Towards Successful Deployment of Wellbeing Sensing Technologies: Identifying Misalignments across Contextual Boundaries

Jina Suh
Microsoft Research
Redmond, USA
jinsuh@microsoft.com

Javier Hernandez
Microsoft Research
Redmond, USA
javierh@microsoft.com

Koustuv Saha
University of Illinois at Urbana-Champaign
Urbana, USA
ksaha2@illinois.edu

Kathy Dixon
Microsoft
Redmond, USA
kdixon@microsoft.com

Mehrab Bin Morshed
Georgia Institute of Technology
Atlanta, USA
mehrab.morshed@gatech.edu

Esther Howe
University of California, Berkeley
Berkeley, USA
howe@berkeley.edu

Anna Kawakami
Carnegie Mellon University
Pittsburgh, USA
akawakam@andrew.cmu.edu

Mary Czerwinski
Microsoft Research
Redmond, USA
marycz@microsoft.com

***Abstract*—Affective computing technologies have emerged as promising tools for supporting mental health and wellbeing but face deployment challenges in work and nonwork contexts due to misalignments in boundary preferences, data ownership, values and incentives, wellbeing definitions, and power dynamics. This paper presents a case study on the deployment of a just-in-time emotional support agent in the workplace, highlighting the five categories of misalignments that undermine the successful deployment of personal mental health support systems across these contextual boundaries. The identification and analysis of these misalignments contributes to a deeper understanding of the complexities and challenges faced when implementing affective computing systems for holistic mental health support. By emphasizing the importance of considering these misalignments in future research and development, this paper aims to contribute to the ongoing discourse on the effective and ethical deployment of affective computing technologies across work and nonwork contexts.**

***Index Terms*—wellbeing, sensing, workplace, context, deployment**

I. Introduction and Background

Mental health disorders affect a significant portion of the global population, with many individuals unable to access adequate professional support due to a variety of factors, such as limited availability of mental health professionals, financial constraints, and social stigma [1]–[3]. Digital mental health solutions, such as mobile applications and online teletherapy platforms, have emerged as promising tools to bridge this gap [4], but they often face limitations in terms of personalization, engagement, and effectiveness in addressing mental health needs [5]–[7]. In this context, affective computing can be harnessed to support individuals in managing their mental wellbeing by providing measures of affective states and behaviors as well as personalized, context-aware, and timely interventions [8]–[11].

In recent years, there has been an increased emphasis on holistic, person-centered strategies for managing mental health [12]–[14] that take a comprehensive approach to considerating all aspects and contexts to understand and treat the whole person [15]. However, affective computing systems face several challenges when it comes to weaving these contexts together and providing a holistic mental health care. For example, obtaining a holistic understanding of an individual requires gathering and analyzing data from multiple sources and contexts, such as personal, work, and social environments. Collecting data from various streams raises concerns about privacy and consent, while not having access to all relevant information about the context in which emotions and behaviors occur may lead to inaccurate interpretations and unhelpful assumptions. To provide holistic mental health care, affective computing systems need to be integrated with existing support systems, but such integration requires complex coordination and collaboration across stakeholders and environments. Although it may seem plausible to address these challenges with technical solutions given adequate resources and sophisticated technological advancements, the fundamental issue lies in the social, conceptual, and contextual boundaries and the practical and ethical concerns that arise when crossing these boundaries.

These boundaries have become increasingly blurred and complicated due to the digital exhaust generated by prolific and ubiquitous technologies, as well as the growing interest of companies and organizations in mining this data for their own purposes [16]–[19]. In the workplace setting, the heightened digitization of the workplace [20]–[22], alongside the rise of remote and hybrid work prompted by the pandemic, led to a surge in corporate interest in using affective computing and passive sensing technologies to support workplace wellbeing [23]–[26]. These technologies are often motivated by the ease of unobtrusive data collection and intentions of improving personal wellbeing. However, despite recent efforts to guide the ethical use of these technologies [27], [28], they often fail to address socio-technical and ethical complexities that arise with deployment in the real-world settings where individuals

979-8-3503-2745-8/23/$31.00 ©2023 IEEE

cannot be completely decoupled from their environment [29]–[33]. This crossing the contextual boundaries of work and nonwork can introduce a range of complexities that challenge the potential impact and in-practice desirability of wellbeing sensing technologies.

In this paper, we present a case study on the deployment of a just-in-time (JIT) emotional support agent in the workplace to highlight challenges and concerns surrounding contextual boundaries. Grounded in four user studies that evaluated the design, effectiveness, and feasibility of the system in the work and nonwork contexts, we consider the perspectives of various stakeholders to present five misalignments that undermine successful deployment of personal mental health support systems across the contextual boundaries. We propose that future research efforts in affective computing should strive for holistic mental wellbeing support and carefully consider these dimensions of misalignments.

II. Refinement of JIT Emotional Support Agent

With the goal of helping individuals manage their stress, we designed an emotional support agent that used passive sensing technologies to nudge people to engage in stress-reduction micro-interventions when the system detected higher than average stress levels. The system consists of (1) a passive-sensing software installed on their desktops that captures contextual, behavioral, and physiological data about the user, (2) a user interface (e.g., chatbot, mobile or desktop apps) that checks in with the user, facilitates the consumption of micro-intervention activities, and collects subjective stress levels, and (3) a service that incorporates all data streams and runs an algorithm that infers the user's stress level and the opportune moment to intervene.

The development and refinement of the just-in-time emotional support agent involved multiple studies. We start with Study 1 that introduces the concept of our system without a specific context in mind, and the study revealed potential benefits as well as privacy concerns across different contexts of use. In Studies 2 and 3, we deploy such a system specifically targetting the work context to probe how a system situated at work can interact with personal wellbeing concepts such as stress. In Study 4, we characterize the challenges of deploying such a system at work from multiple stakeholders. In this section, we briefly describe each study and summarize uncovered challenges of deploying a personal wellbeing sensing and intervention system in the workplace setting.

A. Study 1: Co-refinement of Emotional Support Agent

The goal of the first study was to understand the feasibility of an emotional support agent for individuals and to identify opportunities for refinement of design. The concept of a JIT emotional support agent was tested in 4 focus group sessions with 4 participants in each session. The focus group participants were a mix of technology enthusiasts and information workers. They were shown high-fidelity mockups of the system to collaboratively discuss and refine the features of the concept and how the system might be integrated into their daily lives.

Overall, the participants expressed that the concept was highly relevant. They saw immediate benefits of de-escalating a short-term issue through intervention as well as long-term benefits of training to help with personal growth and emotion regulation. Getting a full picture of oneself was highly sought after, as one participant expressed: "For me, my biggest thing is the interoperability... I think it also gives you an entire picture, too... mental health is just as important as your physical health and it all kind of works together." Participants were willing to provide feedback to the system so that it adapts to them individually, but there were general concerns about the intrusiveness of camera or audio-based sensing invading other family members' lives (e.g., children at home) or being used against them (e.g., in a court of law). Some participants were willing to provide access to their work and personal calendars in order to help the system detect stress and intelligently plan interventions (e.g., taking a break after consecutive meetings). On the other hand, others highlighted limited access to work calendar and emails as a barrier.

B. Study 2: Stress Prediction through Passive Sensing

The second study [34] aimed to build and deploy a fully functional passive sensing software that captures contextual, behavioral, and physiological signals from webcams, computer applications and activities, keyboard and mouse, email, and calendar. We deployed this passive sensing software specifically in the work context because work is a major source of stress and to understand how these passively captured signals relate to stress experienced at work. The passive sensing software was a custom Windows desktop app that captured and aggregated data from common workplace tools such as a webcam, keyboard, or mouse. Webcam video feeds were analyzed for facial expressions and physiological states using non-contact techniques [35], while peripheral devices contributed signals such as keystroke speed and mouse movements. Application usage and device-independent telemetry were also measured. The system integrated with Outlook for email and calendar metrics, with computed data sent to a cloud-based orchestrator via Azure Service Bus. The software was installed on 50 information workers' work computers (desktops or laptops) and collected their subjective daily and momentary stress ratings for four weeks. Using this data, we built and evaluated a machine-learning model that can accurately infer moments of high stress.

We found that passive sensors are highly effective in detecting triggers and manifestations of workplace stress (F1-score of up to 78%). Even though facial expressions (e.g., Facial Action Units) were the most predictive features of stress, webcams were perceived as the least comfortable in the context of workplace stress monitoring. Participants reported varying levels of comfort about sharing sensed data or personal stress levels with human resources (HR) and their managers, highlighting intricate power dynamics. Certain subsets were okay to share given specific circumstances. For example, sharing "meetings and engagement data" was okay with direct managers, sharing stress level data was okay with HR "to promote wellbeing," and sharing stress level data was okay with "chosen medical provider to better diagnose workplace-related stress factors and solutions."

C. Study 3: JIT Stress-Reduction Micro-Intervention System

The goal of the third study [36] is to deploy one instantiation of an emotional support agent to evaluate its effectiveness and feasibility in the same work context. In a third study, we designed and developed a chatbot to deliver JIT micro-interventions. The system leveraged data streams from the above passive sensing software to infer the user's stress level. The stress level, computed every 30 seconds, is an average of five components identified as stress sources in prior work: received emails [37], daily meetings [38], time into the day [39], facial expressions [39], and heart rate [40]. When the inferred stress level is above a threshold, the system sent a chat message to the user asking to engage in an intervention.

The micro-interventions were based on components of Cognitive Behavioral Therapy (CBT) and Dialectical Behavioral Therapy (DBT) and designed to take under five minutes. They were comprised of either a short video, a single-turn text prompt, or a brief therapeutic conversation with the chatbot under three types of interventions – 'Get my mind off work', 'Feel calm and present', and 'Think through my stress.' In a four-week study, we deployed the chatbot to the work computers of 86 information workers; 43 participants received JIT nudges and the other 43 manually scheduled the interventions in the calendar. We then captured the system usage data to evaluate the usage and effectiveness of the interventions as well as to determine tailoring variables that improve the likelihood of engagement and effectiveness.

We found that our micro-interventions, especially 'Think through my stress' ($\Delta\bar{x}$=-0.41), were effective in reducing momentary stress ($\Delta\bar{x}$=-0.34 overall out of a 5-point stress scale across all types, $t(1084)$=18.113, $p \ll 0.001$). 77% of those that manually scheduled interventions sought out automated nudges, and 70% of JIT participants expressed tht JIT interventions were a good reminder to help people take time out of their day, especially when stressed: "It made me think more on my stress and help me with better work-life balance." JIT participants further suggested improvements to timing and frequency of the nudges to have some level of control because frequent nudges were disruptive of focus, which indicates potentially conflicting goals between productivity and stress management. Intervention preferences greatly varied across participants, highlighting that a one-size-fits-all intervention approach may not work. The fact that sensing was performed on the work desktop and that the study was mostly taking place during work hours added some confusion around work-nonwork boundaries. For example, one participant noted that "I don't normally open work computer on the weekends," and another noted that "I didn't understand if I was supposed to be rating my work stress or my personal stress, my work resources or my personal resources." As we have seen from both Study 1 and Study 2, participants in Study 3 were also concerned about their privacy, given the video and audio-based data collection. In addition, they desired that the conversational texts generated while interacting with the chatbot should be kept for themselves.

D. Study 4: Multi-stakeholder Perspectives on Wellbeing Sensing at Work

Finally, in a fourth study [33], we aimed to understand the potential harms and benefits of deploying such a system at work. We conducted storyboard-driven interviews with 33 participants across three stakeholder groups (1. organizational governors—*decision-makers including managers, leaders, HR, and legal personnel*, 2. AI builders—*researchers, designers, and developers of wellbeing technologies*, and 3. worker data subjects—*end-user workers on whom the envisioned wellbeing sensing would be applied*). The storyboards allowed the depiction of various stakeholders, contexts, and instantiations of wellbeing sensing and intervention technologies, eliciting feedback on the concerns and desires about wellbeing technologies in the workplace.

We found cascading impacts of deploying such technologies across individual, interpersonal, and organizational layers as well as potential harms arising from ambiguous and misaligned notions of wellbeing. At the *individual layer*, while such technologies can promote individual wellbeing and self-reflections, recommendations based on over-simplification or over-generalization of wellbeing definitions can be unhelpful or even harmful. At the *interpersonal layer*, the systems may promote proactive and positive manager-worker relationships to support worker wellbeing, but managers may collect data that propagates productivity-centric wellbeing metrics or unhealthy social comparison or lack sufficient contextual factors to fully interpret the data. At the *organizational layer*, such technologies could enable organization-wide policies that support worker wellbeing, but there are potentials to misuse the data for incentives that do not support worker wellbeing and to thwart transparency.

III. Challenges due to Misalignments across Work and Nonwork Contexts

Across our studies, we envisioned and explored the use of personal wellbeing sensing and intervention technologies and how such technologies could be deployed and used in workplace contexts. Through our studies, we learned that the deployment of such technologies, regardless of who owns and controls the deployment, has inherent challenges in crossing contextual boundaries and supporting mental health holistically. We conducted a thematic analysis to organize these challenges into five categories of misalignments across contexts that must be considered when deploying affective computing and wellbeing sensing technologies for holistic, person-centered mental health support. For this discussion, we focus specifically on contextual boundaries of work and nonwork.

A. Misalignment in Boundary Preferences

Over the years, through decades of research in work and nonwork boundaries, interfaces, separation, and integration, we know that the preferences for these boundaries vary across individuals [41], [42]. Some seek a complete separation of work and nonwork, while others integrate nonwork activities into work throughout the day. These boundary preferences are also influenced by the flexibilities and policies that vary across organizations [43]–[45]. We also have learned that work bleeds into nonwork (e.g., spillover effects [46]) and nonwork bleeds into work (e.g., childcare challenges during the pandemic [47]). In other words, personal wellbeing is not only influenced by the workplace culture and organizational structure but also the desires and motivations of the individual and their relationships. This phenomenon is best described by the social-ecological model [48].

In the context of affective computing technologies, when wellbeing sensing technologies are deployed in the workplace, it necessarily takes a stance that personal wellbeing is now in the purview of workplace and organizational wellbeing. In fact, this workplace wellness perspective has existed since the advent of Employee Assistance Programs (EAPs) in the 1950s and labor and workers' rights movements of the 1800s [49]. However, traditional workplace wellbeing programs like EAPs are fundamentally different from affective computing and sensing technologies in terms of their approach to data collection and maintaining boundaries. Participation in the EAPs (i.e., professional psychotherapy) is typically infrequent in comparison to continuous sensing and physically or temporally decoupled from work, and its confidentiality is protected by federal law in the US [50]. On the other hand, the premise of sensing technologies requires that they are ubiquitous, unobtrusive, and always on [13], [51], making it difficult to separate data from work or nonwork.

From a whole-person mental health point of view, wellbeing sensing should seamlessly weave in and out of work and nonwork contexts, as Study 1 participant pointed out about which context our emotional support agent should be optimized for: "There's no work Henry and home Henry. There's just Henry (name replaced)." As many of our participants across the study reported, interoperability between personal and work data streams is necessary to paint the full picture of the person, but the separation of the data sources is also necessary to understand the root cause of the symptoms. Others found it uncomfortable and somewhat "creepy" (Study 4) to disclose personal wellbeing situations at work at the granularity that sensing technologies can: "We cross the threshold from being someone's employer to being more than that (Study 4)." With sensing, boundaries that could traditionally be set according to personal preferences become incredibly difficult to maintain. Addressing these challenges is essential as it raises the question of who has the power to set the boundary where and the strengths of these boundaries in the context of affective computing technologies.

B. Misalignment in Data Ownership

As individuals navigate in and out of work and nonwork contexts and interact with digital tools, they leave digital breadcrumbs that can be mined to obtain a better understanding of that person. How much you can separate digital footprints about work from nonwork may depend on the strength of the work and nonwork boundary separation and integration, but we all know that tending to personal needs, whether it is taking a biobreak or running personal errands, frequently happen at work. According to the United States Bureau of Labor Statistics, the average number of hours worked per week is around 34.4[1], and the average number of hours of sleep per day is 6.8 hours in the US[2]. This means that approximately 40% of their waking hours are spent at work, and a significant portion of their personal data is generated within the workplace. However, questions regarding the ownership, access, and rights to this data remain unresolved and contentious, complicating the integration of data streams across work and nonwork contexts.

One key challenge in this regard is determining the extent to which data generated at work can be considered personal data. While certain types of data, such as health records and personal communications, may be more easily classified as personal data, other types may fall into a gray area. For example, our emotional support agent leveraged data from computer activities (e.g., keyboard and mouse usage, desktop windows and browser activity), facial expressions from webcam, and behavioral and workplace demand signals from email and calendar, with no real way to differentiate personal activity from work activity. In addition, employees are increasingly being asked to use their personal devices for work with bring-your-own-device (BYOD) policies, making it even more difficult to separate the data streams [52]. Even if wellbeing sensing technologies are brought in as personal tools, rich information about your work context (e.g., your daily meeting schedule, your general affect in email communications) may be off limits to these technologies. This ambiguity complicates the process of integrating data streams, as the privacy, security, and confidentiality concerns associated with personal data may not apply uniformly across all types of data.

Furthermore, the ownership of data generated at work is often contested, with both employees and organizations claiming rights to this information. Our participants across the studies expressed that they want seamless integration of work and personal data across devices, seeking to get access and control the data generated about them. Participants in Study 2 sought control over the types or granularity of data and with whom to share that data. But, organizations often maintain data generated by their employees as a valuable asset that should be protected and utilized for their benefit. This conflict over data ownership can lead to tensions between employees and organizations, impeding the effective integration of data streams across contexts, which is paramount to holistic mental health support.

In addition to the legal and ethical complexities surrounding data ownership, there are practical challenges that may hinder the integration of data streams. For example, integrating data from different contexts may require sharing sensitive personal information across data security boundaries of various

[1]https://www.bls.gov/news.release/empsit.t18.htm#ces_table2.f.p

[2]https://news.gallup.com/poll/166553/less-recommended-amount-sleep.aspx

companies (e.g., Fitbit, Outlook, iOS), potentially exposing individuals to privacy risks. Strict data-sharing agreements and technical safeguards that meet the local and federal regulations must also be put in place, further complicating the process. The misalignment in data ownership between employees and organizations, combined with the complex legal, ethical, and practical challenges associated with integrating data streams across work and nonwork contexts, presents a significant obstacle to the realization of holistic mental health care through affective computing technologies.

C. Misalignment in Values and Incentives

The deployment of affective computing technologies for mental health support in the workplace necessitates an examination of the potential misalignment in values and incentives between individuals, managers, and organizations. While individuals may prioritize their personal wellbeing and mental health, arguments for workplace wellbeing programs are typically associated with increasing productivity and quality, reducing cost due to absenteeism and healthcare spending, and improving profitability [53]. For example, organizations may deploy wellbeing sensing tools with the ultimate goal of optimizing employee performance and efficiency. Measuring wellbeing primarily through the lens of productivity assumes that people "live to work" (Study 4), and that their mental health is in service to their work performance. However, individuals may not share this perspective, instead valuing a balance between work and personal life and seeking to use these tools to improve their overall wellbeing beyond the workplace. This tension between wellbeing and productivity was clearly observed in our studies where many Study 3 participants mentioned difficulties of incorporating self-care activities into their work contexts and expressed skepticism about integrating our tools into their professional lives, fearing it might lead to decreased productivity.

This divergence in values and incentives can lead to unhealthy behaviors and paradoxes due to the increased visibility of employee behaviors [19]. For example, employees may feel pressured to maintain high levels of productivity at the expense of their mental health, as the sensing tools make their performance more transparent to themselves, managers, and colleagues. This increased visibility makes social comparisons easier, leading to heightened fear of missing out and increased stress, burnout, and other negative mental health outcomes.

The potential misuse of data generated by sensing tools can exacerbate these tensions. The deployment of wellbeing sensing technologies in the workplace is often paid for by the organization. As one Study 4 participant pointed out, it may be that "organizations are incentivized to maximize profit to get more out of workers," and that organizations may want to recoup the cost of deploying such technologies somehow. Organizations may be tempted to use the information collected to rank employees, make promotion and compensation decisions, or even penalize those who do not meet specific productivity or wellbeing benchmarks [53]–[55]. Bringing your own personal wellbeing sensing tools can be challenging due to psychologically unsafe workplace culture [56]. The misalignments in values and incentives embedded in the deployment of these technologies can ultimately erode the trust between employees and their employers, as well as between colleagues, rendering these tools unhelpful in supporting holistic mental health.

D. Misalignment in Wellbeing Definition

The implementation of wellbeing sensing technologies in the workplace not only raises concerns about differing values and incentives but also highlights the challenges associated with defining and measuring wellbeing. Wellbeing is a multifaceted and highly individualistic concept, which makes it difficult to establish a one-size-fits-all definition or set of metrics that can accurately capture and address the unique needs and perspectives of an individual. In addition, emotion recognition technologies, commonly used in wellbeing applications, face challenges such as the evolving theory of human emotions, difficulties in labeling emotions, and lack of representative and generalizable data [27].

Across our studies, we found that stress is idiosyncratic, personalized stress models outperform generic ones and that participants desired to define what stress means for themselves. These findings suggest that understanding and addressing individual stressors and coping mechanisms may be crucial for optimizing wellbeing in the workplace. However, balancing the need for personalization with the desire to scale out solutions can be challenging, as tailoring technologies to individual needs requires additional effort and resources, even beyond development costs. For example, one Study 1 participant explained that they would be willing to provide feedback to the system to improve its stress detection algorithm, but they would stop using the system if they could not learn after 10 failures. Many of the Study 3 participants who used our just-in-time emotional support agent for four weeks found nudging to be disruptive, even though these nudges are necessary to fine-tune the system at the beginning.

Besides technical and user experience challenges of addressing the diversity of wellbeing definitions, the multiplicity of wellbeing definitions is further complicated in the organizational contexts, where a general notion of wellbeing is often necessary to make organization-wide decisions on policy (e.g., meeting-free Fridays). In an organizational context, the definition of wellbeing may be influenced by various stakeholders, including leaders, managers, analysts, human resources personnel, AI developers, system builders, and individual employees. The challenge lies in determining whose definition of wellbeing should be prioritized and how to ensure that the sensing technologies cater to diverse needs and preferences.

For instance, a manager may view working beyond the typical 9-to-5 schedule as unhealthy, while a working parent might consider taking a break to attend to family responsibilities and resuming work later in the day as a healthy way to balance work and life. A worker may temporarily override their wellbeing needs to be immersed in a work project that they are passionate about, but the organization may reassign that project to another employee to distribute the workload evenly for collective wellbeing sake. These differing perspectives can create tensions in deploying and using wellbeing sensing technologies, as employees may feel that their personal wellbeing strategies are being judged or misunderstood. When the wellbeing definitions of an individual misaligns with the wellbeing definitions encoded in the wellbeing sensing technologies by their organization or developers, one Study 4 participant expressed that they may feel "gaslit by this technology."

E. Misalignment in Power

Across our studies, we found that personal emotional support tools can empower individuals to be more aware of their emotions and regain control over their lives. Just-in-time capability especially helped with escalated situations to "take steps to actually change it back...before it goes into that whole other cycle (Study 3)." One Study 1 participant envisioned that it could give them a convenient cover: "Interrupt me before I get into a fight in the sense that oh wait a minute I got to take this phone call."

However, addressing mental health issues often requires more than just surface-level solutions or tools. Intervention strategies for workplace stress are commonly grouped into three categories: primary, secondary, and tertiary [57]–[60]. Primary strategies involve organizational-level changes (i.e., a culture shift) [57]. Secondary strategies, which are the most common, target individuals experiencing stress and aim to detect and reduce their stress. Tertiary prevention typically involves Employee Assistance Programs (EAPs). Although primary strategies may be necessary for long-term benefits, because they are challenging to implement, the focus often shifts towards secondary strategies [60].

Wellbeing sensing technologies fall in this secondary category, with the promise to empower individuals to manage their stress responses. Employees might feel empowered by using wellbeing sensing technologies to gain insights into their mental health and wellbeing. As one Study 4 participant said, this knowledge could enable them to speak up and advocate for change: "I think half the battle is bringing up the issue." Unfortunately, some of our Study 4 participants also deemed them as "putting band-aids" and "palliative care to make people feel better about the situation." That is because, while these wellbeing sensing technologies can help alleviate some symptoms, they might not be sufficient to address the deeper, systematic issues that contribute to stress and other mental health challenges. Meaningful change often comes from shifting organizational culture, priorities, and processes, which requires a concerted effort from both employees and management.

One aspect to consider is the power dynamics in the workplace. Despite the benefits of these wellbeing tools, the existing power structures at work might make it difficult for employees to trust the organization with their personal information or to believe that their concerns will be taken seriously. Furthermore, the same data collected by these technologies can potentially be used to maintain the status quo and perpetuate unhealthy work cultures, ultimately harming employee wellbeing. If the system is deployed and the data is collected by the employer, it is crucial to examine whether the consent was obtained meaningfully given the power asymmetry [27], [61].

In cases where wellbeing sensing technologies are mandated by employers, they can be perceived as intrusive and undermining individual autonomy. This can lead to a sense of disempowerment among employees, who may feel that they have no control over their own wellbeing data or the decisions that affect their mental health. On the other hand, employees who bring their own technologies to work may encounter difficulties in seamlessly integrating these tools with work data, further contributing to feelings of powerlessness. Therefore, introducing affective computing and sensing technologies into the work context, whether by employers or by employees, may work against mental health goals due to an imbalance in power and autonomy.

IV. Discussion and Conclusion

Our case study on the deployment of a just-in-time emotional support agent in the workplace highlights challenges of respecting the appropriate and individualized boundary preferences, giving the rights to the data generated about them, aligning the system design and organizational goals and wellbeing definitions with their own personal values and wellbeing definitions, and being sensitive to the power dynamics across work and nonwork contexts. We argue that designing and implementing affective computing systems for holistic mental health support and deploying them across work and nonwork contexts must be coupled with careful consideration of contextual boundaries and the various dimensions of misalignments presented in this paper.

To navigate the complex landscape of contextual boundaries and address the challenges associated with misalignments, a multifaceted approach that considers the perspectives and needs of various stakeholders, including employees, managers, and organizations, is necessary. This involves deliberating with multiple stakeholders about the appropriate notions of "wellbeing" that should be targeted through data and sensing technology. It is important to establish a sustainable pipeline and process for incorporating these diverse perspectives through policy and regulation in centralized way to address the social, conceptual, and ethical concerns that arise when crossing contextual boundaries. Such process should also allow tailored and decentralized approach to cater to the needs of individuals and teams.

Moreover, it is crucial to consider how technology can shift the responsibility for wellbeing challenges across individuals and organizations and the allocation of power to make meaningful changes at individual or organizational levels. By reflecting on who is responsible for worker wellbeing, who is doing what work to support worker wellbeing, and the gains and losses that come with quantification, we can better align the deployment of affective computing technologies with the needs and values of stakeholders. Providing individuals with the power to freely decide whether, when, and with whom they want to share wellbeing data is an important step towards addressing concerns about privacy and autonomy. This necessitates the development of processes that give appropriate ownership of data and that provide meaningful consent for using the wellbeing technology and for collecting data that respects individual choices and preferences.

Finally, to provide holistic mental health support, affective computing technologies should be integrated with existing support systems and complemented by organizational and cultural changes that foster a healthy work environment. This requires a shift in focus from solely addressing individual stress responses to tackling systemic issues that contribute to mental health challenges. By considering these dimensions of misalignments and working towards a more holistic, person-centered approach to mental health support, affective computing technologies have the potential to significantly improve the wellbeing of individuals across work and nonwork contexts.

In conclusion, affective computing technologies are promising avenues for supporting mental health and wellbeing across various contexts, but their successful deployment requires careful consideration of the contextual boundaries and the numerous challenges that arise due to misalignments in boundary preferences, data ownership, values and incentives, wellbeing definitions, and power dynamics. By addressing these misalignments and striving for a holistic approach to mental health support that incorporates the perspectives of multiple stakeholders, future research in affective computing can help bridge the gap between digital mental health solutions and the diverse needs of individuals.

Ethical Impact Statement

The deployment of affective computing technologies for mental health support raises ethical concerns, such as misalignments in boundary preferences, data ownership, and power dynamics. Addressing these concerns requires a multifaceted approach considering various stakeholders' perspectives and interdisciplinary research. Ensuring meaningful stakeholder involvement, privacy, autonomy, and informed consent is crucial for mitigating potential harms. Integrating these technologies with existing mental health support systems necessitates addressing systemic issues and fostering organizational and cultural changes that align with stakeholders' needs and values. By acknowledging study limitations and considering ethical concerns, affective computing research can contribute to the development of ethically responsible digital mental health solutions catering to diverse needs.

References

[1] B. G. Druss, T. Bornemann, Y. W. Fry-Johnson, H. G. McCombs, R. M. Politzer, and G. Rust, "Trends in mental health and substance abuse services at the nation's community health centers: 1998–2003," *American Journal of Public Health*, vol. 98, no. Supplement_1, pp. S126–S131, 2008.

[2] J. P. Jameson and M. B. Blank, "Diagnosis and treatment of depression and anxiety in rural and nonrural primary care: National survey results," *Psychiatric Services*, vol. 61, no. 6, pp. 624–627, 2010.

[3] R. Mojtabai, M. Olfson, N. a. Sampson, B. Druss, P. S. Wang, K. B. Wells, H. a. Pincus, and R. C. Kessler, "Barriers to mental health treatment: results from the WHO World Mental Health surveys." *Psychological Medicine*, vol. 41, no. 8, pp. 1751–1761, 2011.

[4] K. Woodward, E. Kanjo, D. J. Brown, T. M. McGinnity, B. Inkster, D. J. Macintyre, and A. Tsanas, "Beyond mobile apps: a survey of technologies for mental well-being," *IEEE Transactions on Affective Computing*, vol. 13, no. 3, pp. 1216–1235, 2020.

[5] A. Thieme, D. Belgrave, and G. Doherty, "Machine learning in mental health: A systematic review of the hci literature to support the development of effective and implementable ml systems," *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 27, no. 5, pp. 1–53, 2020.

[6] M. Neary and S. M. Schueller, "State of the field of mental health apps," *Cognitive and Behavioral Practice*, vol. 25, no. 4, pp. 531–537, 2018.

[7] J. Borghouts, E. Eikey, G. Mark, C. De Leon, S. M. Schueller, M. Schneider, N. Stadnick, K. Zheng, D. Mukamel, D. H. Sorkin *et al.*, "Barriers to and facilitators of user engagement with digital mental health interventions: systematic review," *Journal of medical Internet research*, vol. 23, no. 3, p. e24387, 2021.

[8] A. Sano, P. Johns, and M. Czerwinski, "Designing opportune stress intervention delivery timing using multi-modal data," in *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2017, pp. 346–353.

[9] Y. S. Can, B. Arnrich, and C. Ersoy, "Stress detection in daily life scenarios using smart phones and wearable sensors: A survey," *Journal of biomedical informatics*, vol. 92, p. 103139, 2019.

[10] A. Alberdi, A. Aztiria, and A. Basarab, "Towards an automatic early stress recognition system for office environments based on multimodal measurements: A review," *Journal of biomedical informatics*, vol. 59, pp. 49–75, 2016.

[11] L. M. Vizer, L. Zhou, and A. Sears, "Automated stress detection using keystroke and linguistic features: An exploratory study," *International Journal of Human-Computer Studies*, vol. 67, no. 10, pp. 870–886, 2009.

[12] K. Saha, T. Grover, S. M. Mattingly, V. Das Swain, P. Gupta, G. J. Martinez, P. Robles-Granda, G. Mark, A. Striegel, and M. De Choudhury, "Person-centered predictions of psychological constructs with social media contextualized by multimodal sensing," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 5, no. 1, pp. 1–32, 2021.

[13] S. Mirjafari, K. Masaba, T. Grover, W. Wang, P. Audia, A. T. Campbell, N. V. Chawla, V. D. Swain, M. D. Choudhury, A. K. Dey, and et al., "Differentiating higher and lower job performers in the workplace using mobile sensing," *Proc. ACM IMWUT*, 2019.

[14] O. Rudovic, J. Lee, M. Dai, B. Schuller, and R. W. Picard, "Personalized machine learning for robot perception of affect and engagement in autism therapy," *Science Robotics*, vol. 3, no. 19, 2018.

[15] N. Mead and P. Bower, "Patient-centredness: a conceptual framework and review of the empirical literature," *Social science & medicine*, vol. 51, no. 7, pp. 1087–1110, 2000.

[16] J. Yang, O. Arif, P. A. Vela, J. Teizer, and Z. Shi, "Tracking multiple workers on construction sites using video cameras," *Advanced Engineering Informatics*, vol. 24, no. 4, pp. 428–434, 2010.

[17] J. Muckell, Y. Young, and M. Leventhal, "A wearable motion tracking system to reduce direct care worker injuries: An exploratory study," in *Proceedings of the 2017 International Conference on Digital Health*, 2017, pp. 202–206.

[18] P. Robles-Granda, S. Lin, X. Wu, G. J. Martinez, S. M. Mattingly, E. Moskal, A. Striegel, N. V. Chawla, S. D'Mello, J. Gregg *et al.*, "Jointly predicting job performance, personality, cognitive ability, affect, and well-being," *IEEE Computational Intelligence Magazine*, vol. 16, no. 2, pp. 46–61, 2021.

[19] P. M. Leonardi and J. W. Treem, "Behavioral visibility: A new paradigm for organization studies in the age of digitization, digitalization, and datafication," *Organization Studies*, vol. 41, no. 12, pp. 1601–1625, 2020.

[20] K. Zickuhr, "Workplace surveillance is becoming the new normal for us workers," *Washington Center for Equitable Growth. https://equitablegrowth. org/research-paper/workplace-surveillance-is-becomingthe-new-normal-for-us-workers/. Institute for Research on Labor and Employment University of California, Berkeley*, vol. 2521, pp. 94 720–5555, 2021.

[21] M. Finnegan, "The new normal: When work-from-home means the boss is watching," https://www.computerworld.com/article/3586616/the-new-normal-when-work-from-home-means-the-boss-is-watching.html, 2020, accessed: 2022-09-07.

[22] M. Barbaro, "The rise of workplace surveillance: Is your productivity being electronically monitored by your bosses?" https://www.nytimes.com/2022/08/24/podcasts/the-daily/workplace-surveillance-productivity-tracking.html?, 2022, accessed: 2022-09-07.

[23] V. Das Swain, K. Saha, G. D. Abowd, and M. De Choudhury, "Social media and ubiquitous technologies for remote worker wellbeing and productivity in a post-pandemic world," in *2020 IEEE Second International Conference on Cognitive Machine Intelligence (CogMI)*. IEEE, 2020, pp. 121–130.

[24] A. Bleakley, D. Rough, J. Edwards, P. Doyle, O. Dumbleton, L. Clark, S. Rintel, V. Wade, and B. R. Cowan, "Bridging social distance during social distancing: exploring social talk and remote collegiality in video conferencing," *Human–Computer Interaction*, vol. 37, no. 5, pp. 404–432, 2022.

[25] D. Russell, C. Neustaedter, J. Tang, T. Judge, and G. Olson, "Videoconferencing in the age of covid: How well has it worked out?" in *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–2.

[26] C.-L. Yang, N. Yamashita, H. Kuzuoka, H.-C. Wang, and E. Foong, "Distance matters to weak ties: Exploring how workers perceive their strongly-and weakly-connected collaborators in remote workplaces," *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. GROUP, pp. 1–26, 2022.

[27] J. Hernandez, J. Lovejoy, D. McDuff, J. Suh, T. O'Brien, A. Sethumadhavan, G. Greene, R. Picard, and M. Czerwinski, "Guidelines for assessing and minimizing risks of emotion recognition applications," in *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2021, pp. 1–8.

[28] D. C. Ong, "An ethical framework for guiding the development of affectively-aware artificial intelligence," in *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2021, pp. 1–8.

[29] H. Kaur, D. McDuff, A. C. Williams, J. Teevan, and S. T. Iqbal, ""i didn't know i looked angry": Characterizing observed emotion and reported affect at work," in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–18.

[30] V. Das Swain, L. Gao, W. A. Wood, S. C. Matli, G. D. Abowd, and M. De Choudhury, "Algorithmic power or punishment: Information worker perspectives on passive sensing enabled ai phenotyping of performance and wellbeing," in *CHI Conference on Human Factors in Computing Systems*, 2023.

[31] S. Corvite, K. Roemmich, T. Rosenberg, and N. Andalibi, "Data subjects' perspectives on emotion artificial intelligence use in the workplace: A relational ethics lens," *Proceedings of the ACM on Human-Computer Interaction*, 2022.

[32] D. A. Adler, E. Tseng, K. C. Moon, J. Q. Young, J. M. Kane, E. Moss, D. C. Mohr, and T. Choudhury, "Burnout and the quantified workplace: Tensions around personal sensing interventions for stress in resident physicians," 2022.

[33] A. Kawakami, S. Chowdhary, S. T. Iqbal, Q. V. Liao, A. Olteanu, J. Suh, and K. Saha, "Sensing wellbeing in the workplace, why and for whom? envisioning impacts with organizational stakeholders," *Proceedings of the ACM on Human-Computer Interaction (CSCW)*, 2023.

[34] M. B. Morshed, J. Hernandez, D. McDuff, J. Suh, E. Howe, K. Rowan, M. Abdin, G. Ramos, T. Tran, and M. Czerwinski, "Advancing the understanding and measurement of workplace stress in remote information workers from passive sensors and behavioral data," in *2022 10th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2022, pp. 1–8.

[35] X. Liu, J. Fromm, S. Patel, and D. McDuff, "Multi-task temporal shift attention networks for on-device contactless vitals measurement," *NeurIPS*, 2020.

[36] E. Howe, J. Suh, M. Bin Morshed, D. McDuff, K. Rowan, J. Hernandez, M. I. Abdin, G. Ramos, T. Tran, and M. P. Czerwinski, "Design of digital workplace stress-reduction intervention systems: Effects of intervention type and timing," in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–16.

[37] G. Mark, S. T. Iqbal, M. Czerwinski, P. Johns, A. Sano, and Y. Lutchyn, "Email duration, batching and self-interruption: Patterns of email use on productivity and stress," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 2016, pp. 1717–1728.

[38] G. Mark, S. Iqbal, and M. Czerwinski, "How blocking distractions affects workplace focus and productivity," in *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*, 2017, pp. 928–934.

[39] D. McDuff, E. Jun, K. Rowan, and M. Czerwinski, "Longitudinal observational evidence of the impact of emotion regulation strategies on affective expression," *IEEE Transactions on Affective Computing*, 2019.

[40] N. Hjortskov, D. Rissén, A. K. Blangsted, N. Fallentin, U. Lundberg, and K. Søgaard, "The effect of mental stress on heart rate variability and blood pressure during computer work," *European journal of applied physiology*, vol. 92, no. 1, pp. 84–89, 2004.

[41] T. D. Allen, E. Cho, and L. L. Meier, "Work–family boundary dynamics," *Annu. Rev. Organ. Psychol. Organ. Behav.*, vol. 1, no. 1, pp. 99–121, 2014.

[42] E. E. Kossek, M. N. Ruderman, P. W. Braddy, and K. M. Hannum, "Work–nonwork boundary management profiles: A person-centered approach," *Journal of Vocational Behavior*, vol. 81, no. 1, pp. 112–128, 2012.

[43] J. B. Olson-Buchanan and W. R. Boswell, "Blurring boundaries: Correlates of integration and segmentation between work and nonwork," *Journal of Vocational behavior*, vol. 68, no. 3, pp. 432–445, 2006.

[44] A. Foucreault, A. Ollier-Malaterre, and J. Ménard, "Organizational culture and work–life integration: A barrier to employees' respite?" *The International Journal of Human Resource Management*, vol. 29, no. 16, pp. 2378–2398, 2018.

[45] Y. Bogaerts, R. De Cooman, and S. De Gieter, "Getting the work-nonwork interface you are looking for: The relevance of work-nonwork boundary management fit," *Frontiers in Psychology*, vol. 9, p. 1158, 2018.

[46] J. G. Grzywacz, D. M. Almeida, and D. A. McDonald, "Work–family spillover and daily reports of work and family stress in the adult labor force," *Family relations*, vol. 51, no. 1, pp. 28–36, 2002.

[47] E. E. Kossek, T. L. Dumas, M. M. Piszczek, and T. D. Allen, "Pushing the boundaries: A qualitative study of how stem women adapted to disrupted work–nonwork boundaries during the covid-19 pandemic." *Journal of Applied Psychology*, vol. 106, no. 11, p. 1615, 2021.

[48] R. Catalano, *Health, behavior and the community: An ecological perspective*. Pergamon Press New York, 1979.

[49] M. Rucker, "The interesting history of workplace wellness," *Retrieved from*, 2016.

[50] R. P. Maiden, "Employee assistance program evaluation in a federal government agency," *Employee Assistance Quarterly*, vol. 3, no. 3-4, pp. 191–203, 1988.

[51] R. Wang, F. Chen, Z. Chen, T. Li, G. Harari, S. Tignor, X. Zhou, D. Ben-Zeev, and A. T. Campbell, "Studentlife: assessing mental health, academic performance and behavioral trends of college students using smartphones," in *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*, 2014, pp. 3–14.

[52] A. Scarfo, "New security perspectives around byod," in *2012 Seventh International Conference on Broadband, Wireless Computing, Communication and Applications*. IEEE, 2012, pp. 446–451.

[53] I. Ajunwa, K. Crawford, and J. Schultz, "Limitless worker surveillance," *California Law Review*, 2017.

[54] A. Rosenblat, T. Kneese, and D. Boyd, "Workplace surveillance," *Open Society Foundations' Future of Work Commissioned Research Papers*, 2014.

[55] L. D. Introna, "Workplace surveillance, privacy and distributive justice," *ACM SIGCAS Computers and Society*, vol. 30, no. 4, pp. 33–39, 2000.

[56] M. F. Dollard, *Psychosocial safety climate: A lead indicator of workplace psychological health and engagement and a precursor to intervention success*. Routledge, 2012.

[57] C. L. Cooper and S. Cartwright, "An intervention strategy for workplace stress," *Journal of psychosomatic research*, vol. 43, no. 1, pp. 7–16, 1997.

[58] S. D. Sidle, "Workplace stress management interventions: What works best?" *Academy of Management Perspectives*, vol. 22, no. 3, pp. 111–112, 2008.

[59] S. Reynolds, "Psychological well-being at work: is prevention better than cure?" *Journal of psychosomatic research*, vol. 43, no. 1, pp. 93–102, 1997.

[60] K. M. Richardson and H. R. Rothstein, "Effects of occupational stress management intervention programs: a meta-analysis." *Journal of occupational health psychology*, vol. 13, no. 1, p. 69, 2008.

[61] S. Chowdhary, A. Kawakami, M. L. Gray, J. Suh, A. Olteanu, and K. Saha, "Can workers meaningfully consent to workplace wellbeing technologies?" *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT)*, 2023.