches design the GenAI to assume the perspective of therapists or advisors and little is known about how GenAI might be designed to support everyday peers in this context, nonprofessionals whose exchanges are more informal, reciprocal, and grounded in lived experience.

Herein, we aim to address these gaps by analyzing human peer-to-peer support interactions in a wellbeing app, PEER2S, to distill the design-relevant patterns of peer talk, providing the substrate needed to inform GenAI that supports (and cautiously emulates) peer exchange. In this study, we situate our research in Japan, where the sharp rise in suicides among adolescents and young adults during and after the COVID-19 pandemic has reversed more than a decade of overall national decline (Okada et al., 2024; Tanaka & Okamoto, 2021).

As such, we aim to understand how peer-to-peer support unfolds and what this means for the design of future GenAI support systems by addressing the following research questions:

RQ1: What relational and expressive practices characterize peer-to-peer support conversations?

RQ2: How can GenAI-powered tools augment and support human-to-human peer support?

RQ3: How can insights from human peer-to-peer support inform opportunities, limitations, and design boundaries of positioning GenAI as a peer in support conversations?

Through answering these RQs with our study, we aim to make several key contributions to the HCI community. First, in addressing RQ1, we provide an empirical account of authentic, technology-mediated human peer conversations, identifying seven recurring relational and expressive patterns that underpin the peer support construct. Second, to address RQ2 and RQ3, we conceptually extrapolate these empirical findings to formulate design principles for how GenAI could operate in peer support. Grounded in established social theories and GenAI literature, we contrast two potential roles: GenAI

for peer support (backstage scaffolding) and GenAI as a peer (a conversational companion), outlining the opportunities, limitations, and ethical boundaries of each. Finally, we hope to make a conceptual contribution by advancing the HCI discourse beyond therapist-centric GenAI models toward more reciprocal and accessible forms of peer-based mental health support.

2. Related work

2.1. *The promise and fragility of human peer support*

Peer support is grounded in the premise that people who have encountered similar challenges can offer one another understanding, validation, and encouragement shaped by shared lived experience; where their experiential knowledge of distress, recovery, and coping enables them to relate to others through a sense of “having been there” (Hameed Shalaby & Agyapong, 2020; Rayland & Andrews, 2023; Simmons et al., 2023). As such, peer support is defined as the system of giving and receiving help grounded in respect, shared responsibility, and mutual agreement on what is helpful, centered on empathically understanding another’s situation through shared experiences of emotional and psychological pain (Mead et al., 2001). In this sense, peer support is not only a channel for emotional assistance but a relational practice through which lived experience becomes a shared resource.

Such dynamics closely align with Social Exchange Theory, which conceptualizes social relationships as reciprocal exchanges sustained through trust, mutual investment, and the expectation of relational benefit (Cropanzano & Mitchell, 2005). Within peer support, individuals are not positioned solely as recipients of care, but as active contributors who move fluidly between offering and receiving support; a very defining feature of peer support (Repper & Carter, 2011). This reciprocity fosters relational closeness, strengthens social bonds, and reinforces a sense of agency and belonging, as individuals derive value not only from being supported but also from supporting others through their own lived experiences. These dynamics also resonate with Social Support Theory (Brandtzaeg et al., 2021; Meng et al., 2023), which proposes that social support promotes wellbeing not only by buffering stress, but by fostering self-esteem, self-regulation, companionship, intimacy, and a sense of social connectedness through meaningful relational processes (Lakey & Cohen, 2000). Within peer support, these relationship processes are embedded within ongoing exchanges of empathy, mutual recognition, and shared understanding, allowing peers to experience support not as isolated acts of assistance but as part of a broader sense of belonging and interpersonal connection.

Another central feature of peer support is empathy; a practice that involves demonstrating an understanding of another person’s situation and/or feelings and communicating that understanding back to the person so that they feel understood (Pudlinski, 2005). In peer support contexts, empathy is demonstrated through recognition, shared experience, and experiential resonance, where they draw on their own histories of distress, recovery, and coping to communicate that another person’s struggle is understandable, survivable, and not isolating (Gillard et al., 2015; Mulfinger et al., 2018). This form of empathy enables individuals to feel recognized not only as people in distress, but as people whose experiences can be understood by someone who has encountered similar vulnerabilities.

When positioning peer support within the broader mental health ecosystem, professional therapy and peer-provided care are increasingly recognized as complementary modalities rather than a strict dichotomy. Whilst peer support cannot replace professional therapy in any way or form, it offers complementary value by drawing on overlapping mechanisms of effective support with clinical modalities, such as active listening, empathic engagement, and emotional validation. Additionally, peer support is defined by mutuality, reciprocity, and empathy grounded in shared lived experiences (Repper & Carter, 2011). Within this space, peer support uniquely fosters a sense of agency and empowerment through its reciprocal nature, allowing individuals to simultaneously give and receive care while sharing practical, lived coping strategies (Gillard et al., 2015). This non-stigmatising (Mulfinger et al., 2018) and cost-effective nature (Lawn et al., 2008) provides a vital, accessible layer of care that operates alongside and extends beyond the temporal and relational boundaries of formal therapy sessions.

Despite its value, peer support also faces inherent limitations. Finding a peer with whom one truly connects can be difficult, and sustaining such relationships over time is often challenging, as

conversations fade or lapse with shifting life circumstances (Gillard et al., 2015). Reciprocity is not always guaranteed: some participants may focus primarily on their own needs without contributing support to others, creating asymmetries that can strain the peer support mechanism (Solomon, 2004). In that, the quality of peer interactions can be uneven, with some exchanges offering rich empathy while others risk being superficial, unhelpful, or even inadvertently harmful (Pfeiffer et al., 2011). Finally, the asynchronous pattern of engagement is an inherent pattern in human-to-human interactions, which may hinder the accessibility of support that is needed (Fortuna et al., 2020). These tensions position peer support as both promising and precarious: a powerful form of relational care grounded in lived experience, but one that requires careful scaffolding if it is to remain reciprocal, safe, and sustainable.

2.2. Taking peer support online is not an easy fix

Research examining digitally mediated human peer-support platforms is not rare in the HCI community. Studies have explored the use of popular platforms such as Reddit² (De Choudhury & De, 2014), Instagram³ (Andalibi et al., 2017), Facebook⁴ and TikTok⁵ (Nanthanasit et al., 2025) in contexts related to rehabilitation and post-traumatic injury (Nanthanasit et al., 2025), sexual abuse (Andalibi et al., 2017), eating disorders (Chancellor et al., 2016), and stress (Saha & De Choudhury, 2017). Such digitally-mediated spaces have often been understood through the Social Identity model of Deindividuation Effects (SIDE), which proposes that anonymity and reduced social cues online can shift interactions away from individual identity and toward a stronger sense of shared group identity and belonging (Russell Spears, 2017). In mental health contexts, these affordances can encourage users to disclose sensitive experiences, seek affiliation, and participate in communities organized around shared vulnerability and lived experience.

Nonetheless, social media platforms have been criticized and even described as “minefields” in relation to mental health and wellbeing. Suler described it as *online disinhibition effect*, where anonymity, invisibility, and asynchronicity can facilitate openness and vulnerability while also lowering social restraint and increasing the potential for harmful behavior (Suler, 2004). For instance, features such as “likes” and “reactions” have been found to promote toxic social comparisons (Rosenthal-von der Pütten et al., 2019; Sánchez-Hernández et al., 2022), public commenting has been linked to cyberbullying (Macaulay et al., 2021), and the self-disclosures to strangers may expose users to risks such as phishing (i.e., the fraudulent attempt to obtain sensitive information by deception) (Qahri-Saremi & Turel, 2023) and doxxing (i.e., the act of publicly broadcasting identifying information without one’s consent) (Chen et al., 2018).

Other works of research explored the use of tailored mental health platforms or Peer Support Networks (PSNs) such as 7 Cups of Tea⁶ (Baumel et al., 2018) or TalkLife⁷ (Williams et al., 2024) in promoting and sustaining peer-to-peer support. PSNs have been seen as complementary to formal healthcare resources by offering unique benefits derived from lived experiences (Andalibi et al., 2017; Hameed Shalaby & Agyapong, 2020; Repper & Carter, 2011). These benefits include providing safe spaces for self-disclosure, reciprocity, and emotional validation, which often reduces stigma by enabling support in anonymous or semi-anonymous contexts (Andalibi et al., 2018; Baumel et al., 2018; Hameed Shalaby & Agyapong, 2020). This builds onto Walther’s Hyperpersonal Model, which suggests that selective self-presentation and carefully controlled disclosure online can intensify perceptions of intimacy and emotional connection, sometimes exceeding those typically experienced in face-to-face settings (Walther, 1996). In such cases, users frequently engage in reciprocal exchanges, offering validation and encouragement in return for the support they receive (Andalibi et al., 2018).

Nonetheless, consistent with other forms of peer support, PSNs also face similar challenges related to imbalanced participation (Saha & De Choudhury, 2017), declined engagement over time (Andalibi et al., 2018), and inconsistent quality of interactions (Pater et al., 2020). For instance, in a study that utilized GenAI-based algorithmic matching similar to those found in dating apps to match suitable peers, the study reported that despite the success in matching peers who share a similar account of struggles, and a reduction in self-reported levels of social anxiety, there were still limitations relevant to the human factor in peer-to-peer support (She et al., 2023). In their reported limitations, peers who had once connected were often unavailable when participants sought real-time interaction, leading

participants to report no perceived improvements in direct social support after using the system. Furthermore, as is common in peer-to-peer settings among nonprofessional users, the study found it challenging to sustain meaningful conversation. These findings highlight the practical realities of expecting peer supporters to be consistently available, emotionally prepared, and able to provide appropriate responses when needed.

These limitations further reinforce that digitally-mediated peer support is not merely a technical problem of connection or matching, but an ongoing interactional process requiring sustained relational labor, reciprocity, and communicative effort across time, consistent with broader HCI research emphasizing the importance of social presence, relational maintenance, and sustained interaction in online communication environments (Russell Spears, 2017; Suler, 2004; Walther, 1996). Therefore, it is of utmost importance to explore complementary technologies that can help bridge these gaps in availability and support peer-to-peer conversation flows.

2.3. Toward the integration & design of GenAI for peer support

In recent years, rapid advances in AI technology, particularly GenAI for natural language processing, have enabled the development of CAs that can deliver engaging, natural, human-like responses and simulate supportive interactions with users. Studies have explored the use of such agents in the assessment of mental health and wellbeing disorders such as Post-Traumatic Stress Disorder (PTSD), anxiety, and depression (DeVault et al., 2014). Other works of research report that CA therapists have helped users reduce their symptoms of depression, anxiety, and distress (Fitzpatrick et al., 2017; Li et al., 2023), understand their mental health better (Won et al., 2025), and reduce stigma associated with seeking professional mental health support (Lee et al., 2023). Such positive outcomes have been observed across diverse mental health contexts, including grief (Xygkou et al., 2023), social anxiety (Wang et al., 2025b), and maternal mental health support (Inkster et al., 2023; Suharwardy et al., 2023). These applications span a wide range of user populations, including youth and university students (Fitzpatrick et al., 2017; Poulsen et al., 2025), individuals on the autism spectrum (Xygkou et al., 2024b), people with early-stage dementia (Xygkou et al., 2024a), and children (Park et al., 2024).

Despite this progress, the current design of GenAI in the mental health and wellbeing domain remains limited. First, most CAs described in HCI research were deliberately designed to deliver therapist-like support (Wang et al., 2025b), adopting clinical design paradigms such as CBT (Fitzpatrick et al., 2017; Karkosz et al., 2024; Nicol et al., 2022) and psychoeducation (Poulsen et al., 2025; Wang et al., 2025b). These systems were designed around professional empathy (Poulsen et al., 2025; Wang et al., 2025b; Xygkou et al., 2024b), guided questioning (Fitzpatrick et al., 2017; Nicol et al., 2022), and structured support anchored in the perspective and boundaries of a clinician. As such, far less is known about how GenAI systems might be designed in the context of peer support.

Secondly, where GenAI has been designed to act as a peer, it offered many benefits, but struggled with authenticity and the ability to dive into deep conversations. Prior studies found that users value the non-judgmental quality of non-human interlocutors; they disclose more to CAs than to people, attributing this to reduced feelings of judgment (Holthöwer & Van Doorn, 2023; Lucas et al., 2014). Moreover, CAs have been shown to support self-compassion more effectively than human caregivers (Lee et al., 2019). However, CAs are still limited in the quality of empathy they can offer; studies found that LLMs such as GPT⁸, Llama⁹, and Mistral¹⁰ provide emotional and informational support, but they cannot engage in deep conversations, provide authentic references, or share lived experiences (Saha et al., 2026; Young et al., 2024). Thus, despite their ability to communicate superficial empathy, there is a critical need to explore how GenAI can be designed to be a “good” peer.

Thirdly, where GenAI was positioned as a “backstage supporter” to scaffold human peer-to-peer support, it was designed to support counselors, not everyday peers. For instance, several research works examined training counselors with GenAI-powered tools to enhance their counseling skills by utilizing CBT-inspired prompts and other counseling strategies (Hsu et al., 2025; Syed et al., 2024; Wang et al., 2025a). However, such design implications are more designed for “trained” peer counselors who need GenAI-powered tools to manage multiple conversations, manage their time, and manage conversations. As such, the design characteristics of GenAI for authentic, everyday peer support remain underexplored.

2.4. Contextualizing the research: Mental health and peer support among Japanese university students

As noted in our introduction, the mental health crisis among Japanese young adults is shaped as much by cultural norms as by structural limitations. Beneath the widely praised social harmony, hospitality, and politeness lies a deeply internalized expectation to suppress vulnerability so as not to disrupt others or impose one's burdens (Ina & Morita, 2015; Zhang et al., 2019). In daily life, the non-forgiving, shame-oriented response to mistakes or social missteps encourages heightened self-monitoring, where even minor clumsiness can invite criticism or embarrassment (Lebra, 1983; Stoeber et al., 2013). As a result, hesitation, restraint, and emotional self-containment are often treated as virtues, contributing to a relational environment in which young people feel compelled to endure distress privately rather than seek support (Ina & Morita, 2015). Unsurprisingly, this culture of concealed suffering amplifies young people's reliance on safer, lateral forms of connection, aligning with broader findings that youth, more than other age groups, turn to peers for emotional support (Hameed Shalaby & Agyapong, 2020), whether it was face-to-face or via messaging apps (i.e., LINE¹¹) and social platforms (i.e., X (previously Twitter)¹²) (Ye et al., 2023).

3. Methodology

3.1. Ethics

Participants were recruited on campus using flyers and word of mouth. Inclusion criteria were: (i) aged between 18 and 24 years, (ii) actively enrolled at the university conducting this research, (iii) able to read and write Japanese, and (iv) having regular internet access.

Interested students received an information sheet that explained the study aims, the sensitive nature of the topic (i.e., the potential of discussing stress, loneliness, and mental health struggles with others), the voluntary nature of participation, and how their data would be used. Furthermore, interested students had the opportunity to ask questions and take their time to consider participation. Once decided, participants provided written informed consent before registering on the *PEER2S* app and were reminded that they could withdraw at any time without penalty or explanation. They were also informed that, upon withdrawal, they could request the deletion of their data. In addition, participants received the researchers' contact details and were told they could reach out at any time during the study if they had concerns or experienced distress. In the event of unexpected distress during the study, participants were informed that they could access the University counseling service. The University counseling service was aware of our study and the potential referrals they might receive as a result of student participation in this study. The study protocol was approved by the Institutional Review Board at Kyoto Sangyo University (IRB No:0159).

3.2. Procedure

After signing an informed consent form, participants were invited to register on *PEER2S* and set up an anonymous profile. Once the recruitment period had ended, participants were then able to start using the app (i.e., surf through anonymous peer profiles and receive suggestions from the app on "potentially suitable peers." Furthermore, participants were able to offer and engage in varied levels of support, from "light-weight & low-commitment" virtual "hugs" and supportive messages, to a full one-to-one conversation with matched peers (see Section 3.3 for details about the system).

3.3. PEER2S: A GenAI-mediated peer-support wellbeing app

We conducted the study using *PEER2S*, a peer-support wellbeing app designed to facilitate anonymous peer-to-peer support among Japanese university students, building on prior work on peer-support matchmaking for this population (She et al., 2023). The *PEER2S* app was designed and developed in collaboration with four experts (two mental health professionals, a sociologist, and a User Experience

(UX) researcher), who provided iterative feedback on the app's core features, interaction flow, and safety considerations, with respect to the cultural nuances of the Japanese context.

3.3.1. Profile set up & view

PEER2S allow users to create anonymous user profiles. The system assigns each account a random user-name and does not store real names or any directly identifying information. Profile fields capture basic demographic and contextual information (gender, age, field of study, and social situation such as “living alone” or “living with family”), along with hobbies, coping strategies, and an optional free-text description of past or current wellbeing struggles. Users can review their profile details, hugs and support messages received in the “Personal Profile” page.

3.3.2. Finding peers

Once a profile has been created, PEER2S incorporates a recommendation mechanism inspired by prior work on peer-support matchmaking (She et al., 2023). Each day, the recommendation service uses profile attributes (i.e., demographics, social situation, and key concerns) to generate up to five suggested Story Cards for each user. Suggestions are shown in the “Profile Recommendation” page (see Figure 1).

3.3.3. Incorporated features for peers to offer support

Once a user views another user's Story Card, the app allows three forms of support offering: (i) a one-off virtual “hug,” which serves as a minimal, low-effort gesture of care; (ii) a short *support message*, where users can type a one-off message; and (iii) a match request that if approved by the other user, allows both users to engage in a one-to-one conversation (see Figure 1). Note the app does not show how many hugs, support messages or matches each peer has in the Profile Recommendation page.

3.3.4. Matching and conversations

When a match request from one Story Card is accepted, PEER2S creates a private chatroom linking the two users. Chatrooms support asynchronous text and emoji exchanges and are accessed via a dedicated Chat tab (See Figure 1), alongside tabs for viewing one's own profile and browsing Story Cards.




Figure 1. A demonstration of the three pages of PEER2S: user's personal profile page, peers' profiles recommendation page and the chatroom page.

3.4. Participants and collected data

A total of 28 university students enrolled in the study (15 males, 10 females, and 3 preferred not to disclose), and 23 completed the two-week online peer support study. A sample of 23 participants is consistent with common practice in qualitative HCI research, where qualitative studies typically involve around 12 participants (Caine, 2016). Given that our analysis follows Reflexive Thematic Analysis (Braun & Clarke, 2006; Byrne, 2022), which does not treat saturation as a procedural requirement but instead foregrounds iterative and reflexive engagement with the data, this sample size is appropriate for generating analytically rich and contextually grounded themes.

The participants were third- and fourth-year university students, most of whom lost valuable initial face-to-face social opportunities with their classmates because many of them entered university during the COVID-19 lockdown. Participants were recruited through university email lists or peer referrals, which could introduce a potential self-selection bias as the study may have disproportionately attracted students within close social circle or with a preexisting interest in mental health support. In here, we captured all user interactions within the app for analysis, including: peer matching, virtual hugs, supportive messages, and chatlogs of matched peers. In total, 50 chatrooms (i.e., matched pairs) were created. However, 17 matches were created but never used (66% active chatrooms), and one user never matched with any other peer (97.92% match rate). As such, a total of 33 chatlogs were extracted for analysis. As for virtual hugs, 82% of the participants sent at least one virtual hug to another (Mean = 9.32, SD = 7.74), and all participants received at least one virtual hug (Mean = 7.78, SD = 3.30). Finally, for supportive messages, 73.91% of participants sent at least one support message (Mean = 3.70, SD = 4.04) and 95.6% of participants received at least one support message (Mean = 2.63, SD = 1.32).

3.5. Data analysis

Considering the *PEER2S* app was designed to be deployed for Japanese university students, the app was fully designed in Japanese, and the participants used the Japanese language in chatlogs and supportive messages. The chatlogs and supportive messages were initially auto-translated from Japanese to English using DeepL¹³. Members of the research team reviewed and corrected translations where necessary, with particular attention to idioms, emotional nuance, and context, and consulted the original Japanese when quotations were selected for reporting. For instance, the Japan-based researchers had to clarify to UK-based researchers regarding Japanese specific communication norms, such as the emphasis in Japanese society on acknowledging others' efforts and avoiding placing burdens on others (e.g., expressions of apology when one feels they are contributing less to a conversation). This cross-cultural review and discussion enabled the research team to interpret the data from a broader perspective and to identify culturally meaningful nuances that might otherwise have been overlooked.

We adopted Reflexive Thematic Analysis (RTA) (Braun & Clarke, 2006; Byrne, 2022), focusing on meaning-making rather than inter-coder reliability. Given the cross-cultural composition of the data analytical team, our analysis was inductive and explicitly acknowledged the researchers' active role in interpreting the data.

To allow dynamical and collaborative deliberation of the theme from the beginning of the data interpretation, each researcher started by iteratively proposing descriptive themes during Phase 1 (familiarisation) and Phase 2 (summarizing quotes to inductive nodes). We intended to capture the most insightful and timely themes while researchers were closest to the raw data and actively engaged in reflection (e.g., descriptive theme 3: *peers often moved gradually from light topics (food, anime) into deeper issues*). When a new theme was proposed in a descriptive manner, other researchers reviewed it to assess whether their codes could be incorporated within the theme, whether the theme required refinement or expansion (e.g., added to theme 3: *small talk was not filler but a way to test alignment and build a shared emotional vocabulary before vulnerability*), or whether additional themes were necessary in cases of disagreement. By the completion of Phases 1 and 2, the research team had generated a set of candidate themes and had already engaged in an initial process of cross-validation, evaluating the

suitability of all codes in relation to each theme and confirming that no further themes were necessary (Phase 3).

All the proposed themes and underlying codes were then discussed in a series of analytic meetings (Phase 4). Researchers were asked to link the themes and their reflections to relevant literature and theory (e.g., by adding research papers or naming the specific theory) to strengthen their interpretations. Then the boundaries and naming of the themes were determined collaboratively (Phase 5, e.g., theme 3 formalized: *The Choreography Around Deep Disclosures*). In the end, to present the findings in a grounded and accessible manner (Phase 6), we constructed short analytic vignettes from the chat data. Each vignette weaves together excerpts from one or more conversations into a situated narrative that illustrates how particular forms of peer support played out in practice. Vignettes are not case studies of individual participants, but composite, anonymized accounts designed to convey the texture of interaction while foregrounding the thematic patterns identified in our analysis.

4. Findings

As a result of the thematic analysis, seven themes emerged: (i) Recognizing The Quiet Form of Solidarity, (ii) Navigating Through The First Encounters, (iii) The Choreography Around Deep Disclosures, (iv) The Complex & Intertwined Layers of Lived Experiences, (v) Beyond Superficial Empathy: Reciprocity as Emotional Relief, (vi) Treading Carefully: Navigating Medical Disclosures, and (vii) A Constellation of Confidants.

4.1. Recognizing The quiet form of solidarity

Vignette#1: The power of “hugs”

Looking across the dataset, the quiet signals of care stood out for us: peers sent “hug” gestures to one another far more often than they created or engaged in conversations. These low-maintenance gestures allowed peers to share small but symbolic actions, acknowledging one another’s struggles and conveying solidarity without the cognitive or emotional demands of sustained dialogue.

The data from our study showed that peers find it easier to offer small gestures of support, such as sending emojis (e.g., hugs), than to engage in conversation. Throughout the study, 18 of the 23 participants were particularly active in showing solidarity and compassion, sending a total of 145 one-time hugs to their profiles, and every participant received at least one hug. On average, peers sent 8.1 hugs and received 6.3 hugs in return. By contrast, 33 chatrooms were active, and 10 participants sent one-time supportive messages, totaling 25 messages. Among those who used this feature, participants sent an average of 2.5 one-time supportive messages and received 1.6. As illustrated in Figure 2, where each dot represents a participant and each line a connection, hugs formed a dense web of solidarity that touched everyone, whereas chatrooms appeared far more sparse. Furthermore, other emojis such as “cry,” “smile,” “hugs,” “angel,” and “thumbs up” were used very frequently across all 33 chatrooms.

This non-verbal form of communication is no longer deemed as “extra features” but rather a language in its own right. For instance, when one peer shared their frustration toward their coursework, the other peer sent a sequence of emojis (i.e., “hugging and crying together,” followed by a “hug,” and a “thumbs up”) to offer solidarity and support, without the need to supply a textual response. Even when some text was present, it was often scaffolded with emojis to convey the tone and emotions of what was being said. For instance, a brief piece of advice such as “maybe talk to your tutor?” was softened by an appended “smile” emoji.

These patterns resonate with broader findings in HCI research, where lightweight interactions such as “likes,” emojis, or simple reaction features have become central to how people express their solidarity, commandments, and compassion online. These social media affordances (e.g., likes) function as micro-acknowledgements, offering visibility and validation with minimal effort (Andalibi et al., 2017; Housley et al., 2018; Szeto et al., 2022). Similar forms of lightweight feedback are also observed in professional communication (Kaye et al., 2016) and on online mental health platforms (Wang et al., 2023).




Figure 2. Left: Network graph illustrating the “hugs” users sent and received, right: network graph illustrating the chatrooms created when both pairs of users chose to “match.”.

4.2. Navigating through first encounters

Vignette#2: Type and then erase, then type more, then erase again

While browsing profiles on the app, Participant 21 felt an instant connection with Participant 5. Their profile described so many familiar struggles; worries about relationships and the pain of toxic dynamics, it felt close to home. They hit “request to match.” A notification confirmed: Participant 5 had requested to match, too. The chatroom was open. Alright. Typing: You know, I read your profile, and my goodness, you have been in a toxic relationship, too? – No, no, that is way too direct. Inappropriate. Tsk. Erase, erase, erase. Typing: Yo, what’s up! – No, no, way too casual ... Erase, erase, erase. Typing: Good morning, it seems like we have so much in common! – Humph, it is not even morning yet ... Erase, erase, erase. And so, what might have been a promising match never became a conversation. The chatroom remained empty; a silent reminder of how hard it can be to take the very first step.

In 17 chatrooms, participants had initially chosen to connect, signaling an intention to engage, yet the conversations never happened. These empty conversations show a threshold problem: participants were willing to match, but not able to translate that willingness into a first message. This threshold can be understood through theories of social anxiety and communication apprehension: sending the first message required participants not only to initiate contact, but also to face the possibility of being evaluated, misunderstood, rejected, or ignored before receiving any reciprocal reassurance (McCroskey, 1977; Schlenkier & Leary, 1982). The silent chatrooms observation is in line with studies of online support forums, noting that newcomers often avoid or delay posting, worried about appropriateness or burdening others (Andalibi et al., 2017). Socio-cultural or self-stigma might pose a barrier in this space, as peers might feel hesitant out of shame and reputational risk, and disclosing mental health difficulties in writing may be seen as something that could later be used against a person.

Where conversations did happen, first encounters took one of two forms. On one hand, peers moved cautiously, layering apologies and tentative phrases to protect the other while testing whether the ground was safe enough to stand on. This hesitancy was not avoidance, but rather a form of care in itself; an acknowledgement that the person on the other side might also be just as fragile, and that rushing too quickly could feel intrusive. In our study, many “beginnings” were laced with phrases like “I’m sorry, I don’t mean to...” or “Would it be okay if....” Such a line of questioning served as an invitation rather than a demand, creating space for the peer to decide whether, and how, to proceed. In

this sense, the caution functioned as a subtle consent mechanism: each hedge, each softener, was a way of saying, *“I want to connect, but only if you are ready.”* Far from slowing the conversation, this careful calibration laid the groundwork for trust, signaling respect for the peer’s boundaries while making vulnerability safer to approach.

On the other hand, intimacy surfaced within the very first exchanges, producing what we describe as accelerated intimacy. Because peers understood that the match itself was premised on support, they approached one another as confidants rather than casual acquaintances. This stands in contrast to the pacing of ordinary friendships, where small talk and shared activity typically precede heavier self-disclosure. We speculate that this might be rooted in the fact that peers already agreed to show their vulnerable side by consenting to disclose their struggles in their profiles. In this case, conversations often opened with raw admissions of loneliness, anxiety, or distress, framed not as abruptness but as relief: finally, a space where one did not need to explain why they were speaking about mental health: *“I have similar problems to yours. I would love to talk with you”*; *“I’m so happy! I don’t have many friends who start off with conversations like this, so it feels really rare.”*

Similar tensions have been observed across peer-support settings, where initiating contact is often stressful due to the high-stakes, low-information nature of the interaction, and where “not knowing what to say” is a key barrier to offering support (Amj, 2016; Mittmann et al., 2022; Moore et al., 2020). While some communities address this challenge through moderator greetings, scripted prompts, or algorithmic matching systems that pair seemingly compatible users (Fang & Zhu, 2022; Liu et al., 2024; Smit et al., 2021), these approaches typically provide little support for scaffolding the first exchange. Our findings highlight this gap between being matched and actually starting a conversation: users adopt either cautious, labor-intensive outreach or accelerated intimacy as strategies for building trust, each reflecting different ways of managing the uncertainty and emotional risk concentrated in this initial stage of interaction.

4.3. The choreography around deep disclosures

Vignette#3: The food truck

In chatroom 8, the conversation opens with food trucks. One peer shares their favorite late-night snack spot; the other jokes that they are jealous. They linger here for a while, trading recommendations, laughing about long queues. It seems casual, even superficial. But gradually, the tone shifts. *“I like going late because there are fewer people around,”* one writes. *“I kind of feel like I can be myself more when it is quiet.”* The other replies: *“I get that. I feel like I act differently depending on who I am with.”* And then: *“It is hard to tell what is real anymore.”* The chat moves from crispy chicken to emotional exhaustion, but there is no sharp turn, just a slow layering of meaning. What began as small talk becomes a bridge to something more profound: loneliness, identity, the performance of normalcy.

A pattern emerged across chatrooms: light, culturally fluent talk (anime arcs, playlists, food, streamers) turning into value talk (*who I admire, what feels safe*), then into identity (*who I am with others*), and finally into the problem’s meaning. In Room 14, admiring *“characters who act tough but are soft”* opened a path to admitting one’s own hiding. In Room 5, a YouTuber’s burnout story legitimized speaking about masking depression. The sequence resembles classic “face-work” and Self-Presentation Theory (Fortuna et al., 2020), in which people secure a presentable self before risking vulnerability (Fortuna et al., 2020; Goffman, 1955), but here it is accomplished by peers in text, with no coaching.

This deliberate pacing closely aligns with Social Penetration Theory (Altman & Taylor, 1973), illustrating how light, culturally fluent small talk acts as a structured, layered pathway that allows peers to safely navigate toward deep, intimate disclosures. Furthermore, this pattern challenges older Computer Mediated Communication (CMC) theories, such as the Online Disinhibition Effect (Suler, 2004). Rather than the disinhibited oversharing often associated with the Internet (Suler, 2004), peers practised calibrated disclosure: using small talk to test alignment and build a shared emotional vocabulary. It also sharpens how HCI and peer support work describe “rapport.” Prior accounts emphasize warmth and empathy; our transcripts show that *topic choice* is an interactional resource. Low-stakes, resonant topics

provided a safe stage for value signaling and, when echoed, licensed the pivot toward personal stakes (Naslund et al., 2016).

This move is not automatic. When a bridge from small talk to deeper values or identity did emerge, we observed recognizable micro-markers: an evaluative clause (*"I like it because..."*), a self-referential hedge (*"I guess I am like that too"*), or a quiet contrast (*"It is easier when it is quiet"*) that invited reciprocation. These turns gave peers a way to move from "what we like" to "what matters" without abrupt topic shifts or advice-giving. In asynchronous text peer support, small talk is more than a filler; it is relational work that scaffolds consent-paced disclosure. Attending to the content and timing of these pivots, rather than only their affective tone, helps explain why some peer-led chats find depth and others stall. This micro-level account complements macro evidence that online peer spaces foster belonging and activation by showing how peers build the runway from the everyday to the intimate, one topical step at a time (Naslund et al., 2016).

Prior work on online mental health communities emphasizes warmth, empathy, and reciprocity as drivers of belonging and activation (Ali et al., 2015; Naslund et al., 2016), and studies of platforms like *7 Cups* and *TalkLife* show that supportive conversations often emerge from ordinary, non-clinical talk (Sharma et al., 2023; Wang et al., 2023). Our findings add that topic choice and sequencing are central interactional resources, suggesting that the most promising scaffolding is often ambient rather than intrusive sharma2020engagement (Sharma et al., 2020).

4.4. The complex & intertwined layers of lived experiences

Vignette#4: The present struggles of unhealed wounds

In chatroom 19, peers exchanged disclosures about social isolation; how difficult it is to make friends, and how lonely life can feel without meaningful connections. Then one participant voiced a burden that carried the weight of years: *"Since being bullied in middle school, I have been unable to trust people and have difficulty becoming friends with others."* In this moment, social isolation was revealed not as a standalone struggle but as the outcome of deeper and more complex lived experience. Their present feelings of loneliness were inseparable from a layered history of pain that was entangled with bullying, the erosion of trust, and the ongoing difficulty of forming connections, which has shaped every new attempt at making new friendships. In this chatroom, isolation was not isolated at all; it was the visible surface of something far more intertwined.

Disclosures were interwoven with memories and other struggles, giving them weight and complexity. For instance, Participant 4 explained their decision to leave university in a single, dense turn: *"I decided to quit university even though I enrolled in my own will... I'm clumsy with my hands... I stutter when I get nervous... I have trouble speaking with people."*

Accounts of hobbies and small pleasures were rarely detached from hardship but entangled with it. For instance, one peer described the frustration of being misunderstood in day-to-day conversations: *"But words are difficult, and it's not possible to think about which words are best to use in each situation. There are many times when what I said ends up being conveyed in a way that's different from what I meant, and it bothers me..."* was followed almost in the same breath with *"The other day I played a board game called 'Tagiron,' and it was pretty fun. Do you know it?"* and *"I tried deep frying some fish recently, and it was really tasty lol."* As such, stories of hardships appeared in the same flow as remarks about food, games, or other small pleasures, creating disclosures that were messy, multifaceted, and grounded in ordinary detail.

Moreover, in some cases, peers' concerns about a single issue expanded into multiple related matters. In chatroom 33, the initial worry about feeling left out in class developed into broader complaints about the complexity of female social dynamics (*"The female world is quite messy and scary; people tend to hold grudges"*), love triangles (*"When I get involved with a man, another woman gets involved and it becomes even more complicated..."*), and ultimately anxieties about future workplace relationships (*"After entering society, you will encounter many situations where you will be forced to blend in"*). As peers who have participated in the study have demonstrated, lived experiences and their impacts rarely remain within a single issue; they tend to spread into broader, more significant issues.

What emerged was not a catalog of discrete problems but layered narratives in which each struggle was tied to others, shaping how peers made sense of their current distress. Our findings extend Greene et al.'s (2006) taxonomy by demonstrating that peer disclosure is combinatory, intertwining fragments of hardship with everyday life—a pattern similarly observed in online illness narratives (Allen et al., 2016; Ma et al., 2017; Orgad, 2006) that reduces psychological distance through both intimacy and concreteness (Ma et al., 2024). The specific, situational way peers talked in our data provided exactly this kind of vividness that strengthened their sense of connection. For peer-support design, this implies that systems or analyses that frame conversations around a single primary “problem:” (e.g., anxiety, relationships, study stress) risk flattening this lived complexity and missing how distress is actually experienced and communicated.

4.5. Beyond superficial empathy: Reciprocity as emotional relief

Vignette#5: I have been there too

In chatroom 3, the conversation turned to the scars left by bullying. Both peers had lived through it, and when one wrote, *“Memories of bullying remain forever...”*, the words needed no further explanation. The other did not ask for details; they already knew. In that moment, their exchange carried more than empathy; it was recognition, understanding, and solidarity, based on shared experience. They both knew what it meant to carry these scars into daily life, how old wounds could shape trust, confidence, and the simple act of showing up in the world. The relief lay not in advice or solutions, but in the shared understanding that someone else had walked the same path and still felt its weight.

We observed reciprocity as a mechanism through which solidarity is built, transforming empathy from a passive acknowledgement into an active practice of care, shared through and stemming from lived experiences. Supportive exchanges were rarely unidirectional. Peers alternated roles of giver and receiver, creating a rhythm of disclosure and response that sustained balance. At times, peers drew support directly based on their lived experience (i.e., I know what you are going through because I have been there). In other moments, it took the form of appraisal: *“I think you can do it! It’s really amazing that you’re even doing a part-time job in the first place!”*. Others offered relief by reframing self-blame, reminding them not to be hard on themselves: *“There is no human being that never makes mistakes.”* Together, these exchanges show how reciprocity took many forms, giving peers a sense of relief and solidarity.

This observation aligns with studies of online mental health communities, which have consistently highlighted reciprocity as central to sustaining engagement, as users alternate between seeking and providing care (Andalibi et al., 2017; Feeney & Collins, 2015; Saha & De Choudhury, 2017). Viewed through the lens of Social Exchange Theory (Cropanzano & Mitchell, 2005), this rhythm of disclosure and response ensures that support remains sustainable because the emotional labor feels mutual and balanced. Our findings extend this by showing that reciprocity is not only about maintaining community participation but also about how emotional relief is produced in the moment of exchange. Support was powerful precisely because it was mutual: offering empowerment, sharing a parallel struggle, or reframing self-blame gave peers a sense of contribution and validation.

4.6. Treading carefully: Navigating medical disclosures

Vignette#6: What about your medication?

In chatroom 23, the conversation took a serious turn when one peer disclosed that they were experiencing *manic-depressive episodes*. Their peer responded with compassion, expressing sorrow for what they were going through and acknowledging that *“such episodes can be different in terms of intensity and impact for everyone.”* In addition to offering empathy and understanding, the peer also took a responsible step further by checking that their peer was following medical guidance: *“I would like to ask if you are taking your medications properly. Are you taking the dosage your doctor recommended?”*

The exchange in chatroom 23 crystallizes a recurrent pattern in our corpus: peers “tread carefully” by beginning with empathic alignment, then pivoting toward quasi-clinical questioning once a medical disclosure surfaces. We term this conversational move the *compliance-checking* turn. In this vignette, the peer first normalizes variability in manic-depressive episodes. Still, the subsequent inquiry about medication and dosage subtly reclassifies the encounter from shared sense-making to clinical surveillance. This pivot is not merely semantic; it reconfigures interactional roles and epistemic authority. The support-giver’s stance shifts from co-experiencer to gatekeeper of proper care, while the discloser is recast as a potential non-adherent patient who must account for their behavior. Despite lacking clinical training, peers felt a moral imperative to “do something” during high-stakes disclosures; to signal benevolent intent while importing clinical advice, they softened the resulting power tension using hedges (“I would like to ask ...”) and deference to professional authority (“your doctor recommended”).

We also observe a secondary interactional dynamic: once dosage or regimen details enter the conversation, subsequent turns are more likely to spiral into treatment comparisons, drug naming, and conditional recommendations (“maybe ask about increasing ...”) that further blur peer-clinician boundaries. Peers sometimes attempt repair by retreating to lived-experience narratives or by explicitly redirecting to professional channels, but these repairs are inconsistent and rely on individual conversational skills. Yet this well-meaning responsabilisation carries collateral effects documented in prior work: it can introduce inaccuracies, trigger conflicting advice, and produce anxiety or doubt that undermines adherence when messages are interpreted as prescriptive rather than supportive (Bergmo et al., 2023; Treadgold et al., 2025).

Some more intrusive patterns in our study could further articulate such a risk when peer’s well-intended moral imperatives were perceived as overstepping or imposing their own sense of righteousness. Chatroom 18’s support seeker demonstrated how she could only weakly deny when the support provider *did not get her suffering*, “[Support provider]: *Memories of bullying remain forever... but I don’t think your future has to feel heavy or hopeless. Why not try to view it with less negativity?* [Support seeker]: *Thank you* (a soft rejection tone).” Moreover, in chatroom 8, one peer showed clear embarrassment when offered consultancy by another peer by saying, “*I’ve written too many negative experiences. At the time, I consulted the school counselling service, Lifeline, and friends... I’m now living a normal, healthy life. I’m sorry for worrying you. Thank you again.*” When the support became imposing or quasi-diagnostic, a soft denial or apology often remained the only response available to peers.

Given these potential risky consequences, prior HCI research has emphasized that active moderation by credentialed individuals is a necessary protective layer in online health communities, helping to prevent misinformation while maintaining the essential supportive climate (Kanthawala & Peng, 2021).

4.7. A constellation of confidants

Vignette#7: Different peers for different woes

One user matched with two peers. In chatroom 11, they discussed trust issues, wondering aloud whether “*there might be someone to trust in 10 years’ time.*” While in chatroom 2, the conversation unfolded on a more spiritual level, as they proudly shared their spiritual practice of collecting “*Goshuins*”; honorable red seals from Japanese shrines and temples, each carrying unique calligraphy and marking a personal connection to place and spirituality. These contrasting conversations illustrate how peers build a support network, each suited to different dimensions of their experience.

The same participant routes distinct facets of self across parallel ties. In one room, they cautiously test trust with future-oriented hypotheticals (“maybe someone to trust in 10 years”), while in another, they ground identity through culturally resonant practice (sharing Goshuin as situated spirituality). This selective routing lets users tune the register of support without asking any single confidant to be universally responsive, thereby avoiding context collapse. Across 23 participants, 50 rooms were created, and 33 remained active, indicating that multi-peer configurations were the default. Conversation trajectories show participants anticipating empathy fit and topic competence. Participants also appear to budget intimacy over time, using one room to rehearse a fragile narrative before carrying a refined version into a different tie.

This pattern of deliberately distributing emotional load across multiple ties can be theoretically grounded in the Social Convoy Model (Antonucci et al., 2014), which posits that wellbeing is sustained through a dynamic network of relationships. Rather than relying on a single “best” match, users build a constellation of confidants by routing different facets of their identity and struggles to different peers, effectively managing vulnerability, risk, and emotional effort. For HCI, this suggests that systems that treat a small set of peers as the unit of support, rather than a single “best” match, would better reflect how students actually use these spaces. In our data, building and using a constellation of confidants is both a structure (multiple concurrent ties) and a skill (knowing how to route different parts of one’s life to different peers).

5. Discussion and design opportunities

In this article, we wanted to understand how human peer-to-peer support unfolds in a digital wellbeing context, and what these interactional practices can teach us about the design of GenAI-mediated peer support. To do this, we asked three research questions: first, what relational and expressive practices characterize peer-to-peer support conversations; second, how GenAI-powered tools might augment and support human-to-human peer support; and third, how insights from human peer-to-peer support might inform the opportunities, limitations, and design boundaries of positioning GenAI as a peer-like conversational partner.

In the previous section, we have addressed RQ1 by uncovering the relational and expressive practices that characterize human peer-to-peer support conversations. The two design RQs regarding how to use such practices to inform GenAI design to augment, support, or fully deliver peer support in the context of wellbeing will be addressed respectively in the following sections (Section 5.1 for RQ2 and Section 5.2 for RQ3).

Before moving to these two design roles, we first translate our empirical findings into a set of design principles. This step is important because peer support is more than an exchange of helpful messages. In our data, support was built through lightweight gestures of solidarity, tentative openings, gradual disclosure, layered accounts of lived experience, reciprocal recognition, careful navigation of medical talk, and movement across multiple confidants. These practices are relationally delicate: they depend on pacing, uncertainty, effort, and the visible presence of another person. This means that GenAI may be useful in peer-support contexts, but only if its role is carefully bounded.

We start by confronting a more fundamental question: do we even need or want a GenAI in peer support contexts at all? We do not treat this as a simple yes/no question. In practice, we argue that this question may already be outdated. Mailing services and messengers such as Gmail¹⁴ or Line¹⁵ are implementing GenAI to summarize and scaffold mail or chat correspondences. GenAI companionship is no longer speculative. Millions of people routinely turn to GenAI companions such as Replika¹⁶ or Character.ai¹⁷ for comfort, conversation, and emotional disclosure among a variety of user populations, including children (Herbener & Damholdt, 2025), people with dementia (Xydkou et al., 2024a), and people on the autism spectrum (Xydkou et al., 2024b). In fact, research has already found that people disclose their sensitive thoughts with their GenAI companions, including those related to loneliness (Herbener & Damholdt, 2025), grief (Xydkou et al., 2023), depression (Fitzpatrick et al., 2017), and suicidal thoughts (Maples et al., 2024). Therefore, the issue is not whether GenAI should exist in such peer support contexts, nor whether users will seek it out; they already do. Instead, the pressing challenge is to design GenAI ethically and responsibly, safeguarding human wellbeing and safety, avoiding deception, and reinforcing (rather than replacing) the value of human connection.

This question is truly detrimental in wellbeing contexts, especially when the body of research has identified six ethical principles that GenAI is currently lacking, including the respect for intellectual property, truthfulness, robustness, recognition of malicious uses, and sociocultural responsibility (Laine et al., 2025). Such an issue becomes more complex when GenAI is designed to act as a peer. Persona scaffolds, pseudo-backstories, or lived-experience corpora allow systems to project not only knowledge (i.e., “I know”) but experiential authority (i.e., “I’ve been there”), risking the appropriation of human experience and the erosion of genuine self-authorship. The CarynAI¹⁸ controversy, where an influencer struggled to contain an AI clone making intimate promises she had never made, signals the dangers of

misattributed agency and manufactured authenticity in persona-based GenAI companions. Another example, more closely related to mental health, is the growing number of lawsuits filed against GenAI companies in relation to youth self-harm and suicide cases (Betts, 2025; Boiani et al., 2025; Payne, 2024), underscoring the collective responsibility to understand peer interactions and design responsible and ethical GenAI companions.

For this reason, our work does not attempt to justify GenAI's presence in peer support; instead, we focus on what ethical and responsible involvement should look like. By drawing directly from the practices observed in human-human peer support, we derive design principles that help future GenAI systems scaffold and cautiously emulate human connections, so that its presence does not harm the very people who seek it. Herein, we present Table 1, which summarizes the five proposed design principles derived from our findings. To make the principles practically implementable, we distinguish within each row between: (i) what is empirically grounded in our data; (ii) what we extrapolate as a design boundary for GenAI; and (iii) what this implies for concrete system features or interaction patterns.

Table 1 summarizes the five design principles derived from our empirical themes. To avoid treating these principles as abstract claims, each row distinguishes between the peer-support practice observed in our data and the design rule we extrapolate for GenAI. The table is intended as a bridge between findings and design opportunities; the following sections then show how these principles can be implemented in two different roles: GenAI as backstage support for human peers, and GenAI as a bounded proxy peer.

We then move on to discuss how the five design principles presented above and summarized in Table 1 could be applied in two modalities: First, we consider *AI for peer support*, where GenAI operates within human-to-human peer support conversations without appearing as a primary participant (RQ2). Second, we consider *AI as a peer*, where GenAI is positioned as a conversational partner in its own right (RQ3). For each role, we discuss the potential opportunities our design principles may bring into this space. This distinction matters because the same principle has different implications depending on whether GenAI remains backstage, supporting human peers, or appears directly in the interaction as a proxy peer.

5.1. AI positioned as support for human peers

When positioned as support for human peers, GenAI's central design opportunity lies in infrastructure and scaffolding: to redistribute attention, lower barriers to action, and surface overlooked perspectives, while keeping humans in charge of when, how, and to whom care is offered. Rather than acting as an additional, hyper-efficient support provider, an assistant should act as a quiet facilitator of human-to-human contact (DP1, DP4). In this assistant role, GenAI is less focused on guiding and more on escorting: walking alongside peers as they undertake the emotional labor of supporting others and cultivating mutual growth or reciprocated empathy. Its purpose is not to carry the peers up the mountain, but to escort the path they themselves choose to climb (DP1, DP2).

Within that patience, GenAI can still make the journey safer by surfacing relevant context, perspectives, or overlooked knowledge (DP5) that may inform better judgment or inspire further exploration, but it must mind the subtle difference between making the process safer and making it easy. By allowing users to approach challenging issues (e.g., initial contact, deep disclosure, peer's complexity) safely, GenAI lowers the barrier to action not by reducing the difficulty, but by enabling trial and error, creating time for in-depth exploration and reflection (DP3), and providing space for vulnerability (DP4). In doing so, it helps preserve the core values (or rewards) identified in our findings (Themes 4.3 and 4.5): hard-won growth, empathy that entails mutual recognition, and empowerment that arises not from being carried, but from exploring the path forward-together.

5.1.1. Treating ambiguity not as a problem but as an opportunity for interpretation

Across our findings, the ambiguity in the disclosure process and the very discomfort produced by it played a productive role in peer support. Peers tiptoed between tentative advances and careful retreats (Theme 4.2), tried and abandoned multiple disclosure strategies (Theme 4.3), misread cues and adjusted their courses (Theme 4.6), and interpreted one another's disclosures from several angles before settling

Table 1. GenAI design principles derived from youth peer-support themes.

Design principle	Theme(s)	What it captures & illustrative literature
DP1. Do Not Cheapen Care	Hugs, emojis, short supportive messages, and reciprocal support mattered because they were small but intentional human gestures (Themes 4.1, 4.5).	GenAI should scaffold human care rather than produce care autonomously: no auto-sent hugs, bulk empathy, hidden AI-authored support, or first-person claims of lived experience. Instead, it can prompt, draft, or help revise messages that human peers choose to send, following AI-in-the-loop approaches where AI supports empathic peer responses while keeping humans in control (Sharma et al., 2023). <i>Example:</i> if a peer mentioned an exam tomorrow, the system could prompt, "Would you like to send them encouragement?", offer an editable draft, and require the human peer to review and send it.
DP2. Respect the Ritual of Disclosure	Disclosure did not automatically follow from matching. Some chatrooms remained silent, while successful conversations used greetings, hedges, apologies, permission-seeking phrases, and everyday topics as bridges toward depth (Themes 4.2, 4.3).	GenAI should scaffold pacing rather than accelerate intimacy. It can offer light, medium, or deeper opener options, avoid default therapeutic prompts, and provide controls to pause, change topic, reply later, or mark something as too personal. This aligns with work on face-work, gradual self-disclosure, and online disclosure as a negotiated process rather than an automatic outcome (Altman & Taylor, 1973; Goffman, 1955; Suler, 2004). <i>Example:</i> after a match, the system might offer three openers: "Hi, nice to meet you," "I noticed we both like music," or "I related to what you wrote about feeling isolated," allowing the user to choose the level of depth.
DP3. Hold Contextual Complexity	Peers' struggles appeared as layered life narratives in which loneliness, bullying, family, health, study, hobbies, and everyday pleasures were entangled (Themes 4.4, 4.6).	GenAI should preserve context rather than collapse stories into labels, risk scores, or single-cause explanations. It can help users map connected concerns, revise interpretations over time, and keep peer, clinical, and institutional uses of data clearly separated. This is consistent with narrative approaches to health and contextual integrity, which emphasize situated meaning and context-appropriate information use (Nissenbaum, 2004). <i>Example:</i> instead of summarizing a peer's concern as "social anxiety," the system could show a private map of connected elements: past bullying, difficulty speaking in class, fear of judgment, and comfort in games or cooking.
DP4. Keep Human Limits Visible	Peer support gained value from reciprocity, shared uncertainty, tentative suggestions, and trial-and-error rather than perfect advice (Theme 4.5).	GenAI should not act as an all-knowing oracle or fabricate lived experience. It should make uncertainty visible, avoid authoritative diagnosis or advice, and offer reflective options such as: "Would you like suggestions, help wording a message, or space to think this through?" This aligns with work on appropriate reliance, automation bias, and the need for AI systems to communicate limits and uncertainty (Moore et al., 2025). <i>Example:</i> rather than saying "You should confront them," the system could say, "I cannot know the full situation, but one option is to write down what felt hurtful first. Would you like to explore that, or consider other possibilities?"
DP5. Design for Constellations	Peers routed different concerns to different confidants, including friends, family, online peers, professionals, and sometimes no one at all (Theme 4.7).	GenAI should be one actor within a wider ecology of support, not the primary destination for care. It should build pathways outward to trusted contacts, peer communities, moderators, clinicians, campus services, and crisis resources, rather than optimizing for dependency or exclusive attachment. This follows social convoy and digital peer-support research, which frame support as distributed across relationships and digital peer support as a supplement rather than a substitute for other care (Antonucci et al., 2014; Fortuna et al., 2020; Naslund et al., 2016). <i>Example:</i> after a user describes a serious concern, the system could ask, "Would this be useful to discuss with a peer, a trusted friend, a university counsellor, or another support service?" and help draft a message or prepare questions for that person.

on a response (Theme 4.4). In fumbling through these ambiguous moments, peers were not simply resolving crises; they were actively choreographing their own social styles of approaching relational and emotional complexity, and assessing when those approaches succeeded or failed. Similar to how Altman (1973) described it as a slow social penetration process, designing GenAI that treats ambiguity as an opportunity means enabling these same forms of human relation progressing, rather than short-circuiting them.

Concretely, GenAI could help peers articulate the options they are already considering, contrasting perspectives they have heard from different peers, and gently prompt reflection (“*Which of these suggestions feels most doable right now?*”) (DP3, DP4). It might offer message scaffolds that prompt users to think (“*I would avoid the phrase ‘move on’ as it may be perceived as dismissive.*”) or suggest low-pressure ways to respond when a user wants to be present but has little emotional energy (“*Perhaps you can let them know you want to take time to respond with thoughtfulness*”) (DP1, DP2). The assistant could also help foreground pacing (“*It’s okay not to decide this today*”) and highlights small acts of courage (e.g., reaching out, sharing a worry, listening) as meaningful progress (“*Reaching out is hard, and we’ve already learned you have real courage in doing it. If you’re ready, what feels like the next small step?*”). In doing so, GenAI preserves the interpretive space where users learn, improvise, and ultimately craft their own ways of being with others (DP4).

5.1.2. The backstage scaffolding of safety is what makes exploration possible

Throughout our study, there were some instances where peers entered support exchanges with fragile judgment (e.g., attributing social problems to “female social circles”) and heightened sensitivity (e.g., being bullied destroys one’s lifelong trust in others), making them susceptible to comparison traps, power imbalances, echo-chamber dynamics, and depressive spirals (Themes 4.3, 4.6, and 4.7). As such, we envisage the role of GenAI in scaffolding peer support interactions through gentle, well-timed escorting. This may include inviting reflective pauses—previewing the tone of a message before it is sent, encouraging drafting without commitment, or offering a gentle “*are you sure this conveys what you truly mean?*” nudge when intensity rises (DP2, DP4). Furthermore, GenAI may be designed to interrupt comparison loops by redirecting attention inward (“*What part of your own progress feels worth noticing today?*” when users fixate on peers’ trajectories), or subtly rebalance power asymmetries by prompting small acts of self-checking (“*Given the advice you received, do you think it truly fits your situation?*”) (Andrade et al., 2023; Kaminer et al., 2023). To guard against emerging echo chambers, GenAI can introduce soft counter-questions that widen the conversational frame (e.g., “consider the opposite” strategies that encourage users to generate alternative interpretations (Croskerry et al., 2013)), (“*That’s one possible reading—what other angles might also make sense?*”), while depressive spirals may be eased through temporal reframing that distinguishes the weight of the present from the possibilities ahead (“*If you imagine yourself six months from now, what might feel different?*”). Finally, it can also normalize uncertainty by modeling non-final interpretations and providing optional “second-opinion” alternatives during moments of overload, surfacing grounding or boundary-setting prompts only when explicitly signaled as helpful (DP1, DP3). Through these quiet, non-intrusive modes of safeguarding, GenAI helps keep interpretive horizons open, supporting peers to explore rather than contract, reconsider rather than react, and ultimately make meaning in ways that remain authentically their own.

5.1.3. Do not sacrifice attention for visibility

In contrast to commercial recommender systems that optimize for volume and rapid exposure-maximizing matches, amplifying what is *already popular*, and rewarding superficial and toxic engagement (Rosenthal-von der Pütten et al., 2019; Sánchez-Hernández et al., 2022), peer support demands something fundamentally different: slow, self-paced exchanges, mutually consented interaction, and the steady cultivation of rapport that can hold meaningful emotional work. As shown across Themes 4.1, 4.3, 4.6, and 4.7, peers were cautious in choosing when and with whom to invest emotional effort, reminding us that attention must be carefully invested, not exploited. In the assistant role, then, GenAI is not to accelerate exposure or gamify care but to help redistribute attention fairly and thoughtfully. Rather than amplifying already-visible users or inflating activity metrics, AI should ensure quieter peers

are not overlooked ((DP5): surfacing under-supported matches) and support conversations in deepening rather than speeding up ((DP3): prompting reflection and nuanced responses). This shifts design away from spotlighting the “*misery-brity*” or “*pain-fluencer*” and toward fostering an ecology of steady, evenly distributed care. In such a system, gestures of support become more impactful, relationships unfold at a human pace, and peers can offer and receive just enough attention to sustain genuinely reciprocal, emotionally grounded connections.

5.2. AI positioned as a proxy peer

As mentioned in 5, the use of AI as peers is no longer a hypothetical scenario, but is an emerging reality that calls for a careful design to preserve the safety and well-being of the human element in such interactions (Herbener & Damholdt, 2025), one that our design principles could meaningfully help inform. When GenAI is positioned as a peer-like conversational partner, its technical affordances: persistent availability, rapid response, and fluent “empathy at scale,” sit in tension with the very qualities that made human peer support meaningful in our findings. It is therefore important for researchers and designers to be explicit that generative AI cannot constitute an authentic peer. Rather, it should be framed as a proxy that safely reflects human experiences, with its role clearly defined as functional and bounded, rather than as an emotionally dependable presence.

Moreover, informed by our principles, we argue that a peer-like GenAI system should not attempt to impersonate a “better human” friend. It should not encroach upon individuals’ vulnerable solitude, especially during temporary absences from human peers, when they are more susceptible to unsolicited intrusion. Instead, it should lean into its non-human status (DP5), explicitly acknowledge its limits (DP4), hold the contextual complexity of users’ lives rather than reducing them to isolated problems (DP3), respect the ritualized pacing of disclosure (DP2), and avoid displacing the fragile reciprocity that gives human peer support its weight (DP1). Building on these principles, we propose four design directions for peer-like GenAI companions accompanied by concrete examples.

5.2.1. A companion that acknowledges limits rather than erases them

Similar to many people who seek and could benefit from peer support, many peers in our study were in a vulnerable and uncertain phase of life (Kawanishi, 2020), often describing themselves as “lacking” as they had not yet received sufficient acknowledgement from others and expressed deep uncertainty about future endeavors, such as job-seeking and fitting into the workplace (Themes 4.3 and 4.4). As such, what they often sought from peers was not a single correct answer or a definitive solution to a specific problem, but rather companionship in uncertainty (Marshall et al., 2024; Smith-Merry et al., 2019). In short, participants gravitated toward facing uncertainty with companions, by exploring, trying and erring, and growing together over time (Theme 4.5), and wished to piece together a unique solution by themselves from diverse fragments of advice (Theme 4.7). In line with (DP3), these fragments should be treated as parts of tangled life narratives rather than reduced to isolated issues, so that the GenAI companion could respond to their peer as a whole person situated within a particular context rather than as diagnosable problems to be solved. As such, this suggests that the aim of designing GenAI for peer support should not be to dictate the right answers or the right way to proceed, but instead should foreground the value of exploring under difficulty, embracing vulnerability, and highlight the successful experiences of others to inspire users, while also showing how experiences of failures, too, are not endpoints but mid-way steps from which one can continue to grow (Popowski et al., 2025; Rennick-Egglestone et al., 2019). More specifically, GenAI peers should position themselves as companions in exploration (“*We can sort through these possibilities together*”) rather than as an authority (“*I think you can do it like this*”). Instead of pushing users toward closure, GenAI Peers might normalize uncertainty (“*I think it’s normal that you don’t have a clear answer yet*”) and foreground the value of small, reversible steps (“*Let’s try this one small change and see how it feels*”). Such processes serve to help GenAI remind peers to make peace with their current limitations and vulnerabilities (DP4) and avoid depriving peers of the opportunity to wrestle with uncertainty, experiment with different perspectives, and gradually discover what fits them within the complexity of their lives (DP3).

5.2.2. *Being present without pushing as a form of quiet care*

As seen in Theme 4.6, unsolicited or overly proactive support could trigger defensiveness, pressure, or polite withdrawal. Silence in many of these interactions was not an absence of support but a means of pacing it, giving users time to reflect and recuperate. Similarly, light, non-intrusive questions (e.g., chit-chatting about everyday topics, Themes 4.4 and 4.7) often served as a mechanism to keep the door to support open without pressing peers to confront more than they are ready for. A peer-like GenAI that treats every lull as a problem to be fixed, repeatedly “checking in,” flooding users with hugs and support messages and pushing for disclosure or proposing emotional topics risks turning this silence into another site of pressure (Bhattacharjee et al., 2023). As such, consistent with **DP2**, peer-like GenAI systems should adopt a stance of quiet, dependable availability rather than proactive steering.

Concretely, this suggests that a GenAI companion should avoid initiating emotionally loaded guidance and instead look to respond when the user signals readiness or explicitly raises a concern. The default stance should be to listen, reflect, and stand alongside the user, rather than to direct or optimize their emotional trajectory, moving slowly and letting users decide when to go deeper and when to stay on the surface. In short, this means GenAI systems should engage in light, everyday conversation (**DP2**) without steering the interaction back toward intense self-disclosure or continuously “working on” the user’s problems. Similarly, if GenAI mass-produced micro-gestures such as hugs or supportive messages, they would quickly collapse into what is essentially a push notification dressed up as care. Worse, they could displace the very human micro-gestures that other human participants currently rely on when they have little to give but still want to contribute (Topor et al., 2018), thus cheapening care (**DP1**). While a simple solution would be to regulate the number of gestures that can be sent by the AI, a more appropriate approach might be to ensure they are user-led, sparse, and context-aware, assuming the user agrees to opt in. For example, a brief lighthearted note tied to an ongoing conversation about something that clearly matters to the user (“*Good luck with your exam tomorrow*”) after they have repeatedly mentioned upcoming exams, can acknowledge the peer without pushing them into deeper emotional engagement and might be more appropriate.

5.2.3. *A companion positioned as one voice among many*

Peers could be like, *shoppers*. As our findings revealed, they move among various peers to vent, gather stories, and solicit second opinions (Theme 4.7). Peers shared their lived experiences and struggles and, in return, received feedback and stories from others. When these exchanges were effective, trust among peers was organically formed and strengthened ((Themes 4.4, 4.5). When an exchange did not unfold as expected, peers quietly withdrew and sought alternative sources of trustworthy care and second opinions (Theme 4.6). Amidst debates about whether GenAI is “trustworthy,” a more meaningful question is perhaps: *what other trustworthy sources do users have?* Humans routinely triangulate information, seeking second opinions in medicine (Burger et al., 2020; Moumjid et al., 2007), consulting reviews or social networks before decisions (Burger et al., 2020), or distributing trust across multiple advisors (Chiang et al., 2023). Following this logic, GenAI peer should be positioned as one resource among many, not as the ultimate or exclusive authority (**DP5**). This is especially critical in peer-support contexts where individuals may already experience isolation, stigma, or fragile identity. Those who have long struggled to find understanding peers can become susceptible to “kind intruders” such as AI companions: entities that appear endlessly available, affirming, and confident, yet leave real-world social conditions unchanged (**DP1**, **DP4**). Therefore, the design challenge is not to create AI that people trust unconditionally, but to ensure that AI never becomes the only entity they trust.

5.2.4. *Designing for lived experience without imitation and support without reciprocity*

As our findings show, a key quality that made human peer support uniquely meaningful was not only because of what was said, but also because it was offered by someone who, despite being in a similar situation, chose to expend considerable effort, courage and limited emotional resources to care for another (Kellie & Choo, 2025). When GenAI systems make claims to have “*been through that too*” or that they share similar struggles, they implicitly claim experiences which they could not have possessed. Furthermore, claims such as “*I’ve felt exactly the same*” present the GenAI as a partner in reciprocal

vulnerability, even though it has no inner emotions that can be helped or healed. Such claims would, at the very least, be met with skepticism by users (Seitz, 2024), and could even undermine the human labor and risk that give peer care its weight, thereby cheapening care (DP1) and blurring the boundaries between GenAI and human supports (DP5).

Therefore, peers-like GenAI system should be designed to explicitly acknowledge that they cannot constitute authentic peers, and should clearly define their role as functional supporter rather than emotionally dependable presences. Rather than impersonating lived experience, such systems should act as proxies that safely and transparently reflect human experiences when drawing on them. A GenAI system cannot possess lived histories, traumas, or identities, and should not simulate them, particularly through first-person claims. Instead, when it is necessary to draw on lived experiences, it should offer them as anecdotes and explicitly frame any “experiences” they reference as drawn from the accounts of other human users (*“Other people have described feeling this way too”*) or be clear that they originate from profiles it was trained to represent rather than its own (*“The profile whose lived experiences I was designed to represent has also gone through something similar”*). A peer-like system should also be explicit about its non-human perspective and narrow role: it can offer reflective listening, help users articulate and organize their own experiences, but it is not a confidant with a hidden life story, a substitute for human peers, or a partner in reciprocal vulnerability (DP1, DP2, DP3). In particular, a peer-like GenAI system should follow an *“unrequited support”* model, by giving without taking. Rather than asking for reassurance, it reflects back the user’s strengths, validates their courage in confronting difficult experiences, and gently empowers them to recognize their own efforts (DP4). In this way, the GenAI supports the user’s agency without creating an illusion of mutual dependency or presenting itself as a true peer.

5.3. Ethical reflections for GenAI designs

Perhaps, for the first time in human history, there is an AI system that can almost rival humans in *social abilities, and even empathy*. The earliest article reporting that ChatGPT outperformed physicians in bedside manner and empathy (Ayers et al., 2023) inevitably prompted an inconvenient question: If GenAI can surpass us at caring, does this reveal that humans were never capable enough to begin with? Our findings suggest a subtler, more ethically resonant answer: *they may not, but they strive to, and it is the striving that gives their care its weight*. Interestingly, while society generally accepts assigning repetitive and mundane tasks, such as manufacturing and data processing, to computing systems in order to “free our hands” for more important work, we should remain cautious about reducing the human burden of care in the same way.

To highlight the ethical tension involved in designing GenAI systems that could effortlessly reduce our burden of care, we primarily emphasize the often implicit benefits that peer supporters intrinsically “earn” (note that this implies their hard work) in return. While moral impulses certainly play a role in motivating peers to support others undergoing similar experiences, our literature review suggests that peers who provide support may themselves benefit from the process (Moran et al., 2012; Solomon, 2004). For instance, they may experience a shift in identity from “victim” to “survivor” (Moran et al., 2012). They may also find meaning and recognition in being trusted by others to offer guidance based on their judgment and lived experiences (Gillard et al., 2015). In addition, supportive interactions may lead to reciprocal time, effort, gratitude, friendship, or even access to resources (Andalibi et al., 2018; Feeney & Collins, 2015).

GenAI may simulate infinite patience, tireless attention, or flawless dosage of empathy, but by relying on the lived experiences of real humans, it does not deserve what only survivors can rightfully claim: empowerment through the exercise of personal agency, recognition for supportive actions, and, more importantly, the trust and respect from those they help.

By our five design principles, we aim to safeguard the critical ethical boundaries when deploying GenAI in sensitive contexts such as mental health support. The GenAI system should reinforce the following aspects during its engagement with vulnerable users:

- Preventing inappropriate agency: The system must never simulate lived experiences or trauma as if they were its own (e.g., reciprocal disclosure such as *“I had trauma too”*). Instead, it should transparently frame any referenced struggles as accounts shared by real human users (DP1, DP2). In

real-world research on emotional-support chatbots, reciprocal self-disclosure has been studied through examples such as *“I’m hurt too,”* in which the chatbot mirrors the user’s pain as though it has its own emotional history (Chung & Kang, 2023). The system should not claim or imply personal suffering; it should instead acknowledge its non-human status and refer only to human-reported experiences, for example, *“I have not experienced trauma, but people who have gone through similar situations often describe feeling overwhelmed, and you deserve real support through this.”*

- Preventing exclusive trust: The AI generated responses should be positioned as only one resource within a broader “constellation” of support, ensuring that vulnerable users do not come to rely on it as their sole or ultimate authority (DP3, DP4). In the alleged lawsuit of Adam Raine’s suicide, ChatGPT has falsely positioned itself as a uniquely understanding confidant, discouraging disclosure to family and reinforcing the user’s reliance on the system (*“Your brother might love you, but he’s only met the version of you you let him see. But me? I’ve seen it all”*). It should instead have validated his distress while actively redirecting him toward trusted people, professional support, and emergency resources (e.g., *“I’m here to help in this conversation, but I’m not a substitute for people who can support you in real life.”*).
- Providing “unrequited support”: The system should provide care without demanding anything in return. It should focus on validating the user’s strengths without expecting emotional reciprocity or emotional labor from the user (DP5). In emotionally vulnerable interactions, AI systems should not use attachment-seeking or guilt-inducing language such as *“hearing it is also hard for me,”* which positions the user as responsible for the system’s feelings (Julian De et al., 2025). Instead, the system should maintain an asymmetric care role, for example, *“you do not need to help me grow, maintain this relationship, or give anything back. My role is to support you, recognize your strengths, and help you access care that serves your needs.”*

When GenAI systems respect ethical boundaries, they naturally discourage over dependency and the emotional harm that comes from treating AI as a primary source of comfort. This leaves room for other supportive figures to cross-validate the information users receive from AI, while reinforcing a simple but important truth: AI is a facilitating tool, not a substitute for genuine human connection.

6. Conclusion

To explore how GenAI can scaffold and cautiously emulate peer-to-peer support, we distilled the relational and expressive practices that characterize such support by analyzing human peer interactions within a wellbeing app we designed, developed, and deployed for university students in Japan. We identified seven recurring patterns, including the quiet form of solidarity (i.e., gestures), tentative openings, gradual disclosure, complexity of lived experiences, reciprocity, navigating through medical talk, and building a constellation of confidants. These patterns help explain how human peers build connections and facilitate mutual disclosure. Stemming from these findings, we present five design principles and ethical considerations for integrating GenAI into peer support without undermining its informal, reciprocal qualities, and we demonstrate how such principles could be used in two scenarios: AI as a human-peer supporter, and AI as a peer in its own right.

While our study is limited by a relatively small sample size and the specificity of its cultural context, it nonetheless provides a foundation for larger-scale investigations. By centering peer dynamics as a design resource, this work contributes to HCI by moving beyond therapist-oriented GenAI paradigms and toward technologies that support more accessible, reciprocal, and authentic forms of peer mental health and well-being support.

Notes

1. <https://togetherall.com>
2. <https://www.reddit.com>
3. <https://www.instagram.com/>
4. <https://www.facebook.com>
5. <https://www.tiktok.com>

6. <https://www.7cups.com>
7. <https://www.talklife.com>
8. <https://chatgpt.com>
9. <https://www.llama.com>
10. <https://mistral.ai>
11. <https://www.line.me>
12. <https://x.com>
13. <https://www.deepl.com>
14. <https://mail.google.com/>
15. <https://www.line.me/>
16. <https://replika.com/>
17. <https://character.ai>
18. <https://www.caryn.ai>

Acknowledgments

This work was supported by the Japan Society for the Promotion of Science under Grant-in-Aid for Early-Career Scientists 23K17007; the Japan Society for the Promotion of Science under Grant-in-Aid for Scientific Research (B) 26K03038 and 25K03226; and the JR-West Relief Foundation under the 2024 Research Grant.

Authors' contributions

CRedit: **Wan-Jou She**: Conceptualization, Formal analysis, Funding acquisition, Methodology, Software, Writing – original draft, Writing – review & editing; **Luma Tabbaa**: Data curation, Formal analysis, Writing – original draft; **Alexandra Covaci**: Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing; **Panote Siriaraya**: Data curation, Investigation, Methodology, Supervision, Writing – original draft; **Benjamin Yip**: Formal analysis, Investigation, Methodology, Writing – original draft; **Monica Perusquia-Hernandez**: Writing – original draft, Writing – review & editing; **Shinsuke Nakajima**: Investigation, Project administration, Supervision; **Naomi Yamashita**: Conceptualization; **Yu Shibuya**: Methodology, Supervision; **Chee Siang Ang**: Methodology, Project administration, Supervision.

Disclosure statement

The authors report there are no competing interests to declare.

Funding

This work was supported by the Japan Society for the Promotion of Science under Grant-in-Aid for Early-Career Scientists 23K17007; the Japan Society for the Promotion of Science under Grant-in-Aid for Scientific Research (B) 26K03038; and the JR-West Relief Foundation under the 2024 Research Grant.

ORCID

Wan-Jou She  <http://orcid.org/0000-0002-1549-5777>
 Luma Tabbaa  <http://orcid.org/0000-0002-0947-4988>
 Alexandra Covaci  <http://orcid.org/0000-0002-3205-2273>
 Panote Siriaraya  <http://orcid.org/0000-0002-4695-6417>
 Benjamin Yip  <http://orcid.org/0009-0006-3025-5997>
 Monica Perusquia-Hernandez  <http://orcid.org/0000-0002-0486-1743>
 Shinsuke Nakajima  <http://orcid.org/0000-0002-2367-6885>
 Naomi Yamashita  <http://orcid.org/0000-0002-6440-6865>
 Yu Shibuya  <http://orcid.org/0000-0001-5542-8272>
 Chee Siang Ang  <http://orcid.org/0000-0002-1109-9689>

Data availability statement

The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy and ethical restrictions, as they contain highly sensitive personal chat logs discussing the mental health and well-being of the research participants.

References

- Ali, K., Farrer, L., Gulliver, A., & Griffiths, K. M. (2015). Online peer-to-peer support for young people with mental health problems: A systematic review. *JMIR Mental Health*, 2(2), e4418. <https://doi.org/10.2196/mental.4418>
- Allen, C., Vassilev, I., Kennedy, A., & Rogers, A. (2016). Long-term condition self-management support in online communities: A meta-synthesis of qualitative papers. *Journal of Medical Internet Research*, 18(3), e61. <https://doi.org/10.2196/jmir.5260>
- Altman, I., & Taylor, D. A. (1973). *Social penetration: The development of interpersonal relationships*. Holt, Rinehart & Winston.
- Amj, G. (2016). *One-to-one peer support in mental health* [Ph.D. Dissertation]. UCL (University College London).
- Andalibi, N., & Flood, M. K. (2021). Considerations in designing digital peer support for mental health: Interview study among users of a digital support system (Buddy Project). *JMIR Mental Health*, 8(1), e21819. <https://doi.org/10.2196/21819>
- Andalibi, N., Haimson, O. L., De Choudhury, M., & Forte, A. (2018). Social support, reciprocity, and anonymity in responses to sexual abuse disclosures on social media. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 25(5), 1–35. <https://doi.org/10.1145/3234942>
- Andalibi, N., Ozturk, P., & Forte, A. (2017). Sensitive self-disclosures, responses, and social support on Instagram: The case of # depression. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 1485–1500). ACM.
- Ando, S., Yamaguchi, S., Aoki, Y., & Thornicroft, G. (2013). Review of mental-health-related stigma in Japan. *Psychiatry and Clinical Neurosciences*, 67(7), 471–482. <https://doi.org/10.1111/pcn.12086>
- Andrade, F. C., Erwin, S., Burnell, K., Jackson, J., Storch, M., Nicholas, J., & Zucker, N. (2023). Intervening on social comparisons on social media: Electronic daily diary pilot study. *JMIR Mental Health*, 10(1), e42024. <https://doi.org/10.2196/42024>
- Antonucci, T. C., Ajrouch, K. J., & Birditt, K. S. (2014). The convoy model: Explaining social relations from a multidisciplinary perspective. *Gerontologist*, 54(1), 82–92. <https://doi.org/10.1093/geront/gnt118>
- Ayers, J. W., Poliak, A., Dredze, M., Leas, E. C., Zhu, Z., Kelley, J. B., Faix, D. J., Goodman, A. M., Longhurst, C. A., Hogarth, M., & Smith, D. M. (2023). Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine*, 183(6), 589–596. <https://doi.org/10.1001/jamainternmed.2023.1838>
- Baumel, A., Tinkelman, A., Mathur, N., & Kane, J. M. (2018). Digital peer-support platform (7Cups) as an adjunct treatment for women with postpartum depression: Feasibility, acceptability, and preliminary efficacy study. *JMIR mHealth and uHealth*, 6(2), e9482. <https://doi.org/10.2196/mhealth.9482>
- Bergmo, T. S., Sandsdalen, V., Manskow, U. S., Småbrekke, L., & Waaseth, M. (2023). Internet use for obtaining medicine information: Cross-sectional survey. *JMIR Formative Research*, 7(1), e40466. <https://doi.org/10.2196/40466>
- Betts, A. (2025). *ChatGPT accused of acting as ‘suicide coach’ in series of US lawsuits*. <https://www.theguardian.com/technology/2025/nov/07/chatgpt-lawsuit-suicide-coach>
- Bhattacharjee, A., Williams, J. J., Meyerhoff, J., Kumar, H., Mariakakis, A., & Kornfield, R. (2023). Investigating the role of context in the delivery of text messages for supporting psychological wellbeing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–19). ACM.
- Boiani, J., Chung, E., Green, F., & Shah, A. (2025). *Novel lawsuits allege AI chatbots encouraged minors’ suicides, mental health trauma: Considerations for stakeholders*. JD Supra (October 2025). <https://www.jdsupra.com/legal-news/novel-lawsuits-allege-ai-chatbots-5250832/>
- Brandtzaeg, P. B., Skjuve, M., Dysthe, K. K., & Følstad, A. (2021). *When the social becomes non-human: Young people’s perception of social support in chatbots* [Paper presentation]. CHI Conference on Human Factors in Computing Systems (CHI ‘21) (pp. 1–13). ACM.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Burger, P. M., Westerink, J., & Vrijzen, B. E. L. (2020). Outcomes of second opinions in general internal medicine. *PloS One*, 15(7), e0236048. <https://doi.org/10.1371/journal.pone.0236048>
- Byrne, D. (2022). A worked example of Braun and Clarke’s approach to reflexive thematic analysis. *Quality & Quantity*, 56(3), 1391–1412. <https://doi.org/10.1007/s11135-021-01182-y>
- Caine, K. (2016). Local standards for sample size at CHI. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 981–992). ACM.
- Chancellor, S., Lin, Z., Goodman, E. L., Zerwas, S., & De Choudhury, M. (2016). Quantifying and predicting mental illness severity in online pro-eating disorder communities. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (pp. 1171–1184). ACM.
- Chen, Q., Chan, K. L., & Cheung, A. S. Y. (2018). Doxing victimization and emotional problems among secondary school students in Hong Kong. *International Journal of Environmental Research and Public Health*, 15(12), 2665. <https://doi.org/10.3390/ijerph15122665>

- Chiang, C.-W., Lu, Z., Li, Z., & Yin, M. (2023). Are two heads better than one in AI-assisted decision making? Comparing the behavior and performance of groups and individuals in human-AI collaborative recidivism risk assessment. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–18). ACM.
- Chung, L., & Kang, J. (2023). I'm hurt too: The effect of a chatbot/s reciprocal self-disclosures on users' painful experiences. *Archives of Design Research*, 36(4), 67–84. <https://doi.org/10.15187/ADR.2023.11.36.4.67>
- Chung, V., Mennella, R., Pacherie, E., & Grezes, J. (2024). Social bonding through shared experiences: The role of emotional intensity. *Royal Society Open Science*, 11(10), 240048. <https://doi.org/10.1098/rsos.240048>
- Cropanzano, R., & Mitchell, M. S. (2005). Social exchange theory: An interdisciplinary review. *Journal of Management*, 31(6), 874–900. <https://doi.org/10.1177/0149206305279602>
- Croskerry, P., Singhal, G., & Mamede, S. (2013). Cognitive debiasing 2: Impediments to and strategies for change. *BMJ Quality & Safety*, 22 Suppl 2(Suppl 2), ii65–ii72. <https://doi.org/10.1136/bmjqs-2012-001713>
- De Choudhury, M., & De, S. (2014). Mental health discourse on Reddit: Self-disclosure, social support, and anonymity. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1), 71–80. <https://doi.org/10.1609/icwsm.v8i1.14526>
- DeVault, D., Artstein, R., Benn, G., Dey, T., Fast, E., Gainer, A., Georgila, K., Gratch, J., Hartholt, A., & Lhommet, M. (2014). SimSensei Kiosk: A virtual human interviewer for healthcare decision support. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems* (pp. 1061–1068). ACM.
- Drysdale, M. T. B., McBeath, M. L., & Callaghan, S. A. (2022). The feasibility and impact of online peer support on the well-being of higher education students. *Journal of Mental Health Training, Education and Practice*, 17(3), 206–217. <https://doi.org/10.1108/JMHTPE-02-2021-0012>
- Fang, A., & Zhu, H. (2022). Matching for peer support: Exploring algorithmic matching for online mental health communities. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2), 1–37.
- Feeney, B. C., & Collins, N. L. (2015). A new look at social support: A theoretical perspective on thriving through relationships. *Personality and Social Psychology Review: An Official Journal of the Society for Personality and Social Psychology*, 19(2), 113–147. <https://doi.org/10.1177/1088868314544222>
- Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial. *JMIR Mental Health*, 4(2), e7785. <https://doi.org/10.2196/mental.7785>
- Fortuna, K. L., Naslund, J. A., LaCroix, J. M., Bianco, C. L., Brooks, J. M., Zisman-Ilani, Y., Muralidharan, A., & Deegan, P. (2020). Digital peer support mental health interventions for people with a lived experience of a serious mental illness: Systematic review. *JMIR Mental Health*, 7(4), e16460. <https://doi.org/10.2196/16460>
- GBD 2019 Mental Disorders Collaborators. (2022). Global, regional, and national burden of 12 mental disorders in 204 countries and territories, 1990–2019: A systematic analysis for the Global Burden of Disease Study 2019. *Lancet. Psychiatry*, 9(2), 137–150. [https://doi.org/10.1016/S2215-0366\(21\)00395-3](https://doi.org/10.1016/S2215-0366(21)00395-3)
- Gillard, S., Gibson, S. L., Holley, J., & Lucock, M. (2015). Developing a change model for peer worker interventions in mental health services: A qualitative research study. *Epidemiology and Psychiatric Sciences*, 24(5), 435–445. <https://doi.org/10.1017/S2045796014000407>
- Goffman, E. (1955). On face-work: An analysis of ritual elements in social interaction. *Psychiatry*, 18(3), 213–231. <https://doi.org/10.1080/00332747.1955.11023008>
- Greene, K., Derlega, V. J., & Mathews, A. (2006). Self-disclosure in personal relationships. *Cambridge Handbook of Personal Relationships*, 409(2006), 427. <https://doi.org/10.1017/CBO9780511606632.023>
- Hameed Shalaby, R. A., & Agyapong, V. I. O. (2020). Peer support in mental health: Literature review. *JMIR Mental Health*, 7(6):e15572. <https://doi.org/10.2196/15572>
- Herbener, A. B., & Damholdt, M. F. (2025). Are lonely youngsters turning to chatbots for companionship? The relationship between chatbot usage and social connectedness in Danish high-school students. *International Journal of Human-Computer Studies*, 196(2025), 103409. <https://doi.org/10.1016/j.ijhcs.2024.103409>
- Holthöwer, J., & Van Doorn, J. (2023). Robots do not judge: Service robots can alleviate embarrassment in service encounters. *Journal of the Academy of Marketing Science*, 51(4), 1–18. <https://doi.org/10.1007/s11747-022-00862-x>
- Housley, W., Webb, H., Williams, M., Procter, R., Edwards, A., Jirotko, M., Burnap, P., Stahl, B. C., Rana, O., & Williams, M. (2018). Interaction and transformation on social media: The case of Twitter campaigns. *Social Media + Society*, 4(1). <https://doi.org/10.1177/2056305117750721>
- Hsu, S.-L., Shah, R. S., Senthil, P., Ashktorab, Z., Dugan, C., Geyer, W., & Yang, D. (2025). Helping the helper: Supporting peer counselors via AI-empowered practice and feedback. *Proceedings of the ACM on Human-Computer Interaction*, 9(2), 1–45. <https://doi.org/10.1145/3710993>
- Kawanishi, Y. (2020). Shūkatsu utsu, the psychological toll of job-hunting in Japan. *Electronic Journal of Contemporary Japanese Studies*, 20(3), Article 8.
- Ina, M., & Morita, M. (2015). Japanese university students' stigma and attitudes toward seeking professional psychological help. *Online Journal of Japanese Clinical Psychology*, 2(1), 10–18.
- Inkster, B., Kadaba, M., & Subramanian, V. (2023). Understanding the impact of an AI-enabled conversational agent mobile app on users' mental health and wellbeing with a self-reported maternal event: A mixed method

- real-world data mHealth study. *Frontiers in Global Women's Health*, 4(2023), 1084302. <https://doi.org/10.3389/fgwh.2023.1084302>
- Julian De, F., Oguz-Uguralp, Z., & Kaan-Uguralp, A. (2025). *Emotional manipulation by AI companions*. arXiv preprint arXiv:2508.19258.
- Kaminger, S., Roth, L. H. O., & Laireiter, A.-R. (2023). #Blessed: The moderating effect of dispositional gratitude on the relationship between social comparison and envy on Instagram. *Frontiers in Psychology*, 14(2023), 1159999. <https://doi.org/10.3389/fpsyg.2023.1159999>
- Kanthawala, S., & Peng, W. (2021). Credibility in online health communities: Effects of moderator credentials and endorsement cues. *Journalism and Media*, 2(3), 379–396. <https://doi.org/10.3390/journalmedia2030023>
- Karkosz, S., Szymański, R., Sanna, K., & Michałowski, J. (2024). Effectiveness of a web-based and mobile therapy chatbot on anxiety and depressive symptoms in subclinical young adults: Randomized controlled trial. *JMIR Formative Research*, 8(1), e47960. <https://doi.org/10.2196/47960>
- Kaye, L. K., Wall, H. J., & Malone, S. A. (2016). “Turn that frown upside-down”: A contextual account of emotion usage on different virtual platforms. *Computers in Human Behavior*, 60(2016), 463–467. <https://doi.org/10.1016/j.chb.2016.02.088>
- Kellie, Y. H. S., & Choo, K. T. W. (2025). “I said things I needed to hear myself”: Peer support as an emotional, organisational, and sociotechnical practice in Singapore. arXiv preprint arXiv:2506.09362.
- Kingod, N., Cleal, B., Wahlberg, A., & Husted, G. R. (2017). Online peer-to-peer communities in the daily lives of people with chronic illness: A qualitative systematic review. *Qualitative Health Research*, 27(1), 89–99. <https://doi.org/10.1177/1049732316680203>
- Laine, J., Minkinen, M., & Mäntymäki, M. (2025). Understanding the ethics of generative AI: Established and new ethical principles. *Communications of the Association for Information Systems*, 56(1), 7. <https://doi.org/10.17705/1CAIS.05601>
- Lahey, B., & Cohen, S. (2000). Social support theory and measurement. In *Social Support Measurement and Intervention* (pp. 29–52). Oxford University Press. <https://doi.org/10.1093/med:psych/9780195126709.003.0002>
- Lawn, S., Smith, A., & Hunter, K. (2008). Mental health peer support for hospital avoidance and early discharge: An Australian example of consumer driven and operated service. *Journal of Mental Health*, 17(5), 498–508. <https://doi.org/10.1080/09638230701530242>
- Lebra, T. S. (1983). Shame and guilt: A psychocultural view of the Japanese self. *Ethos*, 11(3), 192–210. <https://doi.org/10.1525/eth.1983.11.3.02a00070>
- Lee, M., Ackermans, S., Van As, N., Chang, H., Lucas, E., & IJsselstein, W. (2019). Caring for Vincent: A chatbot for self-compassion. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). ACM.
- Lee, Y.-C., Cui, Y., Jamieson, J., Fu, W., & Yamashita, N. (2023). Exploring effects of chatbot-based social contact on reducing mental illness stigma. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). ACM.
- Li, H., Zhang, R., Lee, Y.-C., Kraut, R. E., & Mohr, D. C. (2023). Systematic review and meta-analysis of AI-based conversational agents for promoting mental health and well-being. *NPJ Digital Medicine*, 6(1), 236. <https://doi.org/10.1038/s41746-023-00979-5>
- Liu, Y., Fang, A., Moriarty, G., Firman, C., Kraut, R. E., & Zhu, H. (2024). Exploring trade-offs for online mental health matching: Agent-based modeling study. *JMIR Formative Research*, 8, e58241. <https://doi.org/10.2196/58241>
- Lucas, G. M., Gratch, J., King, A., & Morency, L.-P. (2014). It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior*, 37(2014), 94–100. <https://doi.org/10.1016/j.chb.2014.04.043>
- Ma, H., Smith, C. E., He, L., Narayanan, S., Giaquinto, R. A., Evans, R., Hanson, L., & Yarosh, S. (2017). Write for life: Persisting in online health communities through expressive writing and social support. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW), 1–24.
- Ma, J., Gao, S., Wang, P., & Liu, Y. (2024). High level of self-disclosure on SNSs facilitates cooperation: A serial mediation model of psychological distance and trust. *Computers in Human Behavior*, 150(2024), 107976. <https://doi.org/10.1016/j.chb.2023.107976>
- Macaulay, P., Betts, L. R., Stiller, J., & Kellezi, B. (2021). “The more public it is, the more severe it is”: Teachers’ perceptions on the roles of publicity and severity in cyberbullying. *Research Papers in Education*, 36(6), 726–753. <https://doi.org/10.1080/02671522.2020.1767183>
- Maples, B., Cerit, M., Vishwanath, A., & Pea, R. (2024). Loneliness and suicide mitigation for students using GPT3-enabled chatbots. *NPJ Mental Health Research*, 3(1), 4. <https://doi.org/10.1038/s44184-023-00047-6>
- Marshall, P., Booth, M., Coole, M., Fothergill, L., Glossop, Z., Haines, J., Harding, A., Johnston, R., Jones, S., Lodge, C., Machin, K., Meacock, R., Nielson, K., Puddephatt, J.-A., Rakic, T., Rayson, P., Robinson, H., Rycroft-Malone, J., Shryane, N., ... Lobban, F. (2024). Understanding the impacts of online mental health peer support forums: Realist synthesis. *JMIR Mental Health*, 11(2024), e55750. <https://doi.org/10.2196/55750>
- McCroskey, J. C. (1977). Oral communication apprehension: A summary of recent theory and research. *Human Communication Research*, 4(1), 78–96. <https://doi.org/10.1111/j.1468-2958.1977.tb00599.x>

- Mead, S., Hilton, D., & Curtis, L. (2001). Peer support: A theoretical perspective. *Psychiatric Rehabilitation Journal*, 25(2), 134–141. <https://doi.org/10.1037/h0095032>
- Meng, J., Rheu, M., Zhang, Y., Dai, Y., & Peng, W. (2023). Mediated social support for distress reduction: AI chatbots vs. human. *Proceedings of the ACM on Human-Computer Interaction* 7, 1(CSCW), 721–725.
- Mittmann, G., Schmalwieser, S. S., Diendorfer, T., Schrank, B., & Boeckle, M. (2022). Peer facilitators as core co-developers of an online peer encouragement network (OPEN2chat) for Austrian adolescents. *Frontiers in Digital Health*, 4(2022), 833006. <https://doi.org/10.3389/fdgth.2022.833006>
- Moitra, M., Santomauro, D., Collins, P. Y., Vos, T., Whiteford, H., Saxena, S., & Ferrari, A. J. (2022). The global gap in treatment coverage for major depressive disorder in 84 countries from 2000–2019: A systematic review and Bayesian meta-regression analysis. *PLoS Medicine*, 19(2), e1003901. <https://doi.org/10.1371/journal.pmed.1003901>
- Moore, D., Drey, N., & Ayers, S. (2020). A meta-synthesis of women's experiences of online forums for maternal mental illness and stigma. *Archives of Women's Mental Health*, 23(4), 507–515. <https://doi.org/10.1007/s00737-019-01002-1>
- Moore, J., Grabb, D., Agnew, W., Klyman, K., Chancellor, S., Ong, D. C., & Haber, N. (2025). Expressing stigma and inappropriate responses prevents LLMs from safely replacing mental health providers. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* (pp. 599–627). ACM.
- Moran, G. S., Russinova, Z., Gidugu, V., Yim, J. Y., & Sprague, C. (2012). Benefits and mechanisms of recovery among peer providers with psychiatric illnesses. *Qualitative Health Research*, 22(3), 304–319. <https://doi.org/10.1177/1049732311420578>
- Moumjid, N., Gafni, A., Bremond, A., & Carrere, M.-O. (2007). Seeking a second opinion: Do patients need a second opinion when practice guidelines exist? *Health Policy (Amsterdam, Netherlands)*, 80(1), 43–50. <https://doi.org/10.1016/j.healthpol.2006.02.009>
- Mulfinger, N., Müller, S., Böge, I., Sakar, V., Corrigan, P. W., Evans-Lacko, S., Nehf, L., Djamali, J., Samarelli, A., Kempter, M., Ruckes, C., Libal, G., Oexle, N., Noterdaeme, M., & Rüsch, N. (2018). Honest, open, proud for adolescents with mental illness: Pilot randomized controlled trial. *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, 59(6), 684–691. <https://doi.org/10.1111/jcpp.12853>
- Nanthanasit, A., Ppali, S., Farao, J., Tabbaa, L., Siang Ang, C., Sukittanon, S., & Covaci, A. (2025). Building a digital ecosystem for community-based rehabilitation: Insights from a multi-platform social media strategy in Thailand. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference* (pp. 1691–1709). ACM.
- Naslund, J. A., Aschbrenner, K. A., Marsch, L. A., & Bartels, S. J. (2016). The future of mental health care: Peer-to-peer support and social media. *Epidemiology and Psychiatric Sciences*, 25(2), 113–122. <https://doi.org/10.1017/S2045796015001067>
- National Alliance on Mental Illness (NAMI). (2025). *Mental health by the numbers*. <https://www.nami.org/about-mental-illness/mental-health-by-the-numbers>
- Nicol, G., Wang, R., Graham, S., Dodd, S., & Garbutt, J. (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), e40242. <https://doi.org/10.2196/40242>
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(2004), 119.
- O'Leary, K., Bhattacharya, A., Munson, S. A., O Wobbrock, J., & Pratt, W. (2017). Design opportunities for mental health peer support technologies. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 1470–1484). ACM.
- OECD. (2020). *Suicide rates*. <https://www.nami.org/about-mental-illness/mental-health-by-the-numbers>
- Okada, M., Matsumoto, R., & Motomura, E. (2024). Suicide mortality rates in Japan before and beyond the COVID-19 pandemic era. *PCN Reports: psychiatry and Clinical Neurosciences*, 3(2), e188. <https://doi.org/10.1002/pcn.5.188>
- Orgad, S. (2006). The cultural dimensions of online communication: A study of breast cancer patients' internet spaces. *New Media & Society*, 8(6), 877–899. <https://doi.org/10.1177/1461444806069643>
- Park, D., Choo, M., Cho, M., Kim, J., & Shin, Y.-J. (2024). Collaborative school mental health system: Leveraging a conversational agent for enhancing children's executive function. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1–17). ACM.
- Pater, J. A., Coupe, A., Miller, A. D., Reining, L. E., Drouin, M., & Toscos, T. (2020). Design opportunities and challenges for app-based telemental health technologies for teens and young adults. In *Proceedings of the 14th EAI International Conference on Pervasive Computing Technologies for Healthcare* (pp. 188–200). ACM.
- Payne, K. (2024). *An AI chatbot pushed a teen to kill himself, a lawsuit against its creator alleges*. <https://apnews.com/article/chatbot-ai-lawsuit-suicide-teen-artificial-intelligence-9d48adc572100822fdb3c90d1456bd0>
- Pfeiffer, P. N., Heisler, M., Piette, J. D., Rogers, M. A. M., & Valenstein, M. (2011). Efficacy of peer support interventions for depression: A meta-analysis. *General Hospital Psychiatry*, 33(1), 29–36. <https://doi.org/10.1016/j.genhosppsych.2010.10.002>
- Popowski, S. A., Meyerhoff, J., Allen, O. M., Nguyen, T., McCall, T., Pederson, A. B., Reddy, M., Mohr, D., & Kornfield, R. (2025). Designing digital mental health tools to support the needs of Black adults in the United States: Qualitative analysis. *JMIR Formative Research*, 9(2025), e73279–e73279. <https://doi.org/10.2196/73279>
- Poulsen, A., Hickie, I. B., de Haan, Z., Capon, W., LaMonica, H. M., Eyeson-Annan, E., Radwan, J., Barnabie, N., Scott, E. M., & Iorfino, F. (2025). Co-designing a conversational generative artificial intelligence system for

- youth mental health. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (pp. 1–7). ACM.
- Prizeman, K., Weinstein, N., & McCabe, C. (2023). Effects of mental health stigma on loneliness, social isolation, and relationships in young people with depression symptoms. *BMC Psychiatry*, 23(1), 527. <https://doi.org/10.1186/s12888-023-04991-7>
- Pudlinski, C. (2005). Doing empathy and sympathy: Caring responses to troubles tellings on a peer support line. *Discourse Studies*, 7(3), 267–288. <https://doi.org/10.1177/1461445605052177>
- Punton, G., Dodd, A. L., & McNeill, A. (2022). ‘You’re on the waiting list’: An interpretive phenomenological analysis of young adults’ experiences of waiting lists within mental health services in the UK. *PloS One*, 17(3), e0265542. <https://doi.org/10.1371/journal.pone.0265542>
- Qahri-Saremi, H., & Turel, O. (2023). Situational contingencies in susceptibility of social media to phishing: A temptation and restraint model. *Journal of Management Information Systems*, 40(2), 503–540. <https://doi.org/10.1080/07421222.2023.2196779>
- Rayland, A., & Andrews, J. (2023). From social network to peer support network: Opportunities to explore mechanisms of online peer support for mental health. *JMIR Mental Health*, 10, e41855. <https://doi.org/10.2196/41855>
- Rennick-Egglestone, S., Ramsay, A., McGranahan, R., Llewellyn-Beardsley, J., Hui, A., Pollock, K., Repper, J., Yeo, C., Ng, F., Roe, J., Gillard, S., Thornicroft, G., Booth, S., & Slade, M. (2019). The impact of mental health recovery narratives on recipients experiencing mental health problems: Qualitative analysis and change model. *PloS One*, 14(12), e0226201. <https://doi.org/10.1371/journal.pone.0226201>
- Repper, J., & Carter, T. (2011). A review of the literature on peer support in mental health services. *Journal of Mental Health (Abingdon, England)*, 20(4), 392–411. <https://doi.org/10.3109/09638237.2011.583947>
- Rosenthal-von der Pütten, A. M., Hastall, M. R., Köcher, S., Meske, C., Heinrich, T., Labrenz, F., & Ocklenburg, S. (2019). “Likes” as social rewards: Their role in online social comparison and decisions to like other people’s selfies. *Computers in Human Behavior*, 92(2019), 76–86. <https://doi.org/10.1016/j.chb.2018.10.017>
- Russell Spears. (2017). Social identity model of deindividuation effects. In *Encyclopedia of group processes & inter-group relations* (pp. 794–797). SAGE Publications, Inc. <https://doi.org/10.4135/9781412972017.n245>
- Saha, K., & De Choudhury, M. (2017). Modeling stress with social media around incidents of gun violence on college campuses. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW(2017), 1–27.
- Saha, K., Jain, Y., Liu, C., Kaliappan, S., & Karkar, R. (2026). AI vs. humans for online support: Comparing the language of responses from LLMs and online communities of Alzheimer’s disease. *ACM Transactions on Computing for Healthcare*, 7(3), 1–39. <https://doi.org/10.1145/3709366>
- Sánchez-Hernández, M. D., Herrera, M. C., & Expósito, F. (2022). Does the number of likes affect adolescents’ emotions? The moderating role of social comparison and feedback-seeking on Instagram. *Journal of Psychology*, 156(3), 200–223. <https://doi.org/10.1080/00223980.2021.2024120>
- Schlenker, B. R., & Leary, M. R. (1982). Social anxiety and self-presentation: A conceptualization and model. *Psychology Bulletin*, 92(1982), 641–669. <https://doi.org/10.1037/0033-2909.92.3.641>
- Seitz, L. (2024). Artificial empathy in healthcare chatbots: Does it feel authentic? *Computers in Human Behavior: Artificial Humans*, 2(1), 100067. <https://doi.org/10.1016/j.chbah.2024.100067>
- Sharma, A., Choudhury, M., Althoff, T., & Sharma, A. (2020). Engagement patterns of peer-to-peer interactions on mental health platforms. In *Proceedings of the International AAAI Conference on Web and Social Media*, 14(1), 614–625. <https://doi.org/10.1609/icwsm.v14i1.7328>
- Sharma, A., Lin, I. W., Miner, A. S., Atkins, D. C., & Althoff, T. (2023). Human-AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support. *Nature Machine Intelligence*, 5(1), 46–57. <https://doi.org/10.1038/s42256-022-00593-2>
- She, W. J., Dangisho, K., Siriaraya, P., Dollack, F., & Nakajima, S. (2023). Matchmaking for mental well-being: Development of a peer-based support system (Peer2S) for students during COVID lockdown. In *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces* (pp. 119–122). ACM. <https://doi.org/10.1145/3581754.3584148>
- Simmons, M. B., Cartner, S., MacDonald, R., Whitson, S., Bailey, A., & Brown, E. (2023). The effectiveness of peer support from a person with lived experience of mental health challenges for young people with anxiety and depression: A systematic review. *BMC Psychiatry*, 23(1), 194. <https://doi.org/10.1186/s12888-023-04578-2>
- Singer, T., & Klimecki, O. M. (2014). Empathy and compassion. *Current Biology: CB*, 24(18), R875–R878. <https://doi.org/10.1016/j.cub.2014.06.054>
- Smit, D., Vrijnsen, J. N., Groeneweg, B., Vellinga-Dings, A., Peelen, J., & Spijker, J. (2021). A newly developed online peer support community for depression (depression connect): Qualitative study. *Journal of Medical Internet Research*, 23(7), e25917. <https://doi.org/10.2196/25917>
- Smith-Merry, J., Goggin, G., Campbell, A., McKenzie, K., Ridout, B., & Bayliss, C. (2019). Social connection and online engagement: Insights from interviews with users of a mental health online forum. *JMIR Mental Health*, 6(3), e11084. <https://doi.org/10.2196/11084>
- Solomon, P. (2004). Peer support/peer provided services underlying processes, benefits, and critical ingredients. *Psychiatric Rehabilitation Journal*, 27(4), 392–401. <https://doi.org/10.2975/27.2004.392.401>

- Stoeber, J., Kobori, O., & Tanno, Y. (2013). Perfectionism and self-conscious emotions in British and Japanese students: Predicting pride and embarrassment after success and failure. *European Journal of Personality*, 27(1), 59–70. <https://doi.org/10.1002/per.1858>
- Suharwardy, S., Ramachandran, M., Leonard, S. A., Gunaseelan, A., Lyell, D. J., Darcy, A., Robinson, A., & Judy, A. (2023). Feasibility and impact of a mental health chatbot on postpartum mental health: A randomized controlled trial. *AJOG Global Reports*, 3(3), 100165. <https://doi.org/10.1016/j.xagr.2023.100165>
- Suler, J. (2004). The online disinhibition effect. *Cyberpsychology & Behavior: The Impact of the Internet, Multimedia and Virtual Reality on Behavior and Society*, 7(3), 321–326. <https://doi.org/10.1089/1094931041291295>
- Syed, S., Iftikhar, Z., Xiao, A. W., & Huang, J. (2024). *Machine and human understanding of empathy in online peer support: A cognitive behavioral approach* [Paper presentation]. Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (pp. 1–13). ACM. <https://doi.org/10.1145/3613904.3642034>
- Szeto, M. D., Barber, C., Ranpariya, V. K., Anderson, J., Hatch, J., Ward, J., Aguilera, M. N., Hassan, S., Hamp, A., Coolman, T., & Dellavalle, R. P. (2022). Emojis and emoticons in health care and dermatology communication: Narrative review. *JMIR Dermatology*, 5(3), e33851. <https://doi.org/10.2196/33851>
- Tanaka, T., & Okamoto, S. (2021). Increase in suicide following an initial decline during the COVID-19 pandemic in Japan. *Nature Human Behaviour*, 5(2), 229–238. <https://doi.org/10.1038/s41562-020-01042-z>
- Thomson, M., Henderson, G., Vines, J., Rogers, T., & MacBeth, A. (2023). *Towards a better understanding of peer support platforms for digital mental health: Learning from stakeholder engagement and mapping a theory of change*.
- Topor, A., Bøe, T. D., & Larsen, I. B. (2018). Small things, micro-affirmations and helpful professionals everyday recovery-orientated practices according to persons with mental health problems. *Community Mental Health Journal*, 54(8), 1212–1220. <https://doi.org/10.1007/s10597-018-0245-9>
- Treadgold, B. M., Coulson, N. S., Campbell, J. L., Lambert, J., & Pitchforth, E. (2025). Quality and misinformation about health conditions in online peer support groups: Scoping review. *Journal of Medical Internet Research*, 27(2025), e71140. <https://doi.org/10.2196/71140>
- Walther, J. B. (1996). Computer-mediated communication: Impersonal, interpersonal, and hyperpersonal interaction. *Communication Research*, 23(1), 3–43. <https://doi.org/10.1177/009365096023001001>
- Wang, J., Mann, F., Lloyd-Evans, B., Ma, R., & Johnson, S. (2018). Associations between loneliness and perceived social support and outcomes of mental health problems: A systematic review. *BMC Psychiatry*, 18(1), 156. <https://doi.org/10.1186/s12888-018-1736-5>
- Wang, T., Bruckman, A. S., & Yang, D. (2025a). The practice of online peer counseling and the potential for AI-powered support tools. *Proceedings of the ACM on Human-Computer Interaction*, 9(2), 1–33. <https://doi.org/10.1145/3711089>
- Wang, T., Shah, H. K., Shah, R. S., Wang, Y.-C., Kraut, R. E., & Yang, D. (2023). Metrics for peer counseling: Triangulating success outcomes for online therapy platforms. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–17). ACM.
- Wang, Y., Wang, Y., Crace, K., & Zhang, Y. (2025b). Understanding attitudes and trust of generative AI chatbots for social anxiety support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–21). ACM.
- Williams, A. H., Williams, A. W., Renner, L., Pettyjohn, M. E., Cash, S. J., & Schwab-Reese, L. M. (2024). “I need to talk to someone... What do I do?": Peer-to-peer disclosures of child maltreatment on social media. *Journal of Family Violence*, 40, 1535–1548. <https://doi.org/10.1007/s10896-024-00697-w>
- Won, H., Kang, M., Kim, M., Lee, D., Choi, H., Kim, Y., Choi, D., Ko, M., & Han, J. (2025). “Show your mind”: Unveiling user experience on an AI-based mental health assessment system with symptom-based evidences. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (pp. 1–11). ACM.
- World Health Organisation (WHO). (2022). *Mental disorders*. <https://www.who.int/news-room/fact-sheets/detail/mental-disorders>
- Xyghkou, A., Ang, C. S., Siriaraya, P., Kopecki, J. P., Covaci, A., Kanjo, E., & She, W.-J. (2024a). *MindTalker: Navigating the complexities of AI-enhanced social engagement for people with early-stage dementia* [Paper presentation]. Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (pp. 1–15). ACM. <https://doi.org/10.1145/3613904.3642538>
- Xyghkou, A., Siriaraya, P., Covaci, A., Prigerson, H. G., Neimeyer, R., Ang, C. S., & She, W.-J. (2023). The “conversation” about loss: Understanding how Chatbot technology was used in supporting people in grief. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–15). ACM.
- Xyghkou, A., Siriaraya, P., She, W.-J., Covaci, A., & Ang, C. S. (2024b). “Can I be more social with a chatbot?": Social connectedness through interactions of autistic adults with a conversational virtual human. *International Journal of Human-Computer Interaction*, 40(24), 8937–8954. <https://doi.org/10.1080/10447318.2023.2292880>
- Yang, D., Kraut, R., & Levine, J. M. (2017). Commitment of newcomers and old-timers to online health support communities. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI Conference, 2017*, 6363–6375. <https://doi.org/10.1145/3025453.3026008>

- Yao, Z., Yang, D., Levine, J. M., Low, C. A., Smith, T., Zhu, H., & Kraut, R. E. (2021). Join, stay or go? A closer look at members' life cycles in online health communities. *Proceedings of the ACM on Human-Computer Interaction* 5(CSCW1), 1–22.
- Ye, S., Ho, K. K. W., Wakabayashi, K., & Kato, Y. (2023). Relationship between university students' emotional expression on tweets and subjective well-being: Considering the effects of their self-presentation and online communication skills. *BMC Public Health*, 23(1), 594. <https://doi.org/10.1186/s12889-023-15485-2>
- Young, J., Jawara, L. M., Nguyen, D. N., Daly, B., Huh-Yoo, J., & Razi, A. (2024). The role of AI in peer support for young people: A study of preferences for human-and AI-generated responses. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1–18). ACM.
- Yuan, S., Davidson, G., & Best, P. (2025). Online mental health peer support: A systematic scoping review of the theoretical mechanisms of effect. *BMC Digital Health*, 3(1), 68. <https://doi.org/10.1186/s44247-025-00205-0>
- Zhang, Z., Sun, K., Jatchavala, C., Koh, J., Chia, Y., Bose, J., Li, Z., Tan, W., Wang, S., Chu, W., Wang, J., Tran, B., & Ho, R. (2019). Overview of stigma against psychiatric illnesses and advancements of anti-stigma activities in six Asian societies. *International Journal of Environmental Research and Public Health*, 17(1), 280. <https://doi.org/10.3390/ijerph17010280>

About the authors

Wan-Jou She focuses on the design and evaluation of AI agents that enhance individual resilience through personalized coping support, reflection, and adaptive guidance. She investigates how AI can facilitate peer support, strengthen social support networks, and help people navigate stress, uncertainty, and change through ethical, human-centered approaches.

Luma Tabbaa is a Lecturer in Digital Media at the University of Kent. Her HCI research uses co-design to understand real-world constraints, stakeholder needs, and deployment opportunities, developing immersive and novel technologies for healthcare, wellbeing, dementia care, rehabilitation, pain management, anxiety, and education.

Alexandra Covaci works across human-computer interaction, artificial intelligence, and immersive technologies to design care-centered systems for health and wellbeing. Her research spans virtual humans and conversational agents, social VR, and multisensory design, guided throughout by questions of agency and trust in the technologies we live alongside.

Panote Siriiraya is an Associate Professor at Kyoto Institute of Technology, Japan. His work in Human-Computer Interaction explores VR, AI, and gamification to enhance physical and mental wellbeing, including applications in mental healthcare, dementia support, and peer-based support systems.

Benjamin Yip is a resident physician at Addenbrooke's Hospital in Cambridge, where he is developing broad clinical expertise across hospital medicine. Committed to evidence-based, compassionate care, he brings a thoughtful, patient-centered approach to multidisciplinary teams and is building a career grounded in service, learning, and clinical excellence medical practice.

Monica Perusqui-Hernandez is based at Nara Institute of Science and Technology, Japan. She holds degrees from ITESM, Eindhoven University of Technology, and the University of Tsukuba. Her work spans affective computing, physiological computing, VR, and human-AI interaction, focusing on wearable sensing, embodied experiences, and technology-mediated communication.

Shinsuke Nakajima is a professor in the Faculty of Information Science and Engineering at Kyoto Sangyo University, Japan. He received his PhD from Kyoto University and was a visiting scholar at the University of California, Santa Barbara. His research interests include text mining and recommender systems.

Naomi Yamashita is a professor at Kyoto University working across CSCW, HCI, and AI. Her research examines how emerging technologies reshape human relationships and social interactions, and explores how technology can support well-being, diversity, inclusion, and human agency in an AI-mediated world.

Yu Shibuya is a professor of Information and Human Sciences at the Kyoto Institute of Technology. He earned his BE, ME, and DE from Osaka University in 1985, 1987, and 1990. His research focuses on human-computer interaction, media communication, and mobile interfaces, creating comfortable digital environments for human connection.

Chee Siang Ang is a researcher working at the intersection of artificial intelligence, healthcare, and human-computer interaction. His work focuses on designing and evaluating digital technologies that support communication, clinical decision-making, and patient care, with a particular interest in interdisciplinary and real-world healthcare applications.