

MIND: Empowering Mental Health Clinicians with Multimodal Data Insights through a Narrative Dashboard

Ruishi Zou*
Columbia University
New York, New York, USA
ruishi.ray.zou@outlook.com

Shiyu Xu*
Columbia University
New York, New York, USA
shiyuxu@umich.edu

Margaret E Morris
Information School
University of Washington, Seattle
Seattle, Washington, USA
margiem@uw.edu

Jihan Ryu
Hamilton-Madison House
New York, New York, USA
jihani.ryu@gmail.com

Timothy D. Becker
New York-Presbyterian Hospital/
Weill Cornell Medicine
New York, New York, USA
tdb2002@med.cornell.edu

Nicholas Allen
Center for Digital Mental Health
University of Oregon
Eugene, Oregon, USA
nallen3@uoregon.edu

Anne Marie Albano
Columbia University
New York, New York, USA
aa2289@cumc.columbia.edu

Randy Auerbach
Department of Psychiatry
Columbia University
New York City, New York, USA
rpa2009@cumc.columbia.edu

Dan Adler
Information Science
Cornell University
New York, New York, USA
dadler@infosci.cornell.edu

Varun Mishra
Northeastern University
Boston, Massachusetts, USA
v.mishra@northeastern.edu

Lace M. Padilla
Khoury College of Computer Science
Northeastern University
Boston, Massachusetts, USA
l.padilla@northeastern.edu

Dakuo Wang
Northeastern University
Boston, Massachusetts, USA
d.wang@northeastern.edu

Ryan Sultan
Columbia University
New York, New York, USA
rs3511@cumc.columbia.edu

Xuhai Xu
Department of Biomedical
Informatics
Columbia University
New York City, New York, USA
xx2489@columbia.edu

Abstract

Advances in data collection enable the capture of rich patient-generated data: from passive sensing (e.g., wearables and smartphones) to active self-reports (e.g., cross-sectional surveys and ecological momentary assessments). Although prior research has demonstrated the utility of patient-generated data in mental healthcare, significant challenges remain in effectively presenting these data streams along with clinical data (e.g., clinical notes) for clinical decision-making. Through co-design sessions with five clinicians, we propose MIND, a large language model-powered dashboard designed to present clinically relevant multimodal data insights

for mental healthcare. MIND presents multimodal insights through narrative text, complemented by charts communicating underlying data. Our user study ($N=16$) demonstrates that clinicians perceive MIND as a significant improvement over baseline methods, reporting improved performance to reveal hidden and clinically relevant data insights ($p < .001$) and support their decision-making ($p = .004$). Grounded in the study results, we discuss future research opportunities to integrate data narratives in broader clinical practices.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Applied computing** → **Health informatics**.

Keywords

Clinical Dashboard, Multimodal Data, Narrative Visualization, Natural Language Processing, Mental Healthcare

ACM Reference Format:

Ruishi Zou, Shiyu Xu, Margaret E Morris, Jihan Ryu, Timothy D. Becker, Nicholas Allen, Anne Marie Albano, Randy Auerbach, Dan Adler, Varun

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3790529>

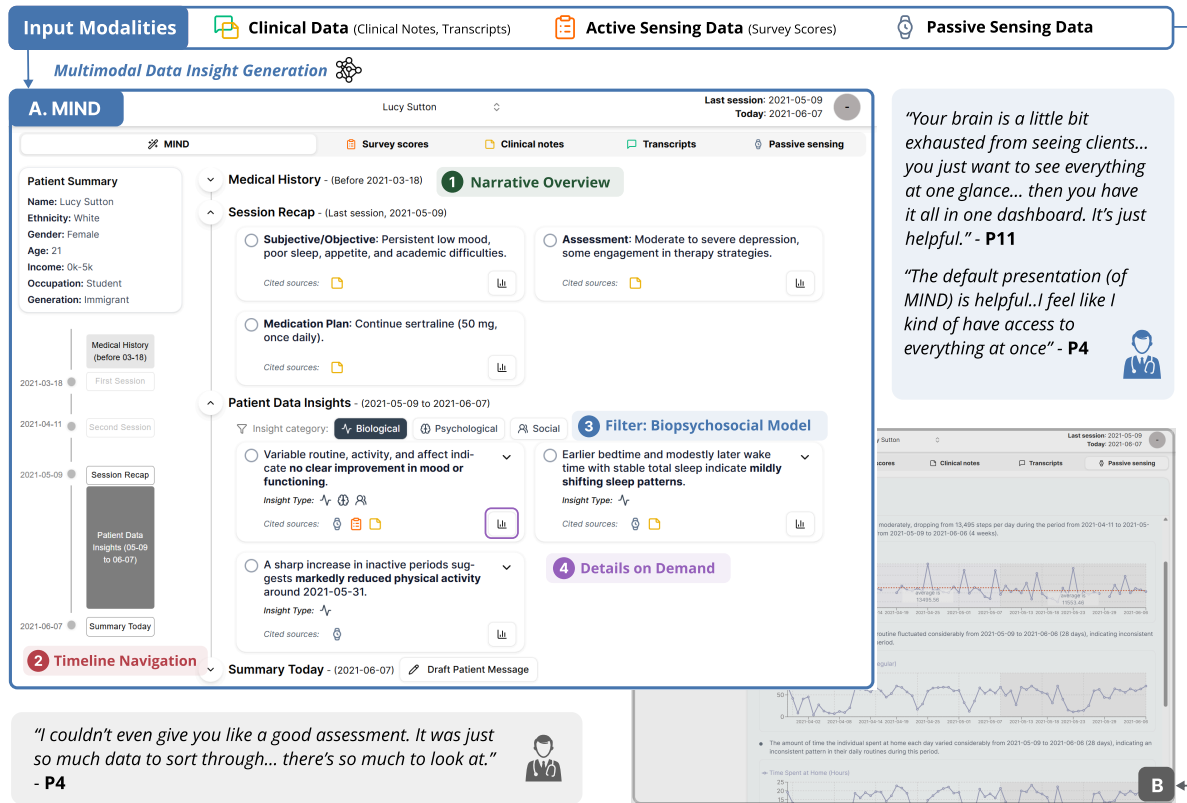


Figure 1: A comparison between the proposed narrative dashboard MIND (A) with a data collection dashboard design (B), showing the data from the same hypothetical patient. MIND presents computationally generated multimodal insights in a narrative way ①, allowing clinicians to better understand multimodal patient data. Clinicians can also easily navigate through timeline ②, filter by insight category ③, and expand details on demand ④.

Mishra, Lace M. Padilla, Dakuo Wang, Ryan Sultan, and Xuhai Xu. 2026. MIND: Empowering Mental Health Clinicians with Multimodal Data Insights through a Narrative Dashboard. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 26 pages. <https://doi.org/10.1145/3772318.3790529>

1 Introduction

Recent advances in technology have enabled patient-generated data to be captured at scale [22, 39]. For instance, mobile phones and wearable devices can passively capture patients’ biological, psychological, and behavioral markers longitudinally, while digital methods such as ecological momentary assessment (EMA) enable frequent self-reports of patients’ states and contexts. While traditional clinical data sources such as electronic health records (EHRs) remain widely used, patient-generated data are increasingly being applied in mental health assessment and care [23, 65, 69, 95]. Specifically, they afford clinicians an alternative and complementary perspective on patients’ daily lives outside clinical settings—one that is difficult to obtain in current practice.

Prior research has investigated three broad data modalities related to mental healthcare [4]: 1) **clinical data** (e.g., clinical notes,

and transcripts), and two types of patient-generated data: 2) **active sensing data** (e.g., survey scores), and 3) **passive sensing data** (e.g., sleep, steps). One of the straightforward ways of leveraging these data modalities is to build machine learning (ML) models to “detect/predict” [5] mental health measures (e.g., stress [25], depression [102]). However, recent evidence suggests these models barely generalize across datasets [2, 62, 104], limiting their real-world applicability. Moreover, building detection/prediction ML models hardly provides clinicians with meaningful insights for treatment decision-making in real-world practices. This leads to another thread of research that takes a more interpretive approach by viewing these data modalities as complementary perspectives to support more informed clinical decision-making [85]. However, without clinical abstraction or contextual explanation, the inherent complexity of multimodal data (e.g., data format: textual clinical data vs. numerical sensing data; temporal frequency: daily passive sensing data vs. weekly active sensing data vs. monthly per-session clinical data) creates significant challenges for clinicians to synthesize and reference in practice [7, 10, 28].

One prevalent solution to present multimodal data at scale is through clinical dashboards (e.g., [100]), a technique that promises

rich data information “at a glance” [31]. Researchers have developed dashboards that accommodate clinical, active sensing, and passive sensing data in separate views (Fig. 1.B) [57, 100, 107]—a design pattern coined as **data collection dashboard** by Bach et al. [8]. While they accommodate large volumes of information, data collection dashboards are designed with minimal curation of what and how the data is presented, leaving mental health clinicians to forage for relevant data and synthesize insights themselves. Given the time-sensitive nature of clinical practice, presenting multimodal data through data collection dashboards may impede their use in clinicians’ daily workflow, undermining their potential to empower mental health clinicians to better track patients’ progress and optimize treatment outcomes with insights from multimodal data.

To address these gaps and empower mental health clinicians to obtain patient insights from multimodal mental health data, we argue for an alternative approach—designing a **narrative dashboard** [8], a type of dashboard that presents multimodal data insights in a coherent “story”. As a manifestation, we propose a proof-of-concept system **MIND** (Fig. 1.A)—short for **M**ultimodal data **I**ntegrated **N**arrative **D**ashboard, a computationally curated dashboard that displays multimodal data insights to clinicians through data-driven narratives [76].

To ensure the curation aligned with clinicians’ needs, we motivated MIND by co-designing with five mental health clinicians from varying backgrounds. Through an iterative process, we refined both the design and computation aspects of MIND and summarized four design goals: intuitive, text-first, adaptive, and insightful (§3). Then, we implemented MIND (§4), an interactive interface that displays multimodal insights through data-driven narratives to support various mental healthcare workflows. Organized with a combination of linear (chronological insight organization) and parallel (by Biopsy-chosocial [27, 72] model) narrative structures, we used cards as the key means to accommodate multimodal data insights, allowing users to switch between two information levels (Fig. 6). To assemble the information used in MIND, we designed a hybrid computation pipeline (§5) that pairs the narrative capabilities of large language models (LLMs) with the rigor of rule-based data exploration to curate clinically relevant insights to assist mental health clinicians in better utilizing insights from multimodal patient data. Moreover, we designed the pipeline in a modular way to support future expansion or personalization of the system to more specific workflows.

We evaluated MIND via a user study (§6) with 16 licensed mental health clinicians comparing MIND with a non-narrative data collection dashboard baseline. Our results reveal that clinicians perceive MIND can reveal hidden and clinically relevant data insights ($p < .001$) and support their decision-making ($p = .004$). Additionally, clinicians perceive MIND as superior to the data collection dashboard baseline in presentation efficiency ($p = .006$) and cohesiveness ($p = .004$), as well as multimodal data integration ($p < .001$).

In summary, our main contributions are as follows:

- We describe clinicians’ information needs through a co-design study with 5 mental health experts, identifying key concerns when they engage with mental health dashboards.
- We propose MIND, a narrative dashboard curated by a computation pipeline that assists mental health clinicians in obtaining multimodal patient data insights. We open-source

our insights generation pipeline and interface design at: <https://github.com/sea-lab-space/MIND>

- We evaluate MIND via a user study with 16 licensed mental health clinicians using simulated patient encounter scenarios. We demonstrate how careful computation curation of clinical narratives can empower clinicians in understanding multimodal patient data.

Our research demonstrates how multimodal data narratives, paired with carefully designed computation and dashboard design, can empower clinicians to obtain richer information for clinical reasoning. We envision that such narrative approaches can extend beyond mental health and engage other stakeholders (e.g., patients, caregivers) to reshape data-driven decision support across healthcare domains (§7).

2 Related Work

We review three lines of work at the intersection of ubiquitous computing, health informatics, and visualization to situate our research. Specifically, we review 1) what data modalities are used in mental healthcare (§2.1), 2) how multimodal data are currently applied in mental healthcare (§2.2), and 3) design patterns and practical challenges of clinical dashboards presenting multimodal data (§2.3).

2.1 Data Modalities in Mental Healthcare

The digitization of healthcare over the past decades has gradually transformed healthcare towards more data-driven [34]. Advancements at the intersection of ubiquitous computing and health informatics now enable the collection of multimodal data for mental healthcare. Those multimodal data types could be categorized into clinical data, active sensing data, and passive sensing data [4]. In this work, we focus on four specific **data source types** that are either widely accessible in EHRs or increasingly prevalent in mental healthcare contexts: 1) *clinical notes*, 2) *clinical transcripts*, 3) *survey scores*, and 4) *passive sensing data*.

Clinical notes and transcripts are the two text-based clinical data source types we consider in this study. *Clinical notes* can be retrieved from EHRs [11] and are critical to the workflow of mental healthcare practice. These notes include clinician-driven, unstructured source texts that reflect clinicians’ assessments and decision-making information about patients. *Clinical transcripts* that contain all conversation details between patients and providers are gradually becoming accessible to more clinicians [12]. Typically collected through automatic speech recognition, clinical transcripts provide a complementary source to assist clinical documentation.

The other prevailing patient-generated data source types are survey scores and passive sensing data, two time-series, numerical data sources we consider in this study. *Survey scores* are patient self-report mental health measurements collected through active engagement by patients. Those measurements are typically validated methods for evaluating patients’ symptoms or functioning, such as PHQ for depression [52, 59] and GAD for anxiety [51, 86]. Those scores are collected using methods such as cross-sectional surveys or ecological momentary assessments. In addition to actively collected survey scores, sensors embedded in mobile devices (e.g., smartphones, wearables) collect *passive sensing data*, a type of data collected with little human effort (e.g., sleep, activity). While

those modalities are not direct indicators of mental health conditions, recent research has been exploring data-driven methods to infer mental health symptoms [21].

Given the complementary nature of the four data types, we investigate how their integration can empower mental health clinicians. In the following sections, we first review the current practices and challenges in leveraging multimodal data, then motivate our approach in integrating multimodal data for mental health clinicians.

2.2 Leveraging Multimodal Data for Mental Healthcare

Research has explored multiple ways to leverage multimodal mental healthcare data through both automated and human-centric approaches. For instance, prior work has used transcripts to generate structured clinical notes [36, 58, 74]. To reduce the complexity of these notes, studies have also proposed automated summarization methods [6, 64] and visual analytics systems [88–90], enabling clinicians to review and interpret information at scale. One prevailing paradigm for utilizing patient-generated data is to develop models that “detect/predict” [5] mental health measures such as mood [56, 67], stress [25, 40, 109], and depression [93, 102, 103]. However, recent evidence suggests that such models often fail to generalize across datasets [2, 62, 104] and are typically centered on enabling low-burden methods for continuous symptom monitoring outside the clinic [3, 43]. Few have discussed how patient-generated multimodal data can be used inside of clinics and linked to clinical notes and transcripts, to provide clinicians with insights that support in-session treatment decision-making beyond prediction [5].

In this work, we make an initial attempt to thread together all four data source types as complementary information to support mental health clinicians. Inspired by recent research [28, 39, 47], we propose a LLM-powered pipeline that translates multimodal data into data insights. We further embed the data insights in a narrative dashboard to facilitate clinicians’ consumption and reasoning with multimodal text insights.

2.3 Designing Dashboards for Clinical Practice

Effectively communicating the rich information embedded in health data is an active research topic [84]. Dashboards, “*a visual display of data used to monitor conditions and/or facilitate understanding*” [73, 99], emerged as a promising solution because of its potential to allow clinicians to grasp complex information “*at a glance*” [8, 31]. Specifically, we focus on clinical dashboards [24], a subset of healthcare dashboards focusing on informing clinicians with relevant data with the aim to improve patient care quality (*e.g.*, [70]).

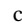
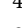
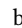
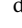
Researchers have designed clinical dashboards for multiple data types (see reviews [24, 38, 81, 114]), including clinical notes [88–90, 92], time-series/event sequence data [26, 45, 70, 71, 77, 78, 80, 94], domain-specific collection of multiple data types [30, 44, 48, 100, 107]. While these dashboards effectively support targeted workflows, many remain single-modal (*e.g.*, focusing only on notes) or adopt a data collection dashboard design pattern—dashboards that present high volumes of information with minimal curation of what and how the data is presented [8]. For example, CardioAI [100] integrates passive sensing with voice assistant data to summarize cardiotoxicity risk, yet presents the two modalities in separate modules.

Such layouts expose all relevant information but risk increasing clinicians’ cognitive burden, particularly in time-sensitive clinical environments where mental capacity is already stretched [108, 110].

To address this limitation, we explore an alternative paradigm—designing narrative dashboards [8]. Drawing from narrative visualization research [76], we integrate a “storyline” that threads together multimodal data. Perhaps the most relevant research to our work is from Zhang et al. [113], which proposed using the Five Ws (who, when, what, where, and why) from journalistic reporting to visualize patient information. We extend this perspective by applying linear and parallel narrative techniques to communicate multimodal mental health insights for clinicians. Further, building on calls to center text in clinical visualization [88], we leverage the expressive potential of LLM-generated insights (as evidenced in §2.2) to communicate multimodal data. We expect the combination of multimodal data, algorithmic insight curation, and narrative presentation to provide mental health clinicians with a complementary view they can leverage in their daily workflow [5].

3 MIND Design

While prior works discussed clinical dashboard designs in multiple clinical settings (§2.3), less is known about mental health clinicians’ needs in making sense of multimodal data across clinical and patient-generated data [57] and how to design systems to support this task. To execute our research in a principled way, we adopted a process following the design study methodology [75]. We chose this methodology because it’s proven for developing visualization systems that address real-world tasks, and has been widely applied in HCI and health-related systems design (*e.g.*, [45, 53]).

To characterize the needs of mental health clinicians, we conducted a co-design study with five mental health experts. We focus on the four clinically meaningful data types identified in §2.1: 1) clinical notes , 2) clinical transcripts , 3) survey scores , and 4) passive sensing data . In the following, we first describe our co-design process (§3.1). Then, we report insights from expert feedback (§3.2) and outline four design goals (§3.3) that guided the development of MIND.

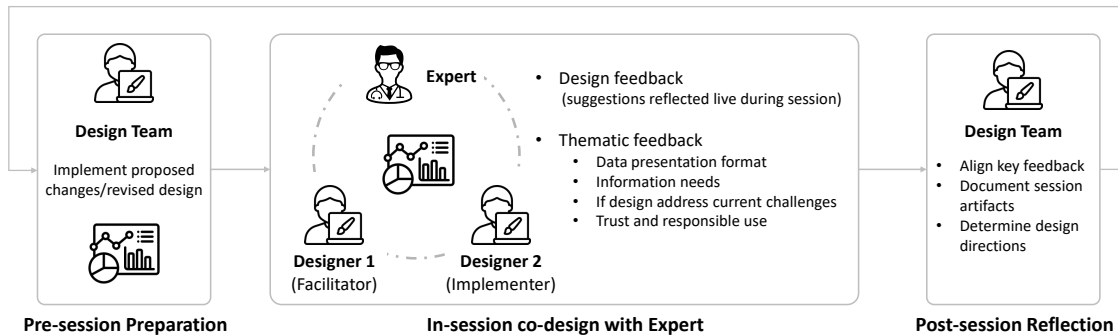
3.1 Co-designing with Experts

We engaged five domain experts (E1-E5) in psychiatry and psychology (Tab. 1) to bridge the mental health domain with the technical areas and collaboratively design MIND [75]. We recruited experts through a combination of convenience and purposive sampling [29] to ensure they represented perspectives from diverse mental health clinical settings (*e.g.*, academic medicine, digital health innovation, private practice, and global mental health). We conducted multiple co-design sessions with experts spanning over two months to ensure MIND targets real-world clinical needs. We envisioned MIND to be mainly used for patient case review at multiple stages in the clinical workflow, often in situations where decision time is limited, *e.g.*, preparation (patient recap prior to the clinical encounter) and consultation (decision-making during clinical sessions).

Before the co-design process, we reviewed prior literature and created an initial dashboard. Then, in co-design sessions (approximately 1 hour each), we followed an unstructured, exploratory

Table 1: Experts participated in the co-design process.

E#	Gender	Profession	Area of Expertise	Years of Experience
E1	Female	Clinical Psychologist	Technology for mental health, Researcher	16-20 years
E2	Male	Psychiatrist	Digital mental health	7-10 years
E3	Male	Psychiatrist	Mental health, Substance abuse	7-10 years
E4	Male	Research Psychologist	Digital mental health, Depression, Researcher	20+ years
E5	Male	Psychiatrist	ADHD, Cannabis use	11-15 years

**Figure 2: Illustration of the procedure of the co-design process.**

interview process (Fig. 2) and invited experts to propose and refine prototypes together alongside designers on a version of the prototype (updated based on prior co-design sessions). We scheduled each expert individually for focused design sessions, and two authors (the designers) of the paper participated in each session. During each session, one author (*i.e.*, Facilitator) facilitates the discussion, and the other takes notes and implements real-time updates to prototypes (*i.e.*, Implementer). Besides detailed suggestions, we elicit feedback on broader topics such as their preferred data presentation format (*i.e.*, text, charts) for different tasks (*e.g.*, overview, fact-checking), information needs (*i.e.*, what needs to be shown on the dashboard), constraints in current dashboards/information systems, and how to promote trust for clinicians using an AI-powered dashboard. We instantiate larger revisions between sessions to reflect the broader thematic suggestions.

Throughout the co-design, we created seven early iterations of MIND by rapid prototyping in Figma¹ and developing interactive interfaces through programming. Fig. 3 illustrates four design checkpoints that marked the progression of design (see Appendix A for prototype screenshots). MIND evolved from a chart-heavy dashboard design towards a more text-centered, narrative dashboard design. The earliest prototypes (Fig. 3.C1, C2) relied on stacked or tiled text-chart pairings that clinicians found visually overwhelming and lacked sufficient logic. Specifically, C1’s heavy emphasis on charts remains hard for mental health clinicians to interpret. While C2 improved upon C1 by integrating more text, it remains scattered and without a clear focal point. Building on C1 and C2, C3 (Fig. 3.C3) introduced a narrative structure that grouped text descriptions of data insights. Welcoming this change, clinicians continued to request clearer navigation and categorization given

the large number of insights. The final checkpoint, C4 (Fig. 3.C4), improved upon C3 by providing components that enable easier user control in navigating text descriptions of data insights. Moreover, after we finalized the implementation, we additionally ran two pilot sessions with E1 and E3 with the full-functioning MIND dashboard and finalized the design.

3.2 Design Process Findings

Synthesizing feedback across the co-design process, we distill takeaways regarding mental health experts’ needs in viewing multimodal data. We refer to design checkpoints (C#) to provide context to these findings.

3.2.1 Clinicians’ workflows differ across specialties, but common ground exists at the conceptual level. Clinicians proposed multiple workflows working with multimodal data. We found common approaches that clinicians use to make sense of multimodal data. In earlier iterations (C1, C2, C3), experts often proposed adding time-related information about the timeframe used to derive insights, indicating that they make sense of a patient’s progress by looking at a sequence of events. Another common approach is through categorization, first proposed as a feature by E2. Early prototypes (C2, C3) took E2’s suggestion by organizing information by thematic topics (*e.g.*, sleep, activities, emotions). While experts appreciate the feature, they did not reach consensus on which topics should be included because of differences in clinical practice. To provide a common ground for all clinicians, E5 suggested adopting recognized conceptual frameworks such as the Biopsychosocial model [27, 72], an established framework in the mental health domain, as a familiar baseline for categorization. Although not all experts used the model in practice, they found it easy to understand and appropriate for navigating insights(C4). Such feedback suggested that we

¹<https://www.figma.com/>

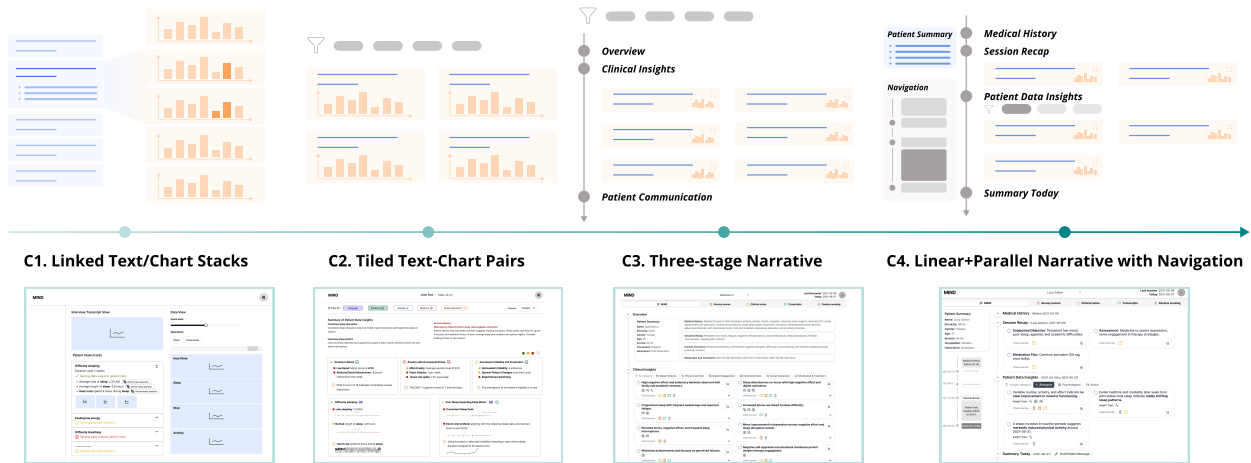


Figure 3: Four checkpoint designs (C1-C4) showing the co-design progression incorporating experts’ ideas and feedback. The top row illustrates the design pattern, and the bottom row presents the prototype artifacts that instantiate each pattern. Information layout becomes more structured, while charts become more closely integrated with text content with each iteration. We present expanded views of the prototype artifacts for C1 (Fig. 13), C2 (Fig. 14), and C3 (Fig. 15) in the Appendix.

should organize information through workflow-agnostic structures or established mental health frameworks.

3.2.2 Clinicians prefer text in most cases but value visual aids for navigation and information seeking. Early prototypes were primarily designed around charts (C1), but several experts (E1, E3, E5) emphasized that text is the preferred format to convey insights. While charts were considered a powerful way of conveying information, they were less familiar to mental health clinicians and were often treated as secondary information aids. E4 suggested presenting textual insight descriptions as structured summaries (e.g., bullet points), and using charts only as a form of providing supplementary information on demand to avoid cluttering the interface. Regarding the textual insight description, E1 stressed the importance of using clear, simple language rather than confusing technical terminology. E2 and E5 also recommended using icons and color-coding as information cues. Such feedback suggested that we should make text a central, navigable part of the dashboard, while allowing for detail-on-demand via charts.

3.2.3 Clinicians value simple interfaces given time constraints in clinical practice. Clinical workflows are highly time-sensitive, particularly during pre-session preparation. E5 noted that clinicians often only have about five minutes to review and recap a patient’s information, and E2 emphasized that back-to-back sessions leave little time for a deep review. Such reports echoed prior findings (e.g., [88]) that demonstrate a typically 2-10 minutes preparation stage note reviewing time and extended it into the mental health-care settings. In light of the time constraints, prototypes (C1, C2) were frequently critiqued as overwhelming, with experts describing the excessive amount of text and charts as contributing to cognitive overload. Particularly, interfaces that displayed too much raw data made it difficult for clinicians to retrieve the most relevant insight, especially under a short timeframe. Such feedback suggested that

we should display a selected subset of information at first glance to avoid cognitive overload.

3.2.4 Clinicians acknowledge AI could be a useful tool, with caveats in their role and reliability. During the co-design process, clinicians expressed nuanced expectations for how AI should be integrated into the dashboard. While they agreed that AI could be useful for summarizing subtle data patterns that might otherwise be overlooked (E2), they cautioned against diagnostic or prescriptive uses. For example, E3 emphasized that AI-generated suggestions should be phrased in a non-diagnostic manner, while E4 supported AI recommending potential activities based on patient symptoms, but only with clinicians’ approval. Experts also raised concerns about AI-generated content impacting or even replacing their professional judgment, noting that auto-generated insights risked preventing clinicians’ cognitive engagement. To leverage the benefit and maintain clinicians’ agency, they proposed several design ideas to keep the transparency and traceability of the design-generated insights should link back to raw data sources so clinicians could verify and challenge the system when necessary. Overall, clinicians envisioned AI as an assistive tool that surfaces insights and suggestions while they preserve the agency in diagnosis and shared decision-making with patients. Such feedback suggested that we should generate clinically relevant insights in a reliable, traceable way to foster trust in the dashboard.

3.3 Design Implications

By summarizing takeaways through co-designing with clinicians, we derived four design guidelines:

DG1 Intuitive: Design intuitive and at-a-glance clear information presentation. MIND should allow clinicians to easily locate and retrieve relevant clinical information to accommodate the time-sensitive clinical workflow.

DG2 Text-first: Apply text-first design to enable efficient information uptake. MIND should prioritize concise and simple textual summaries over jargonistic content or complex visualizations, and apply visual cues to improve the navigation of the text.

DG3 Adaptive: Structure information to adapt to different information needs. MIND should provide detailed information on demand to supplement the text-based data insights for fact-checking and trust-building.

DG4 Insightful: Generate reliable, clinically relevant data insights. The pipeline that generates clinician-facing content should ensure the generation of data insights relevant to clinicians' workflow, while keeping the process reliable and traceable. Besides, constraints should be applied to the pipeline to ensure insights are concise, clinically appropriate, and free of ambiguous or overly technical phrasing.

4 MIND Interface

To validate the design goals summarized through the co-design process, we created MIND, a multimodal data-integrated narrative dashboard. Following the visual information-seeking mantra [83], “*overview first, zoom and filter, details on demand*”, we designed MIND with two information levels. Accordingly, L1 (§4.1, Fig. 4) narrates multimodal data insights (DG4)—data patterns interpreted to describe a patient's condition—with established narrative structures (DG1) and text descriptions (DG2). Meanwhile, L2 (§4.2, Fig. 6) displays data facts, or “*patterns, relationships, or anomalies extracted from data under analysis*” [17] that support the data insights for trust building and data exploration (DG3).

4.1 L1: Narrating Multimodal Data Insights

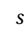
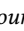
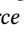
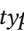


MIND's L1 interface consists of two components: Patient Background for basic information and navigation, and Narrative Overview to show the data insights. Specifically, we employed two narrative structures (*i.e.*, linear and parallel storytelling) and sought to strike a balance between information communication and discovery (*i.e.*, balancing user- and author-driven data stories) [8, 54, 55, 76, 106].

The first information point clinicians see is the **Patient Background** situated at the left of the interface, which presents 1) a summary of the patient's demographic data and 2) a clinical timeline that illustrates the series of mental health treatment sessions conducted with the clinician. From the clinical timeline, clinicians can navigate across the dashboard—each button on the timeline corresponds to a section in the Narrative Overview component. We use the components in this column as the primary hub for “*messaging*” [76], communicating both the summary clinical information and the overall organization of the dashboard.

The main content of MIND lies in **Narrative Overview**, situated at the right of the interface. To ensure the dashboard is intuitive and efficient (DG1, DG2), we took an “*author-driven*” approach [76] by choosing chronology to bind the story that enforces a specific order in presenting the data insights. Specifically, laid out following the *linear narrative* (*i.e.*, in the exact chronological order of how events unfold), clinicians can make sense of the patient's data by

reading sequentially through the four sections (S1-S4): Medical History (Fig. 4.1, the narrative's “*setting*”), Session Recap (Fig. 4.2, the narrative's “*setting*”), Patient Data Insights (Fig. 4.3, the narrative's “*rising-climax*”), and Summary Today (Fig. 4.4, the narrative's “*resolution*”). Such a sequence also closely aligns the key narrative stages identified by Yang et al. [106], who applied Freytag's pyramid narrative structure [32] to data stories. While research in narrative visualization has shown the potential of using anachronies (*i.e.*, present events in an order that deviates from the exact chronology) to improve user engagement in reading data stories [54], experts in our co-design session championed the timeline design, which we argue is more common [42] and anticipated in clinical practice.

Regarding interaction, we expand the Session Recap section by default and collapse other sections to focus on clinical notes insights—a common starting point for patient review for mental health clinicians. Clinicians can click the button next to the section titles or the button on the timeline to expand \vee or collapse \wedge each section. Following the typical design of a narrative dashboard, we keep every insight on the same page and display overflow content in the Narrative Overview component with scrolling [8]. In the following, we first introduce the visual primitive that contains the data insights. Then, we break down the four sections in the Narrative Overview and explain the rationale behind the design.

4.1.1 Visual Primitive: Insight Card. We present multimodal data insights within **Insight Cards** (Fig. 5). Specifically, we lay out key insight information: *textual insight description*, the *insight types* (biological \mathcal{A} , psychological \mathcal{P} , and social \mathcal{S}), and the *four data source types* (*i.e.*,    ) used to generate the insight in the center. Both insight types and data source types are coded in icons, and the data source types are further color-coded to highlight the multimodality of the data insight. Two clickable buttons, the *selector*  and the *drill-down* , link the data insight with the Summary Today section in L1 and the drill-down panel in L2, respectively.

4.1.2 S1: Medical History. The Medical History section contains summative descriptions of a patient's medical history beyond the ongoing mental health treatment sessions. We included the patient's past physical and psychological history with time markers indicating the onset of those symptoms (Fig. 4.1).

4.1.3 S2: Session Recap. Session Recap contains a three-piece summary from a past session note (Fig. 4.2). Drawing on the structure of a widely applied note format, SOAP (Subjective, Objective, Assessment, Plan) [72], we present summaries in three insight cards with prefixes: Subjective/Objective, Assessment, and (medication) Plan. We merge the S and O components because both capture observations of the patient's state. We keep the Plan section strictly limited to medication-related matters. Clinicians can add the insight to the S4: Patient Communication Section or also drill down to see the raw notes through the drill-down button (see §4.2).

4.1.4 S3: Patient Data Insights. Building on the session recap, Patient Data Insights (Fig. 4.3) shows “*what happened since the last visit*” [88]. Within the section, we employ the Biopsychosocial model [27, 72]—a framework that breaks down mental health causes into biological, psychological, and sociological factors—to structure data insights through *parallel narrative* (*i.e.*, multiple separated stories linked to a common theme). Specifically, we selected the

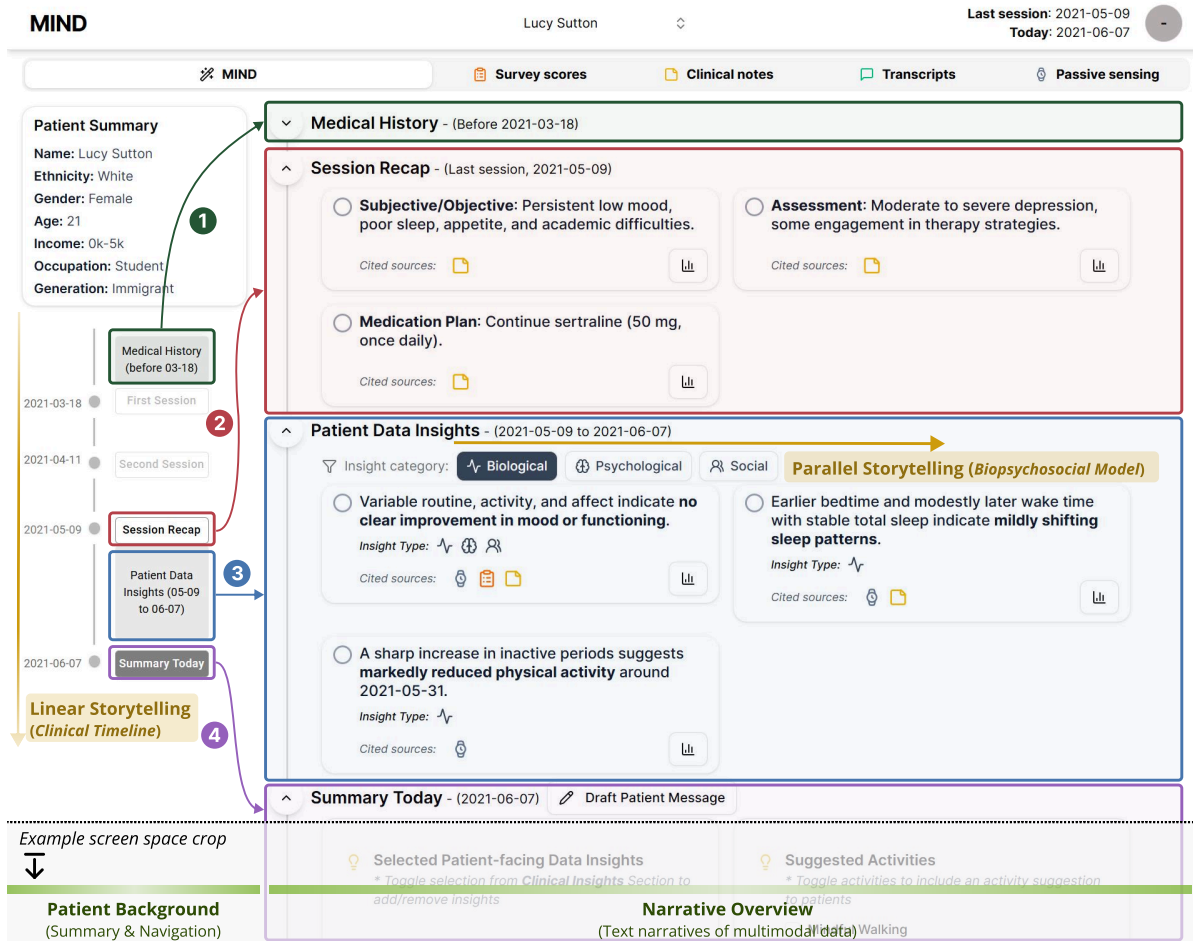


Figure 4: The L1 overview interface of MIND. MIND communicates multimodal data insights through a combination of linear and parallel storytelling techniques. On the outer level, MIND narrates insights using a linear four-part structure: Medical History (collapsed) ①, Session Recap ②, Patient Data Insights ③, and Summary Today ④, with overflowing content accessible through scrolling ↓. Within the Patient Data Insights section, MIND further applies three parallel narrative threads based on the Biopsychosocial model for clinician-guided information faceting.




Figure 5: Insight card as visual primitive. Three sections: insight description, insight types, and data source types, present the multimodal insights. The selector and drill-down buttons links the card with other parts of the MIND.

Biopsychosocial model because it is 1) suggested by experts in our co-design study, 2) widely known by mental health clinicians [60], and 3) carves an appropriate amount of “sub-stories” (three) for

clinicians to explore. Each “story” is about one of the three facets in the Biopsychosocial model [72]. Clinicians are given the affordance to select which “storyline” they want to see with a filter function with three options: Biological \mathcal{A} , Psychological \mathcal{B} , and Social \mathcal{R} .

Within each parallel story, we laid out insight cards in a grid-like, two-column setting. We argue that current capabilities of computation can not decide “the most important” insight for clinicians. As a heuristic, we ordered data insights that used more data sources before single-modal data insights, and left the final prioritization and interpretation to clinicians (*i.e.*, designing “user-driven” data stories [76]). Regarding data insight description, we employed a two-part template: 1) factual descriptions of the patient’s condition (*e.g.*, “earlier bedtime”) plus 2) how it implies the patient’s condition (*e.g.*, “shifting sleeping patterns”, bolded to foster efficient reading). Similar to the Session Recap section, clinicians can add the insight to the S4: Patient Communication Section or drill down to see the raw data facts that support the insight.

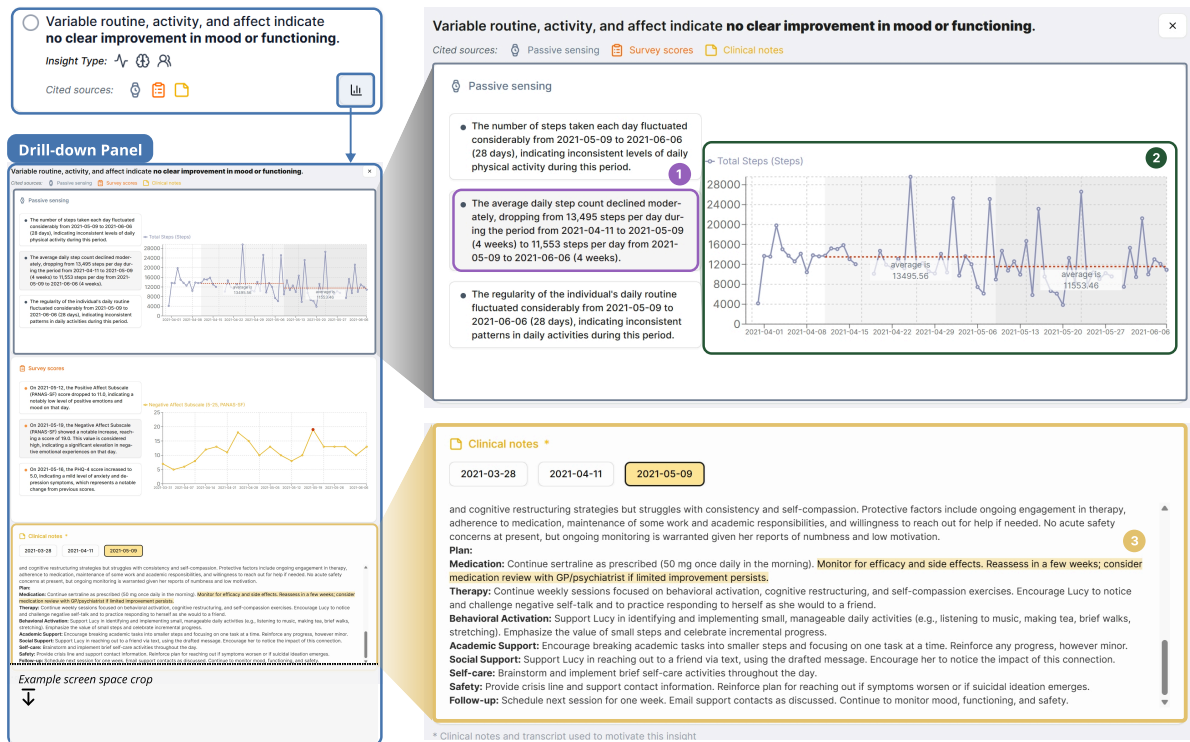


Figure 6: Drill-down panel for a patient data insight. Users can open the panel by clicking the drill-down button in L1. Each section provides detailed information for clinicians to explore underlying data, and each ❶ summarized bullet point corresponds with a ❷ chart supporting the data fact. ❸ Text highlight shows the text piece most relevant to the data insight.

4.1.5 S4: Summary Today. The Summary Today section (Fig. 4.❶) offers an open-ended conclusion to the Session Recap (S2) and Patient Data Insight (S3) by showing all selected data insights to provoke post-session patient communication. Based on the context, clinicians can select several suggested activities and use the *Draft Patient Message* function to create a text message that integrates patient data insights and suggested activities for communication.

4.2 L2: Data Facts on Demand

The L2 drill-down panel included detailed displays of data facts that support the data insights in L1. Clinicians can click on the drill-down button from the insight cards (Fig. 6). Inside the drill-down panel, we display the data insight description on top, followed by the data source types. Subsequently, we categorize data facts by their data source type: passive sensing, active sensing (survey scores), clinical notes, and clinical transcripts. Additionally, we put numerical data first to give clinicians an immediate, quantitative view, then textual data source types as supporting evidence.

To illustrate the data facts from passive and active sensing data sources, we pair each data fact (Fig. 6.❶) with a chart (Fig. 6.❷). We created visual elements on the charts (e.g., average value, trend line) to clearly demonstrate the key data attribute described in the data fact. Moreover, we chose simple chart types to minimize the learning curve from clinicians while following best practices from the visualization community (e.g., using a linear y-axis scale; using

the most appropriate chart type: line charts to display data trend/bar charts to show differences and extremes) to accurately reflect the underlying data and mitigate potential data-driven biases.

To show insights inspired by clinical notes and transcripts, we highlight the text that contributes to each insight to enable efficient skimming (Fig. 6.❸). We also show the full history of clinical data to support further backtracking, while highlighting the current date for quick navigation to the most recent date. Moreover, the drill-down panel provides a critical guardrail, enabling clinicians to rigorously ascertain the correctness of the presented insights.

4.3 Implementation

We implement MIND using a typical web stack, using React² as the UI framework. To render the charts used in MIND, we used Recharts³ for its built-in interactive data visualizations. For the content, the system consumed a JSON data structure passed from the computation pipeline (§5).

5 MIND Computation Pipeline

To curate the content for MIND, we implemented a proof-of-concept pipeline to transform multimodal data into reliable, clinically relevant data insights (DG4). In this section, we first describe how

²<https://react.dev/>

³<https://recharts.org/>

we formulate the computation tasks and utilize the multimodal data (§5.1). Then, we provide descriptions of the design of each module in our pipeline (§5.2, §5.3). Last, we introduce the technical implementation details of the pipeline (§5.4).

5.1 Insight Discovery Tasks, Data Usage, and Method

First, we define the types of insights needed and the time points or intervals associated with their computation. For each upcoming session, we consider generating information for two sections in the L1 Narrative Overview component: 1) session recap (*i.e.*, what was discussed last session, Fig. 4.📄), and 2) patient data insights (*i.e.*, how the patient’s condition has evolved between the sessions, Fig. 4.📄). We define three time points: 1) *today* (*i.e.*, the upcoming session date), 2) *last session* (the previous session date), and *the retrospective sessions* (the session before the last session). Fig. 7 shows how each data modality contributes to session recap and patient data insights.

First, we discover session recap insights only using the last session’s clinical notes 📄 (Fig. 7.📄) because they are clinician-driven narratives that reflect what they deemed important about the patient. Meanwhile, we decided not to include sensing data 📄📄 for the session recap because they are too remote for the current date.

Second, we discover patient data insights using all four data source types 📄📄📄📄. For clinical notes and transcripts, we use them to inform data insight discovery (*e.g.*, analyze sleep data if mentioned in clinical notes). We refer to this process as *guided discovery* (Fig. 7.📄). In parallel, relying solely on sensing data 📄📄, we designed an *exploratory discovery* (Fig. 7.📄) process to capture “anomalies” [57] that may be neglected by structured discovery.

Within our pipeline, all modules described are driven by a large language model (LLM) agent unless explicitly stated. Specifically, we followed best practices in building LLM applications (*e.g.*, using zero-shot Chain-of-Thought prompting [49], prompt-chaining techniques [101]) to optimize the performance of our pipeline.

5.2 Insight Discovery for Session Recap

We generate the Session Recap section with **Summarizer**. Summarizer takes in raw clinical notes and transforms them into a modified SOAP structure specified in §4.1. For each insight item, we constrain the LLM to generate insights of no more than 12 words

to keep them succinct. Moreover, we instruct the LLM to provide evidence on which pieces of clinical notes are used to create the insight summary to support the text highlight feature (§4.1).

5.3 Insight Discovery for Patient Data Insight

We employed a hybrid approach that combines rule-based and LLM-based methods. Through three modules **Analyzer** (§5.3.2), **Synthesizer** (§5.3.3), and **Narrator** (§5.3.4), the pipeline takes the four data source types 📄📄📄📄 as input and outputs data insight-fact pairs. The system was architected using a pipes-and-filter software architecture (Fig. 8) to minimize dependencies across modules. We also indicate whether each component module is rule-based ⚙️ or LLM-based 🧠 with icons.

5.3.1 Insight Generation Framework. Taking inspiration from automatic visualization research, we employ the data fact taxonomy [17] for a principled approach in managing the data fact discovery process. Specifically, the data fact taxonomy provides formal definitions to categorize data insights (*e.g.*, value, distribution, difference) in visual analytics systems. Recent research has successfully applied this taxonomy in several computational approaches to extract insights from multimodal datasets (*e.g.*, [16, 82, 96, 115]).

While the original fact taxonomy included 12 data fact types, we consider a subset of five data fact types most relevant to our task: *comparison*, *trend*, *outlier*, *extreme*, and *difference*. Specifically, the first three data fact types measure data attributes across one or multiple time periods (Δt or $\Delta t_1, \Delta t_2$), whereas the latter two measure data attributes at single time points (*extreme*: t ; *difference*: t_1, t_2). Given the data fact types and their temporal attributes, we define a *fact* as an intermediate layer between raw data and data facts. We formulate a *fact* as a 4-tuple:

$$fact := \{type, Time, Value, attribute\}$$

where *type* is one of the five fact types we selected. Meanwhile, *Time* and *Value* are tuples that define the temporal feature of data facts. Combining *type* and *value*, *attribute* is a derived feature describing the characteristics of the data fact. Tab. 2 documents the Time, Value, and Attribute of the five fact types we consider in this study.

5.3.2 Analyzer. Analyzer (Fig. 8. M1), transforms raw multimodal data into *facts*. To derive data facts grounded with clinical data, we first process clinical notes and transcripts through two agents: Inquirer (Fig. 8. M1.1) and Planner (Fig. 8. M1.2). Analyzer reflects




Figure 7: Data modalities used for generating each insight thread. 📄 Session Recap uses clinical note 📄 from the last session. 📄 The Structured Patient Data Insights thread combines clinical data (last session) 📄 with new patient-generated sensing data 📄📄, using historical data for context. 📄 The Exploratory Patient Data Insights thread analyzes only new patient-generated sensing data 📄📄 since the last session.

Table 2: Template-based data fact generation rules used by the *Analyzer* module. We specify each fact type with its temporal feature (single time point/time duration, $t/\Delta t$), value ($v/\Delta v$), and a value-derived attribute (a). We include the entity related to the data fact ($\{e\}$, e.g., “total sleep”) in the description syntax to provide context for a data fact.

Fact Type	Time ($t/\Delta t$)	Value ($v/\Delta v$)	Attribute (a)	Description Syntax
Outlier	Δt	v	spike/dip	An anomaly was detected for $\{e\}$ on $\{t\}$, which $\{a\}$ to $\{v\}$.
Trend	Δt	-	rise/fall/stable/ cyclic/variable/ no (trend)	The $\{e\}$ showed $\{a\}$ trend from $\{\Delta t_1\}$ to $\{\Delta t_2\}$.
Comparison	$\Delta t_1, \Delta t_2$	$\bar{v}_{\Delta t_1}, \bar{v}_{\Delta t_2}$	increase/decrease/ remained stable	The average $\{e\}\{a\}$ from $\{\bar{v}_{\Delta t_1}\}$ (Δt_1) to $\{\bar{v}_{\Delta t_2}\}$ (Δt_2).
Difference	t_1, t_2	v_{t_1}, v_{t_2}	more/less	The $\{e\}$ was $\{v_{t_1}\}$ on $\{t_1\}$ and became $\{a\}$ at $\{v_{t_2}\}$ on $\{t_2\}$.
Extreme	t	v	max/min	The $\{e\}$ reached its $\{a\}$ value of $\{v\}$ on $\{t\}$.

on clinical data, especially the forward-facing plans, to generate a set of questions aligned with mental health clinicians’ interests. Then, we pass all passive and active sensing data features to the Planner to determine 1) if the proposed question is answerable, and 2) if answerable, what features the Analyzer should investigate. Those filtered questions are then sent to the Discoverer (Fig. 8. M1.3). To ensure that the extracted data facts are grounded in raw data, we designed Discoverer with two complementary mechanisms: 1) rule-based agents for extracting comparison, trend, and outlier data facts, and 2) LLM-based exploratory agents with tool-assisted verification for extreme and difference data facts.

First, taking as input the *guided* plan from the Analyzer and Planner, we designed a series of rule-based agents to discover outlier, trend, and comparison data facts. Because the discovery goal is to reveal any significant changes between sessions, the time intervals are defined as the period from the last session to today.

We applied statistical and heuristic methods to discover data facts from raw data. For comparison, we apply the Mann-Whitney U test to uncover differences in medians between two time periods split by the last session (see Fig. 5.1). For trend detection, we use the Mann-Kendall test to identify monotonic increases or decreases, complemented by autocorrelation analysis for cyclic detection and the coefficient of variation method for stability assessment. For outlier detection, we apply seasonal-trend decomposition via LOESS (locally estimated scatterplot smoothing) and identify anomalies in the residual component using the Median Absolute Deviation rule, classifying them as spikes or dips.

Second, complementing the guided discovery, we used LLMs for *exploratory* discovery on extreme and difference data facts. Grounded by LLMs’ potential in understanding sensing data [28, 39, 47], we provided raw csv to the LLM-based Discoverer (Fig. 8 M1.3) agents to capture data facts. Additionally, to mitigate potential

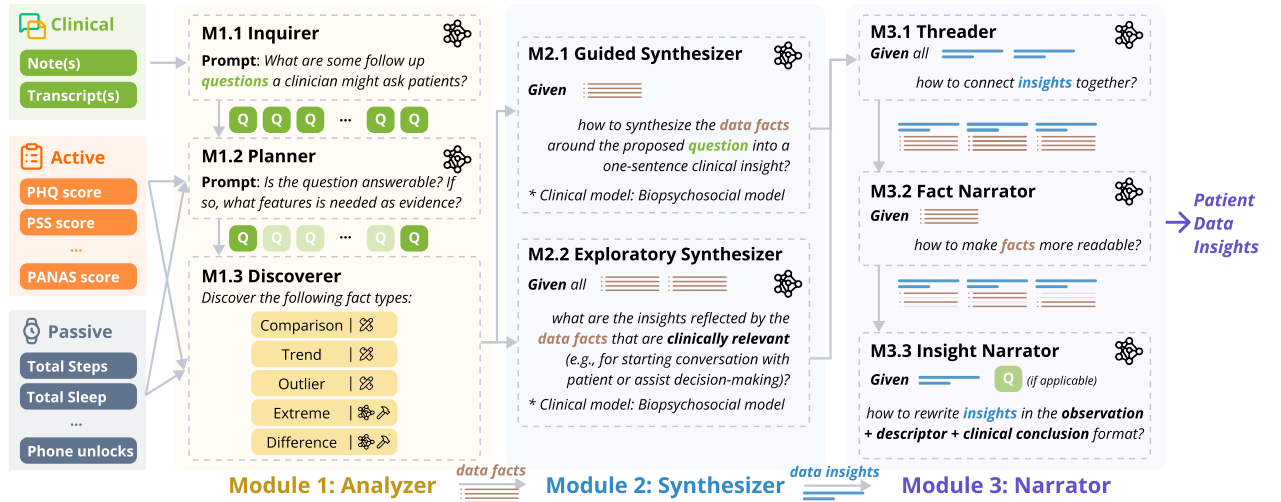


Figure 8: The patient data insight discovery pipeline of MIND contains three modules: *Analyzer* (M1), *Synthesizer* (M2), and *Narrator* (M3). We display the computation method used in each module/submodule using icons, with \otimes indicating rule-based approaches, \textcircled{L} indicating LLM-based approaches, while $\textcircled{L}\textcircled{A}$ indicating tool-using-LLM-based approaches.

hallucination [41], we guardrail the agent with a data verification tool: whenever a data retrieval task is required, the LLM invokes a predefined script to access raw values directly from the dataset, minimizing the instances where the LLMs might “invent” raw data.

Crucially, both the *guided* and *exploratory* discovery processes are agnostic to the underlying active and passive sensing measurements. Designers can tailor the discovery process with a specific set of measurements best suited to specific mental health cases or clinicians. Moreover, both discovery processes are also robust to missing data. Missing data are either handled directly by methods that tolerate incomplete observations (e.g., Mann-Whitney U and Mann-Kendall tests) or linearly interpolated over time when required (e.g., autocorrelation analysis and anomaly detection). We preserve the raw data (i.e., include data missing) in the output to the MIND interface to indicate missing data to users.

After the hybrid LLM/rule-based approach transformed numerical data into fact specifications, we applied a rule-based natural language generation method to convert these data facts into data fact descriptions. Extracted fact specifications fill in the blanks we designed for each data fact type (Tab. 2, Description Syntax column). Rather than relying on code representations (i.e., JSON specification), we prepared data facts as natural language descriptions to fully leverage the LLM’s language understanding capabilities for subsequent insight synthesis. We also specifically opted for this controlled process to ensure that the description of data facts is accurate and not altered during LLM’s natural language generation.

5.3.3 Synthesizer. Taking the guided and exploratory data facts, **Synthesizer** keeps the two types of generation and combines data facts. Guided Synthesizer (Fig. 8. M2) surface data insights from data facts that potentially answer each computable question proposed in **Analyzer**. If significant data facts exist, guided synthesizers only take significant results as input for data insight synthesis. If no significant data facts exist, we instruct the Guided Synthesizer agent to consume all data facts and generate an insight using up to six data facts. Meanwhile, the Exploratory Synthesizer (Fig. 8. M2.2) surface anomalies observed through one large corpus of data facts. Specifically, we instruct the Exploratory Synthesizer to output at least 15 “clinically relevant” insights (i.e., assist clinical decision-making or good material to raise a patient discussion in a session).

To ensure that each data insight could be traced back to data facts, we instructed both synthesizer modules to output the data fact evidence alongside the synthesized data insight, creating data insight-facts pairs. Moreover, we steered both synthesizer agents to consider the Biopsychosocial model during synthesis and labeled each insight with one or multiple categories in the Biopsychosocial model. To keep final data insights succinct and readable, we constrain both synthesizer agents to generate insights of fewer than 15 words using plain language.

5.3.4 Narrator. The **Narrator** module (Fig. 8. M3) combines the data insight-facts pairs from the **Synthesizer** module into clinician-aligned content. First, the **Threader** (Fig. 8. M3.1) combines the insights from the two synthesizers into a coherent narrative. We keep all insights from Guided Synthesizer (M2.1) because they are clinically grounded and use a combination of sensing and clinical data. We add four to six insights from the Exploratory Synthesizer (M2.2) to balance between guided and exploratory insights.

Taking the full patient data insights as input, the **Fact Narrator** (Fig. 8. M3.2) and **Insight Narrator** (Fig. 8. M3.3) optimize content quality. Specifically, the fact narrator processes the template-generated data fact description from **Analyzer** and paraphrases it to make it more readable. While this approach did make the data fact description longer, experts considered the additional details appropriate for describing the data fact. Subsequently, the insight narrator further rewrites the data insights to align with the clinician’s information needs using the template introduced in §4.1.

It is worth noting that despite our LLM generation guardrails, errors might still occur from **Synthesizer** (e.g., synthesizing clinically unrelated data facts) and **Narrator** (e.g., semantic shifts when rewriting data facts/insights). To mitigate this risk, we designed the L2 drill-down panel to be immediately accessible (i.e., the drill-down button) to the data insights, prompting clinicians to verify insights that seemed nonsensical to them.

5.4 Implementation and Internal Validation

We implemented our computation pipeline in Python. The pipeline communicates the generated content with the frontend via a JSON file, analogous to a live backend server. For LLMs, we used OpenAI’s GPT-4.1 API⁴—a non-reasoning model optimized for building agentic applications. To enable tool integration, we used the OpenAI Agents SDK⁵, allowing the LLM to call functions and retrieve raw data. We used packages such as `statsmodels`, `scipy`, and `pymannkendall` for the statistical tests in the rule-based agents.

Following the design study methodology [75], we conducted multiple rounds of “inward-facing” validations of our computational pipeline. Specifically, two mental health experts (E1 and E2) independently reviewed multiple batches of LLM-generated summaries and data insights to ensure the appropriateness of wording and the factuality of the generated insights. Specifically, experts evaluated whether the wording avoided diagnostic claims, whether the language was easy to read, and whether each insight related to the raw data. For instance, reviewing the sentence “*Sleep disturbances co-occur with high negative affect and digital rumination*”, E2 requested the insights to be more descriptive. We thus modified our computation pipeline to favor wording such as “*Variable bedtime and awakening episodes with stable total sleep and wake time indicate fluctuating sleep quality*.” Experts also flagged cases where insights were not clearly grounded in raw data, and we subsequently applied more stringent feature selection requirements. Overall, their feedback helped refine our templates, constraints, and evidence-linking mechanisms. We terminated the validation process after both health experts agreed upon the generated content.




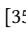
6 User Evaluation









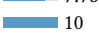




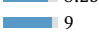






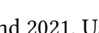
To evaluate if and how well MIND supports mental health clinicians in obtaining insights from multimodal data, we conducted a mixed-method within-subject study with 16 licensed mental health clinicians in a simulated patient review task⁶. We compared MIND

⁴<https://openai.com/api/>

⁵<https://openai.github.io/openai-agents-python/>

⁶Approved by the last author’s institute’s IRB.

Table 3: Demographics of user evaluation participants. We use SV, PS, CN, and CT to denote survey scores , passive sensing data , clinical notes , and clinical transcripts , respectively. We measure participants’ attitude towards AI on a scale of 1 to 10, with a higher score indicating a more positive attitude.

P#	Gender	Practice Focus	Years of Experience	Reported Data Experience				AI Attitude [35] (1-10, negative-positive)
				SV 	PS 	CN 	CT 	
P1	Female	Psychiatry	7-10 years	•	•	•		 6.5
P2	Female	Psychiatry	7-10 years	•	•	•		 7
P3	Female	Psychiatry	7-10 years	•		•	•	 8.75
P4	Male	Counseling, Social Work	10+ years	•	•	•	•	 7.5
P5	Female	Counseling	1-3 years	•	•	•	•	 7.75
P6	Female	Behavior Analyst	10+ years	•		•	•	 10
P7	Male	Clinical Psychology	10+ years	•		•		 9.5
P8	Male	Counseling	4-6 years			•		 5.75
P9	Female	Counseling, Social Work	1-3 years			•		 10
P10	Female	Psychiatry	1-3 years	•	•	•	•	 9
P11	Female	Counseling	4-6 years	•	•	•	•	 8.25
P12	Female	Counseling	1-3 years			•		 9
P13	Male	Psychiatry	10+ years	•	•	•		 9.25
P14	Female	Clinical Psychology	7-10 years	•	•	•	•	 8.5
P15	Female	Clinical Psychology	4-6 years	•	•	•	•	 6
P16	Female	Psychiatry, Social Work	10+ years	•	•	•	•	 6.25
Summary				13	10	16	9	 8.06

with a baseline system FACT that employs the data collection dashboard design that displays the same multimodal data as MIND. We want to investigate:

- **Quality:** How do clinicians evaluate the quality (*e.g.*, clinical relevance) of the generated data insights in MIND?
- **Perception:** How do users perceive the narrative dashboard design of MIND, and how does it impact system usability and perceived workload?

6.1 Participants

We recruited 16 participants (12 self-identified as female, 4 as male) who did not participate in our co-design process through a combination of convenience, purposive, and snowball sampling [29, 33] (Tab. 3). All participants are licensed mental health practitioners with at least 1 year of experience, among whom 5 have over 10 years of experience in mental healthcare. Participants cover various practices, among which 6 reported expertise in psychiatry, 6 in counseling, 3 in clinical psychology, and 3 in social work. Most participants (14 of 16) reported using an EHR system regularly as part of their clinical practice. All participants reported working with clinical notes while using different data modality combinations in their daily practice. Specifically, 7 have prior experience in all four data categories, while 3 have only worked with clinical notes. Using the AI Attitude Scale (AIAS-4) [35], participants reported overall positive attitude in AI ($M=8.06$, $SD=1.42$).

6.2 Materials

To simulate real-world scenarios, we prepared three simulated users: Lucy, Gabriella, and Alison, synthesizing data streams from multiple real-world datasets in collaboration with our experts (E1-E5). We utilized the GLOBEM dataset [105], an open-source longitudinal dataset that includes both passive and active sensing data from

real-world users between 2018 and 2021. Using Beck’s Depression Inventory (BDI-II) [9] in the dataset, we screened out users with moderate or high depression symptoms ($BDI-II \geq 19$) as archetype simulated patients. Then, by manually going through their other EMA survey scores (*e.g.*, PHQ-4 [59], PANAS [98]), we manually sampled three typical instances as simulated patient archetypes.

Meanwhile, since GLOBEM was designed primarily for ML model training and evaluation, we picked a subset of easy-to-understand features for materials preparation and the later evaluation studies. Specifically, we retained all six regularly collected active sensing data features (see Appendix §B.1) and 11 passive sensing data features spanning four categories (Location, Screen, Sleep, and Step, see full list in Appendix §B.2).

Subsequently, sourcing publicly available mental health session tutorial videos, we picked three instances to present to our experts as baselines that we pair with the three patient archetypes. Then, for each patient archetype, we worked with our experts in §3 and paired each archetype with a tutorial session. Subsequently, we created two subsequent mock interview sessions with three clinical notes (see pairing in Appendix §C). Overall, while our setup does not cover all ecological setups (*i.e.*, other note formats or interview strategies exist), we aim to provide a homogeneous baseline across all three simulated patients to minimize the noise caused by different study materials and focus on evaluating the efficacy of the narrative design of MIND.

6.3 Study Design

6.3.1 Conditions. To understand if the computational narrative design of MIND provided benefits for clinicians viewing multimodal patient data, we compare MIND with a baseline interface FACT. Specifically, generated through the same computation pipeline as MIND, FACT displays the information using the data collection

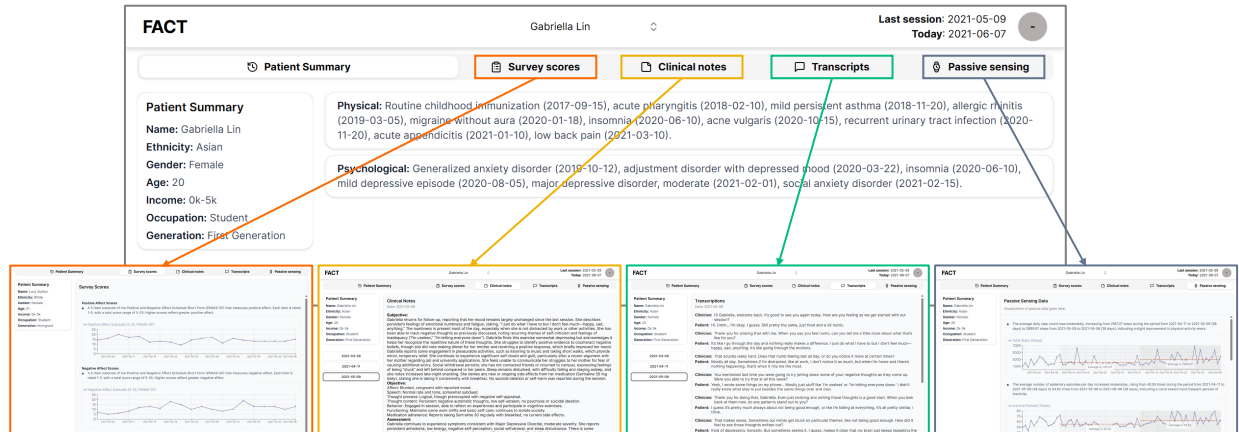


Figure 9: The baseline interface FACT. FACT displays the same multimodal data as MIND, but in four separate views (tabs): Survey scores, Clinical notes, Transcripts, and Passive sensing.

dashboard design pattern, *i.e.*, laying out the four data categories in separate views (Fig. 9). For instance, the view of passive sensing would only contain text descriptions and visualizations about the data facts extracted from passive sensing features. Moreover, we implemented FACT with a similar aesthetic to MIND to minimize the impact of visual appearance on the study results.

We selected FACT over a traditional EHR as the baseline for two primary reasons. Traditional EHR systems might not integrate the full spectrum of data modalities we considered, and such a comparison might conflate the effect of the added data modalities with our narrative design. Additionally, we argue that FACT represents a typical near-future design (see §2.3) integrating multimodal data into existing EHR systems. By controlling the confounding variables between MIND and FACT, we can focus on the effect of the narrative presentation compared to a typical dashboard baseline.

We paired the two conditions (*MIND*, *FACT*) with two simulated patient scenarios (*Lucy*, *Gabriella*) that were similar in characteristics (*e.g.*, complexity, treatment methods). We leave the data for *Alison* for demonstration and system introduction purposes.

6.3.2 Measures. We included the System Usability Scale (SUS) [15] and the NASA-TLX [37] to gauge and compare the usability and clinician’s workload using both systems. Meanwhile, we designed a 10-question, 7-point Likert survey (see Appendix §D) used to capture quantitative user feedback regarding the quality and perception of both systems. Specifically, four questions targeted **quality** of data insights, assessing the clinical relevance of the information, the system’s ability to reveal hidden or complex insights, its effectiveness in supporting clinical decision-making, and the trustworthiness of the presented data. The remaining six questions focused on user **perception**, evaluating aspects of system design such as the ease of interpreting the data, efficiency in reviewing patient information, integration of multimodal data, cohesiveness of the narrative presentation, time-saving potential, and support for effective communication during patient interactions.

6.3.3 Procedures and Tasks. Each study session comprised three stages: a) introduction to the study settings, b) two patient review

tasks using MIND and FACT with a counterbalanced order, and c) a post-study survey and a semi-structured interview to understand participants’ experiences. Upon completion of the study, participants were thanked with a \$50 gift card.

Introduction. We first provided an introduction to the study setting. We informed participants to imagine a hypothetical 4th encounter with a mental health patient whom they had last seen a month ago. We told them that their goal is to recap on the prior sessions using the data presented on the dashboards to recall the patient’s conditions.

Patient Review Tasks. We first gave a brief tutorial on the system (*MIND*, *FACT*) to familiarize them with the interface. Researchers then sent a link to direct participants to the target condition. We counterbalanced the interface condition to control order effect and rotated the patient scenarios (*Lucy*, *Gabriella*) to control the effect of the content on the conditions. For each patient review, we gave participants 6 minutes with the target of giving a verbal summary of the patient’s condition in the end. While the task is goal-oriented with a time limit, we informed participants that they would neither be graded on the number of insights found nor on the quality of the review to foster naturalistic review behavior. Besides, participants were encouraged to follow a think-aloud protocol to allow us to better understand their thinking process.

Post-study Survey and Interview. After participants finished both conditions, participants filled out a survey where they reflected on the two systems’ performance. At the end of each session, we conducted a semi-structured interview to collect qualitative feedback on MIND. Building on their think-aloud feedback, we further asked them questions about their strategies for conducting patient review using both conditions, preference between the two conditions, attitudes towards the AI integration, and limitations and feedback regarding the functionality of MIND. We also asked follow-up questions based on the participants’ responses.




Figure 10: A comparison of information quality between MIND and FACT using a 7-point Likert Scale. Each point in the chart represents a single record from a single participant. We map the values with a divergent color scheme from dark red (1-very low) to dark blue (7-very high). The black dot and line indicate the mean and 95% confidence interval of the metric.

6.4 Findings

Since the study employed a within-subject design, we conducted quantitative analysis between MIND and FACT using the Wilcoxon Signed-Rank test. Data processing and statistical analyses were conducted in R [91], using the `coin` and `rstatix` packages. We report the relevant statistical indicators, including the p -value, Z -value, and the rank-biserial effect size (r_b). Additionally, we used the AI attitude score (1-10) and data familiarity (1-4, *i.e.*, number of reported data experience) as covariates and conducted regression analyses (using the `lmerTest` package) to investigate if they impact the evaluation results.

Meanwhile, we analyzed the qualitative data through an iterative process of qualitative coding and affinity diagramming based on principles of thematic analysis [14, 19]. Two coders independently reviewed a subset of interview transcripts and notes to generate initial codes. Subsequently, in an iterative process, the two coders met for multiple rounds to discuss and refine the coding structure. Finally, the two coders collaboratively surfaced high-level themes from the codes to address our research questions.

6.4.1 Quality: Expert Evaluation. Participants overall evaluated the information presented on MIND as of good quality and significantly better than that presented on FACT (evaluation of DG4, Fig. 10). Specifically, MIND significantly outperforms FACT in providing quality information for clinicians to discover hidden data insights ($M_{MIND}=5.75$ vs. $M_{FACT}=3.93$, $Z=-3.39$, $p<.001$, $r_b=.85$), and in assisting clinical decision-making ($M_{MIND}=5.68$ vs. $M_{FACT}=4.62$, $Z=-2.94$, $p=.004$, $r_b=.74$). Meanwhile, the statistical test does not reveal differences regarding clinical relevance ($M_{MIND}=6.13$ vs. $M_{FACT}=6.13$, $p=.438$) and the trustworthiness of the information ($M_{MIND}=5.93$ vs. $M_{FACT}=5.87$, $p=.922$), with both high average scores suggesting that participants favors the information generated by our computation pipeline. An additional regression analysis using linear mixed effect models (R notation `<Quality_metrics> ~ Condition + AI_Attitude + Data_Familiarity + (1|PID)`, referent FACT condition) revealed no significance on participants' AI attitude

and data familiarity, indicating the results are consistent across populations.

6.4.2 Perception: System Design Evaluation. Participants prefer MIND as the interface for consuming multimodal data with significance in five out of the six measures tested (evaluation of DG1, 2; Fig. 11). MIND significantly outperforms FACT in efficiency in reviewing patient information ($M_{MIND}=5.94$ vs. $M_{FACT}=4.81$, $Z=-2.78$, $p=.006$, $r_b=.70$), integration of multimodal data ($M_{MIND}=5.94$ vs. $M_{FACT}=3.94$, $Z=-3.20$, $p<.001$, $r_b=.80$), cohesiveness of the narrative presentation ($M_{MIND}=5.50$ vs. $M_{FACT}=4.31$, $Z=-2.90$, $p=.004$, $r_b=.73$), time-saving potential ($M_{MIND}=6.25$ vs. $M_{FACT}=4.69$, $Z=-3.20$, $p<.001$, $r_b=.80$), and effective communication during patient interactions ($M_{MIND}=6.38$ vs. $M_{FACT}=5.13$, $Z=-2.94$, $p=.004$, $r_b=.73$). Meanwhile, both MIND and FACT showed positive ratings for ease of interpretation ($M_{MIND}=5.63$ vs. $M_{FACT}=5.13$, $p=.180$), though we reveal no significant difference comparing the two conditions. An additional regression analysis using linear mixed-effect models similar to the above test revealed no significance on participants' AI attitude and data familiarity on their perception of the systems' design.

6.4.3 Perception: System Usability Scale. Overall, MIND showed a comparable system usability to FACT. A breakdown of the system usability subscales (Fig. 12) revealed that participants perceived MIND as a system they would more likely to use frequently ($M_{MIND}=3.94$ vs. $M_{FACT}=2.94$, $Z=-2.68$, $p=.005$, $r_b=.67$) and is better integrated ($M_{MIND}=3.94$ vs. $M_{FACT}=3.19$, $Z=-2.07$, $p=.04$, $r_b=.52$) than FACT. On the other hand, visual analysis revealed that clinicians preferred FACT over MIND in a few aspects, especially its learnability (*i.e.*, no tech support needed $M_{MIND}=3.69$ vs. $M_{FACT}=4.31$, $p=.125$) and simplicity ($M_{MIND}=3.38$ vs. $M_{FACT}=3.81$, $p=.219$). However, the difference is not large enough for significance.

6.4.4 Perception: Workload. Overall, we observed no significant difference across all measures (mental, physical, temporal, performance, effort, and frustration) of workload using the NASA-TLX scale (Fig. 16). We argue that the scores from NASA-TLX reflected




Figure 11: A comparison of the perception of system design between MIND and FACT using a 7-point Likert Scale. Each point in the chart represents a single record from a single participant. We map the values with a divergent color scheme from dark red (1-very low) to dark blue (7-very high). The black dot and line indicate the mean and 95% confidence interval of the metric.




Figure 12: Normalized scores (higher is better) for the adapted SUS (MIND vs. FACT). Bolded items indicate $p < 0.05$.

more on the difficulty of the task itself, and merely integrating the information through narrative did not make the interpretation easier. Such scores suggested that clinicians maintained the agency in making clinical decisions using both MIND and FACT.

6.4.5 Perception: Clinicians’ Feedback. Qualitative data suggested that clinicians found MIND effective for viewing multimodal patient data, though they also suggested several areas for improvement.

Clinicians prefer MIND for its summarization capability, especially under time constraints. Across all 16 participants, clinicians preferred the MIND dashboard over the FACT dashboard. MIND was positively received for its summaries that distilled complex multimodal data into data insights. For example, P1 liked that “it (MIND) does summarize,” whereas FACT was critiqued by having “so much data to sort through” (P4). Under time constraints, clinicians commended MIND because it provides a “quick & dirty snapshot” (P5) for them before entering a patient session.

Clinicians navigated MIND using Session Recap and Timeline as key entry points. Clinicians (P3, P6, P9, P10, P13) identified the

session recap and the timeline as two key entry points for finding insights in the dashboard. P9 commented that “*timeline on the bottom left was really the main way I navigated MIND... and the recap is helpful because it gives you that SOAP format summary.*” P13 liked “*having the timeline visible*” to “orient” and guide “*what to do next.*” P6 remarked that the timeline “*makes it easier to understand how things are progressing, instead of just snapshots*” and that the timeline is “*closer to how I think about patient care.*” Such feedback indicates that the narrative structure functioned effectively to provide clinicians with an intuitive interface for navigating multimodal data.

Verifiability and transparency are keys in fostering trust in MIND. During the interview session, we informed clinicians with a general description of the hybrid insight generation workflow. Clinicians overall commended the quality of the insights, saying that “*(The data insights) seem to be good summaries [of] what I would, in a bird’s-eye view, want to know about how the month went [for the patient]*”. Additionally, when asked what designs helped them trust the system, they (P1, P2, P5, P14) emphasized that the affordance jumping to L2 (Drill-down panel) strengthens their confidence in the insights. P5 commented “*as long as I can check the raw data myself, I’ll generally trust the AI.*” Such a response further demonstrates how MIND successfully embodies DG3. However, we also observed scenarios (P9, P16) where clinicians “overtrust” AI. For example, P9 states that “*I don’t doubt AI ... I trust AI fully.*”

Challenges and concerns around using clinical transcripts in mental healthcare. Several clinicians raised privacy concerns and patients’ discomfort about using verbatim session transcripts as a data source. Some (P1, P7) worried about the confidentiality of patients’ data. P1 commented: “*The idea of having a transcript in somebody’s medical record makes me nervous ... I’m not totally sure that the trade-off and, like, the [benefit and] risk on privacy [is worth it].*” Meanwhile, P7 described that data breach would be “*really life changing for the patients*”. Others (P11, P16) questioned the practicality and validity of integrating clinical transcripts. For example, P16 thinks that transcripts are too wordy and overwhelming to be useful in

daily practice, while P11 questions how the system would handle misalignment between patients' behavioral and linguistic patterns.

Room for Improvement. While we made extensive efforts to prevent the dashboard from overwhelming clinicians, clinicians have differing opinions about MIND's complexity. For example, P4 perceived MIND as “*busy and crowded*” while P8 reported “*too much text*” on the screen. Meanwhile, both P10 and P14 reported excessive clicking interactions navigating the dashboard. In contrast, P7 believes that “*it's (P7's EHR system in real-world practice) not nearly as elegant as this (MIND)*”. The disparity in the perception of MIND's complexity surfaces future needs to personalize MIND to suit different preferences among clinicians.

Functions beyond patient review. We designed the Summary Today section as an open-ended summary that provides a hook for clinicians to communicate with patients after sessions. Clinicians (P6, P9, P11) appreciated its ability to streamline their workflow, while others (P3, P6, P8, P9) proposed imaginative uses for Summary Today. For instance, P8 reflected that “*if I could use this text generation (pointing at the Summary Today section) to write notes, that would be a really helpful tool ... if I could draft note at the end and kind of put these in, that would be helpful for me.*” Future iterations of MIND should include these functionalities to better incentivize clinicians to use the system.

7 Discussion

Our work explored how a narrative dashboard can empower mental health clinicians in understanding multimodal data. While clinicians found data presented in both systems clinically relevant, MIND's narrative design made multimodal data insights more useful in clinical practice. In this section, we reflect on our study and discuss how the narrative dashboard design could benefit broader stakeholders within and beyond mental healthcare. We also discuss the limitations of our research and envision future directions.

7.1 Improving the Curation of Multimodal Data Insights

In this study, we built a hybrid computation pipeline that transforms raw multimodal data into clinically relevant data insights. While our user study demonstrated that our pipeline can generate clinically relevant data insights, we only evaluated its performance on two simulated patient cases. A larger dataset containing heterogeneous real-world corpora to better evaluate the performance of the computation pipeline and, in turn, improve upon the current implementation. For example, research could attempt to design a full LLM-powered agentic workflow (e.g., [1, 18]) or use multimodal models pre-trained with temporal passive sensing data (e.g., [111]). We advocate for future work to propose methods to collect clinician-aligned datasets related to multimodal data to better understand what counts as useful data narratives and improve automatic algorithms in curating multimodal data insights.

Additionally, our current implementation does not explicitly account for the noise inherent within each data modality. For instance, passive sensing data can be inconsistent due to sensor variability, while clinical and active sensing data are subject to noise from imperfect recall and subjective survey scoring [5]. Although our

multimodal approach can partially counteract noise by revealing cross-modal consistencies and inconsistencies, the system does not formally quantify uncertainty—neither in the source data nor in the final curated insights. Consequently, judging the validity and weight of these insights remains a task for the clinician, with MIND serving a supplementary role for synthesizing a data-driven narrative [68]. Future work could therefore develop methods to better quantify the noise in the data and integrate explicit uncertainty cues, whether quantitative (e.g., confidence scores [112]) or qualitative, to help clinicians effectively interpret the derived insights and ensure they are “actionable” for clinical decision-making [5].

7.2 Towards Tailored and Shared Decision-making between Clinician and Patient

MIND includes a Patient Communication section, which presents a potential portal for linking data narratives with patients. However, our tool 1) remains primarily designed for clinicians and the clinical workflow, and 2) represents a design for “common grounds” of mental health practice instead of tailoring for a specific mental health challenge. Future work should therefore seek to adapt MIND to support both clinicians' workflow and the patient-clinician collaborative decision-making process during clinical sessions [63], and afford clinicians and patients tailor the interface to their specific conditions, measurements, data modalities (e.g., patient journals), and treatment plans. One promising approach is to use malleable interfaces, for example, morphing narrative dashboard content into data comics or infographics [97], to better communicate patient/treatment-specific data insights to patients (e.g., [87]). By involving patients in the process, we could foster a greater willingness to share passive sensing data and take active sensing surveys, ultimately creating a positive feedback loop between enhanced data sharing and more data-driven clinical sessions. Furthermore, achieving a shared understanding [61] of multimodal data could promote continuity of care—for instance, by enabling just-in-time notifications based on sensing data to advocate health behaviors and improve treatment adherence [80].

7.3 Narrative Dashboards beyond Mental Health

Our study demonstrated the potential of using narrative dashboards as a means to communicate multimodal data insights to mental health clinicians. We envision that similar narrative dashboards could be extended to domains beyond mental healthcare as a supplement to their current workflow. For instance, cardiologists and oncologists similarly face the challenge of integrating multimodal data sources (e.g., imaging, subjective patient reports, clinical notes) in a short amount of time to inform critical diagnostic and treatment decisions. Future work could seek to design and evaluate domain-specific narrative dashboards as a way to address the unique challenges in integrating multimodal data in other healthcare domains, with the ultimate goal to improve clinical outcomes.

7.4 Ethical Considerations Applying LLM-powered Systems in Health

In addition to the potential to further improve and expand MIND, the integration of an LLM/rule-based hybrid pipeline into MIND raises broader ethical considerations for generative technologies

in clinical workflows. For instance, recent works revealed risks related to bias, sycophancy, and stigmatizing language when applying LLMs in mental health contexts [66], which can cause serious ramifications [13]. While we designed MIND to be clinician-facing, mitigating the direct risk of exposing an underdeveloped algorithm to users, it requires sustained future efforts to understand how computationally curated data insights impact clinicians' decision-making process despite its perceived usefulness. Specifically, content curated by LLMs may be *biased* towards dominant or easily accessible data points, potentially creating a "generative echo chamber" [79] that restricts the breadth of information available to the clinician.

A closely related concern is the *trust* in LLM-powered systems. In our design, MIND employed multiple methods (e.g., showing data source types, supporting drill-down fact-checking) to foster adequate trust towards the data insights. While evaluation demonstrated most clinicians actively check the raw data before making their own judgments, some are more inclined to "overtrust" AI after only a few successes with the generated data insights. Such over-reliance raises concerns regarding the *accountability* of human clinicians working with AI and, without mitigation, may ultimately hamper treatment quality and compromise patient safety.

Additionally, *privacy* is also a concern reflected through our study. Some clinicians expressed reservations about using raw clinical transcripts as a data modality, recognizing that these sensitive records can also serve as potential legal evidence that may put both the patient and the clinician at risk. Subsequently, we advocate for future research to design privacy-preserving data infrastructures to capitalize on their benefits while mitigating inherent limitations.

7.5 Limitations and Future Work

There are several limitations in our study that inform future work opportunities. First, we used homogeneous simulated patient cases in our user study (e.g., SOAP notes, depression-only patients) to treat it as a controlled variable. However, real-world cases could be more heterogeneous (e.g., missing data, diverse note styles), and the quality of the generated insights may vary due to those variances.

Second, we recruited our participants through a combination of convenience, purposive, and snowball sampling, which may not be representative of the demographics of mental health clinicians. For instance, our sample likely over-represented clinicians who are more willing to experiment with new technologies. This is consistent with our population's overall positive attitude towards AI, which might be a contributing factor to why clinicians appreciated the computationally generated data insights.

Third, the evaluation of our computation pipeline might be limited to specific use cases. Although two clinicians reviewed and approved the generated content for the user study, this internal validation may be subjective and does not fully account for challenges using LLMs, such as potential biases during the narrative curation process. Future work should therefore combine expert assessments with large-scale benchmarking to evaluate the pipeline's performance across diverse scenarios and mitigate model risks.

Finally, we conducted our experiments using a controlled lab study. While minimizing confounding variables, it did not fully capture the nuances of actual practice (e.g., strict time constraints, interruptions). Synthesizing the three limitations above, future work

should seek to further diversify the patient data used for evaluation, expand participant demographics through stratified sampling, and observe MIND's use in real-world practice (e.g., through contextual inquiry [46], especially how it might complement current EHR systems) to better understand its applicability and potential for integration into clinical workflows.

8 Conclusion

In this paper, we propose MIND, a multimodal data-integrated narrative dashboard designed to empower mental health clinicians for patient review. Through a co-design study with five mental health experts, we present a data-driven narrative with multimodal data insights based on a hybrid computational pipeline. A user study with 16 licensed clinicians shows that MIND fulfilled our design goals, demonstrating its ability to efficiently reveal hidden, clinically relevant insights and its potential to support clinical decision-making. Clinicians found the narrative structure intuitive and valued the system's emphasis on verifiability through the drill-down panel. MIND demonstrated the feasibility and power of combining careful computation curation with narrative dashboards in presenting multimodal data. We further advocate for future work towards real-world longitudinal evaluations of MIND in clinical practice, as well as applying narrative dashboards towards other healthcare domains and diverse healthcare stakeholders.

Acknowledgment of the Use of AI

We acknowledge the use of Generative AI tools to improve the grammar, style, and readability of this manuscript. These tools were used in our computation pipeline (see details in §5) and text editing, but played no role in the data analysis, interpretation, or generation of the core findings presented.

Acknowledgments

Research reported in this publication was supported in part by the National Cancer Institute of the National Institutes of Health under Award Number R01CA301579. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We also thank the Columbia University Research Stabilization Grant and the Northeastern University Tier-1 Research Seed Grant for supporting this research project. We want to thank all participants for taking part in our study. Additionally, we extend our thanks to the anonymous reviewers for their constructive feedback.

References

- [1] Amirhossein Abaskohi, Amrutha Varshini Ramesh, Shailesh Nanisetty, Chirag Goel, David Vazquez, Christopher Pal, Spandana Gella, Giuseppe Carenini, and Issam H. Laradji. 2025. AgentAda: Skill-Adaptive Data Analytics for Tailored Insight Discovery. doi:10.48550/arXiv.2504.07421 arXiv:2504.07421 [cs]
- [2] Daniel A. Adler, Fei Wang, David C. Mohr, and Tanzeem Choudhury. 2022. Machine Learning for Passive Mental Health Symptom Prediction: Generalization across Different Longitudinal Mobile Sensing Studies. *PLOS ONE* 17, 4 (April 2022), e0266516. doi:10.1371/journal.pone.0266516
- [3] Daniel A. Adler, Fei Wang, David C. Mohr, Deborah Estrin, Cecilia Livesey, and Tanzeem Choudhury. 2022. A Call for Open Data to Develop Mental Health Digital Biomarkers. *BJPsych Open* 8, 2 (March 2022), e58. doi:10.1192/bjo.2022.28
- [4] Daniel A. Adler, Yuewen Yang, Thalia Viranda, Anna R. Van Meter, Emma Elizabeth McGinty, and Tanzeem Choudhury. 2025. Designing Technologies for Value-based Mental Healthcare: Centering Clinicians' Perspectives on Outcomes Data Specification, Collection, and Use. In *Proceedings of the 2025 CHI Conference*

- on *Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3706598.3713481
- [5] Daniel A. Adler, Yuewen Yang, Thalia Viranda, Xuhai Xu, David C. Mohr, Anna R. Van Meter, Julia C. Tartaglia, Nicholas C. Jacobson, Fei Wang, Deborah Estrin, and Tanzeem Choudhury. 2024. Beyond Detection: Towards Actionable Sensing Research in Clinical Mental Healthcare. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 4 (Nov. 2024), 1–33. doi:10.1145/3699755
 - [6] Stergos Afantenos, Vangelis Karkaletsis, and Panagiotis Stamatopoulos. 2005. Summarization from Medical Documents: A Survey. *Artificial Intelligence in Medicine* 33, 2 (Feb. 2005), 157–177. doi:10.1016/j.artmed.2004.07.017
 - [7] Ijeoma Azodo, Robin Williams, Aziz Sheikh, and Kathrin Cresswell. 2020. Opportunities and Challenges Surrounding the Use of Data From Wearable Sensor Devices in Health Care: Qualitative Interview Study. *Journal of Medical Internet Research* 22, 10 (Oct. 2020), e19542. doi:10.2196/19542
 - [8] Benjamin Bach, Euan Freeman, Alfie Abdul-Rahman, Cagatay Turkyay, Saiful Khan, Yulei Fan, and Min Chen. 2023. Dashboard Design Patterns. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (Jan. 2023), 342–352. doi:10.1109/TVCG.2022.3209448
 - [9] Aaron T Beck, Robert A Steer, Gregory K Brown, et al. 1996. Beck depression inventory. (1996).
 - [10] Rasmus H. Birk and Gabrielle Samuel. 2022. Digital Phenotyping for Mental Health: Reviewing the Challenges of Using Data to Monitor and Predict Mental Health Problems. *Current Psychiatry Reports* 24, 10 (Oct. 2022), 523–528. doi:10.1007/s11920-022-01358-9
 - [11] Guthrie S. Birkhead, Michael Klompass, and Nirav R. Shah. 2015. Uses of Electronic Health Records for Public Health Surveillance to Advance Public Health. *Annual Review of Public Health* 36, Volume 36, 2015 (March 2015), 345–359. doi:10.1146/annurev-publhealth-031914-122747
 - [12] Suzanne V Blackley, Jessica Huynh, Liqin Wang, Zfania Korach, and Li Zhou. 2019. Speech Recognition for Clinical Documentation from 1990 to 2018: A Systematic Review. *Journal of the American Medical Informatics Association* 26, 4 (April 2019), 324–338. doi:10.1093/jamia/ocy179
 - [13] Robert Booth and Robert Booth UK technology editor. 2025. Teen Killed Himself after 'Months of Encouragement from ChatGPT', Lawsuit Claims. *The Guardian* (Aug. 2025).
 - [14] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. doi:10.1191/1478088706qp0630a
 - [15] John Brooke. 1996. SUS: A 'Quick and Dirty' Usability Scale. In *Usability Evaluation In Industry*. CRC Press.
 - [16] Qing Chen, Ying Chen, Ruishi Zou, Wei Shuai, Yi Guo, Jiazhe Wang, and Nan Cao. 2024. Chart2Vec: A Universal Embedding of Context-Aware Visualizations. *IEEE Transactions on Visualization and Computer Graphics* (2024), 1–12. doi:10.1109/TVCG.2024.3383089
 - [17] Yang Chen, Jing Yang, and William Ribarsky. 2009. Toward Effective Insight Management in Visual Analytics Systems. In *2009 IEEE Pacific Visualization Symposium*. 49–56. doi:10.1109/PACIFICVIS.2009.4906837
 - [18] Akshat Choube, Ha Le, Jiachen Li, Kaixin Ji, Vedant Das Swain, and Varun Mishra. 2025. GLOSS: Group of LLMs for Open-Ended Sensemaking of Passive Sensing Data for Health and Wellbeing. doi:10.48550/arXiv.2507.05461 arXiv:2507.05461 [cs]
 - [19] Victoria Clarke and Virginia and Braun. 2017. Thematic Analysis. *The Journal of Positive Psychology* 12, 3 (May 2017), 297–298. doi:10.1080/17439760.2016.1262613
 - [20] Sheldon Cohen, Tom Kamarck, and Robin Mermelstein. 1983. A global measure of perceived stress. *Journal of health and social behavior* (1983), 385–396.
 - [21] Vedant Das Swain, Victor Chen, Shrija Mishra, Stephen M. Mattingly, Gregory D. Abowd, and Munmun De Choudhury. 2022. Semantic Gap in Predicting Mental Wellbeing through Passive Sensing. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–16. doi:10.1145/3491102.3502037
 - [22] George Demiris, Sarah J. Iribarren, Katherine Sward, Solim Lee, and Rumei Yang. 2019. Patient Generated Health Data Use in Clinical Practice: A Systematic Review. *Nursing Outlook* 67, 4 (July 2019), 311–330. doi:10.1016/j.outlook.2019.04.005
 - [23] Catherine Dinh-Le, Rachel Chuang, Sara Chokshi, and Devin Mann. 2019. Wearable Health Technology and Electronic Health Record Integration: Scoping Review and Future Directions. *JMIR mHealth and uHealth* 7, 9 (Sept. 2019), e12861. doi:10.2196/12861
 - [24] Dawn Dowding, Rebecca Randall, Peter Gardner, Geraldine Fitzpatrick, Patricia Dykes, Jesus Favela, Susan Hamer, Zac Whitewood-Moore, Nicholas Hardiker, Elizabeth Borycki, and Leanne Currie. 2015. Dashboards for Improving Patient Care: Review of the Literature. *International Journal of Medical Informatics* 84, 2 (Feb. 2015), 87–100. doi:10.1016/j.ijmedinf.2014.10.001
 - [25] Begum Egilmez, Emirhan Poyraz, Wenting Zhou, Gokhan Memik, Peter Dinda, and Nabil Alshurafa. 2017. UStress: Understanding College Student Subjective Stress Using Wrist-Based Passive Sensing. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. 673–678. doi:10.1109/PERCOMW.2017.7917644
 - [26] Jordan J. Elm, Margaret Daeschler, Lauren Bataille, Ruth Schneider, Amy Amara, Alberto J. Espay, Michal Afek, Chen Admati, Abeba Teklehaimanot, and Tanya Simuni. 2019. Feasibility and Utility of a Clinician Dashboard from Wearable and Mobile Application Parkinson's Disease Data. *npj Digital Medicine* 2, 1 (Sept. 2019), 1–6. doi:10.1038/s41746-019-0169-y
 - [27] George L. Engel. 1977. The Need for a New Medical Model: A Challenge for Biomedicine. *Science* 196, 4286 (April 1977), 129–136. doi:10.1126/science.847460
 - [28] Zachary Enghardt, Chengqian Ma, Margaret E. Morris, Chun-Cheng Chang, Xuhai "Orson" Xu, Lianhui Qin, Daniel McDuff, Xin Liu, Shwetak Patel, and Vikram Iyer. 2024. From Classification to Clinical Insights: Towards Analyzing and Reasoning About Mobile and Behavioral Health Data With Large Language Models. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 2 (May 2024), 56:1–56:25. doi:10.1145/3659604
 - [29] Ilker Etikan, Sulaiman Abubakar Musa, Rukayya Sunusi Alkassim, et al. 2016. Comparison of Convenience Sampling and Purposive Sampling. *American journal of theoretical and applied statistics* 5, 1 (2016), 1–4.
 - [30] Hayley I. Evans, Myeonghan Ryu, Theresa Hsieh, Jiawei Zhou, Kefan Xu, Kenneth W. Akers, Andrew M. Sherrill, and Rosa I. Arriaga. 2024. Using Sensor-Captured Patient-Generated Data to Support Clinical Decision-making in PTSD Therapy. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (April 2024), 1–28. doi:10.1145/3637426
 - [31] Stephen Few. 2006. *Information Dashboard Design: The Effective Visual Communication of Data*. O'Reilly Media, Inc.
 - [32] Gustav Freytag. 1895. *Technique of the drama: An exposition of dramatic composition and art*. S. Griggs.
 - [33] Leo A. Goodman. 1961. Snowball Sampling. *The Annals of Mathematical Statistics* 32, 1 (1961), 148–170. jstor:2237615 https://www.jstor.org/stable/2237615
 - [34] David Gotz and David Borland. 2016. 'Data-Driven Healthcare: Challenges and Opportunities for Interactive Visualization. *IEEE Computer Graphics and Applications* 36, 3 (May 2016), 90–96. doi:10.1109/MCG.2016.59
 - [35] Simone Grassini. 2023. Development and Validation of the AI Attitude Scale (AIAS-4): A Brief Measure of General Attitude toward Artificial Intelligence. *Frontiers in Psychology* 14 (July 2023). doi:10.3389/fpsyg.2023.1191628
 - [36] Jiyeon Han, Jimin Park, Jinyoung Huh, Uran Oh, Jaeyoung Do, and Daehee Kim. 2024. AscleAI: A LLM-based Clinical Note Management System for Enhancing Clinician Productivity. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3613905.3650784
 - [37] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Advances in Psychology*, Peter A. Hancock and Najmedin Meshkati (Eds.). Human Mental Workload, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9
 - [38] Danielle Helminski, Jeremy B. Sussman, Paul N. Pfeiffer, Alex N. Kokaly, Allison Ranusch, Anjana Deep Renji, Laura J. Damschroder, Zach Landis-Lewis, and Jacob E. Kurlander. 2024. Development, Implementation, and Evaluation Methods for Dashboards in Health Care: Scoping Review. *JMIR Medical Informatics* 12, 1 (Dec. 2024), e59828. doi:10.2196/59828
 - [39] A. Ali Heydari, Ken Gu, Vidya Srinivas, Hong Yu, Zhihan Zhang, Yuwei Zhang, Akshay Paruchuri, Qian He, Hamid Palangi, Nova Hammerquist, Ahmed A. Metwally, Brent Winslow, Yubin Kim, Kumar Ayush, Yuzhe Yang, Girish Narayanswamy, Maxwell A. Xu, Jake Garrison, Amy Aremnto Lee, Jenny Vafieadou, Ben Graef, Isaac R. Galatzer-Levy, Erik Schenck, Andrew Barakat, Javier Perez, Jacqueline Shreibati, John Hernandez, Anthony Z. Faranesh, Javier L. Prieto, Connor Heneghan, Yun Liu, Jiening Zhan, Mark Malhotra, Shwetak Patel, Tim Althoff, Xin Liu, Daniel McDuff, and Xuhai "Orson" Xu. 2025. The Anatomy of a Personal Health Agent. doi:10.48550/arXiv.2508.20148 arXiv:2508.20148 [cs]
 - [40] Karen Hovsepian, Mustafa al'Absi, Emre Ertin, Thomas Kamarck, Motohiro Nakajima, and Santosh Kumar. 2015. cStress: Towards a Gold Standard for Continuous Stress Assessment in the Mobile Environment. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. Association for Computing Machinery, New York, NY, USA, 493–504. doi:10.1145/2750858.2807526
 - [41] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems* 43, 2 (2025), 1–55.
 - [42] Jessica Hullman, Steven Drucker, Nathalie Henry Riche, Bongshin Lee, Danyel Fisher, and Eytan Adar. 2013. A Deeper Understanding of Sequence in Narrative Visualization. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec. 2013), 2406–2415. doi:10.1109/TVCG.2013.119
 - [43] Thomas R. Insel. 2017. Digital Phenotyping: Technology for a New Science of Behavior. *JAMA* 318, 13 (Oct. 2017), 1215–1216. doi:10.1001/jama.2017.11295
 - [44] Zhihan Jiang, Handi Chen, Rui Zhou, Jing Deng, Xinchun Zhang, Running Zhao, Cong Xie, Yifang Wang, and Edith C.H. Ngai. 2024. HealthPrism: A Visual Analytics System for Exploring Children's Physical and Mental Health

- Profiles with Multimodal Data. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (Jan. 2024), 1205–1215. doi:10.1109/TVCG.2023.3326943
- [45] Jade Kandel, Chelsea Duppen, Qian Zhang, Howard Jiang, Angelos Angelopoulos, Ashley Paula-Ann Neall, Pranav Wagh, Daniel Szafir, Henry Fuchs, Michael Lewek, and Danielle Albers Szafir. 2024. PD-Insighter: A Visual Analytics System to Monitor Daily Actions for Parkinson's Disease Treatment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3613904.3642215
- [46] Holtzblatt Karen and Jones Sandra. 1993. Contextual Inquiry: A Participatory Technique for System Design. In *Participatory Design*. CRC Press.
- [47] Justin Khasentino, Anastasiya Belyaeva, Xin Liu, Zhun Yang, Nicholas A. Furlotte, Chace Lee, Erik Schenck, Yojan Patel, Jian Cui, Logan Douglas Schneider, Robby Bryant, Ryan G. Gomes, Allen Jiang, Roy Lee, Yun Liu, Javier Perez, Jameson K. Rogers, Cathy Speed, Shyam Tailor, Megan Walker, Jeffrey Yu, Tim Althoff, Conor Heneghan, John Hernandez, Mark Malhotra, Leor Stern, Yossi Matias, Greg S. Corrado, Shwetak Patel, Shrivya Shetty, Jiening Zhan, Shruthi Prabhakara, Daniel McDuff, and Cory Y. McLean. 2025. A Personal Health Large Language Model for Sleep and Fitness Coaching. *Nature Medicine* (Aug. 2025), 1–10. doi:10.1038/s41591-025-03888-0
- [48] Yoojung Kim, Eunyoung Heo, Hyunjeong Lee, Sookyoung Ji, Jueun Choi, Jeong-Whun Kim, Joongseek Lee, and Sooyoung Yoo. 2017. Prescribing 10,000 Steps Like Aspirin: Designing a Novel Interface for Data-Driven Medical Consultations. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 5787–5799. doi:10.1145/3025453.3025570
- [49] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large Language Models Are Zero-Shot Reasoners. In *Advances in Neural Information Processing Systems*.
- [50] Kurt Kroenke, Robert L. Spitzer, and Janet BW Williams. 2003. The Patient Health Questionnaire-2: validity of a two-item depression screener. *Medical care* 41, 11 (2003), 1284–1292.
- [51] Kurt Kroenke, Robert L. Spitzer, Janet BW Williams, Patrick O Monahan, and Bernd Löwe. 2007. Anxiety disorders in primary care: prevalence, impairment, comorbidity, and detection. *Annals of internal medicine* 146, 5 (2007), 317–325.
- [52] Kurt Kroenke, Robert L. Spitzer, and Janet B. W. Williams. 2001. The PHQ-9. *Journal of General Internal Medicine* 16, 9 (2001), 606–613. doi:10.1046/j.1525-1497.2001.016009606.x
- [53] Bum Chul Kwon, Min-Je Choi, Joanne Taery Kim, Edward Choi, Young Bin Kim, Soonwook Kwon, Jimeng Sun, and Jaegul Choo. 2019. RetainVis: Visual Analytics with Interpretable and Interactive Recurrent Neural Networks on Electronic Medical Records. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 299–309. doi:10.1109/TVCG.2018.2865027
- [54] Xingyu Lan, Xinyue Xu, and Nan Cao. 2021. Understanding Narrative Linearity for Telling Expressive Time-Oriented Stories. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3411764.3445344
- [55] Bongshin Lee, Nathalie Henry Riche, Petra Isenber, and Sheelagh Carpendale. 2015. More Than Telling a Story: Transforming Data into Visually Shared Stories. *IEEE Computer Graphics and Applications* 35, 5 (Sept. 2015), 84–90. doi:10.1109/MCG.2015.99
- [56] Boning Li and Akane Sano. 2020. Extraction and Interpretation of Deep Autoencoder-based Temporal Features from Wearables for Forecasting Personalized Mood, Health, and Stress. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 2 (June 2020), 49:1–49:26. doi:10.1145/3397318
- [57] Jiachen Li, Xiwen Li, Justin Steinberg, Akshat Choubhe, Bingsheng Yao, Xuhai Xu, Dakuo Wang, Elizabeth Mynatt, and Varun Mishra. 2025. Vital Insight: Assisting Experts' Context-Driven Sensemaking of Multi-modal Personal Tracking Data Using Visualization and Human-In-The-Loop LLM. doi:10.48550/arXiv.2410.14879 arXiv:2410.14879 [cs]
- [58] Steven Y. Lin, Tait D. Shanafelt, and Steven M. Asch. 2018. Reimagining Clinical Documentation With Artificial Intelligence. *Mayo Clinic Proceedings* 93, 5 (May 2018), 563–565. doi:10.1016/j.mayocp.2018.02.016
- [59] Bernd Löwe, Inka Wahl, Matthias Rose, Carsten Spitzer, Heide Glaesmer, Katja Wingenfeld, Antonius Schneider, and Elmar Brähler. 2010. A 4-Item Measure of Depression and Anxiety: Validation and Standardization of the Patient Health Questionnaire-4 (PHQ-4) in the General Population. *Journal of Affective Disorders* 122, 1 (April 2010), 86–95. doi:10.1016/j.jad.2009.06.019
- [60] Ananya Mahapatra and Prerna Sharma. 2024. Biopsychosocial Approach to Mental Health. In *Mental Health Care Resource Book: Concepts and Praxis for Social Workers and Mental Health Professionals*, Meenu Anand (Ed.). Springer Nature, Singapore, 87–100. doi:10.1007/978-981-97-1203-8_5
- [61] Rose McCabe. 2021. When patients and clinician (dis)agree about the nature of the problem: The role of displays of shared understanding in acceptance of treatment. *Social Science & Medicine* 290 (2021), 114208. doi:10.1016/j.socscimed.2021.114208
- [62] Lakmal Meegapapola, William Droz, Peter Kun, Amalia de Götzen, Chaitanya Nutakki, Shyam Diwakar, Prernava Ruiz Correa, Donglei Song, Hao Xu, Miriam Bidoglia, George Gaskell, Altangerel Chagnaa, Amarsanaa Ganbold, Tzolmon Zundui, Carlo Caprini, Daniele Miorandi, Alethia Hume, Jose Luis Zarza, Luca Cernuzzi, Ivano Bison, Marcelo Rodas Britez, Matteo Busso, Ronald Chenu-Abente, Can Günel, Fausto Giunchiglia, Laura Schelenz, and Daniel Gatica-Perez. 2023. Generalization and Personalization of Mobile Sensing-Based Mood Inference Models: An Analysis of College Students in Eight Countries. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 4 (Jan. 2023), 176:1–176:32. doi:10.1145/3569483
- [63] Helena M. Mentis, Anita Komlodi, Katrina Schrader, Michael Phipps, Ann Gruber-Baldini, Karen Yarbrough, and Lisa Shulman. 2017. Crafting a View of Self-Tracking Data in the Clinical Visit. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 5800–5812. doi:10.1145/3025453.3025589
- [64] Rashmi Mishra, Jiantao Bian, Marcelo Fiszman, Charlene R. Weir, Siddhartha Jonnalagadda, Javed Mostafa, and Guilherme Del Fiol. 2014. Text Summarization in the Biomedical Domain: A Systematic Review of Recent Research. *Journal of Biomedical Informatics* 52 (Dec. 2014), 457–467. doi:10.1016/j.jbi.2014.06.009
- [65] 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 (May 2017), 23–47. doi:10.1146/annurev-clinpsy-032816-044949
- [66] Jared Moore, Declan Grabb, William Agnew, Kevin Klyman, Stevie Chancellor, Desmond C. Ong, and Nick Haber. 2025. Expressing Stigma and Inappropriate Responses Prevents LLMs from Safely Replacing Mental Health Providers. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Athens Greece, 599–627. doi:10.1145/3715275.3732039
- [67] Mehrab Bin Morshed, Koustuv Saha, Richard Li, Sidney K. D'Mello, Mumun De Choudhury, Gregory D. Abowd, and Thomas Plötz. 2019. Prediction of Mood Instability with Passive Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 3 (Sept. 2019), 75:1–75:21. doi:10.1145/3351233
- [68] Jodie Nghiemi, Daniel A. Adler, Deborah Estrin, Cecilia Livesey, and Tanzeem Choudhury. 2023. Understanding Mental Health Clinicians' Perceptions and Concerns Regarding Using Passive Patient-Generated Health Data for Clinical Decision-Making: Qualitative Semistructured Interview Study. *JMIR Formative Research* 7, 1 (Aug. 2023), e47380. doi:10.2196/47380
- [69] Jukka-Pekka Onnela and Scott L. Rauch. 2016. Harnessing Smartphone-Based Digital Phenotyping to Enhance Behavioral and Mental Health. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology* 41, 7 (June 2016), 1691–1696. doi:10.1038/npp.2016.7
- [70] S. M. Powsner and E. R. Tuft. 1994. Graphical Summary of Patient Status. *The Lancet* 344, 8919 (Aug. 1994), 386–389. doi:10.1016/S0140-6736(94)91406-0
- [71] Shehjar Sadhu, Dhaval Solanki, Leslie A. Brick, Nicole R. Nugent, and Kunal Mankodiya. 2023. Designing a Clinician-Centered Wearable Data Dashboard (CarePortal): Participatory Design Study. *JMIR Formative Research* 7, 1 (Dec. 2023), e46866. doi:10.2196/46866
- [72] Edward P Sarafino and Timothy W Smith. 2014. *Health psychology: Biopsychosocial interactions*. John Wiley & Sons.
- [73] Alper Sarikaya, Michael Correll, Lyn Bartram, Melanie Tory, and Danyel Fisher. 2019. What Do We Talk About When We Talk About Dashboards? *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 682–692. doi:10.1109/TVCG.2018.2864903
- [74] Benjamin Schloss and Sandeep Konam. 2020. Towards an Automated SOAP Note: Classifying Utterances from Medical Conversations. In *Proceedings of the 5th Machine Learning for Healthcare Conference*. PMLR, 610–631.
- [75] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. 2012. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2431–2440. doi:10.1109/TVCG.2012.213
- [76] Edward Segel and Jeffrey Heer. 2010. Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (Nov. 2010), 1139–1148. doi:10.1109/TVCG.2010.179
- [77] Yuval Shahar and Cleve Cheng. 2000. Model-Based Visualization of Temporal Abstractions. *Computational Intelligence* 16, 2 (2000), 279–306. doi:10.1111/0824-7935.00114
- [78] Yuval Shahar, Dina Goren-Bar, David Boaz, and Gil Tahan. 2006. Distributed, Intelligent, Interactive Visualization and Exploration of Time-Oriented Clinical Data and Their Abstractions. *Artificial Intelligence in Medicine* 38, 2 (Oct. 2006), 115–135. doi:10.1016/j.artmed.2005.03.001
- [79] Nikhil Sharma, Q. Vera Liao, and Ziang Xiao. 2024. Generative Echo Chamber? Effect of LLM-Powered Search Systems on Diverse Information Seeking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 1033, 17 pages. doi:10.1145/3613904.3642459
- [80] Moushumi Sharmin, Andrew Raji, David Epstein, Inbal Nahum-Shani, J. Gayle Beck, Sudip Vhaduri, Kenzie Preston, and Santosh Kumar. 2015. Visualization of Time-Series Sensor Data to Inform the Design of Just-in-Time Adaptive Stress Interventions. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. Association for Computing

- Machinery, New York, NY, USA, 505–516. doi:10.1145/2750858.2807537
- [81] Edna Shenvi, Aziz Boxwala, Dean Sittig, Courtney Zott, Edwin Lomotan, James Swiger, and Prashila Dullabh. 2023. Visualization of Patient-Generated Health Data: A Scoping Review of Dashboard Designs. *Applied Clinical Informatics* 14, 5 (Nov. 2023), 913–922. doi:10.1055/a-2174-7820
- [82] Danqing Shi, Xinyue Xu, Fuling Sun, Yang Shi, and Nan Cao. 2021. Calliope: Automatic Visual Data Story Generation from a Spreadsheet. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 453–463. doi:10.1109/TVCG.2020.3030403
- [83] B. Shneiderman. 1996. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages*. 336–343. doi:10.1109/VL.1996.545307
- [84] Ben Shneiderman, Catherine Plaisant, and Bradford W. Hesse. 2013. Improving Healthcare with Interactive Visualization. *Computer* 46, 5 (May 2013), 58–66. doi:10.1109/MC.2013.38
- [85] Venkatesh Sivaraman, Leigh A Bukowski, Joel Levin, Jeremy M. Kahn, and Adam Perer. 2023. Ignore, Trust, or Negotiate: Understanding Clinician Acceptance of AI-Based Treatment Recommendations in Health Care. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3544548.3581075
- [86] Robert L. Spitzer, Kurt Kroenke, Janet B. W. Williams, and Bernd Löwe. 2006. A Brief Measure for Assessing Generalized Anxiety Disorder: The GAD-7. *Archives of Internal Medicine* 166, 10 (May 2006), 1092–1097. doi:10.1001/archinte.166.10.1092
- [87] Konstantin R. Strömel, Stanislas Henry, Tim Johansson, Jasmin Niess, and Paweł W. Woźniak. 2024. Narrating Fitness: Leveraging Large Language Models for Reflective Fitness Tracker Data Interpretation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3613904.3642032
- [88] Nicole Sultanum, Michael Brudno, Daniel Wigdor, and Fanny Chevalier. 2018. More Text Please! Understanding and Supporting the Use of Visualization for Clinical Text Overview. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3173996
- [89] Nicole Sultanum, Farooq Naem, Michael Brudno, and Fanny Chevalier. 2023. ChartWalk: Navigating Large Collections of Text Notes in Electronic Health Records for Clinical Chart Review. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (Jan. 2023), 1244–1254. doi:10.1109/TVCG.2022.3209444
- [90] Nicole Sultanum, Devin Singh, Michael Brudno, and Fanny Chevalier. 2019. Doccurate: A Curation-Based Approach for Clinical Text Visualization. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 142–151. doi:10.1109/TVCG.2018.2864905
- [91] R Core Team. 2013. R: A language and environment for statistical computing.
- [92] Sanne van der Linden, Rita Sevastjanova, Mathias Funk, and Mennatallah El-Assady. 2023. MediCoSpace: Visual Decision-Support for Doctor-Patient Consultations Using Medical Concept Spaces from EHRs. *ACM Trans. Manage. Inf. Syst.* 14, 2 (Jan. 2023), 15:1–15:20. doi:10.1145/3564275
- [93] Rui Wang, Weichen Wang, Alex daSilva, Jeremy F. Huckins, William M. Kelley, Todd F. Heatherton, and Andrew T. Campbell. 2018. Tracking Depression Dynamics in College Students Using Mobile Phone and Wearable Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1 (March 2018), 43:1–43:26. doi:10.1145/3191775
- [94] Taowei David Wang, Catherine Plaisant, Alexander J. Quinn, Roman Stanchak, Shawn Murphy, and Ben Shneiderman. 2008. Aligning Temporal Data by Sentinel Events: Discovering Patterns in Electronic Health Records. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. Association for Computing Machinery, New York, NY, USA, 457–466. doi:10.1145/1357054.1357129
- [95] Will Ke Wang, Rui Yang, Chao Pang, Karthik Natarajan, Nan Liu, Daniel McDuff, David J Slotwiner, Fei Wang, Matthew McDermott, Xuhai Xu, et al. 2025. Beyond the Clinic: A Large-Scale Evaluation of Augmenting EHR with Wearable Data for Diverse Health Prediction. *arXiv preprint arXiv:2509.22920* (2025).
- [96] Yun Wang, Zhida Sun, Haidong Zhang, Weiwei Cui, Ke Xu, Xiaojuan Ma, and Dongmei Zhang. 2020. DataShot: Automatic Generation of Fact Sheets from Tabular Data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 895–905. doi:10.1109/TVCG.2019.2934398
- [97] Zezhong Wang, Shunming Wang, Matteo Farinella, Dave Murray-Rust, Nathalie Henry Riche, and Benjamin Bach. 2019. Comparing Effectiveness and Engagement of Data Comics and Infographics. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300483
- [98] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology* 54, 6 (1988), 1063.
- [99] Steve Wexler, Jeffrey Shaffer, and Andy Cotgreave. 2017. *The Big Book of Dashboards: Visualizing Your Data Using Real-World Business Scenarios*. John Wiley & Sons.
- [100] Siyi Wu, Weidan Cao, Shihan Fu, Bingsheng Yao, Ziqi Yang, Changchang Yin, Varun Mishra, Daniel Addison, Ping Zhang, and Dakuo Wang. 2025. CardioAI: A Multimodal AI-based System to Support Symptom Monitoring and Risk Prediction of Cancer Treatment-Induced Cardiotoxicity. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–22. doi:10.1145/3706598.3714272
- [101] Tongshuang Wu, Michael Terry, and Carrie Jun Cai. 2022. AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–22. doi:10.1145/3491102.3517582
- [102] Xuhai Xu, Prerna Chikersal, Afsaneh Doryab, Daniella K. Villalba, Janine M. Dutcher, Michael J. Tumminia, Tim Althoff, Sheldon Cohen, Kasey G. Creswell, J. David Creswell, Jennifer Mankoff, and Anind K. Dey. 2019. Leveraging Routine Behavior and Contextually-Filtered Features for Depression Detection among College Students. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 3 (Sept. 2019), 116:1–116:33. doi:10.1145/3351274
- [103] Xuhai Xu, Prerna Chikersal, Janine M. Dutcher, Yasaman S. Sefidgar, Woosuk Seo, Michael J. Tumminia, Daniella K. Villalba, Sheldon Cohen, Kasey G. Creswell, J. David Creswell, Afsaneh Doryab, Paula S. Nurius, Eve Riskin, Anind K. Dey, and Jennifer Mankoff. 2021. Leveraging Collaborative-Filtering for Personalized Behavior Modeling: A Case Study of Depression Detection among College Students. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 1 (March 2021), 41:1–41:27. doi:10.1145/3448107
- [104] Xuhai Xu, Xin Liu, Han Zhang, Weichen Wang, Subigya Nepal, Yasaman Sefidgar, Woosuk Seo, Kevin S. Kuehn, Jeremy F. Huckins, Margaret E. Morris, Paula S. Nurius, Eve A. Riskin, Shwetak Patel, Tim Althoff, Andrew Campbell, Anind K. Dey, and Jennifer Mankoff. 2023. GLOBEM: Cross-Dataset Generalization of Longitudinal Human Behavior Modeling. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 4 (Jan. 2023), 190:1–190:34. doi:10.1145/3569485
- [105] Xuhai Xu, Han Zhang, Yasaman Sefidgar, Yiyi Ren, Xin Liu, Woosuk Seo, Jennifer Brown, Kevin Kuehn, Mike Merrill, Paula Nurius, Shwetak Patel, Tim Althoff, Margaret Morris, Eve Riskin, Jennifer Mankoff, and Anind Dey. 2022. GLOBEM Dataset: Multi-Year Datasets for Longitudinal Human Behavior Modeling Generalization. *Advances in Neural Information Processing Systems* 35 (Dec. 2022), 24655–24692.
- [106] Leni Yang, Xian Xu, XingYu Lan, Ziyang Liu, Shunan Guo, Yang Shi, Huamin Qu, and Nan Cao. 2022. A Design Space for Applying the Freytag's Pyramid Structure to Data Stories. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (Jan. 2022), 922–932. doi:10.1109/TVCG.2021.3114774
- [107] Bingsheng Yao, Menglin Zhao, Yuling Sun, Weidan Cao, Changchang Yin, Stephen Intille, Xuhai Xu, Ping Zhang, Jingzhen Yang, and Dakuo Wang. 2025. More Modality, More AI: Exploring Design Opportunities of AI-Based Multi-modal Remote Monitoring Technologies for Early Detection of Mental Health Sequelae in Youth Concussion Patients. doi:10.48550/arXiv.2502.03732 [cs]
- [108] Jiancheng Ye. 2021. The Impact of Electronic Health Record-Integrated Patient-Generated Health Data on Clinician Burnout. *Journal of the American Medical Informatics Association* 28, 5 (May 2021), 1051–1056. doi:10.1093/jamia/ocab017
- [109] Han Yu and Akane Sano. 2023. Semi-Supervised Learning for Wearable-based Momentary Stress Detection in the Wild. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 2 (June 2023), 80:1–80:23. doi:10.1145/3596246
- [110] Nathan Zeldes and Neil Baum. 2011. Information Overload in Medical Practice. *The Journal of medical practice management: MPM* 26, 5 (2011), 314–316.
- [111] Yuwei Zhang, Kumar Ayush, Siyuan Qiao, A. Ali Heydari, Girish Narayanswamy, Maxwell A. Xu, Ahmed A. Metwally, Shawn Xu, Jake Garrison, Xuhai Xu, Tim Althoff, Yun Liu, Pushmeet Kohli, Jiening Zhan, Mark Malhotra, Shwetak Patel, Cecilia Mascolo, Xin Liu, Daniel McDuff, and Yuzhe Yang. 2025. SensorLM: Learning the Language of Wearable Sensors. doi:10.48550/arXiv.2506.09108 [cs]
- [112] Yunfeng Zhang, Q. Vera Liao, and Rachel K. E. Bellamy. 2020. Effect of Confidence and Explanation on Accuracy and Trust Calibration in AI-assisted Decision Making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 295–305. doi:10.1145/3351095.3372852
- [113] Zhiyuan Zhang, Bing Wang, Faisal Ahmed, I.V. Ramakrishnan, Rong Zhao, Asa Vicedello, and Klaus Mueller. 2013. The Five Ws for Information Visualization with Application to Healthcare Informatics. *IEEE Transactions on Visualization and Computer Graphics* 19, 11 (Nov. 2013), 1895–1910. doi:10.1109/TVCG.2013.89
- [114] Mengdie Zhuang, David Concannon, and Ed Manley. 2022. A Framework for Evaluating Dashboards in Healthcare. *IEEE Transactions on Visualization and Computer Graphics* 28, 4 (April 2022), 1715–1731. doi:10.1109/TVCG.2022.3147154
- [115] Ruishi Zou, Yinqi Tang, Jingzhu Chen, Siyu Lu, Yan Lu, Yingfan Yang, and Chen Ye. 2025. GistVis: Automatic Generation of Word-scale Visualizations from Data-rich Documents. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713881

A Artifacts Used in Co-Design Studies

We provide three checkpoint designs (corresponding to Fig. 3.C1-C3) of MIND (Fig. 13, Fig. 14, and Fig. 15) that mark significant design changes to demonstrate how MIND improved throughout the co-design process.

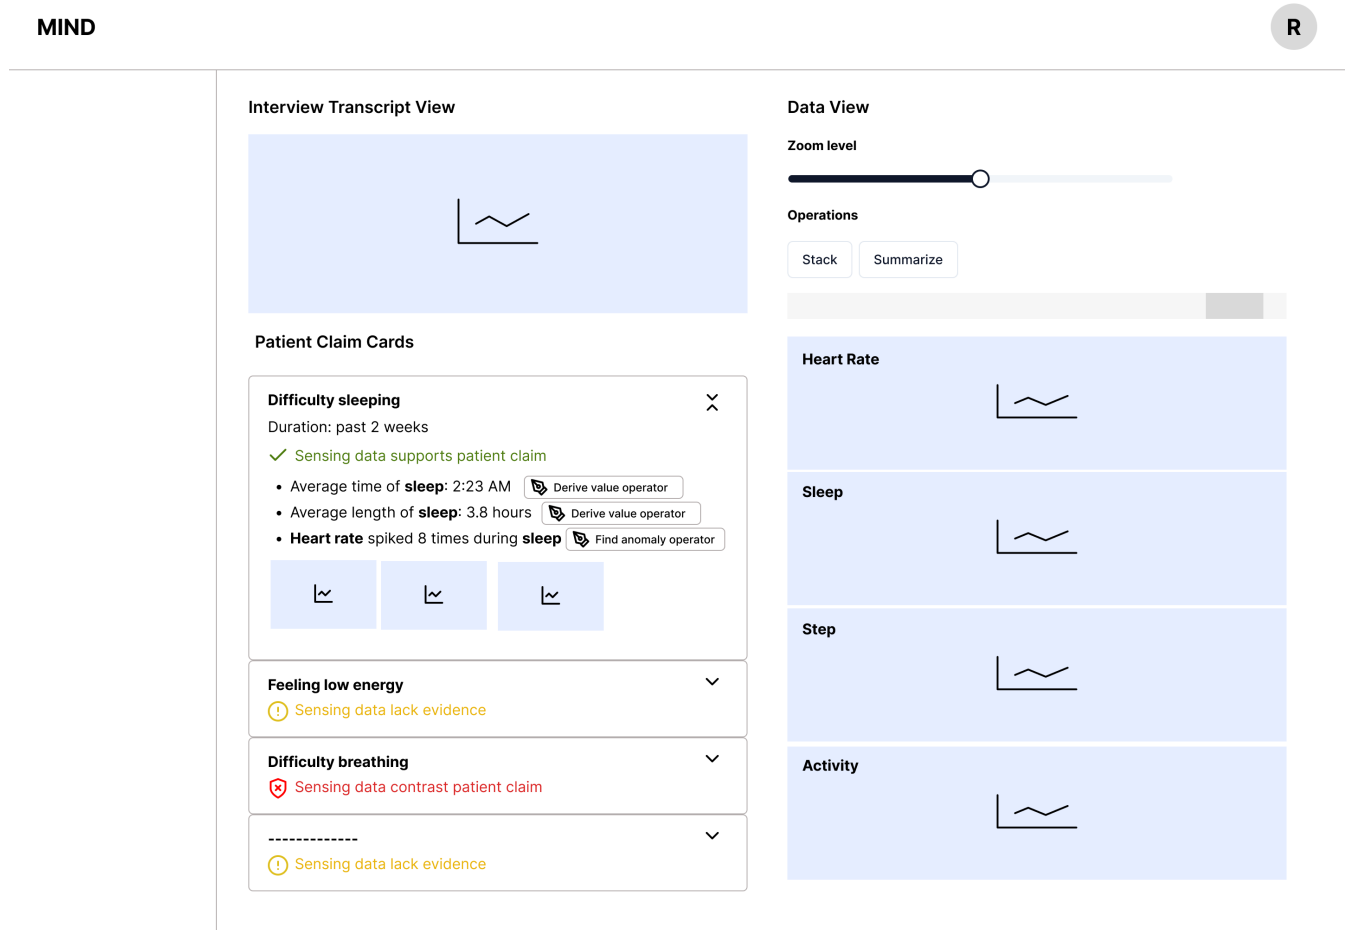


Figure 13: First low-fidelity prototype (Fig. 3.C1). Clinicians pointed out problems such as vague wording, excessive use of graphics, charts being too small within the cards, and an overwhelming dashboard with limited structure.

Filter By: Sleep 🛌 Emotions 😊 Activity 🏃 Medicine 🏠 Social Interaction 🗣️ Passed: 2 Weeks ▾

Summary of Patient Data Insights

Consistent sleep disruption

Consistent sleep disruption observed. Patient reports insomnia and fragmented sleep in session.

Improving sleep pattern

Improved sleep reported and supported by passive data. Patient mentions better rest and fewer interruptions.

⚠️ Contradictory

Discrepancy: Reports better sleep, data suggests otherwise

Patient reports improved sleep, but data suggests ongoing disruption. Sleep quality described as "good" in session, but wearable shows <5 hours average sleep and multiple interruptions nightly. Consider probing further in next session.



✓ Decline in Mood 😞

- Low Mood: Mood score of 4/10
- Reduced Social Interactions: 2 social interactions this week

📌 PHQ-9 score of 18 indicates moderately severe depression.

⊗ Anxiety with Occasional Stress 🔄

- Mild Anxiety: Average anxiety level of 3/10.
- Panic Attacks: 1 per week.
- Heart rate spike: 1 for past week

📌 The GAD-7 supports score of 7 (mild anxiety)

⚠️ Increased Irritability and Frustration 😞

- Noticeable Irritability: 3 instances
- Speech Pattern Changes: fast than usual
- Rapid Screen Switching

📌 The emergence of increased irritability is a new

✓ Difficulty sleeping 🛌

- Late sleeping: 2:23AM

- Normal length of sleep: 3.8 hours

- 📌 Heart rate spiked 8 times during sleep

✓ Poor Sleep Impacting Daily Mood 🛌 😞

- Consistent Sleep Debt

- Direct Link to Mood: aligning with the objective sleep data and lowered mood score (5/10)

- 📌 sharp increase in detected irritability following a night where sleep duration dropped to its lowest point

Figure 14: A representative wireframe of an earlier design artifact of MIND (Fig. 3.C2). At this stage, the design focuses on presenting scattered “data facts” to clinicians without a clear narrative structure.

MIND Gabriella Lin Last Encounter: 2021-05-09 Today: 2021-06-07

MIND Survey scores Clinical notes Transcripts Passive sensing

Overview

Patient Summary
Name: Gabriella Lin
Ethnicity: Asian
Gender: Female
Age: 20
Income: 0k-5k
Occupation: Student
Generation: First Generation

Medical History: Medical Concerns: Mild persistent asthma, allergic rhinitis, migraine, insomnia, acne vulgaris, recurrent UTI, acute appendicitis, low back pain, routine immunizations, acute pharyngitis; Psychiatric Concerns: Generalized anxiety disorder, adjustment disorder with depressed mood, mild and moderate depressive episodes, social anxiety disorder.

Sessions Recap: Persistent low mood, fatigue, negative self-perception, social withdrawal, sleep disturbance, minimal improvement, ongoing self-criticism.

Current Concerns: Emotional numbness, entrenched negative thoughts, difficulty communicating with mother, academic/career pressure, isolation.

Medication and Treatment: 2021-03-28: Sertraline; 2021-04-11: Sertraline; 2021-05-09: Sertraline

Clinical Insights

All category: Sleep Patterns Physical Activity Digital Engagement Emotional State Social Interaction Medication & Treatment

- High negative affect and sedentary behavior observed with family and academic stressors.
- Sleep disturbances co-occur with high negative affect and digital rumination.
- Fragmented sleep with frequent awakenings and reported fatigue.
- Increased phone use linked to sleep difficulty.
- Elevated stress, negative affect, and frequent sleep interruptions.
- Minor improvement in depression scores; negative affect and sleep disruption remain.
- Minimizes achievements and focuses on perceived failures.
- Negative self-appraisal and emotional numbness persist despite therapy engagement.
- Social withdrawal persists, with brief increases in outgoing calls.
- Work shifts and self-care maintained; social isolation persists.
- Consistent medication adherence, no ongoing side effects, slight PHQ-4 improvement.
- Therapy engagement maintained despite moderate depressive symptoms and low positive affect.

Patient Communication Draft Patient Message

Selected Patient-facing Data Insights
* Toggle selection from Clinical Insights Section to add/remove insights

Suggested Activities
* Toggle activities to include an activity suggestion to patients

- Mindful Walking
- Intentional Music Listening
- Gratitude or Positivity Journal
- Gentle Stretching or Yoga
- Low-Pressure Social Connection
- Sleep Wind-Down Routine

Figure 15: The first interactive version of MIND (Fig. 3.C3, scrolling content expanded). At this stage, MIND incorporates the three-act narrative structure, but without a clear timeline structure or narrative inside the *Clinical Insights* section.

B Data Features

B.1 Survey Scores

We document the survey scores used for the mock patients in Tab. 4.

Table 4: Selected Routine Survey Features from GLOBEM [105]

Feature Name	Description
PHQ-4 [59]	A 4-item survey for mental health conditions, including anxiety and depression. Each item is rated 0-3, with a total score range of 0-12, calculated by summing all four items. Interpretation of total scores: Normal (0-2), Mild (3-5), Moderate (6-8), Severe (9-12).
PHQ-4 Anxiety [59]	A subscale of PHQ-4 (equivalent to GAD-2 [51]), consisting of the first two items of PHQ-4. Scores range from 0-6, with scores greater than 3 suggesting anxiety.
PHQ-4 Depression [59]	A subscale of PHQ-4 (equivalent to PHQ-2 [50]), consisting of the last two items of PHQ-4. Scores range from 0-6, with scores greater than 3 suggesting depression.
PSS-4 [20]	A 4-item measure of perceived stress. Each item is rated 0-4, with a total score range of 0-16. Higher scores indicate greater levels of perceived stress, with scores greater than 6 suggesting high stress.
Positive Affect [98]	A 5-item subscale of the Positive and Negative Affect Schedule Short Form (PANAS-SF) that measures positive affect. Each item is rated 1-5, with a total score range of 5-25. Higher scores reflect greater positive affect.
Negative Affect [98]	A 5-item subscale of the Positive and Negative Affect Schedule Short Form (PANAS-SF) that measures negative affect. Each item is rated 1-5, with a total score range of 5-25. Higher scores reflect greater negative affect.

B.2 Passive Sensing Data

We document the passive sensing data features used for the mock patients in Tab. 5.

Table 5: Selected Daily Passive Sensing Features from GLOBEM [105]

Feature Name	Original Unit	Translated Unit	Description
Location			
Distance Traveled	meters	miles	Total distance traveled.
Time Spent at Home	minutes	hours	Time spent at home.
Routine Consistency	0-1	0-100	A score between 0-100: 0 denotes irregular and 100 regular.
Screen			
Phone Unlocks	count	times	Number of unlock episodes.
Total Screen Time	minutes	hours	Total duration of all unlock episodes.
Sleep			
Awakening Episodes	count	times	The number of times a person wakes up during sleep.
Total Sleep	minutes	hours	Total sleep duration.
Bedtime	minutes	hours	Time of first bedtime after midnight.
Wake Time	minutes	hours	Time of last wake time after midnight.
Steps			
Inactive Periods	count	times	Number of sedentary episodes (periods of inactivity).
Total Steps	count	count	Total step count.

C Simulated Patient Cases

We document patient-seed interview link pair in Tab. 6.

D Survey Questions for System Evaluation

Quality

- (1) **Clinical Relevance:** The information presented in the system was clinically relevant.
- (2) **Information Trustworthiness:** I felt I could trust the information provided by the system.
- (3) **Discover Hidden Data Insights:** The system's presentation of data helped me identify important relationships or patterns that were not immediately obvious.

Table 6: Simulated Patient Cases

Simulated Patient Name	P# in GLOBEM	Seed Interview Session Link
Gabriella Lin	963	https://www.youtube.com/watch?v=JKUFWK6iSsw
Lucy Sutton	1044	https://www.youtube.com/watch?v=7LD8iC4NqXM
Alison Daniels	1077	https://www.youtube.com/watch?v=4YhpWZCdiZc

(4) **Assist Clinical Decision-making:** Reviewing the patient’s data in the system increased my confidence in my clinical assessment.

Perception

- (5) **Integration of Multimodal Data:** It was easy to bring together information from different data sources to understand the overall patient situation.
- (6) **Cohesiveness of Presentation:** The system presented the patient’s health information as a clear, cohesive story rather than as separate, unrelated facts.
- (7) **Ease of Interpretation:** I could easily interpret the data presented by the system.
- (8) **Efficiency of Presentation:** Using this system made reviewing patient information more efficient.
- (9) **Patient Communication Support:** Using this system could help facilitate more effective conversations during patient sessions.
- (10) **Patient Review Time Saving:** Using this system could help me save time when reviewing a patient’s condition.

E Supplemental Evaluation Results

We show the NASA-TLX results in Fig. 16.

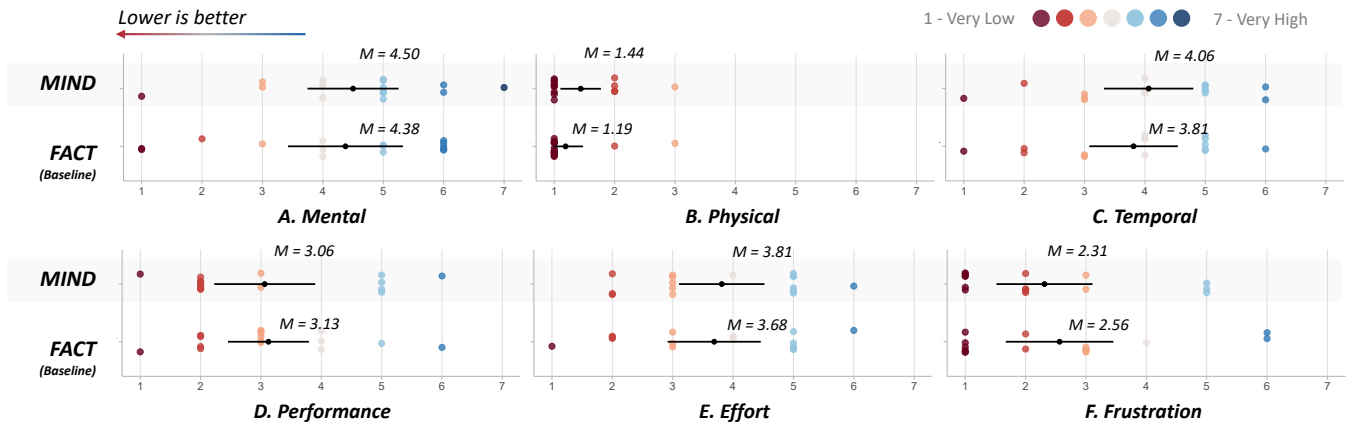


Figure 16: A comparison of perceived workload between MIND and FACT on NASA-TLX. Each point in the chart represents one individual record from one participant. We map the values with a divergent color scheme from dark red (1-very low) to dark blue (7-very high). The black dot and the black line indicate the mean and 95% confidence interval of the metric.