

Deploying AI Methods for Mental Health in Singapore: From Mental Wellness to Serious Mental Health Conditions

Creighton Heaukulani^{1,*}, Ye Sheng Phang¹, Janice Huiqin Weng¹, Jimmy Lee^{2,3} and Robert J. T. Morris^{1,4}

¹MOH Office for Healthcare Transformation, Singapore

²Institute of Mental Health, Singapore

³Lee Kong Chian School of Medicine, Nanyang Technological University, Singapore

⁴Yong Loo Lin School of Medicine, National University of Singapore, Singapore

Abstract

We describe our results from the implementation of machine learning and AI methods in three digital health initiatives serving individuals across the mental health spectrum in Singapore. The first initiative is Project HOPES, which we launched in 2019 for patients with serious mental illnesses. Originating as an observational study on digital phenotypes (collected via smartphones and wrist wearables) of 100 patients with schizophrenia, the tool has now been introduced as a service within a tertiary setting and has expanded to include patients with depression. The strategy dynamically prioritizes patients for early review by care coordinators according to need, which may avoid hospitalizations. Machine learning is used to predict clinical status, i.e., symptoms and functioning, and to predict relapses and other adverse clinical events; the latter can be done at 92% sensitivity and 90% specificity with the available digital biomarkers. The second initiative we describe is *mindline.sg* (www.mindline.sg), a platform for mental wellness in the general population that we created in 2020. Through a public-facing website, we deliver over 800 resources including wellness education, clinically validated self-assessments and triaging, and interactive resources, including an AI chatbot. Launched at the height of the COVID-19 pandemic, with all its attendant stresses, the platform has been visited by somewhere between 10 to 20% of the national population by the end of 2023. The third initiative we describe is *Let's Talk* (<https://letstalk.mindline.sg>), an online peer-support mental health network, which was co-created with youth advocates. The need for this platform was discovered through extensive studies with youth who expressed a desire for human-based support beyond the proliferation of digital solutions. In its first year, the site has been visited by over 80,000 unique users. Trained moderators review content on the site for safety and accuracy, and qualified therapists provide professional support through the free and anonymous *Ask-A-Therapist* service. To scale this service with a growing user base, we have been trialing the use of generative models to aid our therapists in finding relevant resources according to a user's need and to encourage empathetic writing.

Keywords

digital phenotyping, schizophrenia, depression, mental wellness, AI chatbots, large language models, digital health

1. Introduction

Mental Health conditions comprise one of the largest burdens of disease worldwide, especially when measured in Years Lived with Disability (YLDs). Stigma is prevalent in many countries and cultures, which discourages help-seeking. At the MOH Office for Healthcare Transformation and the Institute of Mental Health in Singapore, we adopt a population health approach starting from each end of the mental health spectrum and “working our way in”. On the one hand, we serve the needs of patients with serious mental illnesses, which are defined by the US National Institute of Mental Health (NIMH) as “mental, behavioral, or emotional disorder[s] resulting

in serious functional impairment, which substantially interferes with or limits one or more major life activities.”¹ On the other hand, we serve the mental wellness needs of the population through mental health promotion, where a local agency, the Singapore Association for Mental Health, defines mental wellness as “a positive state of mental health [that] is more than the absence of mental illness,” in which an individual is “able to think, feel and act in ways that create a positive impact on your physical and social well-being.”² Our solutions include digital health tools that increase the continuity of care, shifting care from the hospital/clinic into the home and from healthcare providers into the hands of individuals. In this article, we describe our results from the implementation of machine learning and AI in three digital health initiatives deployed in our strategy, together with the challenges we continually confront as we deploy these tools across the mental health ecosystem in Singapore.

Machine Learning for Cognitive and Mental Health Workshop (ML4CMH), AAAI 2024, Vancouver, BC, Canada

*Corresponding author.

✉ creighton.heaukulani@moht.com.sg (C. Heaukulani);

yesheng.phang@moht.com.sg (Y. S. Phang);

janice.weng@moht.com.sg (J. H. Weng); jimmy_lee@imh.com.sg

(J. Lee); robert.morris@moht.com.sg (R. J. T. Morris)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://www.nimh.nih.gov/health/statistics/mental-illness>

²<https://www.samhealth.org.sg/understanding-mental-health/what-is-mental-wellness/>

1.1. Serious Mental Illnesses: Digital Phenotyping and AI in Schizophrenia and Depression

Patients in Singapore with schizophrenia are usually treated in the specialist or hospital setting. After discharge to the community, it is not uncommon to see relapses. Around 80% suffer at least one relapse within five years of initial remission [1] which often result in emergency room visits or re-hospitalizations. Relapse is extremely disruptive to a patient's life and rehospitalization incurs large costs. Relapsing patients often exhibit psychotic symptoms such as hallucinations, delusions, or disordered thinking. They might display changes in sleep behaviors, mood, social withdrawal and disorganized behaviors. An emerging strategy to detect such changes involves *digital phenotyping*, defined as the "moment-by-moment quantification of the individual-level human phenotype *in situ* using data from personal digital devices" [2]. In 2019, we began the *HOPE-S* observational study [3], for which we developed the *HOPES* digital phenotyping platform described in detail by Wang et al. [4]. As of the end of 2023, the platform has been in continuous operation with patients and clinicians for over four years. Events are currently recorded from the user's smartphone including mobility (derived from obfuscated GPS coordinates), tapping speed on the keyboard, ambient light, screen time, accelerometry, and sociability indices (derived from calls, SMS and *WhatsApp* calls and messaging). Events are also captured from a wrist wearable device measuring heart rate, heart rate variability, activity (through step counts), and sleep (including staging and efficiency).

Over the course of the study, we continuously collected digital phenotyping data from 100 patients with schizophrenia (each patient was followed up for a six-month period), with clinical assessments performed every six weeks to measure symptoms and functioning. The total data collected throughout the trial consists of over 220 million events. We found generally high compliance in wearing of the wrist device (91% of all possible data was successfully collected in the week following enrolment), which required patients to wear the device at all times, including to sleep, and successful data collection from the smartphone (82% of all possible data was collected), which only required patients to not close the App in the background [5]. We note, however, that this high compliance rate was likely aided by a modest inconvenience fee that was provided to the trial patients. Some patients, to whom the study was offered, declined to participate for reasons including privacy or intrusiveness, leaving us with the interesting challenge as to how we can ameliorate and obviate these concerns in the future. In our initial analyses associating clinical status with digital markers, we found that the following (and

other) digital measures were predictive of poor symptoms and functioning: irregular sleep habits (including increased time spent awake in bed and in light stage sleep), decreased steps and GPS mobility, decreased text messages sent, slowed tapping speed, and increased heart rate while asleep, among others [6, 5].

We have also investigated the use of machine learning to predict adverse clinical events, including relapse of psychosis. For these purposes, we defined a relapse as a rehospitalization due to psychosis symptoms or a significant deterioration in clinical status defined using a validated clinician-assessed scale measuring general psychopathology symptoms. Other clinical events include emergency room visits, readmissions due to reasons other than psychosis, and unscheduled clinical visits. For these predictive models, we explored both unsupervised learning approaches to anomaly detection, including time series smoothing and forecasting methods (including the method originally used by Henson et al. [7]) and isolation forests, as well as supervised learning approaches utilizing generalized linear models, random forests, and gradient boosting trees. The models first establish a patient's baseline (on the multivariate data) and attempt to detect deviations from that baseline at the individual-level. Retrospective analyses from our observational trial indicate that we can detect adverse clinical events. The model that performed best varied by situation. For example, when using all digital measures (constituting a very high-dimensional and not very interpretable feature space), the gradient boosting tree performed best. But when we restricted the dataset to the one or two most interpretable features from each sensor (which are studied by clinical staff to explain the model), the isolation forest always performed best. The best model to use therefore depends on the operating mode, and ultimately will be guided by clinical requirements.

Here, we briefly report indicative performance metrics from the gradient boosting tree model. A more thorough report of the methodological descriptions and predictive results will be reported in upcoming publications. A classification setup is used, as studied by Ben-Zeev et al. [8], and sensitivity, specificity, and the harmonic mean score (between sensitivity and specificity) are reported. In our clinical setting, we are willing to accept a reasonable number of false alarms (i.e., lower the algorithm specificity), to keep sensitivity high. A false positive alarm may result in the care coordinator giving the participant a call to check in or sending them an inquiring and supportive text message. We therefore give more emphasis to sensitivity, i.e., being able to sense a deterioration in patient health, even if mild. This is acceptable for our clinical partners who envision a shift from the traditional reactive model of care to a proactive one where early detection and intervention might provide extra support and therefore prevent adverse clinical events. We there-

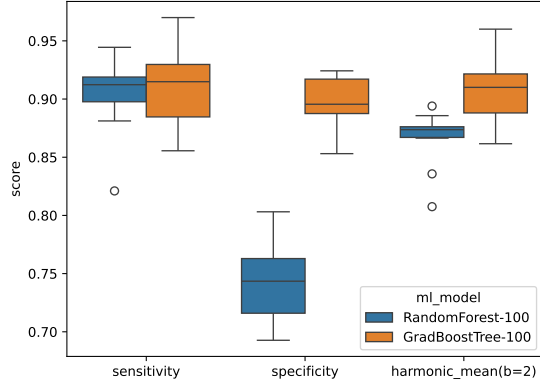


Figure 1: Boxplots of performance metrics for two popular machine learning models on predicting adverse clinical events in schizophrenia, including relapses.

fore explored optimization of the following weighted harmonic mean score:

$$H_{\beta} = (1 + \beta^2) \frac{\text{sensitivity} \times \text{specificity}}{(\beta^2 \times \text{specificity}) + \text{sensitivity}}$$

which has the interpretation that sensitivity is β times as important as specificity (where we note that β can be greater than or less than one depending on whether you value sensitivity or specificity more). This measure is directly analogous to the weighted F -score used in classification. In the experiments that follow, we report H_2 , where sensitivity is considered twice as important as specificity.

Boxplots of H_2 over ten 85%/15% training/testing dataset splits for two machine learning models are shown in fig. 1. The median score (over the test sets) using the gradient boosting tree is $H_2 = 91.0\%$ (91.5% sensitivity, 89.6% specificity). Note that the displayed sensitivity and specificity scores are the median scores for these metrics over the test sets; they do not correspond to the components used to compute the reported H_2 score. This predictive power is high but may be difficult to achieve in a clinical service (versus a controlled study) where compliance to wearable and smartphone data collection without financial incentives could be more challenging. The study by Cohen et al. [9] indicates that this real-world compliance could drop to 50%. Missing and delayed data uploads are common in clinical practice. Experiments on subsets of the features suggest that performance could drop to $H_2 = 85.6\%$ (87.4% sensitivity, 78.7% specificity) with a sparse model containing a subset of measures that are relatively easier to collect (steps, heart rate, accelerometer, screen time, taps in Apps, and tapping speed).

To explore those digital measures that appear most important for prediction, we fit the gradient boosting tree model on the prototypical measures from each fea-




Figure 2: Ranking by feature importance weight in the gradient boosting tree model for the prototypical measures from each sensor.

ture and display the “relative feature importance weights” (the impurity-based feature importance measures implemented by most packages) in fig. 2. In this experiment, we see that tapping speed is considered the most important feature, followed by GPS-based and wrist device measures, including distance travelled, time at home, sleep efficiency score, and number of steps. This example only considers a small subset of the features, however, and we do see such rankings change depending on the experimental setup and the feature set provided.

Having successfully completed the above study, we are now piloting the use of this model as a preventative service at the Institute of Mental Health (a large mental healthcare facility) in Singapore. This service is now being extended to serve both patients with schizophrenia and mood disorders including major depression. For this new *HOPES Clinical Service*, patient-facing components for the smartphone App have been developed to supplement what was purely passive monitoring in the previous observational trial. Supported on both Android and iOS, patients may now interact with the App through *Ecological Momentary Assessments (EMAs)*, wherein they may answer some questions as to how they are feeling at the time of a prompt, and they may additionally record what factors or stressors might be contributing to those feelings. Such data collection is sometimes referred to as *active monitoring*. The EMA responses are transmitted




Figure 3: The clinical care coordinator dashboard in the HOPES service for serious mental illnesses surfaces patients according to R/Y/G triage status, as determined by clinician-defined rules on the digital phenotyping data. Machine learning algorithms may only ever increase the severity of a patient’s status.

to the care coordinators to aid in care and supplement the digital phenotyping data, which together surface patients on a clinical dashboard. Our strategy is to have clinicians in a care monitoring center regularly observe the anomaly detection signals on the dashboard, and with the help of digital tools that can explore a patient’s data the clinical staff may decide to take further interventions, including earlier clinical review for a patient.

The clinical dashboard is designed to achieve the effectiveness of *Intensive Case Management* [10] in the presence of large caseloads. Cost-effective staffing dictates that the caseload of care coordinators (caring for patients that have returned to the community) is prioritized by a positional ordering that is determined by a *Red/Yellow/Green (R/Y/G)* status, shown in fig. 3. A Green (G) status indicates normal behavior and requires no action. A Yellow (Y) status indicates moderate deviation of behavior from normal (or possibly oncoming and escalating symptoms) and triggers automated interventions (described later). A Red (R) status indicates persistent and significant deviation from normal and requires review by the clinical team. To ensure patient safety, it is important that this prioritization be effective and based on an explainable clinical rationale. Hence a rule-based and fully explainable set of criteria has been developed (working closely with the clinicians who oversee the welfare of the patients), which determines whether a patient is flagged Y or R or left in the G state. This explainable property is also desirable to assist a care coordinator when they call the patient: they may explain what they have noticed in the digital markers (e.g., “you don’t seem to have been very active lately” or “it seems you haven’t been sleeping well”).

Patients with detected needs of lower acuity (i.e., those in the Y state) are supported with automated interventions comprised of digital therapeutics delivered through

the HOPES App. These interventions, based on the EMAs as well as the passive data, are timely and individualized, for example, sleep exercises are delivered when patients have poor sleep for two nights or more, and mindfulness exercises are delivered when patients indicate low mood on the EMAs, as just two examples. Such interventions are known as *Ecological Momentary Interventions (EMIs)*, and some of these digital therapeutic exercises are inspired by *cognitive behavioral therapy (CBT)*.

1.2. Our value proposition and strategy

Digital phenotyping enables the continual sensing of needs to drive both automated EMIs and to trigger stepped-up human-based care. The previous standard of care had no such real-time sensing capabilities and has been entirely reactive and episodic. This strategy strengthens the connectivity between care team and patient and augments the mission to increase continuity of care and extend care beyond the clinic into the community. We initially focused on schizophrenia because we anticipated quite strong indicative behavior-based signals. We are now scaling both “up” and “out” by moving into other diagnosed conditions such as depression and into adjacent populations, such as well populations, and into physical health conditions.

1.3. The role of AI and machine learning

Machine learning is used to predict a patient’s clinical status based on the multivariate set of digital markers, i.e., potential predictors of a patient’s current symptom severity and functioning. This may potentially avert the need to come in for clinical visits as is currently required under standard care. Machine learning is also used in the longitudinal prediction of relapses and may raise the color coding of severity and prioritization of the patient

for attention by the care coordinator. We also note that AI is used in a self-contained way in some of the EMI tools in the form of a mental wellness chatbot, which we will describe in the next section on our digital mental wellness platform, *mindline.sg*.

The machine learning models utilized for clinical event prediction include traditional time series smoothing, isolation forests, generalized linear models, random forests, and gradient boosting trees. The methods used for discovering associations between clinical scales and digital markers were multiple linear regression and multi-level/hierarchical linear regression. A detailed description of these methods is beyond the scope of this paper and will appear in upcoming publications.

Explainable signals are important for the clinical service. The multivariate nature of our digital phenotyping data aids clinicians in the interpretation of alerts – for which we provide digital tools to allow clinicians to keep track of patient data. It is important to note that our machine learning-based relapse prediction algorithm is only permitted to *increase patient severity* (from G to Y or from Y to R), which may prompt attention from a care coordinator (see fig. 3). The AI algorithm may never relegate a patient to lower acuity. This is not only for safety, but also aligns with Singapore’s national regulatory guidelines for predictive models in clinical service. As further experience is gained in the detection and management of patients with digital phenotyping, we may be able to move to a wider use of AI algorithms and to allow them to play a more definitive role in the selection of patients for clinical review. As new regulatory guidelines emerge and are navigated, pivots to our design and strategy may occur.

1.4. What’s next?

So far, the bulk of our experience has been with schizophrenia. Our service, however, has expanded to include patients with depression and takes a transdiagnostic approach, which may justify further research trials to develop and refine algorithms for depression in the local context. Additionally, an upcoming research study will evaluate our clinical service. Finally, we are currently expanding our digital phenotyping strategy to mental wellness. In this way, our digital phenotyping and machine learning tools are moving “inward” toward mild severity and well-populations.

2. Mental Wellness in the Community: Tools for the Well and Mild-to-Moderate Needs

Many countries and cultures face persistent challenges to mental health awareness and promotion including stigma toward mental disorders, reluctance to seek help, low mental health literacy, a lack of trained mental health personnel, and underdeveloped mental healthcare ecosystems. The COVID-19 pandemic created a surge in mental healthcare needs, demanding a new impetus to addressing these shortcomings. The pandemic also accelerated the adoption and acceptance of digital health solutions, creating a new opportunity for innovative approaches to address mental healthcare needs.

In June 2020, we launched a Web App, *mindline.sg* (www.mindline.sg), serving as a digital mental health resource website that has grown to include over 800 curated resources, a clinically-validated self-assessment tool for depression and anxiety, and a fully integrated AI chatbot developed for mental health applications from *Wysa*³, a leading partner in the field. More recently, we added curricular and structured learning materials with *Intellect*⁴, another leading partner in digital mental health. The landing page and a therapeutic exercise with the chatbot are shown in fig. 4. The self-assessment tool is comprised of the conventional scales of GAD-7 and PHQ-9 and triages users into *well*, *mild*, *moderate*, and *crisis* levels. The triage status allows us to customize content and recommend appropriate therapeutic exercises. Tailored products for youth and working adults are also provided to the public. Business-to-Business (B2B) customizations and engagements serve a range of ecosystem partners, including workplace partners and educational institutions. Resources on common risk factors (such as financial, employment and caregiver stress) are included, addressing a broad spectrum of determinants of mental health. The platform was developed to be anonymous and to contain authoritative and localised content for various levels of distress, with a focus on wellness and mild needs. Moderate-to-severe needs are primarily served by detection (through the self-assessment triage and AI chatbot) and referral to professional support, which includes counselling centres and emergency 24/7 services, all according to a clinician-designed protocol.

The platform has shown remarkable uptake. The site has been visited by between 10 and 20% of the (targeted) national population. The variability of this estimate is due to the anonymity feature of the site (we can only detect cookie IDs), which is a key feature enabling barrier-free access. If unique users visit from multiple machines or browsers, we may record them more than once. Indeed,

³www.wysa.com

⁴<https://intellect.co>

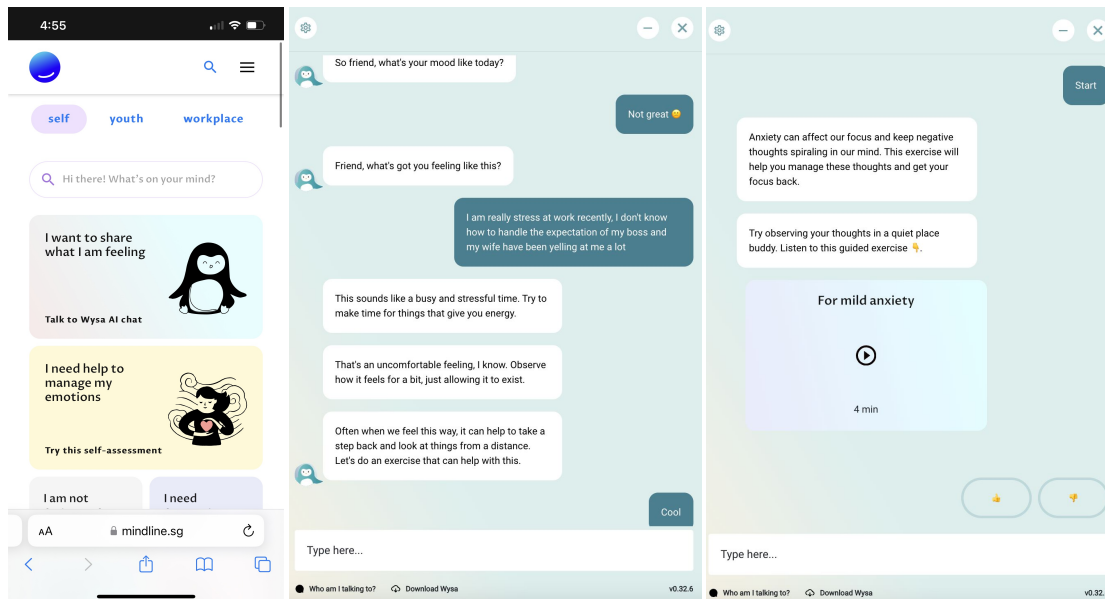


Figure 4: (Left) The *mindline.sg* landing page with the AI chatbot and self-assessment triage tools in the top two panels. (Middle and Right) The *Wysa* AI chatbot directs a user to an exercise inspired by cognitive behavioral therapy.

it is a learning from our implementation that anonymity does limit our ability to evaluate the platform. Another learning has been that successfully scaling the platform required expansive and sustained digital marketing efforts, as well as strategic ecosystem partnerships through the B2B products, and investment in partnership with educational institutions and healthcare providers. About 60% of our user acquisitions come from ads that we post on social media; the next most frequent acquisitions are from search, direct entry of the URL, or use of a QR code from our flyers and posters. Referrals also occur from B2B partner sites. The most popular resources used on the platform are sleep aids. “Mood check-ins” and the self-assessment tool are also popular. The distribution of moods includes “tired”, “unmotivated”, “anxious”, “positive”, “frustrated” and “sad” (in order of decreasing frequency). Scores for GAD-7 and PHQ-9 show a moderate number in a state of crisis: however, we believe that this frequency is affected by anonymous users trying out different answers to the questions to trigger different triage levels and seeing what resources are offered, mainly out of curiosity. We do not view this kind of usage negatively, as we feel that it is important for users to be educated as to what resources are available in times of need, either for themselves or a friend or relative. We have published these results in both process and impact evaluation studies [11, 12].

2.1. Our value proposition and strategy

The goal of the *mindline.sg* platform is to empower individuals in the community to take charge of their own mental health and to provide them the tools they need to offer basic support (“first aid”) to themselves and those around them, all through the ease and convenience of a barrier-free digital solution. This aligns with our strategy to improve population health through digital tools that enable self-empowerment and self-management. Such a strategy transfers some of the care from the system onto the individual and their significant others and moves some care from the clinical setting into the community and the home.

2.2. The role of AI and machine learning

A natural language processing (NLP)-based chatbot from *Wysa* is deployed to engage, triage, chat with, and direct the users to a range of relaxation, mindfulness, and CBT-inspired exercises. The *Wysa* chatbot is designed by a team of psychologists actively involved in patient counseling and has been subjected to numerous studies evaluating effectiveness [13, 14]⁵. Beyond the chatbot, however, we have limited data collected on the site (again, due to the anonymity feature), which limits any machine learning and AI efforts. A mobile App is presently being developed to give users an alternative option, which may

⁵<https://www.wysa.com/clinical-evidence>

be able to better leverage AI to improve user experience and benefit.

2.3. What's next?

The present wellness tool is mainly designed to serve those who are well or have mild conditions, with referral to human-based resources for moderate and crisis cases. We are also exploring the incorporation of clinical adjunct tools such as validated and localized *internet CBT (iCBT)* tools [15], which can be used by mental health-trained primary care providers. The mobile App will enable longitudinal data tracking and clinical management, which in turn may provide new opportunities for AI and machine learning in the service. We have already conducted a feasibility study of the tool in this population [16].

3. Digitally-Enabled Peer and Professional Mental Wellness Support for Youth

After extensive workshops, focus group discussions, and co-design sessions with youth (including those with lived experience), we discovered a desire for social support and meaningful human interactions delivered in a safe online environment. In particular, human-based support is now specifically sought amongst the proliferation of purely digital self-management solutions. We therefore co-created with youth advocates an online peer support network called *Let's Talk* (<https://letstalk.mindline.sg>). The platform soft-launched in October 2022 and has been piloting since. By end of December 2023, the site had received over 80,000 unique visitors (as measured by Google Analytics). This peer-support network's value proposition, over other platforms such as *Reddit*, includes close oversight and management by trained moderators and professional therapists. The moderators and therapists maintain a constructive and supportive atmosphere in the forum; other comparable forums have suffered from trolling, toxicity, spam, and scams. It is noted that some "medically themed" forums are overly commercial or are used by practitioners to advertise their own practices. *Let's Talk* also provides an *Ask-a-Therapist* service where users can pose a question to a panel of qualified professional therapists whom we have engaged; their response is asynchronous but usually occurs within 24 hours. The therapists follow a protocol defined by a *Clinical Advisory Panel* to deal with pressing needs and crisis conditions.

As of December 2023, from among the over 80,000 visitors to the site, over 6,000 have registered an anonymous user account (which is required to post content but is not required to access the forum and read posts). There

have been over 2,800 posts and over 370 *Ask-a-Therapist* questions had been answered. As the platform's user base scales, we anticipate challenges in moderating the platform and responding to questions in a timely manner while maintaining quality. We have therefore started trialing the use of large language models (LLMs) to assist our staff therapists in searching for relevant content from a trusted knowledge base (*mindline.sg*) based on the user's need (inferred from the posted question). We used *GPT-3.5* from *OpenAI*, which we fine-tuned using over 300 question-answer pairs from the *Ask-a-Therapist* service. *Retrieval-augmented generation (RAG)* [17] is used to produce the most suitable resources from the knowledge base, which is indexed from all resources in *mindline.sg*. While the therapist may copy-and-paste recommended resources (and their descriptions) from this tool, they remain fully responsible for the content in their reply. The LLM assistant has helped our staff therapists with close to 30 responses so far, where 88% of those responses have been rated as helpful by the therapist. In fig. 5, we show information on the *Ask-A-Therapist* service and screenshots on how our therapists use the LLM assistant.

3.1. Our value proposition and strategy

Youth mindline is the tailored youth product within the *mindline.sg* platform described in the previous section, and it serves as a companion site to *Let's Talk*. Where *youth mindline* enables self-management and self-empowerment through digital-only "self-help" tools, *Let's Talk* offers a purely human-based form of therapy and engagement. We believe the platform may also address determinants of mental health based on sociability and that users offering support to others has benefits for the helper and the helped. The strategy provides a low-barrier means of accessing professional support and provides this support (which is at the individual level) at scale.

3.2. The role of AI and machine learning

The LLM assistant exemplifies one of our strategies for generative models in healthcare in which an AI agent acts as an assistant to care providers. As we scale, we envision that such agents can dramatically reduce the time to search for a pool of potentially optimal therapies for a patient, client, or user's exhibited needs at the appropriate time, saving the clinician or care provider time. A sufficiently large pool of resources, as is the case in *mindline.sg*, can guard against repetition or a narrow set of recommended resources.

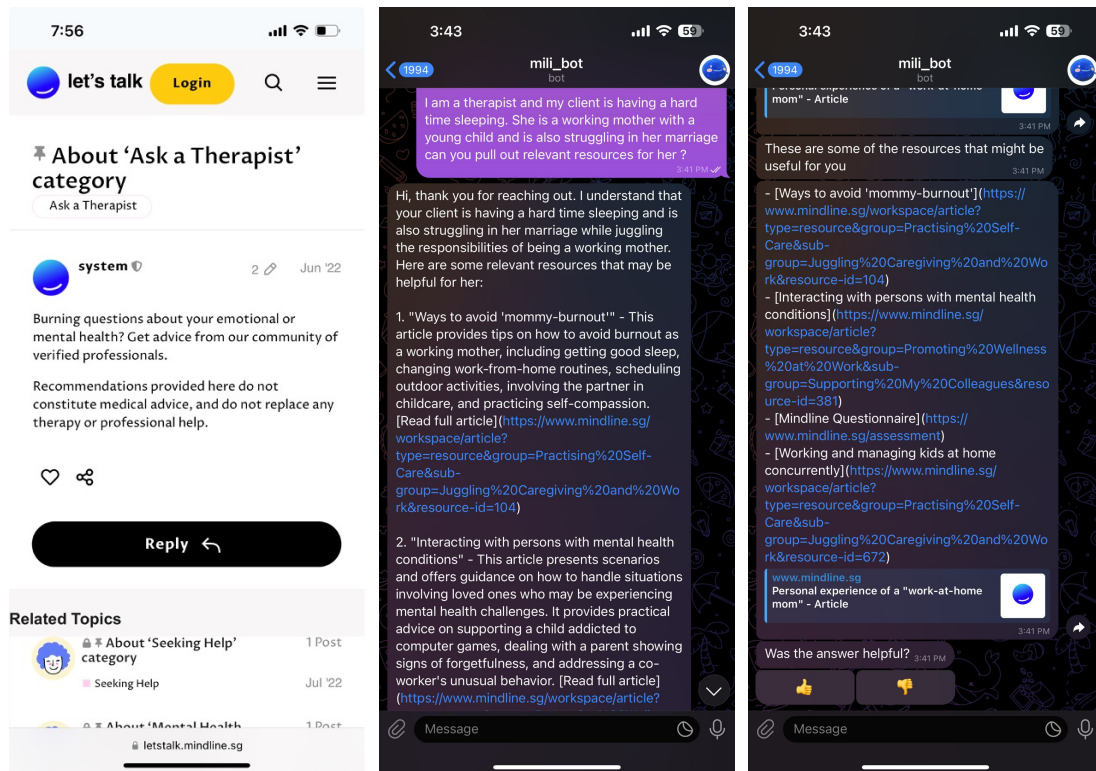


Figure 5: (Left) The Ask-A-Therapist service on the *Let's Talk* digital peer support forum for youth. (Middle) The Telegram interface for the LLM therapist assistant used by *Let's Talk* staff therapists to retrieve therapeutic content relevant for a post. (Right) Feedback is collected from therapists possibly enabling future AI improvement efforts.

3.3. What's next?

Future plans for AI and machine learning in *Let's Talk* include using LLMs to review therapist responses to check for or encourage a more empathetic tone. We may also use LLMs to train our peer supporter volunteers to a standardized level. Finally, LLMs or other NLP techniques can be used to continually assess content on the site to detect toxicity, spam, and misinformation.

4. Conclusion

The initiatives described have shown significant uptake by patients under care or users among our population. They confirm the promise of usefulness of digital and AI tools in providing improved digitally-enabled therapeutics and interventions. These initiatives demonstrate our strategy of starting at extreme ends of the mental health spectrum and working our way in toward the middle, blending interventional tools as we go. Through this strategy, we aim to cover the entire life-course and spectrum of acuity. Along the way, we are incorporat-

ing further care and ecosystem partners, including additional tertiary care partners, primary care providers, allied health, community organizations, peer supporters, workplace partners, and educational partners. We hope to communicate our learnings to others with a similar mission. We also hope to learn from others who are on a similar journey.

Ethical Statement

The authors received IRB approval for the HOPE-S digital phenotyping study described in the first section. The *mindline.sg* and *Let's Talk* services both have *Terms of Use* (available on their websites) that indicate that usage data may be collected and used for research purposes and for service improvement. The authors declare no conflicts of interest.

Responses to Reviewer Comments

We thank the reviewers for their comments. In this published manuscript, we have included succinct references

and further technical details on the machine learning and AI models, including quantitative experimental results. We have added in precise definitions of mental wellness and serious mental illnesses in a new Introduction section. We have made it clear that we are describing results from implementation in the Abstract and Introduction. These changes address all requests by the reviewers.

References

- [1] D. Robinson, M. G. Woerner, J. M. J. Alvir, R. Bilder, R. Goldman, S. Geisler, A. Koreen, B. Sheitman, M. Chakos, D. Mayerhoff, J. A. Lieberman, Predictors of Relapse Following Response From a First Episode of Schizophrenia or Schizoaffective Disorder, *Archives of General Psychiatry* 56 (1999) 241. URL: <http://archpsyc.jamanetwork.com/article.aspx?doi=10.1001/archpsyc.56.3.241>. doi:10.1001/archpsyc.56.3.241.
- [2] J. Torous, M. V. Kiang, J. Lorme, J.-P. Onnela, New Tools for New Research in Psychiatry: A Scalable and Customizable Platform to Empower Data Driven Smartphone Research, *JMIR Mental Health* 3 (2016) e16. URL: <http://mental.jmir.org/2016/2/e16/>. doi:10.2196/mental.5165.
- [3] N. A. Abdul Rashid, W. Martanto, Z. Yang, X. Wang, C. Heaukulani, N. Vouk, T. Buddhika, Y. Wei, S. Verma, C. Tang, R. J. T. Morris, J. Lee, Evaluating the utility of digital phenotyping to predict health outcomes in schizophrenia: protocol for the HOPE-S observational study, *BMJ Open* 11 (2021) e046552. URL: <https://bmjopen.bmj.com/lookup/doi/10.1136/bmjopen-2020-046552>. doi:10.1136/bmjopen-2020-046552.
- [4] X. Wang, N. Vouk, C. Heaukulani, T. Buddhika, W. Martanto, J. Lee, R. J. Morris, HOPES: An Integrative Digital Phenotyping Platform for Data Collection, Monitoring, and Machine Learning, *Journal of Medical Internet Research* 23 (2021) e23984. URL: <https://www.jmir.org/2021/3/e23984>. doi:10.2196/23984.
- [5] Z. Yang, C. Heaukulani, A. Sim, T. Buddhika, N. A. Abdul Rashid, X. Wang, S. Zheng, Y. F. Quek, S. Basu, K. W. Lee, C. Tang, S. Verma, R. J. Morris, Utility of wrist wearable and smartphone-based digital phenotyping in psychosis, Submitted (2023).
- [6] W. Martanto, Y. Y. Koh, Z. Yang, C. Heaukulani, X. Wang, N. A. Abdul Rashid, A. Sim, S. Zheng, C. Tang, S. Verma, R. J. Morris, J. Lee, Association between wrist wearable digital markers and clinical status in Schizophrenia, *General Hospital Psychiatry* 70 (2021) 134–136. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0163834321000098>. doi:10.1016/j.genhosppsy.2021.01.003.
- [7] P. Henson, R. D’Mello, A. Vaidyam, M. Keshavan, J. Torous, Anomaly detection to predict relapse risk in schizophrenia, *Translational Psychiatry* 11 (2021) 28. URL: <https://www.nature.com/articles/s41398-020-01123-7>. doi:10.1038/s41398-020-01123-7.
- [8] D. Ben-Zeev, R. Brian, R. Wang, W. Wang, A. T. Campbell, M. S. H. Aung, M. Merrill, V. W. S. Tseng, T. Choudhury, M. Hauser, J. M. Kane, E. A. Scherer, CrossCheck: Integrating self-report, behavioral sensing, and smartphone use to identify digital indicators of psychotic relapse., *Psychiatric Rehabilitation Journal* 40 (2017) 266–275. URL: <http://doi.apa.org/getdoi.cfm?doi=10.1037/prj0000243>. doi:10.1037/prj0000243.
- [9] A. Cohen, J. A. Naslund, S. Chang, S. Nagendra, A. Bhan, A. Rozatkar, J. Thirthalli, A. Bondre, D. Tugnawat, P. V. Reddy, S. Dutt, S. Choudhary, P. K. Chand, V. Patel, M. Keshavan, D. Joshi, U. M. Mehta, J. Torous, Relapse prediction in schizophrenia with smartphone digital phenotyping during COVID-19: a prospective, three-site, two-country, longitudinal study, *Schizophrenia* 9 (2023) 6. URL: <https://www.nature.com/articles/s41537-023-00332-5>. doi:10.1038/s41537-023-00332-5.
- [10] M. Dieterich, C. B. Irving, H. Bergman, M. A. Khokhar, B. Park, M. Marshall, Intensive case management for severe mental illness, *Cochrane Database of Systematic Reviews* 2017 (2017). URL: <http://doi.wiley.com/10.1002/14651858.CD007906.pub3>. doi:10.1002/14651858.CD007906.pub3.
- [11] J. H. Weng, Y. Hu, C. Heaukulani, C. Tan, J. K. Chang, Y. S. Phang, P. Rajendram, W. M. Tan, W. C. Loke, R. J. T. Morris, Mental Wellness Self-care in Singapore with mindline.sg: A Framework for the Development of a Digital Mental Health Platform for Behaviour Change (Preprint), preprint, *JMIR Preprints*, 2023. URL: <http://preprints.jmir.org/preprint/45761>. doi:10.2196/preprints.45761.
- [12] S. Yoon, H. Goh, X. C. Low, J. H. Weng, C. Heaukulani, Perceived Usability, User Preferences and Impact of a Workplace Digital Mental Wellness Platform “Mindline at Work”: A Mixed Methods Study, preprint, SSRN, 2023. URL: <https://www.ssrn.com/abstract=4579836>. doi:10.2139/ssrn.4579836.
- [13] C. Beatty, T. Malik, S. Meheli, C. Sinha, Evaluating the Therapeutic Alliance With a Free-Text CBT Conversational Agent (Wysa): A Mixed-Methods Study, *Frontiers in Digital Health* 4 (2022) 847991. URL: <https://www.frontiersin.org/articles/10.3389/fdgth.2022.847991/full>. doi:10.3389/fdgth.2022.847991.
- [14] B. Inkster, S. Sarda, V. Subramanian, An Empathy-

- Driven, Conversational Artificial Intelligence Agent (Wysa) for Digital Mental Well-Being: Real-World Data Evaluation Mixed-Methods Study, *JMIR mHealth and uHealth* 6 (2018) e12106. URL: <http://mhealth.jmir.org/2018/11/e12106/>. doi:10.2196/12106.
- [15] S. H. X. Lu, H. A. Assudani, T. R. R. Kwek, S. W. H. Ng, T. E. L. Teoh, G. C. Y. Tan, A Randomised Controlled Trial of Clinician-Guided Internet-Based Cognitive Behavioural Therapy for Depressed Patients in Singapore, *Frontiers in Psychology* 12 (2021) 668384. URL: <https://www.frontiersin.org/articles/10.3389/fpsyg.2021.668384/full>. doi:10.3389/fpsyg.2021.668384.
- [16] Y. S. Phang, C. Heaukulani, W. Martanto, R. Morris, M. M. Tong, R. Ho, Perceptions of a Digital Mental Health Platform Among Participants With Depressive Disorder, Anxiety Disorder, and Other Clinically Diagnosed Mental Disorders in Singapore: Usability and Acceptability Study, *JMIR Human Factors* 10 (2023) e42167. URL: <https://humanfactors.jmir.org/2023/1/e42167>. doi:10.2196/42167.
- [17] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, *Advances in Neural Information Processing Systems* (2020). URL: <https://arxiv.org/abs/2005.11401>. doi:10.48550/ARXIV.2005.11401, publisher: arXiv Version Number: 4.