Societal-Scale Human-Al Interaction Design? How Hospitals and Companies are Integrating Pervasive Sensing into Mental Healthcare

Angel Hsing-Chi Hwang Cornell University Ithaca, NY, USA hh695@cornell.edu Daniel A. Adler Cornell Tech New York, NY, USA daa243@cornell.edu Meir Friedenberg Cornell University Ithaca, NY, USA mdf224@cornell.edu Qian Yang Cornell University Ithaca, NY, USA qianyang@cornell.edu

ABSTRACT

From wearable health tracking to sensor-laden cities, AI-enhanced pervasive sensing platforms promise far-reaching benefits and introduce societal risks. How might designers of these platforms effectively navigate their complex ecology and sociotechnical dynamics? To explore this question, we interviewed designers who undertook this challenge in the mental health domain. They are hospital chief medical information officers and startup founders together striving to create new sensors/AI platforms and integrate them into the healthcare ecosystem. We found that, while all designers aspired to build comprehensive care platforms, their efforts focused on serving either consumers or physicians, delivering a subset of healthcare interventions, and demonstrating system effectiveness one metric at a time. Consequently, breakdowns in patient journeys are emerging; societal risks loom large. We describe how the data economy, designers' mindsets, and evaluation challenges led to these unintended design approaches. We discuss implications for designing pervasive sensing and AI platforms for social good.

CCS CONCEPTS

• Human-centered computing → HCI design and evaluation methods; Empirical studies in HCI.

KEYWORDS

Socio-technical systems, mental health, healthcare ecosystem.

ACM Reference Format:

Angel Hsing-Chi Hwang, Daniel A. Adler, Meir Friedenberg, and Qian Yang. 2024. Societal-Scale Human-AI Interaction Design? How Hospitals and Companies are Integrating Pervasive Sensing into Mental Healthcare. In Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3613904.3642793

1 INTRODUCTION

From sensors monitoring traffic and noise on the road [22] to Natural Language Processing (NLP) models detecting social media users' cognitive states [11], the combination of pervasive sensing and artificial intelligence (AI) offers both exciting opportunities and formidable challenges for human-centered design. A *societal-scale* sensing and AI platform can offer rich insights into a vast user base and inform a wide range of socially beneficial actions. However, considering the example of Facebook (a social media platform that recommends news based on inferred user reading preferences) or Hudson Yards (a "smart neighborhood" that optimizes its operations with pedestrian movement sensing.): Pervasive sensing platforms, when coupled with AI, have also historically triggered protests over surveillance and equity risks and caused real harm. [7, 22, 33, 35].

How might designers of pervasive sensing and AI platforms realize the technology's positive potential while preventing its unintended consequences? This is a difficult question, given such platforms' numerous stakeholders, complex data ecology, and intricate sociotechnical dynamics. Traditional HCI design methods, such as contextual inquiry and rapid prototyping, can be inadequate [15, 42, 59]. HCI researchers coined the term "*scale hack*" to describe how designers made opportunistic modifications of existing methods to cope with these platforms' exceptional complexities [3]. These challenges of creating pervasive sensing and AI platforms for social good have led some in HCI to explore new, scalable design methods [57, 60, 62]. Our research adds to this area of inquiry.

We interviewed nine platform designers from the mental health domain. They are hospital Chief Medical Information Officers (CMIOs) and startup founders striving to create new sensing/AI platforms and integrate them into the larger healthcare ecosystem (as illustrated in Figure 1). We also interviewed 18 relevant stakeholders (patients and physicians) to gain additional perspectives on their efforts. We ask: (1) How do platform designers currently navigate the complexities of designing a pervasive sensing and AI platform? (2) Are there early indications that their approaches effectively harness the technology's promises and account for its risks? What lessons might their approaches offer?

Our interviews revealed that these designers aspired to create all-encompassing sensing and AI platforms by first curating diverse data, then implementing AI models that deliver diverse care interventions, and finally assessing these interventions' superior ability to improve patient well-being. However, they encountered challenges at every step and had to prioritize the data sources they used, the care interventions they delivered, and the evaluation metrics they pursued. In doing so, they increasingly deviated from their goal of creating an all-encompassing care platform. This approach, deemed undesirable even by the designers themselves, also departed from patients' values in important ways. We discuss the lessons this current practice offers for designing pervasive sensing and AI platforms for social good, as well as the research opportunities it opens up.

CHI '24, May 11-16, 2024, Honolulu, HI, USA

^{© 2024} Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA,* https://doi.org/10.1145/3613904.3642793.



Figure 1: A vision of the future mental healthcare, according to prior literature [4, 5, 27, 28, 40]. This vision includes a sensing and AI platform that draws from diverse personal sensing data (the bottom data layer) offers diverse mental health and wellness suggestions (the middle AI layer) and revolutionizes healthcare practices of various kinds (the top health infrastructure layer).

This paper makes two contributions. First, it provides a rare description of the design process of pervasive sensing and AI platforms. It is a valuable reference point for research seeking to improve this practice. Second, this work illuminates many future research topics around addressing the design complexities of societal-scale AI platforms. These topics are important, as smart and connected health/workplaces/homes/cities become increasingly mainstream and AI systems grow ever larger in scale [21, 31–33, 38, 44, 64].

2 RELATED WORK

2.1 Design Complexities

A designer's expertise lies in embracing the full complexities of real-world situations and producing a new, unique solution [51]. Designers generate the solution not by applying a formula or rule, but by drawing upon a *holistic* understanding of the situation.

HCI designers navigate the complexities of human-computer interaction. In order to gain a holistic understanding of those interactions, *without resorting to oversimplification or ignorance*, HCI designers have continually refined and adapted their approaches:

- When people other than computer scientists started using computers, user-centered design methods such as *user personas*, *use scenarios*, and *user journeys* [10, 20, 56]) started scaffolding designers' thinking.
- Later, methods such as *stakeholder mapping* and *value proposition canvas* emerged, aiding designers in coordinating people's

diverse interests and expertise to co-create value [37, 63]. Additionally, concepts such as *value-sensitive design* [19, 37], *participatory design* [12], *speculative design* [17] become increasingly mainstream.

- The rise of AI catalyzed concepts such as *data swim lane* [61], prompting designers to consider the opportunities and risks in data flowing across various users and AI systems;
- As technologies' societal impact grew in importance, the concept of *infrastructure* entered HCI designers' lexicon, drawing their attention to the broader "*organizational, cultural, regulatory, environmental conditions*" that their technologies operate in and influence [41, 57].

This paper investigates how designers navigate the complexities inherent in designing pervasive sensing and AI platforms, with an eye on opportunities for improvement.

2.2 Societal-Scale Sensing and AI Platforms

Several characteristics distinguish pervasive sensing and AI platforms. First, they can collect data from a vast population, whether by sensing individuals (e.g., via smartphones or wearables) or public spaces (e.g., via observing city roads or social media platforms). Second, the AI systems built upon this data connect the platforms' users; how "*people like you*" have interacted with the system previously can impact your experience [59]. Finally, AI systems can extract rich meanings from the data and serve various purposes, both planned and unforeseen. For example, Fitbit's sleep pattern detection has informed users' self-care, aided physicians' neural disorder diagnoses, and even served as evidence of users' mental state in court [6].

Pervasive sensing and AI platforms' broad reach and versatility make them susceptible to unintended consequences. For example, whether it is smart cities, workplaces, homes, or healthcare, the line between pervasive sensing and surveillance is particularly thin [33, 46]. To prevent such unintended consequences, platform designers need to consider all of the aforementioned design complexities:

- Massive user bases, numerous user personas, and use scenarios. Textbook user-centered design methods can be inadequate [10];
- Numerous stakeholder groups and complex stakeholder relationships [3]. Service design methods can be insufficient [62];
- Formidable technical complexities. The platform involves multiple interconnected AI models built upon overlapping yet distinct data streams (as shown in Figure 1), making it difficult, for example, to ensure data representativeness and model fairness.
- Formidable human-AI interaction complexities [54, 59]. Certain outcomes of an AI system (e.g., filter bubbles or misinformation on social media) would not emerge until its user population reaches critical mass [3]. Once materialized, these outcomes can be irreversible [41].
- *Complex infrastructural landscape*, as pervasive sensing and AI platforms span many organizational cultures, regulatory jurisdictions, and cyber-infrastructures [14, 57, 60, 64].

2.3 Designing Societal-Scale AI Systems

How might a designer navigate through the multitude of complexities when determining the functionalities and interactions of a new sensing and AI platform? Existing HCI literature and practitionerfacing textbooks offer no best practice [13, 59]. Reports from industry suggest that designers often resort to ad-hoc, opportunistic approaches. For example, urban designers often decide what sensors to install and what data to collect based on a vague idea of how the data might power AI [35]. 60% of hospital Chief Medical Information Officers (CMIOs; they decide which patient data stream and AI models to include in their hospitals' software system) made these decisions solely based on vendor recommendations; 50% delegated the decisions to decentralized teams [29].

It is noteworthy that recent HCI research has started proposing novel approaches to designing large-scale sociotechnical systems. Some of these approaches center around developing *abstractions* of a system's ecology, allowing designers to grasp all its human, system, and socieotechnical complexities simultaneously and holistically. For example, *infrastructural speculations* offer designers tactics for interrogating the interactions among computing systems, people, social institutions, and political environments.

Other approaches are more incremental. For instance, Norman and Stappers [42] argued that designing complex sociotechnical systems requires a combination of "*incremental muddling through and satisficing*." Grevet and Gilbert [24] proposed *piggyback prototyping*, wherein designers validate a new design's ecological validity and ethical implications before adding it to a social computing platform.

More recent work has adopted, combined, and altered many of these above-mentioned methods through critical lenses [17, 25, 26, 58]. Building on co-design and participatory workshops [12, 52, 55], [26] emphasized the need to incorporate perspectives from underrepresented, marginalized groups. Farias et al. [17] combined speculative design with participatory activities, so that researchers and

domain experts can jointly envision the future of AI development with more diverse communities. Finally, making a twist to conventional world-building, [58] presented *Timelines*, a design activity that purposefully creates space for critical reflections from expert participants.

These approaches are nascent and under-evaluated in design practice. Are they truly effective in designing sensing and AI platforms? We investigate this question in the context of the mental health domain.

2.4 Pervasive Sensing and AI for Mental Health

We chose mental health as our site for studying pervasive sensing and AI platform design for a number of reasons. First, the technologies are maturing and on the cusp of entering the real world. From smart watches monitoring depression to NLP models detecting suicidal ideation on social media, AI models are reaching a level of accuracy that can inform both people's everyday self-care and physicians' medical decisions [11, 53].

Second, the once rigid legal and economic barriers separating consumer and medical devices/data are eroding. In the United States where this research was conducted, federal laws recently started allowing easy health data sharing across commercial apps and Electronic Health Record platforms (EHR) [23, 47]. With user consent, an Apple Watch can now access an EHR's data, and vice versa. Furthermore, insurers are combining insurance policies with mobile health. For example, they offer lower premiums for those exhibiting healthy lifestyles on Fitbit [45]. These policy shifts make possible a truly societal-scale sensing and AI platform that spans across people's daily lives and medical care.

Lastly and importantly, the emerging sensing and AI platforms for mental health encapsulate the great promises, risks, and complexities of designing societal-scale AI systems. Once integrated into the larger healthcare ecosystem, these platforms can significantly broaden the reach of preventative care (Figure 2), improve outpatient monitoring, and alleviate the mental health crisis in the



Figure 2: The promise of pervasive sensing and AI platforms for mental health. While medical care tends to focus on severe illnesses (blue), sensing platforms promise to offer appropriate level of care care continuously (pink) and improve physicians' efficiency capacity (overlap).

United States [1, 2, 8, 27, 34, 39, 40, 48]. Meanwhile, researchers warn of many risks in integrating pervasive sensing and AI systems into healthcare: promoting health surveillance [46], exacerbating inequity [36, 64], and other long-term unintended consequences [9, 18].

3 METHOD

We wanted to understand how designers in the industry have been navigating the complexities of designing pervasive sensing and AI platforms. We wanted to understand the lessons their approaches offer for broader HCI communities.

Towards these goals, we interviewed 9 "*platform designers*", professionals who are in the process of designing new sensing and AI platforms for mental health and integrating them into the larger healthcare ecosystem. They are the Chief Medical Information Officers (CMIOs) of major hospitals implementing new behavioral health initiatives, and startup leaders (founders and executives) (Table 1.) We also interviewed 5 mental health physicians and 13 patients to gain additional perspectives on their efforts.

3.1 Participants

We recruited "*platform designers*" across a wide range of seniority and organizations, through our professional networks and Upwork. With Upwork participants, we reviewed their resumes and did a preinterview to ensure they had relevant platform design experience. Appendix A describes our screening procedure. We also recruited physicians using similar methods.

We recruited mental health patients with diverse experiences using similar recruitment and pre-screening methods. Participants

Participant	Professional Role	Prof. Exp.*				
Platform Designers						
D1	Chief Medical Information Officer (CMIO) of a major hospital system					
D2	Vice President of a wearable health device and service start-up (Business area: Mobile apps that track behavioral data through wearable and mobile sensing devices to reflect real-time mental health states)					
D3	Executive Director of a consumer health tech start-up. (Business area: Integration solutions for clinical and non-clinical data) Formerly: Clinical machine learning engineer.					
D4	CMIO and Director of Healthcare Technology and Clinical System Implementation of a major hospital system	25 - 30				
D5	Founder of a start-up that helps consumer-level sensing devices and apps companies to integrate their products into electronic health record (EHR) systems. (Business area: Integration solutions for cross-hospital, cross-organization technical infrastructure) Formerly: Bioinformatics software developer.					
D6	Consultant for consumer health tech companies (primary clients are those who intend to pitch their products to public health sectors). Formerly: Executive Director for health tech integration at <state> Department of Health.</state>					
D7	Technical Product Manager who manages the development of cloud-based consumer healthcare products and services					
D8	Clinical faculty who leads the implementation of behavioral health programs at a university-affiliated hospital					
D9	Founder of a consumer health tech start-up (Business area: Wearable sensing to track real-time mental health states)					
Physicians						
PH1	Licensed therapist and neuropsychologist who have worked in outpatient, inpatient, and residential treatment settings					
PH2	Licensed therapist and counselor who have worked in outpatient and inpatient settings					
PH3	Licensed therapist and counselor who have worked in outpatient and inpatient settings, specialized in teenager mental healthcare					
PH4	Licensed therapist and counselor who have worked in outpatient and inpatient settings					
PH5	Licensed therapist and counselor with former experience in the public healthcare sector					
Patients						
PA1, PA3, PA6	Patient who receives regular inpatient and outpatient care	-				
PA2, PA4, PA8 PA11, PA12, PA	3, PA9, PA10, Patient who receives regular outpatient care A13	-				
PA5, PA7	Patient who receives regular outpatient care and has professional experience in AI research and development	5 - 10				

Table 1: Study Participants (*Values in Prof. Exp. represent participants' professional experience in years)

in the patient category were receiving mental healthcare services on a weekly/bi-weekly basis. Four of them also receive regular physical health treatments.

3.2 Interviews

We developed two interview protocols: one for interviewing platform designers and another for interviewing physicians and patients. Both protocols were approved by the IRB. Each interview was approximately 60 minutes.

The interviews with platform designers started with inviting them to walk us through step-by-step their recent experience designing their sensing and AI platform. We asked many follow-up questions about their thought processes, external influences, and specific platform design choices. With physicians and patients, we first asked them to share their views of the current mental healthcare ecosystem based on their experiences of providing/receiving care. Next, to help them articulate their preferences for a pervasive sensing and AI platform for mental health, we showed them an illustration of its possible future (similar to Figure 1). From there, we asked participants to share their opinions and their desired changes to this future.

We recorded and transcribed all interviews and analyzed them using affinity diagrams and user journey maps. Both methods helped us to consolidate our understanding of how platform designers navigated the complexities of sensing and AI platforms and made concrete design designs. Appendix D details the themes that emerged from the analyses.

To reduce subjectivity during data analysis, the research team met weekly to discuss and form a working consensus. With preliminary results, we also actively sought feedback from fellow researchers outside our team to identify blind sides that we might have overlooked from the data. Still, we acknowledge that researchers' backgrounds and experiences can influence analyses [49]. Appendix B presents our positionality statement.

4 FINDINGS

Our findings are threefold: (1) Designers aspired to create an allencompassing sensing and AI platform not just for profit, but as a preferred way to improve patient wellness holistically. They took three steps toward this goal: curating diverse data, offering a range of interventions, and demonstrating the efficacy of these interventions in enhancing patients' overall wellness. However, (2) distinctive challenges at each stage led designers to rely heavily on a single data source, implement only a subset of potential interventions, and evaluate interventions using one partial metric at a time. (3) This approach not only deviates from the designers' original vision but also contradicts patients' values.

A note on "patients." There is no definitive line between a mental health patient and everyone else [30]. In our interviews, participants from consumer tech industries referred to receivers of mental wellness/healthcare as *customers* or *users*, while those from medical organizations referred to them as *patients*, even though the populations they refer to overlap broadly. For consistency, we will use the term "*patients*" to refer to those receiving mental health care, including both self-care and professional care.

4.1 Dreams of an All-Encompassing Platform

Platform designers aspired to create all-encompassing sensing and AI platforms that draw from all varieties of training data, deliver all types of mental health care, and improve patient wellness holistically. "We know we need to have connected devices. We know we need to have a mobile application. We know we need to have interventions that are relevant in real time. We know we need to have health coaches that are connecting the dots for people. And so it's a service and technology business." (D2)

Some designers explicitly acknowledged the role of economic incentives in pursuing this aspiration.

"Eventually, every (healthcare) provider wants to be a payer, and every payer wants to be a provider. (That's why) All the wellness, meditation app folks are moving into employee benefit programs (where) you'll see some more clinically oriented solutions. [...] You end up wanting to have a play in all parts of that delivery equation in order to control the economy." (D4)

Importantly, profit is not the only reason. Designers pursued this vision of an all-encompassing sensing/AI platform also because they genuinely believe it can provide more holistic care for patients' mental health. By integrating diverse data, devices, and interventions, physicians could better identify and address the underlying causes of a patient's conditions, rather than treating individual symptoms. This holistic approach is also more likely to keep patients engaged with treatments. "A diabetes patient could not consistently monitor their blood glucose when [they] were depressed or had intense anxiety." (D2). It is worth noting that no platform designers we interviewed excluded physical health from their vision of a holistic mental health platform. "There is no healthcare without mental healthcare," they emphasized.

Towards this vision of fostering patient well-being holistically, platform designers embarked on a three-step journey.

- (1) Curating diverse data. Platform designers worked to curate data from diverse sources: clinical assessments, social media data, patients' digital footprints, and more. Diverse data enabled their machine learning (ML) models to recognize a broader range of mental health conditions. It could also expose errors or biases within certain data sets, thereby improving ML accuracy and reliability. Designers stressed its importance for mental health due to the lack of ground truth. "Oftentimes, you can't do a test that says [mental illness] is present or not" (D1).
- (2) Delivering diverse care interventions with AI. Next, designers worked to create AI applications that, ideally, can deliver various interventions needed to improve a patient's overall wellness. They envisioned these interventions to include day-to-day wellness recommendations (e.g., exercises, diets, and sleep times) and clinical treatments (e.g., pharmaceuticals), mental and physical healthcare, as well as primary, general, and specialized care. Many participants analogized their envisioned platforms as app stores: a one-stop shop where patients could access all clinical and non-clinical care.
- (3) Proving interventions' efficacy in improving patient wellness. Finally, designers worked to persuade key stakeholders that their platforms are more effective at enhancing patients' overall well-being compared to current solutions.

Juxtaposing steps (1) and (2), we can see that the goal of managing patients' holistic well-being has led designers to pursue a platform that drew from all types of data (widening the lower half of Figure 1) while offering a comprehensive set of care interventions (widening the upper half). They aspired to designing pervasive sensing and AI platforms that reaches every corner of the mental health ecology.

4.2 Narrowing Down the Scope at Every Step

In pursuing this ambitious vision, designers encountered distinctive challenges at every step. These challenges forced them to prioritize the data sources they use, the care interventions they deliver, and the evaluation metrics they pursue. In doing so, they gradually veered away from their initial goal of crafting an all-encompassing care platform.

4.2.1 Designing with One Primary Data Source. The platform designers we interviewed sourced their AI training data from one of two sources: Either hospital systems or consumer technology products (e.g., meditation apps, therapy chatbots, or wearable health devices). The difficulties of obtaining data from either source compelled designers to concentrate on a single data source.

Electronic Health Records (EHR) data from hospital systems were most designers' preferred choice, because of the higher data quality. However, accessing EHR data necessitated designers to invest significant time and resources to establish partnerships with hospitals (or other clinical institutions). Furthermore, even designers with EHR data access encountered challenges in retrieving and utilizing it. Several interviewees indicated that leading EHR vendors, notably EPIC, deliberately erected barriers in this process. These vendors wanted to maintain their oligopolistic advantage in EHR data ownership. Therefore, their data-sharing tools were "*extremely old-schooled*" and "*woefully inadequate compared to similar services in other industries*." "*They were not there to get people connected to those systems*." (D5)

Getting data directly from consumer tech users presented its own set of challenges. The quality of such data varied, and this is true across passive sensing data (e.g., wearable heart rate monitoring), users' self-reports (e.g., diet logging), or a combination of both (e.g., therapy chatbot dialogues). As D4 suggested, "*it is even difficult to keep them (patients) wearing their Apple watches.*" As a result, designers found themselves in a perpetual cycle of devising novel strategies to engage users, getting them to maintain consistent interaction with the application or device.

Considering the challenges in obtaining quality data from either source, the platform designers focused on a single source, at least in the initial stages. Moreover, because the initial data source could shape a sensing and AI platform significantly (more on this below), it remains the primary data source, even after the platform matured and gained the capacity to integrate data from multiple sources. In our interviews, even the most established platforms are heavily reliant on a singular data source.

Consequences of designing with one primary data source. The critical role of training data in establishing a large-scale AI platform meant that those who provide data became the primary stakeholders of the platform. Whether they are hospitals or consumer tech users, these stakeholders wield substantial influence over designers'

decisions concerning platform functionalities, model optimization objectives, and interface designs.

Consequently, a dichotomy emerged in platform designs. On the one hand, the platforms predominantly utilizing EHR data placed a strong emphasis on saving physicians' time and cutting down hospitals' expenses. To this end, designers integrated features aimed at, for example, streamlining administrative tasks and shortening the duration of in-person patient consultations.

"If you can get good kind of clinical buy-in and the nurses or the clinicians or the physicians are all like, 'wow, this is really making life a lot easier for me,' then that's a huge win." (D7)

On the other hand, platforms fueled by consumer tech data prioritized improving user engagement and acquiring additional user-generated data. D2 described their approach as "*user-centered service design*." Their design offers a telling example:

> "We asked people what their favorite sports team was, and in between the health messages we would send out score updates or upcoming game updates around their favorite sports team." (D2)

4.2.2 Delivering One or Two Types of Care Interventions. After securing their training data, platform designers began building AI applications that can offer a wealth of mental health and wellness interventions. They started by offering a single type of care, such as providing daily mental wellness suggestions (e.g., exercises, dietary advice, or sleep schedules). Then, they worked to expand their platform by incorporating additional interventions, such as facilitating video consultations with primary care physicians.

However, the fragmented nature of healthcare infrastructure– where various types of care interventions are offered by different hospitals and medical specialties, each covered differently by insurance–often hindered this expansion, for at least two reasons.

Firstly, when designers seek to integrate a new type of care into their platforms, they must first consider: What value can their data uniquely offer to a hospital or department that offers it? Then, they needed to establish partnerships with this hospital or department. This is a time-consuming and competitive endeavor.

For example, Calm (a consumer-level meditation app) is moving into the actual healthcare delivery space. But how do you make the leap? I've downloaded Calm, and I got a meditation app. [...] "How do I connect that (meditation data) to my health care provider, in a way that makes sense to my provider [...] and they can do something with it? There's a lot of challenges around how to connect the wall between medical care and wellness." (D4)

"(When you try to collaborate with a new hospital,) you have to join a waiting list of apps that might be there for years. And there is a waiting list at every hospital." (D5)

Second, designers also needed to integrate their sensor/AI platform into the EHR of the newly partnered hospitals. This involves establishing effective communication between their platform and the application programming interfaces (APIs) of EHR vendors, as well as ensuring that their platform's data (e.g., patients' inputs into Societal-Scale Human-AI Interaction Design?

CHI '24, May 11-16, 2024, Honolulu, HI, USA

the app, AI models' care recommendations, and patient outcomes) are compatible with the hospital's data pipelines.

"Because every hospital did its clinical assessment and treatment differently, with different machines, different values, different sensitivities, different nurses take blood pressures in different ways using different types of blood pressure cuffs. [...] And these variations could create breakdowns everywhere in the system." (D1)

In light of these challenges, all platform designers we interviewed focused on delivering one or two types of care interventions, veering away from their initial vision of delivering comprehensive care. Consider D2's sensing and AI platform for managing Type 2 diabetes as an example. They have successfully expanded the platform from only providing behavioral interventions (e.g., encouraging users to walk 15 minutes after having a pizza or when blood glucose levels are high), to also offering "end-to-end virtual care" (i.e. actual primary care physicians can converse with patients and make prescriptions.) To achieve this expansion, D2's team first established a partnership with Redox (an EHR-app integration vendor), then integrated their platform's data into the partnering primary care clinic's EHR, redesigned the platform's interfaces to suggest its suggestions come from the primary care physicians, and even standardized the costs of these different interventions. Each of these steps was a significant undertaking. However, despite these concerted efforts, the platform could only incorporate two distinct types of care.

Consequences of delivering one or two types of care interven*tions.* Creating a sensing and AI platform that offered only one or two types of care interventions, while beneficial to some extent, resulted in two undesirable consequences. First, the platform required increased efforts from both physicians and patients, yet the care provided did not substantially differ from pre-existing nonsensor-and-AI-based solutions. Without significant patient wellness benefits, both physicians and patients were more likely to become disengaged.

D7 on patient disengagement: "Every time you ask somebody to strap something on their wrist, or to put something around their arm, or to put something up against their forehead is a friction point. It's another device. It's more batteries. It's additional charging. Right there, that's where consumer friction builds up."

D2 on physician disengagement: "They're very pro that (adopting new applications) if it makes them feel closer to their patients, and it makes them feel like they've got actionable data that they can actually use to deliver care. Not just an avalanche of useless data. [...] So many (applications) have come at them that their default response is 'No'. Here comes another thing, another portal, another tab, and the EMR. Here comes another set of data that I have to go through as a provider to figure out what it means."

Secondly, because each sensing and AI platform only provides one or two types of care, patients might "*have to use eight different apps*" to monitor their health and access comprehensive care (D2). Today, patients often find themselves undergoing repetitive medical tests during transitions between different types of care (These transition points include, for instance, upon therapy admission and discharge, during shifts from physical to mental healthcare, and when transitioning from general to specialty care. Figure 3). Both physicians and platform designers that sensing and AI platforms offering limited types of care (e.g., focusing solely on everyday wellness rather than clinical care) might exacerbate these disjointed transitions and necessitate repetitive testing.

"Your heart rate on your Apple Watch is still your heart rate. It is your heart rate on any device. But you now have to do it all over again but with older, slower equipment from the hospitals." (D3)



Figure 3: One physician interviewee's (PH1) illustrated the journey of a patient as they transition between different types of care. Sensing and AI platforms offering limited types of care (e.g., focusing solely on everyday wellness rather than clinical care) might exacerbate these discordant transitions and necessitate repetitive testing.

4.2.3 Demonstrating System Effectiveness One Metric at a Time. After having started delivering care interventions, designers redirected their focus towards evaluating the efficacy of these interventions, to improve the interventions and garner stakeholder buy-in. Their ultimate goal was to prove that their sensor/AI platforms could trigger an upward trajectory in patient wellness (e.g., patients "*didn't need to go to the hospital as much*".) Designers of clinician-facing platforms also aimed to prove the superiority of their platforms over existing solutions or competitors (e.g., It can make physicians "*feel closer to their patients*".) Towards these goals, platform designers conducted clinical trials spanning diverse populations, across various mental wellness or healthcare settings.

However, the benefits of sensing and AI platforms lie in their ability to deliver *diverse* interventions flexibly and improve patient well-being *holistically*. Such benefits are difficult to prove or benchmark. For example, there are no existing benchmarks for assessing the efficacy of integrative mental-and-physical healthcare, because few prior care programs have provided such care. Similarly, there lack of benchmarks or prior clinical trials on patients' *"overall well-being."* Moreover, wellness improvement takes time, yet longitudinal metrics were rare. As D5 acknowledged, *"you can't really assess patients' outcomes within the sales cycle."*

As a result, platform designers resorted to utilizing established, quantitative, and short-term metrics. They assessed the effectiveness of their platform using one metric at a time, with the hope of



Figure 4: Summary of platform designers' approach (§4.2). Designers aspired to create a platform that draws from diverse data (blue), delivers integrative care (purple), and improves patient wellness holistically (pink). However, challenges in securing good data, the fragmented healthcare infrastructure, and the lack of holistic evaluation metrics led designers to narrow their focus at every step. As a result, a dichotomy emerged among platform designs (the two arrows.)

eventually demonstrating its efficiency across all relevant metrics and thereby proving they have improved patient wellness holistically. Designers started by, for example, proving to patients that the platform saved "*this much money on my (health) insurance*", or reduced this many sick leave days taken from work. For clinicians, designers started by proving that "*there weren't all these errors I had to go and correct*" in the EHR system, and then conducting a series of clinical trials; each proving the platform's ability to improve one patient condition.

Consequences of assessing platform efficacy using individual metrics. This one-metric-a-time approach to evaluation led to two downstream effects. Firstly, it prolonged the evaluation process. Out of the nine platform designers we interviewed, only one was able to demonstrate a measurable improvement in patients' overall wellness with one of their AI applications. All other platform designers were still in progress. "*It's this months-and-months-long process.*" (D7)

Secondly, until the individual, partial metrics accumulated and became evidence of the platform's efficacy in improving patients' overall wellness, these metrics often failed to convince platform users of its value. Patients who remained unconvinced by the platform's worth may quit using it.

"We say that we want to improve metrics for patients, like blood pressure, cholesterol, or blood glucose. But they (patients) don't care about any of that. They care about feeling better. They care about things like spending more time with family, and going to their granddaughter's wedding. Those are the kinds of health goals you hear about when you actually interview individuals. We were the ones that were forcing these metrics on them, but that's not what's driving them. What's driving them is their lifestyle." (D2)

Clinicians who remain unconvinced by the platform's worth not only refused to use the platform, but also caused other complications. Platform designers shared that, unless they had secured clinical buy-ins from physicians and nurses, they faced numerous additional metrics imposed by other hospital stakeholders, further slowing down the already prolonged evaluation process.

> "Because hospitals are often very big and bureaucratic, there are a lot of people with their fingers in the pie, all having their own opinions. I've spoken to the nurses, the actual users of the products, which is really great, and (hearing) what does and doesn't work. But then, of course, you also get people along the way who want to jump in. [The] Innovation Department, for example. They might overlook everything the nurses have just said and say, hey, we're Innovation. We're going to start assessing what's best. And then the Accounting Department might step in and say [...] hey, we found this product that's cheaper." (D7)

> "Let's start with the architectural review committee, then the data governance review committee, then the capital committee. And so it's a marathon process. So a lot of people (platform designers) don't make the distance, because just keeping that team engaged for such a long period of time is expensive and it's exhausting." (D5)

Hwang et al.

As such, although all designers aspired to comprehensive care platforms, their design choices ended up serving either consumers/patients or physicians, enhancing one fragment of the healthcare ecosystem or another, and addressing one symptom, one metric at a time.

4.3 Designers & Patients Unknowingly Disagree

We have so far described how platform designers' actions deviated from their own vision of sensing and AI platforms and the breakdowns that have ensued. In this section, we describe the differences between the designers' and the patients' visions of sensing and AI platforms for mental health. While no particular unintended consequences have materialized to date, these differences might also be seen as problematic.

4.3.1 Different Perspectives on the Value of Face-to-Face Patient-Physician Consultation. Physicians, hospitals, and platform designers primarily regarded face-to-face time between physicians and patients as a costly approach to patient data collection and treatment. They highlighted that it is one key cause of mental health physician burnout. Whether designing platforms for patients or physicians, designers aimed to minimize patient-physician face time to reduce costs.

In stark contrast, patients described face-time with physicians as "*ultimate reward*" (PA9). They had undergone day-to-day sensing and logging, interacting with outpatient monitoring AI, and lab tests, for the very purpose of talking to a physician at the end. Facetime was their definition of health *care*.

4.3.2 Different Perspectives on Privacy. Physicians rarely mentioned any privacy considerations when discussing everyday, personal sensing. Instead, they first and foremost cared about data quality. They held "*a fair amount of skepticism*" (PH1) regarding the reliability of data gathered outside clinical settings. Therefore, they worked to collect multiple sources of data, in order to verify an outpatient's personal sensing and draw more robust conclusions regarding their condition.

Similar to physicians' views in some ways, platform designers curated patient data with the grand vision of delivering integrative care and improving patient wellness, rarely articulating how each type of data might improve what type of care for which patients.

In sharp contrast, patients frequently expressed concerns about privacy risks. They conditioned their willingness to share everyday sensing data based on detailed cost-benefit analyses. They asked, for example, how would physicians handle the data I share? What specific improvements in care quality could I myself or my community receive in return? Would I receive more personalized care suggestions if I share this data about my needs, preferences, and mental health states? Based on their past experience, multiple patients expressed concern that clinicians might not grasp the full extent of their illnesses during brief interactions in clinic offices or therapy rooms (P6, P10–P12). They expected that their everyday sensing data would corroborate the severity of their mental health conditions and help them access the level of care they desired.

Patients wanted to manage, filter, and summarize their data according to these specific benefits before sharing it. Some mentioned they wanted to create folders for their shareable mobile sensing data (PA2, PA5, PA6, PA10). Some were only willing to share this data through specific software applications (PA3, PA7, PA8, PA10). Others (PA10, PA6, and PA7) were only willing to share data through laptop or desktop computers because they contained less personal data (e.g., fewer short messages and call logs).

4.3.3 Different Views on Health Fairness and Equity. Platform designers rarely mentioned AI fairness or health equity considerations during the interviews. They categorized users and stakeholders based on the data they contribute, the features they utilize on the platform, or their demographics. D3's efforts to segment user groups according to whether they "self-selected to" use sleep-related features exemplify this approach.

"I need to measure in real time what everyone's doing inside my product, because those people who self-select to certain features become an audience for me. People who landed on the sleep page and asked questions about sleep stuff in my app ... that's my sleep audience. They've already been hand-raisers, and I can automate [sending] sleep reminders to them." (D3)

Multiple platform designers raised the question of whether integrating a sensing and AI platform to the existing mental health ecosystem would mitigate or exacerbate existing health inequities. When we turned this question back to them, all except one platform designer acknowledged that, the patients who already had access to high-quality mental healthcare are also most likely to benefit from the new sensing and AI platforms. The only exception was brought up by D4, who highlighted the potential of combining pervasive sensing, AI, and telehealth platforms to "*bring clinic offices to people who live remotely or have physical disabilities.*" Nevertheless, none of the actions designers took explicitly addressed these risks or potentials for health equity.

In contrast, patients expressed a strong desire for emergent sensing and AI platforms to enhance equal access to mental healthcare. Patients who self-identified as minorities, in particular, wanted these platforms to bring their communities higher-quality care. They were willing to sacrifice some of their privacy and contribute more personal sensing data to help achieve this goal.

"These new technologies may help with reducing biases. As a young woman, I sometimes have trouble getting doctors to really listen to me or take my concerns seriously. So, I wonder if they just see the data, they may be more inclined to believe me – rather than me trying so hard to advocate for my health situation." (PA6)

"There's a lot of potential to learn more about cis-women. Because they fall into the group that would likely buy an Apple Watch, a Fitbit, or other types of tracking devices." (PA10)

5 DISCUSSION

From wearable behavioral sensors to NLP models detecting people's cognitive states, the combination of pervasive sensing and AI holds great promises and conceivable perils. We wanted to aid designers of sensing and AI platforms in effectively leveraging the platforms' promises while preventing their risks. Toward this goal, we uncovered the platform designers' current practices working in the mental health domain. The designers we interviewed worked hard to achieve some social good. In several aspects, their design actions were notably "human-centered." For example, they all sought to create a pervasive sensing and AI platform not just for profit, but to deliver comprehensive care and improve patient wellness holistically. They carried out *user-centered design* and *service design* processes, both when innovating AI systems and services for those who provided training data, as well as in later stages when extending these services to other relevant medical fields. In addition, designers devoted significant effort to establishing collaborations with multiple stakeholders (e.g., hospitals, insurers, EHR vendors, and patients) and bridging disparate segments of the healthcare ecosystem (e.g., primary and specialty care, mental and physical healthcare). These practices offer a positive case study for designing societal-scale AI platforms.

Our findings also revealed several ways in which platform designers' actions deviated from their own design goals. Moreover, their design goals also differed from those of patients in important ways. In what follows, we discuss: What lessons do these divergences offer, such that future sensing and AI platforms are better aligned with designers' and stakeholders' intentions (§5.1)? Are designers' and stakeholders' intentions indeed the best approaches to designing pervasive sensing and AI? (§5.2)? We encourage HCI scholars and designers to join us in contemplating these questions, as these discussions are imperative for advancing HCI design methods into the era of pervasive sensing and AI platforms.

5.1 Embracing the Complexities of Designing Societal-Scale AI Platforms from Day One

Platform designers aspired to create all-encompassing sensing and AI platforms. However, due to three distinctive challenges, their design choices and actions told a different story. We see a clear opportunity for designers to leverage established HCI design methods to tackle all three challenges, crafting sensing and AI platforms more aligned with their vision.

Designing with Diverse Data Sources; Designing for All Stakeholders. When curating data, the nature of the data economy wherein data is power, and quality data embodying knowledge is even greater power—compelled platform designers to prioritize data contributors (i.e., either consumer tech users or clinicians in partnering hospitals) as their primary stakeholders. Designers designed their sensing and AI platforms for the data contributors, sometimes at the expense of deviating from patients' vision of preferred mental healthcare.

Moreover, this data-driven market force was creating a renewed divide between everyday wellness and clinical care, even though the once rigid legal and economic barriers separating consumer and medical devices/data are eroding [23, 47]. To create accurate diagnostic AI models, tech companies and startups must first curate patient data, either from patients directly (e.g., via consumer wellness apps) or from hospitals' EHR data. Consequently, two distinctive types of mental healthcare applications emerged: One treats patients as primary users, optimizing their AI models and interventions for maximal everyday app engagement; the other treats physicians and hospitals as primary users, optimizing time-and cost-efficiency in clinical environments. As a result, even though commercial mental health apps and EHR systems increasingly offer a similar set of services (e.g., chatbots, meditation apps, wearable devices), when a patient enters professional mental healthcare, they nonetheless need to change to a new set of applications and physicians nonetheless need to assess them anew.

Service design methods such as the *value proposition canvas* can potentially assist designers in navigating the intricate stakeholder relationships inherent in a pervasive sensing and AI platform. Moreover, various speculative and participatory design methods [13, 37, 58] can also help align different stakeholders' visions and values.

Central to these methodologies is the recognition that designers should not only account for the immediate users of their systems, but also the indirect users and stakeholders who engage with it through data and data economy [59]. Moreover, designers of a pervasive sensing and AI platform should consider not only its current users, but also anticipate the needs of future users and stakeholders, as the platform grows and AI systems gain new insights from the sensor data. User-centered design methods (e.g., conducting user interviews, and evaluating system prototypes with users) do not automatically guarantee human-centered technology designs [59]. Designing with a narrow group of stakeholders in mind and assessing only their here-and-now experiences may inadvertently lead to designs that harm other stakeholders or result in unintended societal consequences.

Delivering Diverse AI Interventions within Existing Sociotechnical Infrastructure. When implementing AI systems to deliver comprehensive care, the fragmented nature of the healthcare infrastructure—where various types of care interventions often take place across different hospitals and medical specialties—forced platform designers to focus on only one or two types of interventions. After all, sensing and AI platforms extends the existing, clinician-driven healthcare infrastructure, rather than replacing it.

This breakdown echoes the recent research that called for HCI designers to address the broader "organizational, cultural, regulatory, environmental conditions" surrounding computational systems [15, 57]. Methods emerging from this research, such as *Infrastructural Speculations* [57], have the potential to assist designers in effectively situating a pervasive sensing and AI platform within diverse infrastructural contexts.

Evaluating Sensing and AI Platforms Holistically. Designers faced challenges in *evaluating the efficacy of a cascade of data streams and AI models* on patients' overall well-being. The challenge forced designers to employ partial, yet more practical, evaluation metrics. Fairness and equity across patient populations were an afterthought.

Open research questions abound in creating better evaluation metrics for these platforms regarding human-AI interaction. For example, we see a clear opportunity to implement HCI's many AI fairness methods in designing pervasive sensing and AI platforms. Moreover, given patients' privacy mandates and the platforms' complex data ecology, would these AI systems require so much patient daily data to become accurate that they can only be *either* useful *or* privacy-preserving? Considering that certain outcomes of a societal-scale AI platform (e.g., filter bubbles, misinformation on social media) would not emerge until its users reach critical mass [3], how might HCI researchers actively work to anticipate other risks in sensing and AI platforms for mental health?

There is also a critical need for research that incorporates the currently unquantifiable aspects of healthcare into the evaluation of sensing and AI platforms. Metrics of patient-physician relationships and patients' experience of care, among others, have so far remained an afterthought in designing pervasive sensing and AI. While delegating work from over-worked psychiatrists to AI, how can sensing and AI platforms ensure mental health patients still feel cared for? How can we design future AI models and services to ensure patient-psychiatrist rapport, a crucial element in health *care*? These are important future research topics in making health sensing and AI platforms human-centered.

Finally, it is worth noting that these three challenges—the data economy and related stakeholder dynamics, societal-organizational infrastructure, and evaluation challenges of AI assemblages—might be common complexities inherent in various large-scale sensing and AI platforms. If true, platform designers can potentially include these complexities as standard considerations when starting new sensing and AI platforms, and use the HCI methods above to address them. In doing so, they can potentially address a broader spectrum of stakeholders' needs and values and provide more diverse AI interventions. We encourage future research to develop more case studies in sensing and AI platform design, and explore potential new, generalizable design approaches.

5.2 Framing the Problem of Designing Societal-Scale AI Platforms

We have so far discussed research opportunities around expanding designers' current considerations and embracing the full sociotechnical complexities of pervasive sensing and AI platforms. Parallel to these efforts should be methods that enable designers to gain a holistic understanding of these complexities, *without resorting to oversimplification or ignorance*. Designers navigate complex design problems through the process of problem framing and reframing [51]. By offering designers new framings for the complexities surrounding sensing and AI platforms, we can potentially enhance their design actions.

The designers we interviewed did not intentionally frame or delimit the task of designing a pervasive sensing and AI platform. Instead, they assumed a three-step process, with each step addressing a specific aspect of the platform. Illustrated in Figure 4 from the bottom upward, designers addressed the complexities around (1) users and their sensing data, (2) building AI connecting existing healthcare infrastructure, and finally, (3) evaluation. This approach did not account for the technical complexities around building a cascade of interconnected AI models (as shown in Figure 4 as the linked boxes in light blue), the issues of AI fairness and health equity, or the complexities of human-AI co-evolution. These overlooked considerations may have contributed to the disconnect between patients' values and the design of the platform.

How might designers gain a holistic understanding of the human, AI, data, organizational, and infrastructure complexities around sensing and AI platforms, in order to design them in a principled manner? This is an important research question for HCI, as smart and connected health/workplaces/homes/cities become increasingly mainstream and AI systems grow ever larger in scale [21, 31– 33, 38, 44, 64]. We highlight three potential approaches that have shown promise in prior work but have not yet been utilized in designing sensing and AI platforms. Therefore, they merit further study.

Problem Framing and Scoping. Designers can focus on understanding and designing a slice of the pervasive sensing and AI platform; one that serves as a microcosm of the larger platform. For instance, given the pivotal role of patient data in shaping mental health sensing/AI platforms, could a trace-the-data approach enable designers to apprehend the broader societal-technical dynamics of these platforms? Prior research in human-AI interaction has proposed the method of *data swim lane* [61]. We see an opportunity to adapt this method for designing pervasive sensing and AI platforms.

In addition to data, the processes of care transition also have the potential to serve as a microcosm of the broader sociotechnical dynamics, making them a useful focal point for designers. Our interviews revealed that breakdowns most frequently occurred at these transition points: when patients were admitted to or discharged from mental healthcare clinics, moved between mental and physical healthcare facilities, or transitioned from general care to specialty care. Can transition design — the design of processes through which patients (along with their data, the AI systems reliant on that data, their insurance, etc.) undergo care transitions — be a useful way to design the larger, sensing-and-AI-enhanced care ecosystem, considering that these are the very points where breakdowns occur?

Finally, the business model of a sensing and AI platform, along with related laws and policies, could also offer a valuable lens into the platform's design. Consider the issue of scale in platform design. While all platform designers lamented the challenges posed by EPIC, the all-encompassing EHR platform, they also aspire to construct similarly comprehensive platforms themselves. The economic incentive driving this pursuit of scale is understandable. However, from the perspectives of patient outcomes and societal benefits, questions arise: How large should a pervasive sensing and AI platform be? Is there a threshold beyond which it becomes too large? Do our societies truly aspire to a single computational platform capable of delivering all types of care, spanning from mental to physical health, and from everyday self-care to medical treatment? Service design methods, which address value proposition design and corporate policy design questions, can potentially be adapted to address these questions. In summary, designing the push-andpull between business models and policies can be a critical part of designing a socially beneficial sensing and AI platform.

Abstraction and Simulation. Another approach to navigating the complexities of sensing and AI platforms is through simulation, i.e., projecting the possible and likely consequences of a design choice in order to understand if this is a desirable future. Existing HCI design methods serve as valuable reference points for this endeavor. Social simulation and prototyping allow designers to forecast effects at scales [16, 24, 43], and speculative design motivates one to focus on the future and think beyond the capability of existing technology [17]. We encourage future research to apply these

methods to sensing and AI platform design contexts. For instance, what cascading effects might occur when the applications and their user base scale up? What types of new interventions might these applications deliver when new AI capabilities emerge?

To adopt these methods, it is essential to align and move human-AI interaction design, technical setups in the mental health ecosystem, and policy design through joint forces [50, 60]. As technical setups in hospitals lag behind, it becomes critical to offer platform designers a workable sandbox and dataset for development, so that they can effectively test societal-scale AI design *before* deployment in these more dated clinical systems. As our interviewees mentioned, this often requires law enforcement to ensure platform designers attain sufficient support – often through demanding major vendors to release datasets for development. The making of such policies demands insights into what developers and designers need during their working process of model development.

Growth Planning and Staging. In our interviews, platform designers all worked towards creating an all-encompassing sensing and AI platform that reaches all patients and all corners of the healthcare ecology. Interestingly, few actively mapped out a clear path toward this ambitious objective or identified specific milestones. Instead, they often narrowed the focus of their designs when encountering challenges and rarely discussed strategies for scaling back up in the future. For instance, no designer planned what data they need to collect from patients in order to deliver multiple types of care, or how their platforms might accommodate diverse user groups beyond the initial one.

There are open research opportunities around growth staging, that is, assessing and categorizing a pervasive sensing and AI platform's growth trajectory into distinct stages, and designing goals, forms, functions, and value propositions tailored to each stage accordingly. Building upon [42], complex sociotechnical systems cannot be built overnight. Designing them requires a combination of incremental "*muddling through*" and satisficing. This incremental approach to sensing and AI platform design could open up a wealth of design and research opportunities: Could anticipatory data collection and feature planning make the expansion of the platform easier? How might designers set different AI fairness and equity goals as the platform grows larger?

6 LIMITATIONS AND FUTURE WORK

We acknowledge the limitations of our methods and suggest avenues for future research. First, while we identified stakeholder tensions through interviewing with three participant groups separately, we haven't devised plausible solutions to these tensions yet. Bringing these participants together, such as through participatory workshops, could foster discussions on value alignment and even lead to agreed-upon actions to tackle these tensions. We encourage future work to consider methods that can help form communities that might continue working together beyond the research process. Second, as the present work focuses on platform designers' current practices, discussions during our interviews were highly grounded on the current technical capabilities of AI and sensing devices. Future work can consider other speculative approaches that project the longer-term impact of societal-scale AI design based possible technological advances. Finally, while we attempted to incorporate perspectives from direct stakeholders (i.e., patients and physicians), future work should further consider drawing insights from *indirect stakeholders* [19] (e.g., specialists who collect clinical data through lab tests) as they jointly contribute to the future healthcare ecosystem.

ACKNOWLEDGMENTS

This work is partially supported by National Science Foundation under grants IIS-2313077 and IIS-2302977. Angel Hwang is partially supported by Weill Cornell Medicine Seed Grants for Intercampus Collaborative Project "*Improving the Robustness of Mobile Sensing and AI Systems for Mental Health Care*". Qian Yang is supported by Schmidt Futures' AI2050 Early Career Fellowship.

REFERENCES

- Saeed Abdullah, Mark Matthews, Ellen Frank, Gavin Doherty, Geri Gay, and Tanzeem Choudhury. 2016. Automatic detection of social rhythms in bipolar disorder. *Journal of the American Medical Informatics Association* 23, 3 (03 2016), 538–543. https://doi.org/10.1093/jamia/ocv200
- [2] The State Of Mental Health In America. 2023. The State Of Mental Health In America - 2023 Key Statistics. https://www.mhanational.org/issues/state-mentalhealth-america
- [3] Barry Brown, Susanne Bødker, and Kristina Höök. 2017. Does HCI Scale? Scale Hacking and the Relevance of HCI. *Interactions* 24, 5 (aug 2017), 28–33. https: //doi.org/10.1145/3125387
- [4] Stevie Chancellor, Michael L Birnbaum, Eric D Caine, Vincent MB Silenzio, and Munmun De Choudhury. 2019. A taxonomy of ethical tensions in inferring mental health states from social media. In Proceedings of the Conference on Fairness, Accountability, and Transparency. 79–88.
- [5] Stevie Chancellor and Munmun De Choudhury. 2020. Methods in predictive techniques for mental health status on social media: a critical review. NPJ digital medicine 3, 1 (2020), 1–11.
- [6] Nicole Chauriye. 2016. Wearable devices as admissible evidence: Technology is killing our opportunity to lie. Catholic University Journal of Law and Technology 24, 2 (2016), 9.
- [7] Supriya Chinthavali, Varisara Tansakul, Sangkeun Lee, Anika Tabassum, Jeff Munk, Jan Jakowski, Michael Starke, Teja Kuruganti, Heather Buckberry, and Jim Leverette. 2019. Quantification of Energy Cost Savings through Optimization and Control of Appliances within Smart Neighborhood Homes. In Proceedings of the 1st ACM International Workshop on Urban Building Energy Sensing, Controls, Big Data Analysis, and Visualization (New York, NY, USA) (UrbSys'19). Association for Computing Machinery, New York, NY, USA, 59–68. https://doi.org/10.1145/ 3363459.336355
- [8] Tanzeem Choudhury. 2023. 15+ Years in the Making When will the Digital Mental Health Revolution Deliver on its Promise? https://www.csail.mit.edu/event/15-years-making-when-will-digital-mentalhealth-revolution-deliver-its-promise Seimnar at MIT MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) [Accessed: 2024 01 28].
- [9] Sam Clarke and Jess Whittlestone. 2022. A Survey of the Potential Long-Term Impacts of AI: How AI Could Lead to Long-Term Changes in Science, Cooperation, Power, Epistemics and Values. In Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (Oxford, United Kingdom) (AIES '22). Association for Computing Machinery, New York, NY, USA, 192–202. https://doi.org/10.1145/ 3514094.3534131
- [10] Design Council. 2005. The 'double diamond' design process model. Design Council (2005).
- [11] Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering Shifts to Suicidal Ideation from Mental Health Content in Social Media. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 2098–2110. https://doi.org/10. 1145/2858036.2858207
- [12] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2021. Stakeholder Participation in AI: Beyond "Add Diverse Stakeholders and Stir". arXiv:2111.01122 (Nov. 2021). https://doi.org/10.48550/arXiv.2111.01122 arXiv:2111.01122 [cs].
- [13] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The participatory turn in ai design: Theoretical foundations and the current state of practice. In Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization. 1–23.
- [14] Brianna Dym, Namita Pasupuleti, and Casey Fiesler. 2022. Building a Pillowfort: Political Tensions in Platform Design and Policy. Proc. ACM Hum.-Comput.

Societal-Scale Human-AI Interaction Design?

CHI '24, May 11-16, 2024, Honolulu, HI, USA

Interact. 6, GROUP, Article 16 (Jan 2022), 23 pages. https://doi.org/10.1145/3492835

- [15] W. Keith Edwards, Mark W. Newman, and Erika Shehan Poole. 2010. The Infrastructure Problem in HCL In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 423–432. https: //doi.org/10.1145/1753326.1753390
- [16] Daniel A Epstein, Fannie Liu, Andrés Monroy-Hernández, and Dennis Wang. 2022. Revisiting Piggyback Prototyping: Examining Benefits and Tradeoffs in Extending Existing Social Computing Systems. Proceedings of the 25th ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW '22) (2022).
- [17] Pedro Gil Farias, Roy Bendor, and Bregje F. Van Eekelen. 2022. Social dreaming together: A critical exploration of participatory speculative design. In Proceedings of the Participatory Design Conference 2022 - Volume 2. ACM, Newcastle upon Tyne United Kingdom, 147–154. https://doi.org/10.1145/3537797.3537826
- [18] Pedro M. Fernandes, Francisco C. Santos, and Manuel Lopes. 2020. Adoption Dynamics and Societal Impact of AI Systems in Complex Networks. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (New York, NY, USA) (AIES '20). Association for Computing Machinery, New York, NY, USA, 258–264. https://doi.org/10.1145/3375627.3375847
- [19] Batya Friedman, Peter H. Kahn, Alan Borning, and Alina Huldtgren. 2013. Value Sensitive Design and Information Systems. Vol. 16. Springer Netherlands, Dordrecht, 55–95. https://doi.org/10.1007/978-94-007-7844-3_4
- [20] Sarah Gibbons. 2018. Journey Mapping 101. https://www.nngroup.com/articles/ journey-mapping-101/.
- [21] Thomas Krendl Gilbert, Sarah Dean, Nathan Lambert, Tom Zick, and Aaron Snoswell. 2022. Reward Reports for Reinforcement Learning. https://doi.org/10. 48550/ARXIV.2204.10817
- [22] Thomas Krendl Gilbert, Aaron J Snoswell, Michael Dennis, Rowan McAllister, and Cathy Wu. 2022. Sociotechnical Specification for the Broader Impacts of Autonomous Vehicles. arXiv preprint arXiv:2205.07395 (2022).
- [23] William J Gordon and Kenneth D Mandl. 2020. The 21st century cures act: a competitive apps market and the risk of innovation blocking. *Journal of Medical Internet Research* 22, 12 (2020), e24824.
- [24] Catherine Grevet and Eric Gilbert. 2015. Piggyback prototyping: Using existing, large-scale social computing systems to prototype new ones. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. 4047–4056.
- [25] Nava Haghighi, Matthew Jörke, Yousif Mohsen, Andrea Cuadra, and James A. Landay. 2023. A Workshop-Based Method for Navigating Value Tensions in Collectively Speculated Worlds. In Proceedings of the 2023 ACM Designing Interactive Systems Conference. ACM, Pittsburgh PA USA, 1676–1692. https: //doi.org/10.1145/3563657.3595992
- [26] Christina Harrington and Tawanna R Dillahunt. 2021. Eliciting Tech Futures Among Black Young Adults: A Case Study of Remote Speculative Co-Design. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. ACM, Yokohama Japan, 1–15. https://doi.org/10.1145/3411764.3445723
- [27] Derek Hatfield, Lynn McCullough, Shelby HB Frantz, and Kenin Krieger. 2010. Do we know when our clients get worse? An investigation of therapists' ability to detect negative client change. *Clinical Psychology & Psychotherapy: An International Journal of Theory & Practice* 17, 1 (2010), 25–32.
- [28] Nicholas C Jacobson, Kate H Bentley, Ashley Walton, Shirley B Wang, Rebecca G Fortgang, Alexander J Millner, Garth Coombs III, Alexandra M Rodman, and Daniel DL Coppersmith. 2020. Ethical dilemmas posed by mobile health and machine learning in psychiatry research. *Bulletin of the World Health Organization* 98, 4 (2020), 270.
- [29] Sehj Kashyap, Keith E Morse, Birju Patel, and Nigam H Shah. 2021. A survey of extant organizational and computational setups for deploying predictive models in health systems. *Journal of the American Medical Informatics Association* 28, 11 (2021), 2445–2450.
- [30] Corey LM Keyes. 2005. Mental illness and/or mental health? Investigating axioms of the complete state model of health. *Journal of Consulting and Clinical Psychology* 73, 3 (2005), 539.
- [31] Jahoon Koo and Young-Gab Kim. 2021. Interoperability Requirements for a Smart City. In Proceedings of the 36th Annual ACM Symposium on Applied Computing (Virtual Event, Republic of Korea) (SAC '21). Association for Computing Machinery, New York, NY, USA, 690–698. https://doi.org/10.1145/3412841.3441948
- [32] Shashwat Kumar, Debajyoti Datta, Guimin Dong, Lihua Cai, Mehdi Boukhechba, and Laura Barnes. 2022. Leveraging Mobile Sensing and Bayesian Change Point Analysis to Monitor Community-Scale Behavioral Interventions: A Case Study on COVID-19. ACM Trans. Comput. Healthcare 3, 4, Article 37 (Nov 2022), 13 pages. https://doi.org/10.1145/3524886
- [33] Karen Levy. 2022. Data Driven: Truckers, Technology, and the New Workplace Surveillance. Princeton University Press.
- [34] Luona Lin, Assefa Meron, and Karen Stamm. 2023. Practitioners are overworked and burned out, and they need our support. https://www.apa.org/monitor/2023/ 04/psychologists-covid-burnout
- [35] Steve Lohr. 2014. Huge New York Development Project Becomes a Data Science Lab. https://bits.blogs.nytimes.com/2014/04/14/huge-new-york-development-

project-becomes-a-data-science-lab/

- [36] Nicole Martinez-Martin, Henry T Greely, and Mildred K Cho. 2021. Ethical Development of Digital Phenotyping Tools for Mental Health Applications: Delphi Study. JMIR mHealth and uHealth 9, 7 (Jul 2021), e27343. https: //doi.org/10.2196/27343
- [37] Jessica K. Miller, Batya Friedman, and Gavin Jancke. 2007. Value tensions in design: the value sensitive design, development, and appropriation of a corporation's groupware system. In Proceedings of the 2007 International ACM Conference on Supporting Group Work - GROUP '07. ACM Press, Sanibel Island, Florida, USA, 281. https://doi.org/10.1145/1316668
- [38] Jakub Mlynar, Farzaneh Bahrami, André Ourednik, Nico Mutzner, Himanshu Verma, and Hamed Alavi. 2022. AI beyond Deus Ex Machina – Reimagining Intelligence in Future Cities with Urban Experts. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 370, 13 pages. https://doi.org/10.1145/3491102.3517502
- [39] David C. Mohr, Katie Shilton, and Matthew Hotopf. 2020. Digital phenotyping, behavioral sensing, or personal sensing: names and transparency in the digital age. npj Digital Medicine 3, 1 (March 2020), 1–2. https://doi.org/10.1038/s41746-020-0251-5 Number: 1 Publisher: Nature Publishing Group.
- [40] David C Mohr, Mi Zhang, and Stephen M Schueller. 2017. Personal sensing: understanding mental health using ubiquitous sensors and machine learning. *Annual Review of Clinical Psychology* 13 (2017), 23.
- [41] Eric Monteiro, Neil Pollock, Ole Hanseth, and Robin Williams. 2013. From Artefacts to Infrastructures. Computer Supported Cooperative Work (CSCW) 22, 4-6 (Aug 2013), 575-607. https://doi.org/10.1007/s10606-012-9167-1
- [42] Donald A. Norman and Pieter Jan Stappers. 2015. DesignX: Complex Sociotechnical Systems. She Ji: The Journal of Design, Economics, and Innovation 1, 2 (2015), 83–106. https://doi.org/10.1016/j.sheji.2016.01.002
- [43] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 74, 18 pages. https://doi.org/10.1145/3526113.3545616
- [44] Philip WT Pong, Anuradha M Annaswamy, Benjamin Kroposki, Yingchen Zhang, Ram Rajagopal, Gil Zussman, and H Vincent Poor. 2021. Cyber-enabled grids: Shaping future energy systems. Advances in Applied Energy 1 (2021), 100003.
- Shaping future energy systems. Advances in Applied Energy 1 (2021), 100003.
 [45] Inbar Raber, Cian P McCarthy, and Robert W Yeh. 2019. Health insurance and mobile health devices: opportunities and concerns. JAMA 321, 18 (2019), 1767–1768.
- [46] Md Tahmid Rashid, Na Wei, and Dong Wang. 2023. A Survey on Social-Physical Sensing: An Emerging Sensing Paradigm That Explores the Collective Intelligence of Humans and Machines. *Collective Intelligence* 2, 2 (apr 2023), 37 pages. https: //doi.org/10.1177/26339137231170825
- [47] Tera L Reynolds, Meghna Kaligotla, and Kai Zheng. 2022. Investigating the Interoperable Health App Ecosystem at the Start of the 21st Century Cures Act. AMIA Annual Symposium Proceedings 2022 (2022), 942.
- [48] The White House Briefing Room. 2023. Fact Sheet: Biden-Harris Administration Announces New Actions to Tackle Nation's Mental Health Crisis. https://www.whitehouse.gov/briefing-room/statementsreleases/2023/05/18/fact-sheet-biden-harris-administration-announces-newactions-to-tackle-nations-mental-health-crisis/
- [49] Muhammad Sadi Adamu. 2021. Problematising Identity, Positionality, and Adequacy in HCI4D Fieldwork: A Reflection. In 3rd African Human-Computer Interaction Conference: Inclusiveness and Empowerment (Maputo, Mozambique) (AfriCHI 2021). Association for Computing Machinery, New York, NY, USA, 65–74. https://doi.org/10.1145/3448696.3448703
- [50] Hauke Sandhaus, Wendy Ju, and Qian Yang. 2023. Towards Prototyping Driverless Vehicle Behaviors, City Design, and Policies Simultaneously. (2023). https: //doi.org/10.48550/ARXIV.2304.06639
- [51] Donald A Schon and Vincent DeSanctis. 1986. The Reflective Practitioner: How Professionals Think in Action. Taylor & Francis.
- [52] Douglas Schuler and Aki Namioka. 2017. Participatory Design: Principles and Practices (1 ed.). CRC Press. https://doi.org/10.1201/9780203744338
- [53] Anja Thieme, Danielle Belgrave, and Gavin Doherty. 2020. Machine Learning in Mental Health: A Systematic Review of the HCI Literature to Support the Development of Effective and Implementable ML Systems. ACM Trans. Comput.-Hum. Interact. 27, 5, Article 34 (aug 2020), 53 pages. https://doi.org/10.1145/ 3398069
- [54] Margaret L. Traeger, Sarah Strohkorb Sebo, Malte Jung, Brian Scassellati, and Nicholas A. Christakis. 2020. Vulnerable robots positively shape human conversational dynamics in a human-robot team. Proceedings of the National Academy of Sciences 117, 12 (2020), 6370–6375. https://doi.org/10.1073/pnas.1910402117 arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.1910402117
- [55] Axelle Van Wynsberghe and Ângela Guimarães Pereira. 2022. The politics and imaginary of 'autonomous vehicles': a participatory journey. *Humanities and*

Social Sciences Communications 9, 1 (Aug. 2022), 284. https://doi.org/10.1057/ s41599-022-01209-1

- [56] Angela Wilkinson and Roland Kupers. 2013. Living in the futures. Harvard Business Review 91, 5 (2013), 118–127.
- [57] Richmond Y Wong, Vera Khovanskaya, Sarah E Fox, Nick Merrill, and Phoebe Sengers. 2020. Infrastructural speculations: Tactics for designing and interrogating lifeworlds. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. 1–15.
- [58] Richmond Y. Wong and Tonya Nguyen. 2021. Timelines: A World-Building Activity for Values Advocacy. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. ACM, Yokohama Japan, 1–15. https://doi.org/10. 1145/3411764.3445447
- [59] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-Examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376301
- [60] Qian Yang, Richmond Y. Wong, Thomas Gilbert, Margaret D. Hagan, Steven Jackson, Sabine Junginger, and John Zimmerman. 2023. Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3544549. 3573827
- [61] Nur Yildirim, Alex Kass, Teresa Tung, Connor Upton, Donnacha Costello, Robert Giusti, Sinem Lacin, Sara Lovic, James M O'Neill, Rudi O'Reilly Meehan, Eoin Ó Loideáin, Azzurra Pini, Medb Corcoran, Jeremiah Hayes, Diarmuid J Cahalane, Gaurav Shivhare, Luigi Castoro, Giovanni Caruso, Changhoon Oh, James Me-Cann, Jodi Forlizzi, and John Zimmerman. 2022. How Experienced Designers of Enterprise Applications Engage AI as a Design Material. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 483, 13 pages. https://doi.org/10.1145/3491102.3517491
- [62] Daisy Yoo. 2018. Designing with (Political) Complexity: Understanding Stakeholders, Emotion, Time, and Technology in the Case of Medical Aid-in-dying. Ph.D. Dissertation.
- [63] Daisy Yoo, Anya Ernest, Sofia Serholt, Eva Eriksson, and Peter Dalsgaard. 2019. Service design in HCI research: the extended value co-creation model. In Proceedings of the Halfway to the Future Symposium 2019. 1–8.
- [64] Hubert D. Zając, Dana Li, Xiang Dai, Jonathan F. Carlsen, Finn Kensing, and Tariq O. Andersen. 2023. Clinician-Facing AI in the Wild: Taking Stock of the Sociotechnical Challenges and Opportunities for HCL. ACM Trans. Comput.-Hum. Interact. 30, 2, Article 33 (mar 2023), 39 pages. https://doi.org/10.1145/3582430

APPENDIX

A STUDY PARTICIPANT SCREENING

We conducted a pre-screening with platform designer interviewees using the following criteria:

- (1) Their work involved planning, designing, and implementing large AI systems for mental healthcare.
- (2) Their work connected and integrated resources from clinical and non-clinical settings.
- (3) Their work required them to work with multi-stakeholder groups across various healthcare sectors.

A.1 Interviewing Protocol with Platform Designers

With platform designers, we conducted retrospective interviews, asking them to walk us through their recent processes of implementing new devices or applications into practice step-by-step. From there, we unpacked their common motivations, pain points, and ideal scenarios at each stage of their design and work processes. Furthermore, we asked them to elaborate on their approaches to evaluating the success and failure modes of their design.

- (1) Opening
- (2) Tell us one experience moving a new device, or new types of AI (e.g., risk models, outpatient monitoring AI) into practice. What was that process like?

- (3) Questions about **ecosystem**:
 - (a) How do you conceptualize the mental healthcare "ecosystem"?
 - (b) What types of data are currently collected for mental healthcare purposes?
 - (c) What types of technology/tools (e.g., devices, software, systems) are commonly used in support of the current mental healthcare ECOSYSTEM?
 - (d) What types of AI systems are currently used in the mental healthcare ECOSYSTEM, if any?
 - (e) What are some common organizational boundaries in the current mental healthcare ECOSYSTEM?
- (4) Questions about stakeholders:
 - (a) What stakeholders are involved?
 - (b) What are their concerns, needs, and priorities?
 - (c) Are there tensions among these stakeholders' needs?
 - (d) What are the relations and power dynamics in between these stakeholders?
 - (e) What role does stakeholders' experience in the mental healthcare ECOSYSTEM play in your design?
- (5) Questions about societal-scale AI design goals:
 - (a) What are your common goals across your design projects?
 - (b) Based on your design goals, how do you evaluate the success and failure modes of your design respectively?
 - (c) How do you account for stakeholder statistication in your design?
 - (d) How do you account for clinical performance and outcomes in your design?
- (6) Questions about concerns, risk prediction, and management:
 - (a) What risks do you account for your design?
 - (b) How do you assess such risks?
 - (c) What types of risk are deal breakers of a design plan?
 - (d) How do you manage risks around outpatient data or data collected outside of clinical settings?
- (7) Questions about implementation process of societal-scale AI design plans:
 - (a) What are the common stages in your implementation plan?
 - (b) Who typically contributes to an implementation plan? What are their roles?
 - (c) Please walk us through a concrete example of the implementation plan for one of your past or present projects.
- (8) Questions about situating societal-scale AI in mental healthcare:
 - (a) Please describe your ideal vision for the future of mental healthcare.
 - (b) What are some obstacles that prevent this ideal vision?
 - (c) Where and how do the types of design you work on contribute to this vision?
 - (d) Where and how do the types of design you work on bring obstacles to this vision?
 - (e) (If the participant also worked on projects in the physical healthcare area) Which of your prior answers may or may not apply?

A.2 Interview Protocol with Physicians and Patients

With physicians and patients, we first asked them to visualize how the current mental healthcare system operates on a blank digital whiteboard and identified where they see the use of personal sensing and AI systems, if at all. Based on their visualizations, we asked them to elaborate on their struggles with the current healthcare ecosystem and pinpoint where they hope for changes. Next, we provided them with some contextual knowledge on designing and integrating consumer technologies into the mental healthcare ecosystem by walking participants through Figure 1. From there, participants shared their responses toward such visions of future healthcare and, specifically, their willingness to share data to fuel large-scale AI models in such ecosystems.

- (1) Reviewing the current mental healthcare system: To gain an initial understanding of the participants' awareness of the AIs and the current healthcare system, we invited participants to visualize how the current mental healthcare system operates on a template. Specifically, participants would mark whether and where AI applications are used in the present system. During this section, participants would adopt a think-aloud approach, elaborating their present pain points and desires for changes while they map out the current system.
- (2) An item-by-item inquiry of data-centric decisions: To probe users' rationales of data sharing practices with healthcare providers, prior to the study, we asked participants to download and examine their data exported from Google Takeout ¹, encompassing a wide range of user-generated data collected through their Chrome browsers. During the interview, participants were asked whether they would like to share each of the data items with their mental healthcare providers and to elaborate on their decision-making rationales for data sharing and data management.
- (3) Envisioning ideal future healthcare ecosystem: Next, to co-design the future of healthcare with stakeholders, we showed participants Figure 1 and walked them through how future mental healthcare can be accomplished through implementing an ecosystem. Afterward, we invited participants to share their comments and concerns for this future vision, proposing any design recommendations and elaborating on their own ideal of future mental healthcare.

B POSITIONALITY STATEMENT

We formed our research team with two AI/machine learning researchers (one's work focuses on theory-based simulation approaches to study AI fairness and interactions across multiple AI systems; the other specializes in leveraging wearable sensing data to build AI applications for mental healthcare) and two HCI researchers (one focuses on design research and new methods to study large-scale AI systems; the other applies empirical methods to study multi-human, multi-agent interactions at various scales). Per our self-identified demographics, the team consists of two East Asian females and two U.S. White American males. We expect the varying backgrounds of our research team to also contribute to more diverse perspectives throughout the research process.

C STUDY PARTICIPANT DEMOGRAPHICS

Age		Race/Ethnicity		Preferred Pronouns		
Platform Designers						
18 - 29	1	White	5	She/her/hers	4	
30 - 39	4	Black	1	He/him/his	5	
40 - 49	2	Latinx	1			
50+	4	Asian	2			
		Other	0			
Physicians						
18 - 29	2	White	3	She/her/hers	3	
30 - 39	3	Black	1	He/him/his	2	
40 - 49	0	Latinx	0			
50+	0	Asian	1			
		Other	0			
Patients						
18 - 29	5	White	4	She/her/hers	9	
30 - 39	6	Black	2	He/him/his	4	
40 - 49	0	Latinx	1			
50+	2	Asian	4			
		Other	2			

Table 2: Study Participants Demographics.

D THEMES FROM AFFINITY DIAGRAMMING

See Tables 3 and 4 on the next page.

¹See items available for data export on https://takeout.google.com/

Current Design Practices					
AI-driven design considerations	Ecosystem-driven design considerations				
Promises of societal-scale AI	Inherent tensions among stakeholders				
 All-encompassing vision Oversee networked systems Scale-specific qualities of AI Data-driven design goals AI's needs for mass data Two paths to source data (Data-driven dichotomy) Clinical → non-clinical Non-clinical → clinical Data-driven design practices Scoping the platform by individual stakeholder Scoping the platform by specific symptoms Scoping the platform by technical infrastructure 	 Increase vs. reduce doctor-patient face time Patients want to increase doctor-patient face time Physicians want to decrease doctor-patient face time Decision-makers want to decrease doctor-patient face time Targeted treatment vs. comprehensive care Motivations to treat targeted symptoms Motivations to improve overall wellness Cost efficiency of care types 				
Possible Desig	n Breakdowns				
Implications for AI design	Implications for future mental healthcare				
Individual risks	A shift away from patient-centered care				
 Privacy Concerns for pervasive sensing Skepticism for data quality Fairness Optimism with model robustness 	 Fragmentation of healthcare User-engagement drives personalization Fairness and equity as afterthoughts 				
(Unsmooth) transition risks					
 Policy-driven risks Policies define the clinical wall Policies drive unsmooth data transitions Development-related risks No for-development dataset No sandbox Organizational risks 					

Table 3: Themes emerged from interview data analysis.

Perspectives on Mental Healthcare Ecosystem					
Patients' perspectives	Physicians' perspectives				
Ideal vision of mental healthcare	Ideal vision of mental healthcare				
 Efficiency Smooth patient journey Better-quality care Doctor-patient relationship Current status quo and obstacles Neglect for need of care High cost Limited accessibility 	 Efficiency Streamline procedure Current status quo and obstacles High burn-out Repetitive procedures Non-care-related burdens 				
Perspectives on Societal-Scale AI in Mental Healthcare					
Patients' perspectives	Physicians' perspectives				
Concerns	Concerns				
 Pervasive sensing and privacy risks 	 Data as liability 				
Норе	Норе				
Data advocates for careStreamlined patient journey	Streamlined treatment processImproved clinical outcomes				

Table 4: Themes emerged from patient and physician interview data analysis.