Advances in Psychiatry and Behavioral Health 2 (2022) 95–105
ADVANCES IN PSYCHIATRY AND BEHAVIORAL HEALTH

Affective Computing in Psychotherapy

Rahul Khanna, MBBS, FRANZCP[a,*], Nicole Robinson, PhD[b], Meaghan O'Donnell, PhD[a], Harris Eyre, MBBS, PhD[c], Erin Smith[d]

[a]Phoenix Australia, Department of Psychiatry, University of Melbourne, 161 Barry Street, Carlton, VIC 3053, Australia; [b]Electrical and Computer Systems Engineering and Turner Institute for Brain and Mental Health, Monash University, 18 Alliance Lane, Clayton Campus, Victoria 3800, Australia; [c]Brain Capital Alliance, San Francisco, CA, USA; [d]Atlantic Fellow for Equity in Brain Health at the Global Brain Health Institute at University of California, 548 E Crescent Drive, Palo Alto, CA 94301, USA

KEYWORDS

• Digital mental health • Affective computing • Telehealth • Artificial intelligence • Machine learning • Digital psychotherapy

KEY POINTS

- Psychotherapy is evolving away from its roots as an individual, face-to-face interaction to cope with workforce shortages in the face of overwhelming demand and a destabilizing pandemic.
- Affective computing is the branch of computer science focused on detecting, interpreting, processing, and simulating human emotion. It has the potential to overcome the limitations of contemporary approaches to digital therapy.
- Advances in affective computing and adjacent technologies can enhance our understanding and delivery of psychotherapy, though several practical and ethical barriers remain.
- Effective partnerships between system developers and mental health clinicians and consumers will be crucial to widespread clinical adoption.

BACKGROUND

Emerging Trends in Psychotherapy

Psychotherapy, colloquially called the *talking cure*, is an approach to mental health treatment that traditionally used the relationship between a therapist and a client to impart skills or insight to relieve suffering. Psychotherapeutic approaches differ on several dimensions, including by (individual, family, or group), intended aim (personal insight vs direct behavioral change), the centrality of relational elements, duration, and more. Common themes, however, include attention to emotion, thoughts, and behavior. Although formative concepts of psychotherapy trace their roots much further in human history, the psychoanalytic approach pioneered by Freud in the early twentieth century is commonly accepted as the birth of modern psychotherapy [1]. Contemporary approaches represent an evolution or reaction to analysis.

Despite this divergence in foci and content, for most of its history, psychotherapy has primarily been carried out as a face-to-face interaction between a client and a professional mental health care provider. Increasingly, however, the talking cure has evolved into a *communication cure*, representing a broadening of the definition of therapist interaction and the notion of the "therapeutic space." The modality of communication has become increasingly diverse, ranging from face-to-face to telehealth and even text messaging or live chat. The COVID-19 pandemic also has resulted in a global surge in telehealth and other psychotherapy innovations by both increasing demand and breaking prior service models

*Corresponding author, *E-mail address:* Rahul.khanna@unimelb.edu.au

https://doi.org/10.1016/j.ypsc.2022.05.006
2667-3827/22/

in the interest of infection control [2]. Emerging technology-first providers such as Ginger.io, which recently merged with Headspace, use a stepped-care model, whereby a spectrum of modalities is offered based on the need. These evolving modalities have also made it possible for therapy to occur asynchronously. In this approach, the interaction with the therapist can occur asynchronously. The client may communicate their thoughts to a clinician via text, email, or video at any hour, and the therapist would respond when available. Although much less widely used than synchronous telehealth, adoption is growing [3]. Although these approaches bring new advantages to the treatment process, they have been principally motivated by a shortage of qualified therapists. For instance, in the United States, in March 2021, 37% of the population lived in an area experiencing mental health professional shortages. This rises to over 94% in more sparsely populated states [4].

These new modalities increase flexibility and, therefore, access to psychotherapy but have limitations. One aspect of traditional psychotherapy and mental health clinical practice more broadly is the mental state examination (MSE), which is used alongside other clinical assessments and therapeutic modalities during psychotherapy. This examination reflects the clinicians' systematic observation of a client's presentation, including their overt behavior, therapeutic engagement, speech and language use, the coherence and goal-directedness of their thought process, reported mood, visible effect, insight, and judgment. However, an inexact skill, the ability to conduct an MSE is considered a core clinical competency that informs diagnoses and treatment planning. As practice moves further from its face-to-face individual treatment roots, conducting a thorough MSE becomes increasingly challenging as it depends greatly on nonverbal cues lost in videoconferencing. This is one of several spaces where affective computing is demonstrating early promise.

Introducing Affective Computing

Rosalind Picard, widely credited with founding affective computing, defined *affective computing* as "computing that relates to, arises from, and deliberately influences emotion" [5]. Today, it is the discipline of computer science focused on detecting, interpreting, processing, and simulating human emotions (Fig. 1).

The field has several motivations, ranging from improved computer decision-making to enhancing human–computer interactions. As Picard later summarized, emotion "…kept coming up as essential in perception. Emotion biased what we saw and heard. Emotion played major roles not only in perception but also in many other aspects of intelligence that artificial intelligence (AI) researchers had been trying to solve from a cortical-centric perspective. Emotion was vital in forming memory and attention and in rational decision-making. And, of course, emotional communication was vital in human–machine interaction. Emotions influence action selection, language, and whether or not you decide to double-check your mathematical derivations, comment your computer code [or] initiate a conversation." Pertinently for the topic of psychotherapy, Picard reiterates that emotion is a principal "motivating and guiding force in perception and attention" [6].

These diverse aims of affective computing are usually enabled through the combination of data and machine learning. Traditionally, the field focused on interpreting visual and audio data. Examples include using computer vision to interpret facial expressions or audio recordings to infer emotional valence through features such as vocal tone, speed, pitch, prosody, and so forth. Such data points are captured, and a machine learning algorithm is used to "train" software to estimate the probability of a given emotion based on an initial set of previously human-labeled data.

Although speech and facial expressions were the most common data sources of algorithmic training in the early decades of affective computing, in principle any computationally measurable manifestation of an emotional state can be used (see Fig. 1). For example, a clinical bot titled SimSensei uses psychophysiological measures such as heart rate and galvanic skin response to augment vocal features, speech content, and facial expression [7].

Psychophysiological signals have been a mainstay of psychology research since the 1950s but previously had considerable limitations. Chief among them has been the inconvenience, expense, and biasing nature of contact sensors required to capture such data. Once strapped onto a participant's body, they could collect several parameters of the autonomic nervous system, such as temperature, heart rate, and vascular tone. Recent advances in signal processing, however, have led to the emergence of computational psychophysiology techniques. These involve the use of imaging, such as those collected from standard video and thermal cameras, to deduce psychophysiological status through machine learning [8]. Such methods add a rich array of features that can support affective computing.

Affective computing has been improved by the increasing capability of mobile devices, which now

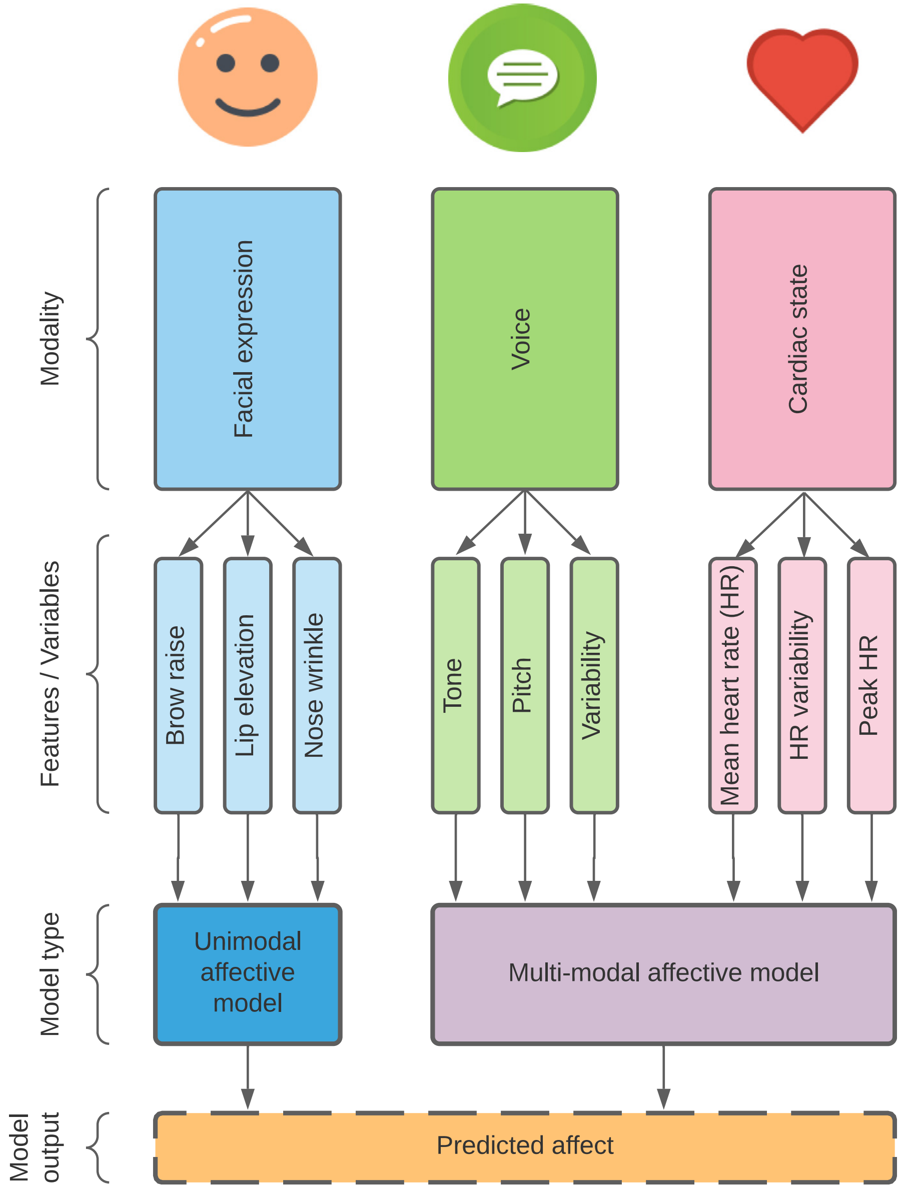

FIG. 1 Affective computing systems commonly consist of affective modalities, derived features, and machine learning models that predict affect.

accompany billions of people throughout the day. Now nearly ubiquitous, smartphones unobtrusively collect rich sensor data spanning multiple modalities, including location, proximity to other devices (and therefore those who carry them), communication behavior, and more. Mobile devices are also capable of increasingly complex computing on their own, reducing the inconvenience and privacy issues associated with the cloud computing previously necessary. Politou and *colleagues* (2017) identified 131 research papers using mobile devices to infer effect [9]. These ranged from unidimensional models like stress levels to more nuanced classifiers capable of distinguishing more distinctive affective states.

Modality, in this context, refers to the data stream(s) used to create the affective computing machine learning software. A unimodal model uses a single source of data, such as voice, to infer effect. A multimodal model can leverage several signal streams at once, as illustrated in Fig. 1. Individual data inputs from a modality that are used by the machine learning algorithm are termed *features*. This approach may make the model more accurate and robust to differences between the training data and the messier world of clinical practice [10]. However, this also adds to the complexity of the model and requires greater volumes of (sometimes intrusive) data and computational power. The cost of acquiring the data and implementing the insights therefore increases. In addition, for unimodal and multimodal approaches, data collection can include passive and active tests that can occur during psychotherapy or in-between sessions. In the context of affective computing for psychotherapy, passive data collective can include information collected from smartphone usage or social media activity, whereas active data collection may involve specific vocal or facial expression analyses during in-person or telemedicine psychotherapy sessions.

CURRENT STATE OF AFFECTIVE COMPUTING APPLIED TO THERAPY

Affective computing is, therefore, a powerful and rapidly evolving discipline. It has the potential to contribute to the full lifecycle of psychotherapy, from advancing our mechanistic understanding of therapy to training clinicians and supporting human-led psychotherapy through autonomous care. A selection of such work follows.

Exploring Underlying Mechanisms

Psychotherapy is a challenge to empirically study. The number of potential variables affecting process and outcome, the complexity of each therapy session, and the subjectivity of human assessors are some of the challenges for systematic comparison and analysis. Affective computer can help advance human's mechanistic understanding of psychotherapy by allowing them to measure affective states more systematically. Such studies featured prominently in Wilshire's (2020) systematic review [11] on interpersonal coordination dynamics in psychotherapy. They note evidence indicating that the affective states of clinicians and their clients covary within a psychotherapeutic interaction. Theoretic constructs, such as the common factors model and polyvagal theory, suggest such covariance may moderate effectiveness. Although some included studies used human-coded data, affective computing is used in an increasing number of studies reviewed. Reviewed studies demonstrated that coordination between the therapist and the client in features such as body movements, physiology, language, and speech patterns can predict outcomes, including symptom improvement, empathy, and therapeutic alliance.

As the investigators note [11], the use of computational methods for such analyses not only allows for the analysis of more longitudinal data but also opens opportunities for automated feedback to improve such coordination, a theme discussed in the following section.

Supporting Human-Delivered Therapy

Affective computing can also support the client's journey of human-delivered therapy, from screening and diagnosis through actual therapeutic content delivery, progress monitoring, and more. These systems have been used to try to screen people who may benefit from therapy (eg, through the affective computing methods to analyze social media posts for the evidence of major depressive disorder or suicide risk [12,13]). Virtual agents like SimSensei have been used to conduct diagnostic assessments, with some suggestions that they may encourage more honest self-disclosure than human assessors [14].

Then, there is a range of work seeking to map digital behavior to emotional states, particularly those of clinical relevance, such as depression and mania. Signals used for such affective systems range from vocal features, as previously described, to more abstract ones like typing behavior. These systems could be used for diagnosis or progress monitoring. Such approaches could enhance measurement-based care (MBC) that is developing an evidence base for enhancing rates and time to recovery in multiple disorders [15]. In traditional MBC, self-report scales are completed on a regular basis to guide therapy. For example, in a protocol developed by Nixon and Elizabeth [16], participants completing a manualized psychotherapy for post-traumatic stress disorder complete a weekly symptom questionnaire. Those who are not showing decreases in questionnaire scores by the fifth session have a reformulation session, whereby the biopsychosocial barriers to recovery are collaboratively reconsidered. Based on this reformulation, the therapist will deviate from the manualized treatment to address barriers (such as substance use) before returning to the core therapy. Replacing or augmenting self-reported measures with affective computing would add a more objective layer to this. Another advantage of such an approach is that it can occur within a client's own environment rather than in the artificial setting of a clinic. This allows for a more naturalistic understanding of the individual's experiential rather than reporting self, which the authors know do not always align [17]. Fig. 2 illustrates an affective computing-based dashboard.

An example of work in this vein comes from Ieso, an online cognitive behavioral therapy (CBT) provider in the United Kingdom. As a large provider of text-based therapy, they were able to leverage a deep learning algorithm to analyze over 90,000 psychotherapy session transcripts to understand the relationship between patient language, therapy content, and outcomes [18,19]. These insights were then distilled into a decision support tool that provides live feedback on the therapist's work within a session, nudging them toward techniques shown to improve outcomes [20]. Using these systems, Ieso was able to achieve recovery rates of 12% and 15% more than the national average in 2021 for depression and generalized anxiety disorder, respectively [21].

Although Ieso's model was limited to text data, the broader concept of live therapist guidance has enormous potential when paired with the latest in affective computing. Although in its infancy, the expansion of teletherapy prompted by the COVID-19 pandemic presents another opportunity to collect the large volumes of data required to build similar models using a broader range of affective signals. Recordings of video therapy sessions could yield not just transcripts but vocal features, facial expressions, and computationally derived physiologic data as well.

Such affective systems could also aid in the training of future psychotherapists. For example, Wilcocks and colleagues demonstrated a psychotherapy simulator to facilitate training in intensive short-term dynamic psychotherapy, helping to train clinicians in a cost-effective way using simulated interactions with the clients. This virtual patient was able to respond to dialogue and display affects to allow trainees to begin practicing skills in a safe environment [22]. Such tools could be paired with Ieso-like clinical guidance systems to present unparalleled training opportunities. Innovative digital training experiences like this would be particularly important in low-income and middle-income countries, where they can help nonspecialists rapidly develop basic psychotherapeutic skills [23].

Delivering Autonomous Therapy

Autonomous therapy systems are stand-alone interventions like chatbots that aim to deliver care without human interaction. They typically use text, audio, and video to teach self-management skills and encourage healthy thought patterns based on CBT principles. By allowing more organic interaction, they present an improvement on online course-like approaches, which linearly deliver therapeutic content. Contemporary systems like *Wysa* or *Woebot* store some persistent information about the user, allowing more tailored therapy delivery and facilitating a sense of connection. In fact, a recent analysis of *Woebot* data showed scores on a therapeutic alliance inventory equal to that of human therapists [24].

Till recently, studies on such tools consistently show that outcomes are enhanced when such programs and paired with a human, even a lay coach or peer, without

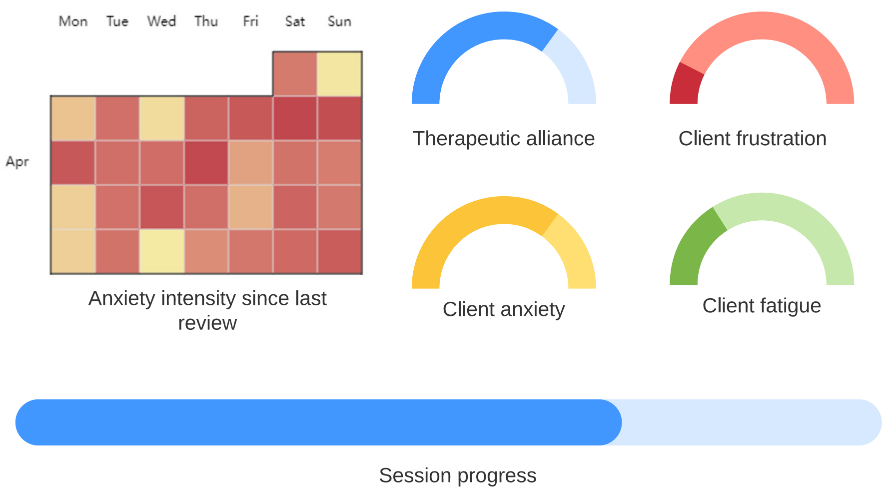

FIG. 2 An affective computing-based dashboard that could support human-delivered therapy.

which both engagement and efficacy markedly drop. Successful therapy depends on a process of learning, which is inherently an emotional process. *Feelings* of curiosity that often trigger learning are reenforced by the pleasure of mastery and sometimes are inhibited by feelings of frustration. A human therapist, coach, or peer can detect such shifts and respond appropriately to keep the client engaged. More mature use of affective computing has the potential to detect and use this nuance to narrow the gap between human-aided and autonomous therapy.

CLINICAL INTEGRATION

Despite the exciting emerging evidence discussed to date, truly clinically integrated solutions such as Ieso's are rare. The strong demand for mental health services and the volume of research and investment activities, however, suggests that they will transform care in the next decade. Ieso themselves raised $53 million in their latest funding round, and Forbes estimated that venture capital funding for United States-based mental health start-ups in 2020 alone exceeded $1.5 billion [25].

Therapists as Key Partners

Ultimately, the successful integration of affective computing into psychotherapy practice will depend on effective partnerships between clinicians and technologists. Most importantly, therapists are best placed to understand and identify improvement opportunities in clinical workflows. This is crucial for an appropriate prioritization of research and investment. Fig. 3 shows an example of workflows that may benefit from affective computing-based applications. Although such outlines seem basic, there will be many nuances between disciplines and therapy types that need broader engagement. Beyond this, there are several roles clinical experts can play in the future of this technology.

High-Quality Data and Expert Evaluation

Another key enabler is the labeled data that forms the backbone of any affective computing system. Such systems perform best when able to "learn" from a large corpus of diverse, high-quality data sets. For clinical applications, this ideally means sensitive recordings of therapy content, of which clinicians and their clients themselves are best positioned as joint custodians. Although clients should be the final arbiter of the data they are comfortable with sharing for system training, this is the best done in dialogue with a clinician who is able to convey the value of their data for improving care.

Further, the existing corpus of data used to train current systems has mostly been coded by the laypeople. That is, nonexperts have provided the human labels for what affect an individual was displaying at the time of data collection. A study that is an exception to this theme illustrates how problematic this can be. Liu and *colleagues* built a robot-assisted rehabilitation therapy for children with autism spectrum disorder [26]. The system was trained to recognize anxiety, engagement, and liking using a mixture of self-reports by children and labeling of videos by parents and clinicians. Interestingly, depending on the pairing (parent, clinician vs child-participant) and the emotional target, the agreement between reporters (measured by kappa score) ranged from 0.15 to 0.63. Given that 1 represents perfect agreement and 0 is the level expected by chance, the study is a reminder that the views of clinically trained and lay observers in effect may differ widely.

It would be impractical to have all such data sets labeled by expert clinicians, but it is important for systems to be able to learn "on the job." That is, systems must be designed to receive and improve based on expert feedback on performance. In this way, affective clinical systems can continuously improve alongside the clinicians that use them.

Appropriate Incentives

The above points around clinician input and improved data will be difficult to achieve without rethinking incentives. Models with the maturity of Ieso's remain rare, and as such early-adopting clinicians of other affective systems may experience inconveniences not immediately offset by benefits. Researchers and enterprises developing and leveraging such tools must, therefore, think creatively about how to appropriately incentivize and remunerate such activities.

Health insurers, governments, and developers of affective clinical systems all stand to gain from scalable, effective clinical innovation. Mental ill-health is associated with reduced productivity and increased physical health costs, on top of the individual human toll and direct health care expenses. All three parties should consider ways to encourage clinicians and patients to adopt novel affective tools, financially and otherwise.

The question of data ownership, and the role of the patients themselves, also needs active thought. Huag wrote a poignant reminder that the clients' perspectives have often been neglected in debates about clinical data ownership, which generally occur between and among

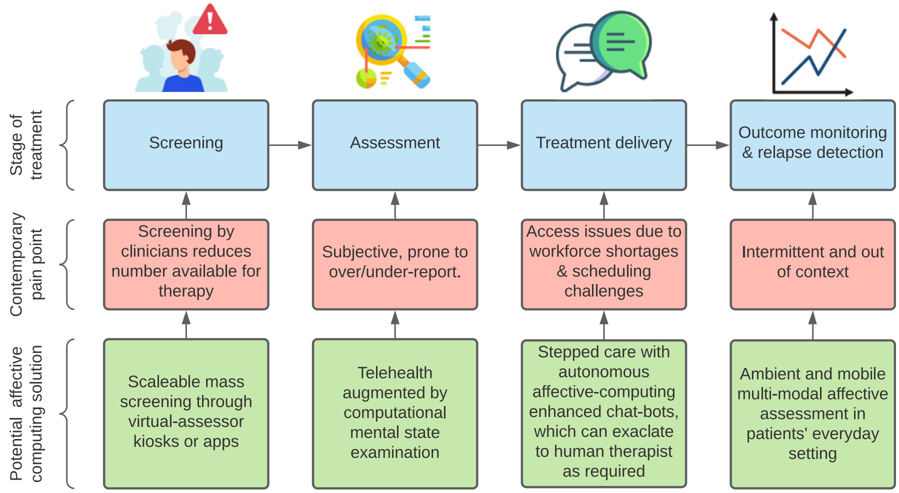

FIG. 3 Example therapy workflows that may benefit from affective computing-based applications.

clinicians, researchers, regulators, and other health enterprise stakeholders. Consumers need to be brought into that data-sharing discussion [27]. Consumers of digital health care products are increasingly skeptical around data privacy and control, including for medically sensitive data. As Huag stresses, however, they can also be savvy about its value, and many are altruistically motivated to share in the interests of other sufferers. Ignoring consumer views can lead to overly permissive and overly restrictive policies and regulations, depending on the level of the paternalism of regulators in a given setting. Both can be a detriment to ethical innovation.

One technology-first path forward may be the use of blockchain-based systems, increasingly termed "Web3." Unlike contemporary "Web2" infrastructure, where data typically reside on institutionally owned servers that are accessed and authored by clients (devices owned by individuals or organizations), Web3 relies on peer-to-peer computing. The technical detail is outside the scope of this article, but the principal distinction between the two for health purposes is data residence and ownership. In a peer-to-peer framework, data remain individually owned. Ownership is tracked on distributed ledgers ("the blockchain"), and protocols enable *smart contracts*. This means that instead of relying on a central authority to manage data and govern use, such requirements are coded directly into the system collecting and sharing the data. Fig. 4 summarizes the structural shift from Web2 to Web3. Odyssey Decentralised Autonomous Organisation (DAO) and other resources are readily available to better understand Web3 principles [28].

Precision Psychology

Affective computing biomarkers may assist in a precision medicine approach to certain psychotherapy modalities. As there are many different types of psychotherapy, affective computing biomarkers may be used to monitor response to psychotherapy and determine which type of therapy a patient is most likely to respond to. Affective computing can be combined with other modalities, such as functional neuroimaging for this decision support [29].

Managing Errors

Another aspect that clinicians, regulators and, other stakeholders must resolve is how to deal with inaccuracies. Affective computing is by no means infallible, and the further a model strays from the context of its training data set and the lower the quality of that data, the less accurate it is likely to be. Particularly problematic are cross-cultural differences in emotional expressions, such as differences in the meaning behind facial expressions, hand gestures, levels of eye contact, physical proximity, and dozens of other variables.

One possible way of managing this is mandating transparency of training data sets for affective computing systems. Much like food labeling, this might necessitate the disclosure of the demographics and context of that data set. It may enumerate which countries the model was trained in and the demographic profiles of participants, including age, gender, ethnicity,

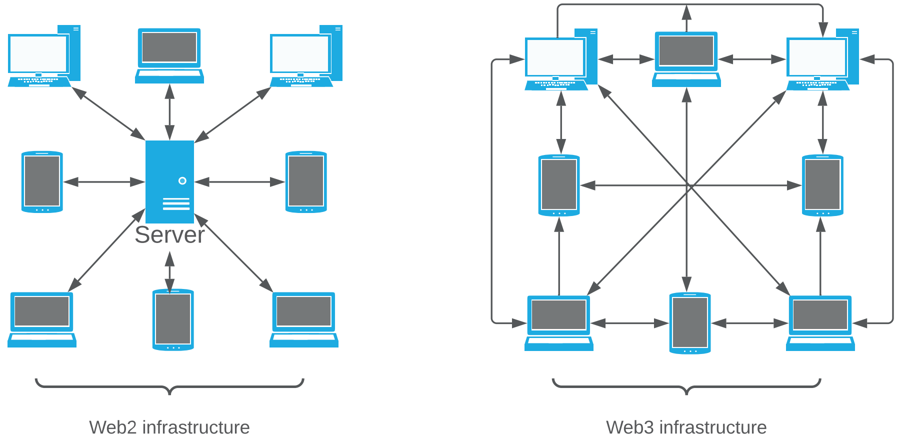

FIG. 4 Depiction of the shift from Web2 to Web3 infrastructure, which may benefit health applications.

and any other dimensions that may influence affective expression. Certain medical conditions or medication may also alter some of the features used in typical models, such as facial scarring for systems relying on expression recognition.

Although there is a formidable challenge, it is important to remember that humans share similar limitations. As Liu's study above reminds us, humans frequently disagree about affective expression. Like our digital counterparts, cultural differences pose particular challenges. Humans, however, have clearer legal standing and greater flexibility to recognize, reflect on, and address their limitations.

THE PATH FORWARD FOR CLINICAL AFFECTIVE COMPUTING

In addition to the above clinical integration questions, the technology has several other challenges that must be considered when developing systems for therapeutic use. Many such issues are considered in the 2021 International Institute of Electrical and Electronics Engineers (IEEE), Global Initiative on Ethics of Autonomous, and Intelligent Systems, which include an article on affective computing [30].

One interesting point posed in this document is the bidirectional relationship between humans and technology. Specifically, just as people shape the systems that accompany them, they are in turn shaped by those they engage with. Tellingly, it seems children as young as 7 months old can assign agency to an autonomous robot [31]. It is uncertain then how generations raised alongside such systems will have their affective experience and expression shaped by such interaction. Even adults may not be immune to the bidirectional influence of affective computing. At a minimum, there may be a shift toward the affective behaviors of the cultural context of the affective system's creators, akin to the impact of colonialism. This is not to say that societal changes are inherently negative, societies are always in a state of evolution. Rather, it highlights the need for cautious implementation and broad investment, where possible, in "home grown" (or at least "home trained") affective systems.

Earlier, the importance of designing systems that clinicians can improve post-implementation was discussed. Similar considerations are important for the therapists' clients who are being analyzed via affective computing. It is ideal for them to also know the outputs of the system. This may serve not only a therapeutic purpose in terms of developing personal insight and reflection but also an ethical duty in terms of transparency and the ability to correct errors. The exact implementation of such feedback, however, must be subject to empirical investigation to detect and address unintended consequences. For example, research suggests the phenomenon of *Zoom fatigue* (fatigue associated with videoconferencing) may partially result from "mirror anxiety" [32]; that is, users having to look at themselves for protracted periods during online interactions.

From a developmental context, we know that children look to carers as a source of truth to learn to understand and regulate their internal emotional states. The accuracy of the carers' regard and the interpretation of the child's effect is, therefore, thought to be a developmental predictor of future psychopathology [33]. It is possible that similar issues may arise with affective systems. The bottom line, as the IEEE concludes, is that stakeholders must be attuned to these risks and seek to mitigate and proactively assess them. As the issue of errors, human actors are not immune to such issues, so it is important to be measured such risks.

In addition, the 2021 "Ethics and governance of artificial intelligence for health: World Health Organization (WHO) guidance" report by the WHO highlights key ethical principles for the use of artificial intelligence in health care, which include protecting autonomy; promoting human well-being, human safety, and the public interest; ensuring transparency, explainability, and intelligibility; fostering responsibility and accountability; ensuring inclusiveness and equity; and promoting artificial intelligence that is responsive and sustainable [34]. These ethical principles must guide the development of affective computing applications for psychotherapy to ensure human rights are upheld and that patient and community interests do not become subordinate to the powerful commercial interest of technology companies or the interests of governments in surveillance and social control [35].

Finally, it is important to start thinking of such systems, as clinical products, and applying similar expectations to those preceding them. This includes publicly acknowledging side effects and somewhat uniquely for more autonomous systems, highlighting their status as *designed and deliberately built* rather than entities in their own right [30]. The IEEE guidelines also suggest that more autonomous systems need to be capable of detecting their use case and modifying accordingly. For example, they should recognize and deal differently with a minor who may not be able to judge the risks and benefits of an affective therapeutic system and give appropriate consent.

SUMMARY

Traditional psychotherapy is inaccessible to many who need it, and the lack of innovative approaches such as affective computing in psychotherapy remains largely absent. The psychotherapy field has begun to innovate to meet this challenge in several ways, including an embrace of telehealth, asynchronous care, and the inclusion of new disciplines like coaches. There have also been attempts at digitizing therapy more comprehensively, such as the use of online treatment companions and chatbots.

In parallel, strides have been made in affective computing. These systems have increased in speed and sophistication, using a range of modalities to interpret, express, and influence emotion. Although there are few such systems already deployed in clinical workflows, many others are demonstrating promise in research and development settings. The billions invested annually speak volumes of the need for such novel techniques in clinical practice.

Work to date demonstrates potential in deepening our mechanistic understanding of psychotherapy, training future therapists, screening those requiring help, assisting in diagnostic assessments, and participating directly in care delivery. This can take the form of supporting human-delivered therapy through clinical decision support and outcome monitoring or delivering sophisticated autonomous treatment.

Maximizing the potential of this emerging technology will depend on effective partnerships with clinicians and consumers. Therapists can contribute by helping prioritize technology to match important clinical challenges and workflows and providing expert feedback on algorithm outputs. They are also key stakeholders in the broader discussion about data ownership. Although Web3 frameworks may facilitate new ways of securing data, consensus and transparency remain important. Consensus is even more crucial when systems fail, and the technology will require the broader health ecosystem, including regulators and insurers, to adapt. The key publications by the WHO and IEEE can help guide the next stages in the broader development and dissemination affective computing systems for psychotherapy.

CLINICS CARE POINTS

- Clinical affective computing systems are beginning to reach clinical application and are attracting substantial consumer, investment, and research interest.
- They have the potential to increase access and quality of psychotherapy, both by supporting human-delivered therapy and supplementing it with autonomous care on a stepped-care model.
- Clinicians should be encouraged to explore adoption into clinical practice while demanding adherence to published guidelines on autonomous systems published by the International Institute of Electrical and Electronics Engineers and World Health Organization from developers.

DISCLOSURE

R. Khanna volunteers as an advisor to Flintworks and More Good Days. He is supported by an Australian

Government Research Training Program (RTP) Scholarship. H. Eyre is an employee of PRODEO LLC. E. Smith has nothing to disclose. N. Robinson is a company director and shareholder with LYRO Robotics.

REFERENCES

[1] Porter R. Madness: a brief history. Oxford University Press; 2021. p. 183–215.
[2] Khanna R, Forbes M. Telepsychiatry as a public health imperative: slowing COVID-19. Aust N Z J Psychiatry 2020;54(7):758.
[3] Chan S, Li L, Torous J, et al. Review of use of asynchronous technologies incorporated in mental health care. Curr Psychiatry Rep 2018;20(10). https://doi.org/10.1007/s11920-018-0954-3.
[4] Hubbard K. Many states face shortage of mental health providers. U.S News World Rep 2021. Available at: https://www.usnews.com/news/best-states/articles/2021-06-10/northeastern-states-have-fewest-mental-health-provider-shortages.
[5] Picard RW. Affective Computing 1995. doi:10.1007/bf01238028.
[6] Calvo R, D'Mello S, Gratch J, et al. The promise of affective computing. Oxford Handb Affect Comput; 2014. p. 1–11. https://doi.org/10.1093/oxfordhb/9780199942237.013.013.
[7] DeVault D, Artstein R, Benn G, et al. SimSensei Kiosk: a virtual human interviewer for healthcare decision support. In: *Proceedings of the 2014 international Conference on autonomous Agents and multi-agent systems*. AAMAS '14. International Foundation for Autonomous Agents and Multiagent Systems; 2014. p. 1061–8.
[8] Cardone D, Pinti P, Merla A. Thermal infrared imaging-based computational psychophysiology for psychometrics. Comput Math Methods Med 2015;2015:1–8.
[9] Politou E, Alepis E, Patsakis C. A survey on mobile affective computing. Comput Sci Rev 2017;25:79–100.
[10] Poria S, Cambria E, Bajpai R, et al. A review of affective computing: from unimodal analysis to multimodal fusion. Inf Fusion 2017;37:98–125.
[11] Wiltshire TJ, Philipsen JS, Trasmundi SB, et al. Interpersonal coordination dynamics in psychotherapy: a systematic review. Cognit Ther Res 2020;44(4):752–73.
[12] Coppersmith G, Leary R, Crutchley P, et al. Natural language processing of social media as screening for suicide risk. Biomed Inform Insights 2018;10:117822261879286.
[13] Chen L, Magdy W, Whalley H, et al. Examining the role of mood patterns in predicting self-reported depressive symptoms. In: 12th ACM conference on web science. WebSci '20. Association for Computing Machinery; 2020. p. 164–73. https://doi.org/10.1145/3394231.3397906.
[14] Lucas GM, Gratch J, King A, et al. It's only a computer: virtual humans increase willingness to disclose. Comput Hum Behav 2014;37:94–100.
[15] Guo T, Xiang Y-T, Xiao L, et al. Measurement-based care versus standard care for major depression: a randomized controlled trial with blind raters. Am J Psychiatry 2015;172(10):1004–13.
[16] Nixon R, Elizabeth M. Does case formulation improve outcomes in cognitive processing therapy for PTSD? a randomized controlled trial. In: ; 2020:151.
[17] Kahneman D, Riis J. Living, and thinking about it: two perspectives on life. Sci Well-being 2005;285–304.
[18] Ewbank MP, Cummins R, Tablan V, et al. Understanding the relationship between patient language and outcomes in internet-enabled cognitive behavioural therapy: a deep learning approach to automatic coding of session transcripts. Psychother Res 2020;31(3):326–38.
[19] Ewbank MP, Cummins R, Tablan V, et al. Quantifying the association between psychotherapy content and clinical outcomes using deep learning. JAMA Psychiatry 2020;77(1):35–43.
[20] Cummins R, Tablan V, Ewbank MP, et al. A tool for gaining insights into psychotherapy. Proc Int World Wide Web 2019;3503–6. https://doi.org/10.1145/3308558.3314128x.
[21] Jee C, Douglas Heaven W. The therapists using AI to make therapy better. MIT Technol Rev 2021. Available at: https://www.technologyreview.com/2021/12/06/1041345/ai-nlp-mental-health-better-therapists-psychology-cbt/.
[22] Wilcocks K, Kapralos B, Kalpin A, et al. A dialogue intervention simulation framework to facilitate psychotherapy training. IEEE 23rd Int Symp Consum Technol 2019;74–5. https://doi.org/10.1109/ISCE.2019.8900976.
[23] Naslund JA, Shidhaye R, Patel V. Digital technology for building capacity of nonspecialist health workers for task sharing and scaling up mental health care globally. Harv Rev Psychiatry 2019;27(3):181–92.
[24] Darcy A, Daniels J, Salinger D, et al. Evidence of human-level bonds established with a digital conversational agent: cross-sectional, retrospective observational study. JMIR Form Res 2021;5(5):1–7.
[25] Jennings K. Venture funding for mental health startups hits record high as anxiety, depression skyrocket. Forbes 2021. Available at: https://www.forbes.com/sites/katiejennings/2021/06/07/venture-funding-for-mental-health-startups-hits-record-high-as-anxiety-depression-skyrocket/?sh=15f1e99e1116.
[26] Liu C, Conn K, Sarkar N, et al. Affect recognition in robot assisted rehabilitation of children with autism spectrum disorder. Proc IEEE Int Conf Robot Autom 2007;1755–60. https://doi.org/10.1109/ROBOT.2007.363576.
[27] Haug CJ. Whose data are they anyway? can a patient perspective advance the data-sharing debate? N Engl J Med 2017;376(23):2203–5.
[28] Odyssey DAO. Intro to Web3. 2022. Available at: https://www.odysseydao.com. Accessed January 29, 2021.

[29] Beam E, Potts C, Poldrack RA, et al. A data-driven framework for mapping domains of human neurobiology. Nat Neurosci 2021;24(12):1733–44.

[30] The IEEE global initiative on ethics of autonomous and intelligent systems. Affective Computing; 2021.

[31] Kamewari K, Kato M, Kanda T, et al. Six-and-a-half-month-old children positively attribute goals to human action and to humanoid-robot motion. Cogn Dev 2005;20(2):303–20.

[32] Fauville G, Luo M, Queiroz ACM, et al. Nonverbal mechanisms predict zoom fatigue and explain why women experience higher levels than men. SSRN Electron J 2021;1–18. https://doi.org/10.2139/ssrn.3820035.

[33] Fonagy P, Gergely G, Target M. The parent-infant dyad and the construction of the subjective self. J Child Psychol Psychiatry 2007;48(3–4):288–328.

[34] World Health Organization. Ethics and governance of artificial intelligence for health. WHO Guidance; 2021.

[35] Buolamwini J, Timnit G. Gender shades: intersectional accuracy disparities in commercial gender classification. Proc Mach Learn Res 2018;81:1–15.