Get To The Point! Problem-Based Curated Data Views To Augment Care For Critically Ill Patients

Minfan Zhang
University of Toronto
Toronto, Canada
zhang.minfan@mail.utoronto.ca

Daniel Ehrmann
Hospital for Sick Children
Toronto, Canada
daniel.ehrmann@sickkids.ca

Mjaye Mazwi
Hospital for Sick Children
Toronto, Canada
mjaye.mazwi@sickkids.ca

Danny Eytan
Hospital for Sick Children
Toronto, ON, Canada
biliary.colic@gmail.com

Marzyeh Ghassemi
MIT
Cambridge, MA, United States
mghassemi@mit.edu

Fanny Chevalier
University of Toronto
Toronto, Canada
fanny@cs.toronto.edu

ABSTRACT

Electronic health records in critical care medicine offer unprecedented opportunities for clinical reasoning and decision making. Paradoxically, these data-rich environments have also resulted in clinical decision support systems (CDSSs) that fit poorly into clinical contexts, and increase health workers cognitive load. In this paper, we introduce a novel approach to designing CDSSs that are embedded in clinical workflows, by presenting problem-based curated data views tailored for problem-driven discovery, team communication, and situational awareness. We describe the design and evaluation of one such CDSS, *In-Sight*, that embodies our approach and addresses the clinical problem of monitoring critically ill pediatric patients. Our work is the result of a co-design process, further informed by empirical data collected through formal usability testing, focus groups, and a simulation study with domain experts. We discuss the potential and limitations of our approach, and share lessons learned in our iterative co-design process.

CCS CONCEPTS

• **Human-centered computing** → *User centered design.*

KEYWORDS

intensive care medicine, clinical decision support systems, visualization, checklist

ACM Reference Format:
Minfan Zhang, Daniel Ehrmann, Mjaye Mazwi, Danny Eytan, Marzyeh Ghassemi, and Fanny Chevalier. 2022. *Get To The Point!* Problem-Based Curated Data Views To Augment Care For Critically Ill Patients . In *CHI Conference on Human Factors in Computing Systems (CHI '22), April 29-May 5, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3491102.3501887

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA

ACM ISBN 978-1-4503-9157-3/22/04…$15.00
https://doi.org/10.1145/3491102.3501887

1 INTRODUCTION

Delivering safe, high-quality medical care in intensive care units (ICUs) is challenging. Patient data—measured or recorded—is often sampled at irregular time intervals, can be redundant, and is prone to contamination by interference and human error [47]. To identify an optimal set of actions, clinicians use their experience to integrate heterogeneous, voluminous, and dispersed data into information that supports decision making. This complex clinical cognitive process must happen efficiently in a context where interruptions are near-constant [28]. The advent of the electronic health record (EHR) has paradoxically added to the cognitive load associated with data integration [14]. While many EHRs incorporate clinical decision support system (CDSS) capabilities to facilitate "meaningful use" [7], embedded CDSSs have had limited success to date.

One potential reason for the limited success of embedded CDSS in EHR systems is that they are poorly integrated into clinical workflows [23, 51]. Increased cognitive load due to usability issues with these CDSSs poses new threats to quality care and patient safety [13, 22], including the introduction of errors caused by fragmented displays and alarm fatigue [41, 51]. Novel strategies are required to ensure that clinician productivity and effectiveness are enhanced, not hindered, by CDSS solutions [6]. An ideal solution would support clinicians' information needs by presenting only relevant data at appropriate times to augment decision making in a manner that does not increase task load, and accounts for complex interruptive and collaborative workflows [28].

Checklists, whose benefits in fostering reproducible medical practice have been documented [52], and data visualization, which is known to support human cognition [39], are compelling potential solutions to the problem of effective presentation of complex information to clinicians. In this work, we identify the characteristics of an ideal solution to inform a checklist- and visualization-based approach to designing CDSS to augment clinical decision making. Clinical decision making is a multi-faceted process of identifying pertinent patient issues in order to propose an appropriate treatment plan. Problem-based approaches are known to provide clinicians with opportunities to use data analysis and metacognitive skills, causal reasoning, systems thinking required for problem-solving in a holistic manner [24, 45]. Additionally, a high-performing patient care team is widely recognized as an essential model for constructing a more patient-centered, coordinated, and effective system of health care delivery [43]. Coordinated team activity in the process of care delivery requires situational awareness to maintain team performance and reduce errors and omissions [40]. Recognizing the importance of these observations, we aim to achieve the following goals: (i) *facilitating problem-based discovery* of relevant patient information, (ii) *reinforcing team communication* around relevant patient problems, and (iii) *improving situational awareness* of patient problems within the busy environment of an ICU.

We propose an approach consisting of **curated data views utilizing clinical problem characterization as a basis for data exploration**. We focus on streamlining information processing and communication by focusing on a checklist of patient problems and providing visual representations of information relevant to the current context within the workflow. We develop different views that support a seamless progression—from an at-a-glance overall assessment of patients' problems for situational awareness and care prioritization, to an ability to support in-depth investigation of a given specific problem in a patient where necessary.

We discuss the design and evaluation of a specific CDSS, *In-Sight*, embodying our generalized approach, which we developed to address the prevalent clinical problem of management of fluids, electrolytes and nutrition (FEN) in critically ill pediatric patients—who experience significant preventable harm as a result of incomplete or inadequate decision making [54]. *In-Sight* served as a technology probe to capture the feasibility and relevance of our approach, while providing a concrete platform to inspire clinicians to think about opportunities to, and implications of changing clinical workflows.

We propose a mixed-methods, multi-stage iterative design research methodology involving clinical experts as collaborators and co-designers, and as external evaluators. Three clinicians part of the research team (co-authors) enabled us to formulate the clinical problem and participated in the co-design of *In-Sight*. Design iterations were also informed by findings from our conducting a formal online usability study of an early prototype with 48 medical staff in the ICU's from four different institutions, and focus group sessions with 12 clinicians, which helped identify gaps in the problem formulation and usability of our prototype. Finally, we performed an ecological simulation using our refined prototype, where 10 intensive care clinicians used the *In-Sight* workflow in a faithful reproduction of the real-world scenario in which the CDSS would be utilized as a part of the multidisciplinary patient rounds.

This paper makes the following contributions: (i) we outline a multi-stage, mixed-methods user-centered approach involving clinicians as core partners in a multi-disciplinary team, to formulate the clinical problem and iterate over CDSS designs, with further insights from evaluations with clinicians external to our research (section 4) (ii) we articulate a set of design goals for CDSS grounded in medical practice and human factors, and propose a general approach to achieve these goals through the combination of checklists and visualization (section 3); (iii) we report on the iterative design of a functional prototype, *In-Sight* (section 5); and (iv) results from a summative ecological study (section 6). Finally, we (v) share the lessons learned and directions for future work (section 7).

2 RELATED WORK

We summarize past work on clinical decision support systems, and the use of checklists and data visualization in healthcare.

2.1 Clinical Decision Support Systems

Clinical decision support systems (CDSSs) are information systems whose goal is to "provide clinicians, patients, or other individuals with knowledge and person-specific information, intelligently filtered or presented at appropriate times to enhance health and healthcare" [31]. CDSSs aid health workers in various aspects of clinical decision making, including diagnosis, prognosis, and treatment decisions. These systems range from simple lists used to solicit advice regarding patient management [12, 20], to systems that organize and display information for patient monitoring [34, 59] and comparison with patient cohorts [19, 57], to sophisticated intelligent systems that recommend the initiation of specific therapeutic interventions [53]; aided by technology including computerized notifications and alerts [33], language processing [26], visualization [34, 61], machine learning [27], and combinations of the above [3, 49].

Our work focuses on supporting diagnosis, articulated as problem lists, in the context of intensive care medicine. Studies describe how incomplete or inaccurate diagnoses contribute to preventable errors [4, 28]. Meta reviews of CDSSs show that these systems optimize decision making regarding treatment (e.g. drug dosing and preventative care), but have performed less well in facilitating diagnosis [51], which is the challenge our work aims to address. Further, while the potential of CDSSs at reducing medical errors and improving patient outcomes has been evident in laboratory studies [18], their adoption in real-life clinical practice is limited [58].

Factors that contribute to this lack of adoption are well documented [18, 30]. One such factor is the poor integration of CDSSs into clinicians' workflow as a result of the mismatch between an over-simplified and idealized conception of work that is linear and localized, and the reality of clinical work, which is interrupted, distributed, interpretative, and collaborative [28]. The healthcare and HCI communities have proposed a number of guidelines and recommendations towards safer, better CDDSs in response to these problems [4, 17, 27, 32, 48, 51, 60]. Our work builds on these guidelines (section 3) and explores the use of clinical problems as checklists combined with data visualization in CDDS.

2.2 Checklists in Medical Settings

A checklist is an organized tool that encompasses a list of action items, tasks, or behaviors arranged in a consistent manner and which serve as cognitive aids to guide users through a set of criteria of consideration for a process [21]. Checklists have become an established tool in medical care to reduce preventable adverse events caused by human error [15], and their widespread adoption has proven effective in improving the quality and safety of care [11].

Checklists are valuable memory aids, which is especially critical in high-intensity fields such as intensive care medicine, where stress, fatigue, and near-constant interruptions can interfere with clinical decision making [11, 21, 52]. Because they are a common reference shared across the multidisciplinary team, checklists act as a catalyst to systematic and standardized care and have been found to foster reproducible medical practice, improve communication, and enhance shared understanding. We leverage the power of checklists, and propose to build lists around patient problems to be systematically checked during daily rounds, further augmented with visualization for problem-based review of patient data.

2.3 Data Visualization in Healthcare

Interactive visualization and visual analytics tools represent a large class of CDDSs. Data visualization is known to be a valuable tool to augment cognition [10], helping people carry out analytical tasks more efficiently by enabling pattern recognition through the concise, structured display of large quantities of information. Since the seminal graphical patient display [38] and Lifelines [35] which were designed to organize and present information of a single patient record to support clinical reasoning, visualization-based CDDSs have evolved to cover a wider range of data types (e.g. vitals, labs, unstructured clinical notes, medical imagery), number of patient records (single vs. cohort), user intents addressed, and analysis support. We refer the reader to prior excellent surveys for a detailed account of the state-of-the-art in this area [39, 56].

Clinical care requires an understanding and evaluation of a patient's history or "clinical trajectory", from past events, to present (and evolving) status, to future prescriptions and prognosis. To facilitate the review and analysis of these multivariate time-based records, a common visualization strategy consists of organizing the clinical data around a timeline [8, 34, 36, 50, 57]. These timelines are often combined into rich interactive dashboard displays that present different facets of the data in synchronized views [16, 17]. Our CDDS, *In-Sight*, falls into this class of systems.

Where our approach differs from prior work is in that we propose to augment an otherwise general dashboard, with a curation mechanism that supports clinicians varying information needs at different stages of their workflow. Our approach is similar in spirit with the notion of semantic zoom visualization found in Midgaard [2] that allows clinicians to reveal the same data at different levels of granularity and abstraction depending on clinical circumstances. We instead propose a curation-based strategy based on the characterization of potential clinical problems in patients, where all important information relevant to the clinicians is highlighted while extraneous data not relevant to the problem at hand is pruned.

3 OUR APPROACH

We are motivated by (a) the need to shift to a human-computer interface paradigm that supports and does not interrupt the clinical workflow [4, 48]. This means that the CDSS is encountered at the point of care, where and when it is needed [32, 60], and slows down decisions making only when necessary [60]. And (b) the need to support and adapt to the physicians' information needs [48]. This adaptation is accomplished by prioritizing and filtering information in a manner that reduces information overload and emphasizes relevant relationships that drive correct clinical inferences [4], and by supporting information sharing and hand-off [28]. Researchers also suggest that clinical decision making should factor in other processes such as situational awareness and problem solving [28].

Drawing on prior work, we articulate the following main goals:

(G1) *Facilitating problem-based discovery*: Making information accessible to clinicians is not sufficient, as clinicians can not afford to spend time retrieving and integrating patient-based information. CDSSs could streamline this process by organizing relevant patient information based on patient problems.

(G2) *Reinforcing team communication around patient problems.* Information gaps often exist between members of the multidisciplinary team involved in decision making. CDDSs that present problem-relevant patient information at the point of care can serve as a support of reference to shared knowledge and shared clinical goals, enabled by a focused discussion of problems and supporting clinical data.

(G3) *Improving situational awareness.* Clinicians need to constantly switch context and re-prioritize in the busy environment of a care unit. CDDSs that allow at-a-glance retrieval of the status of any given patient through ambient, easy-to-process visualization can alleviate the effects of time pressure and interruptions.

We propose a general approach to addressing these goals that consists of curating data views utilizing clinical problem characterization as a basis for data exploration. Curated views, which we discuss in detail below, allow the progression from overall assessment of multiple problems for a given patient, to detailed assessment of individual problems, to in-depth investigation of data relevant to a specific problem. As they proceed with the systematic review of patients' problems, the clinical team records their assessment, by indicating whether their assessment aligns with that of the system or not, from any of the curated views as follows:

The situational-awareness view is the most abstracted view. It displays a standardized checklist of potential problems that clinicians set out to monitor. Items are color-coded to reflect automated detection of the problem and processing status (e.g. if the team has reviewed and acknowledged the problem). This passive view should strive to be simple and readable from a distance, for quick identification of patients' condition, to allow care prioritization and easy context re-acquisition, without explicitly engaging with CDSS.

The at-a-glance view is brought up when the clinical team focuses on a particular patient. This stage assumes active engagement with the CDSS at the bedside. The view presents the same color-coded problem checklist as the problem overview, together with an interactive dashboard display of aggregated patient data providing a detailed holistic view of the patient condition. The dashboard component of this view acts as a traditional general-purpose exploratory data analysis CDDS, in that it should allow the clinicians to perform freeform exploration of patient data through mechanisms such as dynamic queries [46], brushing and linking [9], details on demand, or more advanced visual analytics.

The problem-focused view is triggered when the care team decides to dive deeper into a specific problem from the checklist. Only the charts from the comprehensive dashboard that are relevant to the problem are kept and expanded for improved readability, and further augmented with annotations revealing aspects of the data (e.g. irregularities, exceeding values) explaining why the system identified the presence of the problem.

4 PROBLEM CONTEXT AND METHODOLOGY

Our approach is grounded in co-design and evaluation with domain experts of a CDSS, *In-Sight* (section 5), aimed at supporting the diagnosis of fluid and nutritional anomalies in critically ill pediatric patients. We describe the clinical problem and our methodology.

4.1 The Clinical Problem

Because clinical medicine is practiced by focus area, it is impractical to re-design the entire EHR towards our CDSS goals. Instead, we develop a prototype system to test our novel CDSS design ideas through a technology probe in pediatric nutrition. We selected a prototypical critical care problem to utilize as proof of concept. Management of fluids, electrolytes, and nutrition (FEN) is a daily set of decisions that need to be made for *every* critically ill patient. Information that drives correct decision making relies on the integration of raw data from a variety of domains of the EHR, including nursing observations (such as tabulated intake and output totals, weight, feeding tolerance), select laboratory results (e.g. electrolyte concentrations in the serum, BUN, creatinine) and specific medications (e.g diuretics) that influence patient response. Challenges navigating the EHR to extract these data elements are compounded by the complexity of the interdependent relationships between these data points of interest that are not identified or displayed in the EHR, requiring additional clinician synthesis to identify them [55].

As a result of these challenges, there is significant variability in clinician performance in identifying problems in this domain that is known to be associated with patient harm—an issue that is exacerbated in pediatric populations [54]. Incomplete aggregation of data or incorrect problem formulation leads to inattention to well-known problems like fluid overload [29], fluid creep [42], malnutrition as a result of excessive or inadequate feed administration relative to patient requirements [37] and evolving kidney injury [5]. All of these well-known problems are associated with increased morbidity, length of stay, and mortality in critically ill patients. Importantly, much of this iatrogenic harm could be mitigated by better clinician awareness of relevant patient problems.

We identified this clinical problem as a priority to attempt to address due to (i) its prevalence, (ii) its relevance to every critically ill patient every day, (iii) well-documented harm associated with incorrect or incomplete clinician inferences, and (iv) the complexity and interdependence of data elements that drive correct inference.

4.2 The Task and Context

The clinical problem is addressed for each patient during daily clinical rounds in critical care environments. Rounds are intended to foster group decision making and consolidate team practice [44]. Pertinent information about the patient condition and interval events since the last rounds are reviewed and discussed. The team comes to a consensus patient assessment, which is problem-based before articulating a 'plan' that is intended to address any problems is identified [25]. This can be a cognitively challenging exercise and requires effective integration of large volumes of data in order to arrive at the correct problem formulation [1]. This exercise typically has to be accomplished very efficiently as the rounding team has other patients to assess.

In the traditional rounding format (the current gold standard), an accredited clinician is responsible for reviewing and curating data before rounds and typically records this curated data on a piece of paper that is presented to the rest of the clinical team as a means of helping to identify problems as a team (Figure 5A). To put this into perspective, we gather from a step-by-step internal audit of the workflow using the current EHR solution that 14 steps, 29 different screens, 43 clicks were required to find the predefined data elements necessary for rounds for a single patient, which took 7 minutes to complete.

Table 1: Participants of the Focus Groups*

Focus group #1	Nurse Practitioner (1), Staff (2), Fellow (1)
Focus group #2	Staff (3), Fellow (1)
Focus group #3	Nutritionist (4)

(*) All from the Hospital for Sick Children, Toronto

We attempt to support this task in the context of rounds by developing a CDSS that facilitates accurate, reproducible, and efficient team inference of the relevant problems in the FEN domain for any given critically ill pediatric patient.

4.3 Research Methodology

We employed an agile, multi-stage, mixed-methods approach. A six month-long set of bi-weekly *participatory design sessions* with our multidisciplinary team of 3 computer scientists and 3 clinical domain experts from the Hospital for Sick Children (Canada) at the project outset allowed precise problem and task formulation. Clinical domain experts helped us characterize:

- The key problems that clinical teams ideally need to identify in the FEN domain;
- Information needed to support or refute the existence of the problem(s);
- Data that needs to be aggregated and relationships that need to be highlighted to provide this information;
- The context and constraints of the clinical environment that needed to be considered modifiers in the design.

We used an *iterative multidisciplinary co-design process* where our team of clinicians and computer scientists conducted weekly development sprints and assessments to develop a prototype. Our evolving prototype served as a *technology probe* throughout the process. Three *focus groups* that each involved 4 clinicians external to the project (Table 1) navigating the prototype around defined tasks allowed to solicit further feedback and make refinements.

This allowed us to consolidate an initial prototype of our design that underwent a *formal online usability study* with medical staff in the critical care unit (residents, fellows, staff clinicians, and allied health professionals who were not involved in our research nor focus groups). The 48 participants were drawn from four separate institutions in the USA, Canada, and Israel, and used 3 different EHR solutions (EPIC, Cerner, and Prometheus). This represents a larger sample of critical care practitioners who address these problems in daily rounds. This diversity was important in ensuring that the approach to visualization generalizes beyond a single institution.

We incorporated insights gained through this multi-institutional study—including the introduction of the problem-based checklist to drive data view curation—in a refined mature prototype co-designed with clinicians from our research team during one new round of iterative design. Finally, to gain a more holistic insight into the role of CDSSs embodying our approach in clinical practice, we performed a final assessment in a more *ecologically valid rounding simulation study* using *In-Sight*, where 2 groups of 5 medical staff each (who did not participate in earlier phases of the project), discussed a mock patient as part of their real daily round.

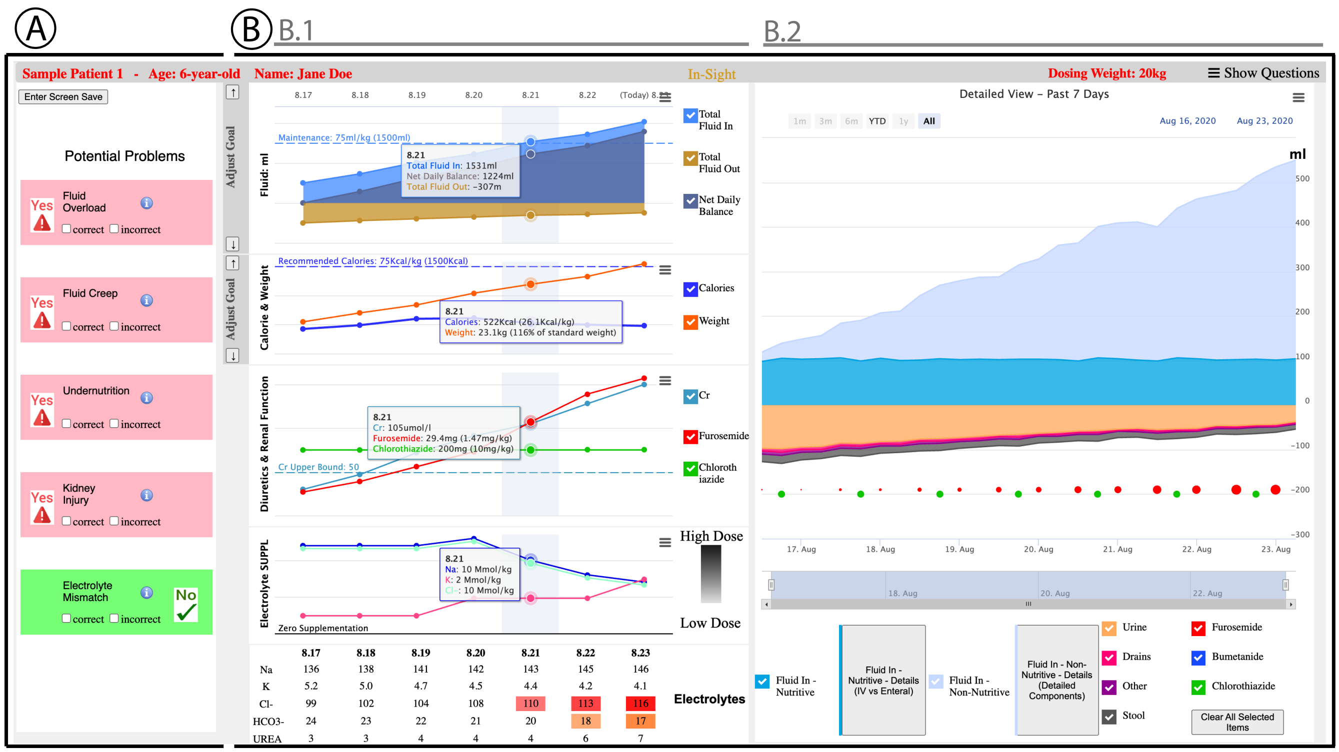

Figure 1: Patient view of In-Sight, comprises of (A) a color-coded checklist of potential problems in a patient, and (B) an interactive dashboard presenting patient integrated data via an (B.1) overview panel displaying five vital aspects of patient FEN status, and (B.2) a detail panel with additional patient data such as sources of fluid in and fluid out and diuretics.

5 IN-SIGHT: DESIGN & WALKTHROUGH

We designed *In-Sight*, a CDSS embodying our approach as a JavaScript web application using the Highcharts library (see https://insight-demo.herokuapp.com/). We first discuss the two core ingredients of *In-Sight*: the dashboard display, and the checklist of problems. We then illustrate their integration into our full-fledged CDSS through a walkthrough of a use case scenario.

5.1 Dashboard Display

The interactive dashboard of *In-Sight* was initially designed independently from the problem checklist, as a traditional visualization-based CDDS integrating patient data to support FEN inferences.

5.1.1 Task Analysis. To better understand tasks and workflows, we used a hierarchical approach whereby the clinician collaborators (i.e. co-authors) outlined the most important and high-level analytical questions and detailed sub-questions. The clinicians also illustrated a step-by-step audit of their current workflow to characterize how data is typically consumed. We identify the following assessment tasks that clinicians perform in our context: (T1) fluid balance and composition, (T2) nutrition received and composition, (T3) growth measures, and (T4) balancing measures.

5.1.2 Design. The *In-Sight* dashboard display (Figure 1B) encompasses an overview panel (Figure 1B.1) and a details panel (Figure 1B.2). The overview panel (B.1) contains views corresponding to the high-level tasks T1-T4, purposefully organized to enable the identification of relationships and context through cross-chart analysis. Seven-day trends are displayed for five vital aspects of patient FEN status, with details for a day revealed in a pop-up when hovering over the charts. Flags (only revealed in the problem-based view in our refined prototype) indicate when a data element is above or below 10% of the desired goal (which can be dynamically adjusted by the care team), or a predefined goal from the medical literature. Our early prototype also included a textual list of the patient's potential problems based on abnormal data at the top of the screen (not shown; see supplemental). The early dashboard was successful as a technology probe, as insights from focus groups and usability evaluation suggest that this information was not prominent enough and could be better leveraged, which inspired using checklists as the core of our problem-based curation approach.

The details panel (B.2) contains elements that were lower in the task hierarchy but nevertheless important for clinical decision making and currently exceedingly difficult to determine in the EHR. Users can adjust the timeline to display fluid details ranging from daily for the past 7 days to every 6 hours for the past 24 hours. All sources of fluid in and fluid out are displayed and can be categorized by nutritive or non-nutritive fluids. All commonly used diuretics are included in the details graph. Hovering over any data element provides details in a pop-up and data elements can be dynamically added/removed directly from the interactive legend.

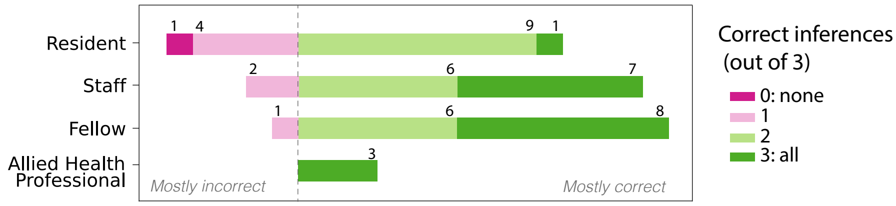

Figure 2: Performance correctly identifying patient problems per participant role.

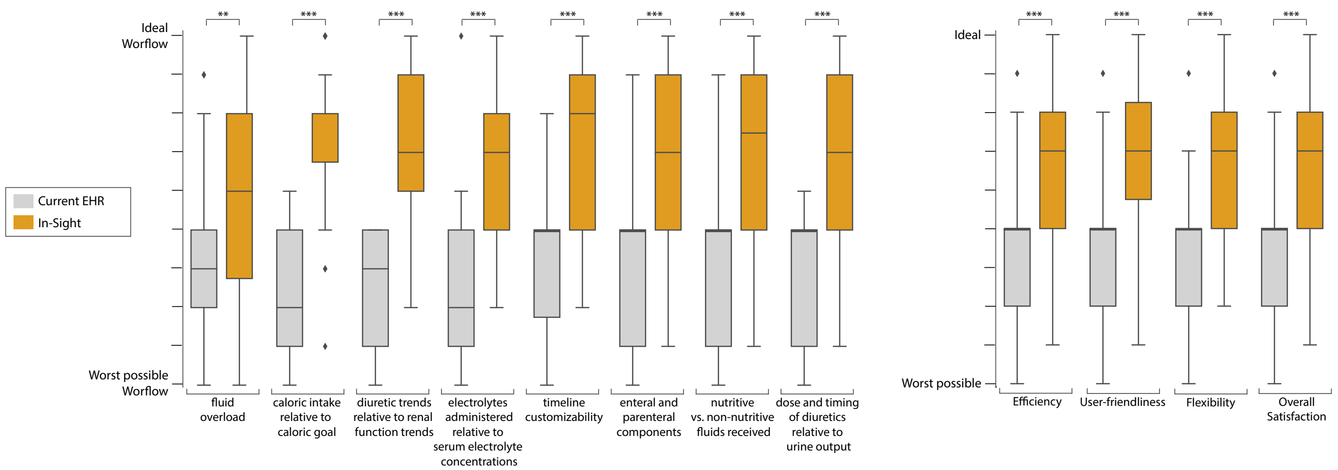

Figure 3: Summary of participants ratings of *In-Sight* and current EHR. Wicoxon rank sum and Kruskall-Wallis tests evaluated for differences between groups (: p < .001, ***: p < .0001).**

5.1.3 Validation. We performed an asynchronous online evaluation of the interactive dashboard component of *In-Sight* with 48 clinical users not involved in the project nor focus groups from four different centers: Hospital for Sick Children (Canada), Children's Hospital of Eastern Ontario (Canada), Boston Children's Hospital (USA), and Rambam Medical Center (Israel). Participants were residents (n=15), fellows (15), staff clinicians (15), and allied health professionals (3), aged 26-58 (median: 36.5) whose experience ranged from 0.5 to 25 years; less than 2 years (n = 14), between 2-4 years (19), between 5-9 years (4), and 10 years and more (11).

We sought to determine whether our dashboard facilitates more accurate and efficient inferences than the EHR in use at these centers (i.e., EPIC (n=30), Cerner (9), Prometheus (8)), and identify usability issues. Participants were presented data from a sample case of a 6-year old hypothetical patient with a number of intended problems. Users were asked to review the patient data using our dashboard, and then complete a series of questions about the patient condition specifically and CDDSs more generally.

Figure 2 shows a breakdown of participants' performance in correctly identifying patient problems. We noted some between-group differences in inference acquisition performance and ratings of the visualization across clinician roles. Residents performed worse (5 out of 15 identified only one or none of the three problems) than fellows (2/15) and staff clinicians (1/15) in the problem formulation task. This group of users has the least clinical experience of those participating in the study. Residents also rated the visualization lower overall than users in other groups.

Figure 3 shows a summary of questionnaire responses. *In-Sight* was rated as significantly superior to the current EHR in every task domain (average of around 2 points higher by all 48 clinicians in a 10-point Likert scale). Ratings were not statistically different between roles or EHR proficiency level for any task. Participants rated *In-Sight* significantly higher than their current EHR in domains of efficiency, user-friendliness, flexibility, and overall satisfaction (average of around 2 points higher).

5.2 Checklist of Problems

5.2.1 Defining the Checklist Items. The problems that we chose to display were prioritized by their prevalence and harm in critically ill pediatric patient populations. For each of the problems prioritized, we created a problem definition supported by either existing medical literature or specific local unit practice as a means of generating rules that permitted automated problem identification in the CDSS. This problem list was reviewed by clinical dietetics teams (dieticians, physicians, and quality champions not involved in the project) to ensure that the problem definition met expert consensus. These definitions are discoverable for clinician users via an infobutton associated with each problem. The problems are not mutually exclusive and can occur asynchronously and repeatedly in a patient's stay in the ICU and are therefore amenable to checklists.

Figure 4: Screenshots of *In-Sight*, at different stages of the workflow. (A, F): the situational awareness view, before the review of the clinical team (A), and after (F). (B): the at-a-glance view, showing a condensed version of the checklist, along with an interactive dashboard of patient data. Clicking on a checklist item allows the clinical team to curate problem-focused views (C, D, E) for closer inspection.

5.2.2 Visual Design. The problem-relevant clinical knowledge is built in the CDDS using computational methods. In the case of *In-Sight*, this knowledge translates to the detection of abnormal data compared to user-defined or clinically-informed data ranges and thresholds. The checklist of problems is color-coded based on the automated identification of potential problems: a red color and "Yes" warning sign indicates that the system flags the presence of a problem, a green color and a "No" checkmark is shown otherwise (Figure 4A). Clinicians manually indicate whether their assessment aligns with that of the system, by confirming that the system is correct or incorrect. The checklist immediately updates to reflect the clinicians acknowledged and reviewed the problem (the background is turned light gray) and to convey assessment by the clinical team: a red stripe on the left side, or a green stripe on the right side, indicating the presence or absence of the problem respectively, according to the clinical team. A trace of the system's recommendation is maintained (i.e. the warning flag and checkmark icons remain visible); see Figure 4F.

5.2.3 Integration With the Dashboard. The problem checklist is presented full-screen, as the situational-awareness view in *In-Sight* (Figure 4A), and compressed on the side of the screen when entering the at-a-glance view (Figure 4B, Figure 1). Clicking on an item of the checklist triggers the problem-focused view for the associated problem, where charts relevant to the problem are curated from the dashboard and magnified for increased readability, and flags are displayed to draw attention to the system-identified abnormal data (e.g. Figure 4C-E).

5.3 Walkthrough of a Use Case Scenario

Tom, a clinician at a pediatric intensive care unit (PICU) is in charge of leading the daily walk rounds involving four other peer healthcare workers. As part of his preparation prior to rounds, Tom has reviewed the data of each patient in the PICU in detail.

5.3.1 Assessment of the Situation of the PICU. As the clinical team walks into the PICU, team members can immediately form an overall idea of each patient's potential problems, by looking at the **situational-awareness view** displayed on the screen at the side of each bedspace (Figure 4A). While his colleagues view this information for the first time today, Tom quickly re-appraises the patients' condition, looking for discrepancies since he last checked. This supports decisions about how to prioritize rounding order (i.e. which patients to evaluate first).

5.3.2 At-A-Glance Assessment of a Given Patient. As he reaches the first patient's bed, Tom interacts with *In-Sight* to switch to the **at-a-glance view**, revealing the interactive dashboard of the patient's integrated data over the past 7 days (Figure 4B), along with a condensed version of the problem checklist. This overview supports a high level review of data pertinent to all of the problems identified. Importantly, this allows rapid high-level verification that some potential problems do not exist (e.g. "Electrolyte Mismatch" in this patient), i.e. are not flagged by the CDSS, not slowing workflow and allowing particular focus on those that are flagged. Tom would like to review the patient's fluid balance which *In-Sight* labeled as being potentially problematic (i.e. "Fluid Overload"), and if the composition of the fluids is appropriate (i.e. "Fluid Creep"). He also wants to understand the patient's nutritional trajectory (i.e. the intake of calories and protein, along with the weight of the patient.) There are other problems of interest to the care team: *In-Sight* also indicates that the patient may suffer from a kidney injury. Tom prompts the team, suggesting that they further investigate the specific potential problem of fluid overload, pointing at the checklist in *In-Sight*.

5.3.3 Problem-Focused Investigation. Tom clicks on the corresponding item in the checklist, which immediately adapts the view to the **problem-focused view** containing data elements that support investigation of "Fluid Overload" (Figure 4C): the top-most chart from the dashboard overview panel is brought into focus, and warning signs are added to the chart, indicating values of the patient's fluid balance (blue portion of the area chart) that exceed desired fluid maintenance threshold (horizontal blue line). This allows Tom to both focus more narrowly on data supporting this specific problem and intentionally draws his attention to the individual data elements that have driven the CDSS to flag that problem according to the built-in clinical rules.

5.3.4 Verification of the Problem. After reviewing this specific problem, Tom and the rounding team agree that the general rule that identifies this problem is applicable to the context of this particular patient. They acknowledge the problem as being 'correct' by clicking the acknowledgment box, which is visually reflected in the system (see top checklist item in Figure 4C,E). This both validates the problem identified by the CDSS and acts as a process and quality measure that permits verification that the rounding team has assessed the problem. In depth assessment of the data may result in scenarios where the rule is not felt to apply to the patient as a result of special medical circumstances, which are not infrequent in critically ill patients; team can then reject the problem.

5.3.5 Completing the Patient Assessment. Tom and the rounding team continue their review of the patient, choosing to focus next on an in depth review of the problem of "Fluid Creep" because they want to identify what sources of fluid intake contribute to the fluid overload they just identified (Figure 4D). This allows them to identify what sources of fluid intake they need to modify to mitigate this problem. Throughout this process they can move from problem specific views back to the at-a-glance view as needed to support their information needs. Problems can be adaptively explored in whatever sequence is most appropriate to the context of the patient for verification and identification of important dependencies.

5.3.6 Communicating the Patient Assessment. After each potential problem has been reviewed and acknowledged (Figure 4E), the rounding team's assessment is complete and the view can be turned back to the situation-awareness view, so that the verified patient problem list is displayed to any other clinicians that interact with the patient in order to ensure consistent information transfer during transitions in care team personnel (Figure 4F).

6 SIMULATION EVALUATION

We conducted an ecological simulation study to better understand the potential and limitations of our CDSS prototype when used at the point of care by intended users within the format of daily bedside patient rounds.

6.1 Study Protocol

We designed a study protocol to facilitate comparisons and contrasts between the traditional rounding format (the current gold standard) and the proposed workflow supported by our CDSS. To that end, two separate rounding groups of clinicians from the Hospital for Sick Children performed simulated rounds in the critical care unit during actual patient rounds. There were 10 participants who were not involved in any other stages of our research (two rounding teams of 5 individuals) that had the same skill mix; a senior staff physician, a clinical fellow physician, the charge nurse, the bedside nurse, and a respiratory therapist. See Table 2.

Table 2: Participants of the Simulation Evaluation*. The two rounding teams included P1-P5 and P6-P10 respectively.

Participant	Role	Experience (years)
P1	Senior staff physician	22
P2	Clinical fellow physician	8
P3	Charge nurse	30
P4	Bedside nurse	5
P5	Respiratory therapist	11
P6	Senior staff physician	8
P7	Clinical fellow physician	6
P8	Charge nurse	12
P9	Bedside nurse	7
P10	Respiratory therapist	27

(*) All from the Hospital for Sick Children (Canada)

Each group performed traditional team rounds on an actual patient using the usual workflow and then simulated team rounds on a synthetic patient in *In-Sight*. The mock patient was the same as used in the online usability evaluation (subsubsection 5.1.3)—which the majority of participants had rated as realistic (n=25) or very realistic (5); and the real patient used in the study was selected by our clinical expert collaborators in such a manner that both patients were of similar medical complexity. The simulation study was conducted at the bedside in the critical care unit during actual patient rounds. The team involved was therefore actively identifying problems at other bed spaces for real patients and enacting treatment plans. This ensured the validity of our ecological simulation.

In both rounding formats, the teams were encouraged to interact with patient data, ask clarifying questions and make observations as indicated to support their information needs. At the conclusion of the evaluation participants in both rounding teams each filled in an anonymous survey about their experience (see supplemental material). As a post-study survey, we performed a thematic analysis of the written open comments and pulled relevant quotes from these surveys, with the survey forms and responses attached in the supplemental materials. The sessions were facilitated and documented by one of the co-authors, but could not be video recorded due to patient privacy concerns. Figure 5 shows pictures taken during the study, using the traditional workflow (A), and with *In-Sight* (B).

6.2 Results

We asked participants to evaluate the ecological validity of our mock patient rounding scenarios by rating their realisticity on a four-point scale (not at all, somewhat, realistic, very realistic). All participants rated either realistic (n=8) or very realistic (n=2).

Overall, all participants (n=10) indicated that they preferred the *In-Sight* format over the traditional format, commenting: *"I like that the monitors can be used for more"* (P8), although this should not be at the cost of adding one more tool to an already overwhelming collection of computerized systems. The senior-most participants (25+ years of experience) pointed: *"We need to be careful of getting too lost in the screens"* (P3), *"Many screens! EPIC, T3, In-Sight. Just not sure"* (P10). We organize the rest of our discussion of results around the goals we set out to achieve with our approach (section 3).

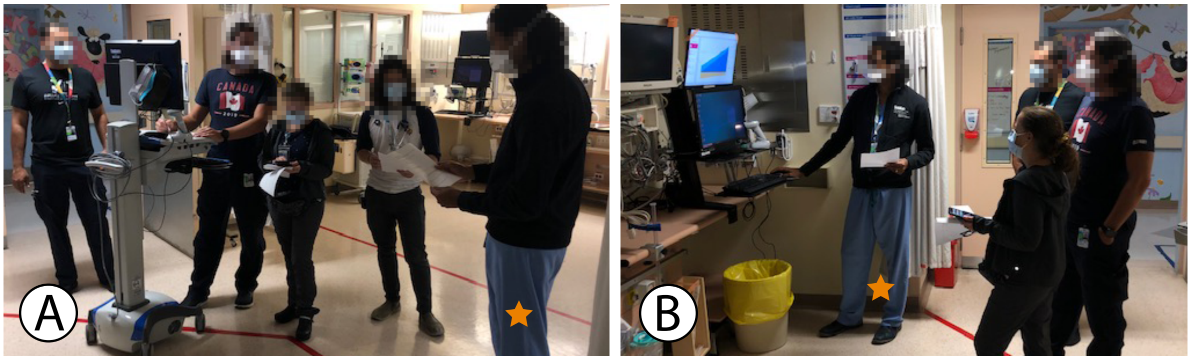

Figure 5: Photographs taken during our ecological study. (A) Traditional rounding format. (B) Using *In-Sight*. The accredited clinician in charge of the rounds is indicated with a star symbol.

6.2.1 Presentation Format of In-Sight Facilitates Problem-Based Discovery (G1). Overall, participants found that the problem communication with *In-Sight* was more efficient than a verbal articulation of problems as typically occurs in traditional rounds. Participants described the CDSS problem characterization as being action-oriented: *"the problem statement makes me want to do something about it."* They felt that problems were clear and precise, allowing *"to quickly get a picture of patients problem list"* (P9) through a *"nice summary of the patient's issues that put them into a clearer context than the traditional speaking rounds"* (P3). This suggests that the combination of the checklist and interactive problem-based curation of views can facilitate the identification of problems and access to the relevant patient data context during the rounds.

Participants found great value in the approach, but some were concerned about adoption, pointing to challenges associated with the feasibility and scalability of the approach: two participants said that they would prefer to use *In-Sight* in a team-based rounding setting *"if it covered more problems"* (P8) and *"if it covered more domains"* (P6). Similarly, P2 would keep with the traditional approach, because *"it covers all domains"*. P3 pointed to the list not being self-sufficient, commenting that *"the bedside clinicians will still need to accurately summarize the patients and formulate their own clear problem lists so that complex patients who don't fit the standard mold can still be readily identified."*

6.2.2 In-Sight Promotes Physical and Mental Team Convergence (G2). Participants generally preferred *In-Sight* as a facilitator of team communication, and we observed our CDSS was associated with a change in team rounding behavior. Team positioning at the bedside in the traditional rounding format is typically dispersed (Figure 5A) with variable participant attention. In contrast, we observed that the *In-Sight* rounding format encouraged team aggregation and joined attention to patient data as a result of the facilitated navigation of relevant patient information (Figure 5B).

To the question of whether they preferred traditional rounds or *In-Sight* facilitated rounds responses were: 7 participants preferred *In-Sight*, 2 preferred a combination of traditional rounds, and 1 traditional rounds. Participants qualified *In-Sight* as more *"collaborative"* (P1), as it helps *"keep attention"* (P4), and makes it *"easy to congregate and look at data together"* (P6). Participants who mentioned they would prefer to keep aspects of the traditional verbal rounds did so because of the limitations of the checklist as opposed to concerns about its potential in supporting communication.

We also noted a change in the qualitative nature of the dialogue between the team members as a result of using *In-Sight*: in traditional rounds, the majority of clarifying questions or comments focused on data values (typically not presented, or only verbally communicated), e.g. *"What were the values of that variable yesterday?"* In contrast, most interventions when using *In-Sight* involved appraisal or interpretation of data (that was explicitly displayed in the CDSS), e.g. *"I would adjust that caloric target for this patient"*, suggesting that instead of processing information from memory, the team was able to focus the discussion on clinical reasoning while relying on the externalized shared knowledge materialized by the list of problems, and further expanded by pulling relevant patient data on demand through drilling down into the dashboard.

These results suggest that *In-Sight* has the potential to improve collaboration between clinicians through enhanced, more focused communication supported by the possibility to easily access and refer to relevant patient data when assessing the patient's problems.

6.2.3 Persistently Visible Problem Lists Anchor Instantaneous Situational Awareness (G3). There was complete consensus in participant responses (n=10) that *In-Sight* can support situational awareness with its views better than the format of the traditional rounds that does not have a comparable mechanism. Our study did not allow us to quantify to which extent this is the case when rounding multiple patients using *In-Sight*, or recovering from interruption. Nonetheless, survey comments are unequivocally pointing to the lack of support with the traditional workflow: *"there is zero awareness with traditional rounds"* (P1). Participants observed that currently, situational awareness relies on either verbal communication between clinicians to understand patient problems or review of rich text notes in the electronic health record (EHR). In contrast, participants appreciated the ability to obtain an at-a-glance identification of patient problems they can come back to re-appraise anytime with our CDSS: *"In-Sight is clear and can be checked back on"* (P2).

Importantly, one participant pointed out that a potential drawback to the persistent display of patient problems at the bedside using *In-Sight* may pose risks to patient confidentiality that would need to be addressed by a careful review of how best to integrate this functionality of the CDSS into clinical workflow.

7 DISCUSSION

Researchers have stressed the need to design CDSS "not only as a functional utility but as an integrated experience", calling for a methodology that more holistically factors in social and physical contexts [60]. Our iterative internal discussions among our multidisciplinary team, multi-center usability evaluation of our initial dashboard, and ecological study at the point of care, together have allowed us to perform a robust assessment of current workflows.

Results from our simulation study suggest that our approach of curating data views utilizing clinical problem characterization as a basis for data exploration has a number of merits regarding our articulated design goals. Clinical experts found the presentation format of *In-Sight* facilitates problem-based discovery (G1); we observed that In-Sight promotes both physical and mental team convergence by facilitating team communication, which was echoed by our participants' comments on improved focused attention and collaboration using our approach (G2); experts also noted that the clarity and visibility of the checklist supports situational awareness, though some improvements are required to preserve patient confidentiality (G3).

In-Sight was successful as a technology probe to further our understanding of the implications of integrating a CDSS that goes beyond presenting aggregated data visually for externalized cognition. Below we discuss important lessons learned from our process and implications on the design of CDSS.

7.1 Designing For User Relevance – Beyond "Yet Another Dashboard"

As noted previously, we found between-group differences in inference acquisition performance and ratings of the visualization by clinician roles (section 5.1.3). The reasons for this are unclear. We hypothesize that the problem formulation task may have been more challenging for the least experienced users and that difficulty and cognitive loading manifested as lower overall ratings. Meanwhile, domain experts in our focus groups stressed the importance of making patient problem identification easier and more direct. These observations influenced how the design evolved.

Reflecting on the discrepancies between different users using our dashboard, we re-centered our focus on the preeminent goal of clinicians, which boils down to a simple question: "what are the patients' problems?" Problem identification is an essential prerequisite for identifying what management decisions and treatment strategies may improve the patient condition. It is intuitive to think that making all of the patient information available to the clinicians in a carefully designed interactive dashboard can support specific analytical tasks such as T1-T4 (subsubsection 5.1.1). And indeed, our results suggest an improvement over current EHRs. Yet, we were taken aback by the fact that it was not enough to support effective problem formulation: interpreting an information display to support problem formulation from a general, multi-purpose dashboard is cognitively demanding and requires expertise. Our study suggests that for best results, technology must go the extra mile in perfectly aligning with clinicians' information needs, by providing shortcuts to curated data views that present only the most important information relevant to the clinical context, while simultaneously pruning extraneous data.

7.2 Standardizing as Competency Scaffolding, Not Process Fossilization

While they are great tools as a mnemonic aid supporting consistency and replicability in care, there are dangers associated with the use of checklists which are well discussed in the literature [11, 21, 52]. One such concern is the prescriptive and micro-managing nature of checklists, which skilled workers see as a direct threat to their competency (e.g. "don't you trust I can think about these myself? It's my job!"). Meanwhile, the reassurance and comfort that checklists imbue compared to having to rely on memory can also yield some to over-rely on them, making it more likely to miss important problems not captured by the list. Participants of our ecological study stressed the need to build more comprehensive problem lists. Building the optimal exhaustive list is utopian, and we reiterate P3's point on the importance of not risking that problems of *"a patient that doesn't fit the standard mold"* pass undetected because of over-reliance on the list.

It is important to note that our goal is not to standardize the process in such a way to suggest a prescriptive, inflexible instruction guide to follow to the dot. Rather, CDDS should strive to act as a practical tool to enable effective competency scaffolding through a shared reference. We made a conscious design choice of not displaying any annotations pointing to anomalous data values in the at-a-glance view of *In-Sight*. This choice was partly to mitigate the information overload and alert fatigue, but was mostly decided in an attempt to prevent priming from computerized annotations. We posit that this view should give clinical reasoning agency back to the clinicians, by fostering exploratory, freeform exploration of the patient data, with minimal guidance from the computer (or with explicitly solicited guidance)—very much like general visual analytics tools do—so as to increase the chances that the experienced clinician captures a singular problem. Now, problems that are easily detectable and well understood should still be fully supported with custom curated data views that highlight that problem, for clinicians to access in a matter of seconds where the general dashboard is too distracting.

Future works should explore how to address this tension between supporting freeform, unguided exploration and providing the best possible custom views for textbook problems.

7.3 That's So Yesterday – Keeping Lists Current

Medicine evolves, and practices change. So should the checklists. CDSS implementing our approach needs to ensure that the checklists are maintained up to date, to reflect the latest advances in medical science. As has been noted in the literature [11, 21, 52], maintaining checklists is challenging in practice, as any update proposal should undergo a complete process of clinical trials to validate their relevance, usability, and impact on patient outcomes. The checklist we developed in this work is a pilot prototype that we do not claim is optimal, nor definitive. The design of a checklist should involve all of the stakeholders. There are opportunities for research to determine how to build general enough lists to cover enough problems, while being specialized enough to be of high relevance to its users, while enabling customization and frequent maintenance of the associated rules.

7.4 Not That Smart, Is It? – Checklist As Partner

Patient conditions, especially in the ICU, can evolve rapidly and serious problems may arise at any time. Meanwhile, many clinical problems are latent ones, for which any noticeable evolution is expected to take several days, or even weeks. This raises the question of when it is best for the computer to (re-)notify a problem in a patient to the care team. On the one hand, there is little value in forcing care workers to re-appraise the same problem at every round, if it is acknowledged as following its course as expected. On the other hand, it is important that the system notifies a change when it detects one, while not overwhelming the clinicians with too many alerts and alarms.

We see an opportunity to expand the capability of the problem checklists to a point where the CDSS is capable of updating the clinical team like an informed peer clinician would. There is enormous potential for future research to integrate more advanced mechanisms to enhance the communication strategy of the CDSS, such as gauging whether the clinical team should be absolutely notified of a change now, or is it more appropriate for it to be caught up on the updated status at the first opportunity when they are physically and mentally available to consume this information. Further, computational methods that leverage clinical knowledge and user-defined input on expected prognosis and projections can help make the system more cognizant of whether the response to a treatment follows course or if something worsens, instead of relying on a strictly localized view of status at a a given instant.

In future work, we also expect that checklists may be able to learn improved rules on the fly, just as clinicians build up expertise through experience. Systems that automatically re-check and verify that their rules are in line with best practices are likely to help with limiting the damage of incorrect past practices.

7.5 Scaling From Ponds to Lakes

Our work tackles a small pond of clinical problem area—managing of fluids, electrolytes, and nutrition in critically ill pediatric patients. As brought up by our participants of the ecological study, resistance to the adoption of a tool like *In-Sight* is to be expected if it does not scale beyond the niche problem and context it is tailored for. We note that the implementation of our approach requires that there is further analysis done to ensure that we can successfully scale the approach when moving to other locations, populations, comorbidity distributions, and available technologies. Now, it is not practical nor desirable for any new unit to start from scratch and build yet another embodiment of our approach that suits their needs. Ideally, a general CDSS building on our approach should be modular and flexible enough to enable seamless integration with existing technology, while allowing customization to specific domain needs and context. We see an opportunity for future research exploring ways of abstracting checklist substrates, where clinical knowledge, care worker preferences, and rules to curate, visualize and annotate data would be easy to define, refine, and redefine by any unit.

7.6 I Know The Problems. Now What?

Our approach focuses on facilitating diagnosis in patients, through the identification and labeling of existing problems. This is only one step in the pipeline: problems need to be acted upon or investigated further; information handoff needs to happen during shifts; treatments, orders, and prognostics need to be recorded; and more. There is an enormous opportunity to address many other parts of the clinical pipeline, e.g., action modeling, input of annotation, justification of decisions, prognostic simulation, etc. How and when to integrate systems like *In-Sight* into clinical practice depends on further evaluation of best practices in deployed visualization and computing systems.

8 CONCLUSION AND FUTURE WORK

There still exist many barriers to designing, evaluating, and deploying CDSSs that effectively and efficiently contribute to enhanced quality and safety of medical care. In this work, our multidisciplinary team of computer scientists and clinicians focuses on solutions that have the potential to address the current limitations of CDSSs in supporting problem-based discovery, team communication, and situational awareness. We propose an approach that combines the power of visual dashboards and problem checklists to effective presentation of relevant patient-based data.

In our work, we designed an entire system from scratch as there is currently no CDSS that aggregates information cohesively for the FEN problem in critically ill pediatric patients. We focused on the FEN visual dashboard, defined the system, prototyped it in the pediatric nutrition setting, discussed limitations and opportunities in focus groups with clinical experts, evaluated an initial prototype with 48 clinical staff in a test study that showed a consistent preference over traditional EHR, and finally performed an ecological simulation of a refined prototype with 10 clinicians.

In future work, we plan to explore augmenting generic CDDS with problem-based curation capabilities, and evaluating machine learning to assist in the process of curating or generating relevant views. Other future work should investigate information obsolescence, by exploring when it is necessary and relevant to refresh the list of problems, how to account for clinicians' knowledge regarding the necessary time before a problem evolves positively, and when data from the latest rounds is no longer up-to-date. These future works should also strive to include domain experts from an extended set of institutions, to assess generalizability of results in clinical centers where the infrastructure and means may differ from that of the centers who we recruited participants from.

We note our approach could be applied to existing visual analytics tools designed to support clinical decision making. By curating data views based on problem characterization, we reduce the cognitive burden associated with presenting too little or too much information at once. There is strong potential in our approach for improving situational awareness and focusing discussion around patient problems in many clinical areas, which should ultimately be formally evaluated through prospective clinical trials focused on care-related and patient outcomes.

ACKNOWLEDGMENTS

We wish to thank the members of the DGP lab for their valuable comments on this work, and all of the participants in our studies for their time and input.

REFERENCES

[1] Shiri Assis-Hassid, Barbara J Grosz, Eyal Zimlichman, Ronen Rozenblum, and David W Bates. 2019. Assessing EHR use during hospital morning rounds: A multi-faceted study. *PloS one* 14, 2 (2019), e0212816. https://doi.org/10.1371/journal.pone.0212816

[2] Ragnar Bade, Stefan Schlechtweg, and Silvia Miksch. 2004. Connecting time-oriented data and information to a coherent interactive visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 105–112. https://doi.org/doi/10.1145/985692.985706

[3] Jason M Baron, Richard Huang, Dustin McEvoy, and Anand S Dighe. 2021. Use of machine learning to predict clinical decision support compliance, reduce alert burden, and evaluate duplicate laboratory test ordering alerts. *JAMIA open* 4, 1 (2021), ooab006. https://doi.org/10.1093/jamiaopen/ooab006

[4] David W Bates, Gilad J Kuperman, Samuel Wang, Tejal Gandhi, Anne Kittler, Lynn Volk, Cynthia Spurr, Ramin Khorasani, Milenko Tanasijevic, and Blackford Middleton. 2003. Ten commandments for effective clinical decision support: making the practice of evidence-based medicine a reality. *Journal of the American Medical Informatics Association* 10, 6 (2003), 523–530. https://academic.oup.com/jamia/article/10/6/523/760582

[5] Rinaldo Bellomo, Claudio Ronco, Ravindra L Mehta, Pierre Asfar, Julie Boisramé-Helms, Michael Darmon, Jean-Luc Diehl, Jacques Duranteau, Eric AJ Hoste, Joannes-Boyau Olivier, et al. 2017. Acute kidney injury in the ICU: from injury to recovery: reports from the 5th Paris International Conference. *Annals of intensive care* 7, 1 (2017), 1–40. https://doi.org/10.1186/s13613-017-0260-y

[6] Hemant K Bhargava and Abhay Nath Mishra. 2014. Electronic medical records and physician productivity: Evidence from panel data analysis. *Management Science* 60, 10 (2014), 2543–2562. https://doi.org/doi/10.1287/mnsc.2014.1934

[7] David Blumenthal and Marilyn Tavenner. 2010. The "meaningful use" regulation for electronic health records. *New England Journal of Medicine* 363, 6 (2010), 501–504. https://doi.org/doi/full/10.1056/nejmp1006114

[8] Alex AT Bui, Denise R Aberle, and Hooshang Kangarloo. 2007. TimeLine: visualizing integrated patient records. *IEEE Transactions on Information Technology in Biomedicine* 11, 4 (2007), 462–473. https://doi.org/10.1109/titb.2006.884365

[9] Andreas Buja, John Alan McDonald, John Michalak, and Werner Stuetzle. 1991. Interactive data visualization using focusing and linking. In *Proceedings of the 2nd conference on Visualization'91*. 156–163. https://doi.org/10.1109/VISUAL.1991.175794

[10] Mackinlay Card. 1999. *Readings in information visualization: using vision to think*. Morgan Kaufmann. https://doi.org/doi/10.5555/300679

[11] Alex Chaparro, Joseph R Keebler, Elizabeth H Lazzara, and Anastasia Diamond. 2019. Checklists: A review of their origins, benefits, and current uses as a cognitive aid in medicine. *ergonomics in design* 27, 2 (2019), 21–26. https://doi.org/10.1177/1064804618819181

[12] James J Cimino, Anthony Aguirre, Stephen B Johnson, and Ping Peng. 1993. Generic queries for meeting clinical information needs. *Bulletin of the Medical Library Association* 81, 2 (1993), 195. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC225762/

[13] Rhonda Collins. 2020. Clinician cognitive overload and its implications for nurse leaders. *Nurse Leader* 18, 1 (2020), 44–47. https://doi.org/10.1016/j.mnl.2019.11.007

[14] Definitive Healthcare, HIMSS Analytics, and Vocera. 2019. *Technology and Clinician Cognitive Overload - Easing the Pain*. Technical Report. https://www.vocera.com/public/pdf/Vocera_HIMSS_Analytics_Survey_Report_2019_UK.pdf

[15] Molla S Donaldson, Janet M Corrigan, Linda T Kohn, et al. 2000. To err is human: building a safer health system. (2000). https://doi.org/10.17226/9728

[16] Anthony Faiola, Preethi Srinivas, and Jon Duke. 2015. Supporting clinical cognition: a human-centered approach to a novel ICU information visualization dashboard. In *AMIA Annual Symposium Proceedings*, Vol. 2015. American Medical Informatics Association, 560. https://pubmed.ncbi.nlm.nih.gov/26958190/

[17] Amy Franklin, Swaroop Gantela, Salsawit Shifarraw, Todd R Johnson, David J Robinson, Brent R King, Amit M Mehta, Charles L Maddow, Nathan R Hoot, Vickie Nguyen, et al. 2017. Dashboard visualizations: Supporting real-time throughput decision-making. *Journal of biomedical informatics* 71 (2017), 211–221. https://doi.org/10.1016/j.jbi.2017.05.024

[18] Amit X Garg, Neill KJ Adhikari, Heather McDonald, M Patricia Rosas-Arellano, Philip J Devereaux, Joseph Beyene, Justina Sam, and R Brian Haynes. 2005. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review. *Jama* 293, 10 (2005), 1223–1238. https://doi.org/10.1001/jama.293.10.1223

[19] Michael Glueck, Peter Hamilton, Fanny Chevalier, Simon Breslav, Azam Khan, Daniel Wigdor, and Michael Brudno. 2015. PhenoBlocks: Phenotype comparison visualizations. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 101–110. https://doi.org/10.1109/tvcg.2015.2467733

[20] Travis R Goodwin and Sanda M Harabagiu. 2016. Medical question answering for clinical decision support. In *Proceedings of the 25th ACM international on conference on information and knowledge management*. 297–306. https://doi.org/10.1145/2983323.2983819

[21] Brigette Hales, Marius Terblanche, Robert Fowler, and William Sibbald. 2008. Development of medical checklists for improved quality of patient care. *International Journal for Quality in Health Care* 20, 1 (2008), 22–30. https://doi.org/10.1093/intqhc/mzm062

[22] Yong Y Han, Joseph A Carcillo, Shekhar T Venkataraman, Robert SB Clark, R Scott Watson, Trung C Nguyen, Hülya Bayir, and Richard A Orr. 2005. Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system. *Pediatrics* 116, 6 (2005), 1506–1512. https://doi.org/10.1542/peds.2005-1287

[23] Marissa A Hendrickson, Genevieve B Melton, and Michael B Pitt. 2019. The review of systems, the electronic health record, and billing. *Jama* 322, 2 (2019), 115–116. https://doi.org/10.1001/jama.2019.5667

[24] Joy Higgs and Mark A Jones. 2008. Clinical decision making and multiple problem spaces. *Clinical reasoning in the health professions* 3 (2008), 3–17.

[25] Oliver Hulland, Jeanne Farnan, Raphael Rabinowitz, Lisa Kearns, Michele Long, Bradley Monash, Priti Bhansali, and H Barrett Fromme. 2017. What's the Purpose of Rounds? A Qualitative Study Examining the Perceptions of Faculty and Students. *Journal of hospital medicine* 12, 11 (2017), 892–897. https://doi.org/10.12788/jhm.2835

[26] Olaronke G Iroju and Janet O Olaleke. 2015. A systematic review of natural language processing in healthcare. *International Journal of Information Technology and Computer Science* 7, 8 (2015), 44–50. https://doi.org/10.5815/ijitcs.2015.08.07

[27] Maia Jacobs, Jeffrey He, Melanie F. Pradier, Barbara Lam, Andrew C Ahn, Thomas H McCoy, Roy H Perlis, Finale Doshi-Velez, and Krzysztof Z Gajos. 2021. Designing AI for Trust and Collaboration in Time-Constrained Medical Decisions: A Sociotechnical Lens. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14. https://doi.org/doi/10.1145/3411764.3445385

[28] Archana Laxmisan, Forogh Hakimzada, Osman R Sayan, Robert A Green, Jiajie Zhang, and Vimla L Patel. 2007. The multitasking clinician: decision-making and cognitive demand during and after team handoffs in emergency care. *International journal of medical informatics* 76, 11-12 (2007), 801–811. https://doi.org/10.1016/j.ijmedinf.2006.09.019

[29] Anna S Messmer, Carina Zingg, Martin Müller, Joel Loic Gerber, Joerg Christian Schefold, and Carmen Andrea Pfortmueller. 2020. Fluid overload and mortality in adult critical care patients—a systematic review and meta-analysis of observational studies. *Critical care medicine* 48, 12 (2020), 1862–1870. https://doi.org/10.1097/ccm.0000000000004617

[30] Mark A Musen, Blackford Middleton, and Robert A Greenes. 2021. Clinical decision-support systems. In *Biomedical informatics*. Springer, 795–840. https://doi.org/10.1007/978-1-4471-4474-8_22

[31] Jerome A Osheroff, Jonathan M Teich, Donald Levick, Luis Saldana, Ferdinand T Velasco, Dean F Sittig, Kendall M Rogers, and Robert A Jenders. 2012. *Improving outcomes with clinical decision support: an implementer's guide*. Himss Publishing. https://www.routledge.com/Improving-Outcomes-with-Clinical-Decision-Support-An-Implementers-Guide/Osheroff-Teich-Levick-Saldana-Velasco-Sittig-Rogers-Jenders/p/book/9780984457731

[32] Jerome A Osheroff, Jonathan M Teich, Blackford Middleton, Elaine B Steen, Adam Wright, and Don E Detmer. 2007. A roadmap for national action on clinical decision support. *Journal of the American medical informatics association* 14, 2 (2007), 141–145. https://doi.org/10.1197/jamia.m2334

[33] Shobha Phansalkar, Judy Edworthy, Elizabeth Hellier, Diane L Seger, Angela Schedlbauer, Anthony J Avery, and David W Bates. 2010. A review of human factors principles for the design and implementation of medication safety alerts in clinical information systems. *Journal of the American Medical Informatics Association* 17, 5 (2010), 493–501. https://doi.org/10.1136/jamia.2010.005264

[34] Catherine Plaisant, Brett Milash, Anne Rose, Seth Widoff, and Ben Shneiderman. 1996. LifeLines: visualizing personal histories. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 221–227. https://doi.org/doi/10.1145/238386.238493

[35] Catherine Plaisant, Richard Mushlin, Aaron Snyder, Jia Li, Dan Heller, and Ben Shneiderman. 2003. LifeLines: using visualization to enhance navigation and analysis of patient records. In *The craft of information visualization*. Elsevier, 308–312. https://pubmed.ncbi.nlm.nih.gov/9929185/

[36] Margit Pohl, Sylvia Wiltner, Alexander Rind, Wolfgang Aigner, Silvia Miksch, Thomas Turic, and Felix Drexler. 2011. Patient development at a glance: An evaluation of a medical data visualization. In *IFIP Conference on Human-Computer Interaction*. Springer, 292–299. https://link.springer.com/chapter/10.1007/978-3-642-23768-3_24

[37] Jan Powers and Karen Samaan. 2014. Malnutrition in the ICU patient population. *Critical Care Nursing Clinics* 26, 2 (2014), 227–242. https://doi.org/10.1016/j.ccell.2014.01.003

[38] Seth M Powsner and Edward R Tufte. 1994. Graphical summary of patient status. *Lancet* 344, 8919 (1994), 386–389. https://doi.org/10.1016/s0140-6736(94)91406-0

[39] Alexander Rind, Taowei David Wang, Wolfgang Aigner, Silvia Miksch, Krist Wongsuphasawat, Catherine Plaisant, and Ben Shneiderman. 2013. Interactive information visualization to explore and query electronic health records. *Foundations and Trends in Human-Computer Interaction* 5, 3 (2013), 207–298. https://doi.org/doi/10.1561/1100000039

[40] Adrian C Rodriguez, Darrell A Lee, and Mary Beth Flynn Makic. 2017. Situational awareness in critical care: An aviation approach to reduce error. *Journal of PeriAnesthesia Nursing* 32, 6 (2017), 650–652.

[41] Keith J Ruskin and Dirk Hueske-Kraus. 2015. Alarm fatigue: impacts on patient safety. *Current Opinion in Anesthesiology* 28, 6 (2015), 685–690. https://doi.org/10.1097/aco.0000000000000260

[42] Jeffrey R Saffle. 2016. Fluid creep and over-resuscitation. *Critical care clinics* 32, 4 (2016), 587–598. https://doi.org/10.1016/j.ccc.2016.06.007

[43] Michelle Schmerge. 2016. Promoting Patient-Centered Team-Based Care: An Advocacy Toolkit. (2016).

[44] Sandeep Sharma, Muhammad F Hashmi, and Rotem Friede. 2020. Interprofessional rounds in the ICU. *StatPearls [Internet]* (2020). https://pubmed.ncbi.nlm.nih.gov/29939553/

[45] Hyoung Seok Shin. 2019. Reasoning processes in clinical reasoning: from the perspective of cognitive psychology. *Korean journal of medical education* 31, 4 (2019), 299. https://doi.org/10.3946/kjme.2019.140

[46] Ben Shneiderman. 1994. Dynamic queries for visual information seeking. *IEEE software* 11, 6 (1994), 70–77. https://doi.org/10.1109/52.329404

[47] Edward H Shortliffe and Michael F Chiang. 2021. Biomedical data: Their acquisition, storage, and use. In *Biomedical informatics*. Springer, 45–75. https://doi.org/10.1007/0-387-36278-9_2

[48] Dean F Sittig, Adam Wright, Jerome A Osheroff, Blackford Middleton, Jonathan M Teich, Joan S Ash, Emily Campbell, and David W Bates. 2008. Grand challenges in clinical decision support. *Journal of biomedical informatics* 41, 2 (2008), 387–392. https://doi.org/10.1016/j.jbi.2007.09.003

[49] 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*. ACM, 422. https://doi.org/doi/10.1145/3173574.3173996

[50] Nicole Sultanum, Devin Singh, Michael Brudno, and Fanny Chevalier. 2018. Doccurate: A Curation-Based Approach for Clinical Text Visualization. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 142–151. http://www.cs.toronto.edu/~fanny/pubs/pdf/doccurate-vast2018.pdf

[51] Reed T Sutton, David Pincock, Daniel C Baumgart, Daniel C Sadowski, Richard N Fedorak, and Karen I Kroeker. 2020. An overview of clinical decision support systems: benefits, risks, and strategies for success. *NPJ digital medicine* 3, 1 (2020), 1–10. https://doi.org/10.1038/s41746-020-0221-y

[52] Øyvind Thomassen, Ansgar Storesund, Eirik Søfteland, and Guttorm Brattebø. 2014. The effects of safety checklists in medicine: a systematic review. *Acta Anaesthesiologica Scandinavica* 58, 1 (2014), 5–18. https://doi.org/doi/10.1111/aas.12207

[53] Thi Ngoc Trang Tran, Alexander Felfernig, Christoph Trattner, and Andreas Holzinger. 2021. Recommender systems in the healthcare domain: state-of-the-art and research issues. *Journal of Intelligent Information Systems* 57, 1 (2021), 171–201. https://doi.org/10.1007/s10844-020-00633-6

[54] Stacey L Valentine, Anil Sapru, Renee A Higgerson, Phillip C Spinella, Heidi R Flori, Dionne A Graham, Molly Brett, Maureen Convery, LeeAnn M Christie, Laurie Karamessinis, et al. 2012. Fluid balance in critically ill children with acute lung injury. *Critical care medicine* 40, 10 (2012), 2883.

[55] J Van der Greef, S Martin, P Juhasz, A Adourian, T Plasterer, ER Verheij, and RN McBurney. 2007. The art and practice of systems biology in medicine: mapping patterns of relationships. *Journal of proteome research* 6, 4 (2007), 1540–1559. https://doi.org/doi/full/10.1021/pr0606530

[56] Qiru Wang and RS Laramee. 2021. EHR STAR: The State-Of-the-Art in Interactive EHR Visualization. In *Computer Graphics Forum*. Wiley Online Library.

[57] 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*. ACM, 457–466. https://doi.org/doi/10.1145/1357054.1357129

[58] Robert L Wears and Marc Berg. 2005. Computer technology and clinical work: still waiting for Godot. *Jama* 293, 10 (2005), 1261–1263. https://doi.org/10.1001/jama.293.10.1261

[59] Adam Wilcox, Spencer S Jones, David A Dorr, Wayne Cannon, Laurie Burns, Kelli Radican, Kent Christensen, Cherie Brunker, Ann Larsen, Scott P Narus, et al. 2005. Use and impact of a computer-generated patient summary worksheet for primary care. In *AMIA Annual Symposium Proceedings*, Vol. 2005. American Medical Informatics Association, 824. https://doi.org/pmc/articles/PMC1560720/

[60] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting intelligent decision support into critical, clinical decision-making processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11. https://doi.org/doi/10.1145/3290605.3300468

[61] Zhilan Zhou, Ximing Wen, Yue Wang, and David Gotz. 2021. Modeling and Leveraging Analytic Focus During Exploratory Visual Analysis. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15. https://doi.org/doi/10.1145/3411764.3445674