Clinical Decision Support Stewardship: Best Practices and Techniques to Monitor and Improve Interruptive Alerts

Juan D. Chaparro[1,2]   Jonathan M. Beus[3,4]   Adam C. Dziorny[5]   Philip A. Hagedorn[6,7]
Sean Hernandez[8,9]   Swaminathan Kandaswamy[3]   Eric S. Kirkendall[8,10,11]   Allison B. McCoy[12]
Naveen Muthu[13,14]   Evan W. Orenstein[3,4]

[1] Division of Clinical Informatics, Nationwide Children's Hospital, Columbus, Ohio, United States
[2] Departments of Pediatrics and Biomedical Informatics, The Ohio State University College of Medicine, Columbus, Ohio, United States
[3] Department of Pediatrics, Emory University School of Medicine, Atlanta, Georgia, United States
[4] Children's Healthcare of Atlanta, Atlanta, Georgia, United States
[5] Department of Pediatrics, University of Rochester School of Medicine, Rochester, New York, United States
[6] Department of Pediatrics, University of Cincinnati, College of Medicine, Cincinnati, Ohio, United States
[7] Division of Hospital Medicine, Cincinnati Children's Hospital Medical Center, Cincinnati, Ohio, United States
[8] Center for Healthcare Innovation, Wake Forest School of Medicine, Winston-Salem, North Carolina, United States
[9] Department of General Internal Medicine, Wake Forest School of Medicine, Winston-Salem, North Carolina, United States

Address for correspondence Juan D. Chaparro, MD, MS, Division of Clinical Informatics, Nationwide Children's Hospital, Columbus, Ohio 43205, United States
(e-mail: juan.chaparro@nationwidechildrens.org).

[10] Department of Pediatrics, Wake Forest School of Medicine, Winston-Salem, North Carolina, United States
[11] Center for Biomedical Informatics, Wake Forest School of Medicine, Winston-Salem NC, United States
[12] Department of Biomedical Informatics, Vanderbilt University Medical Center, Nashville, Tennessee, United States
[13] Department of Pediatrics, University of Pennsylvania Perelman School of Medicine, Philadelphia, Pennsylvania, United States
[14] Department of Biomedical and Health Informatics, Children's Hospital of Philadelphia, Philadelphia, Pennsylvania, United States

Abstract

Keywords
► linical information system use, administration, and maintenance
► clinical decision support
► organizational change management mixed methodologies
► socio-technical aspects of information technology

Interruptive clinical decision support systems, both within and outside of electronic health records, are a resource that should be used sparingly and monitored closely. Excessive use of interruptive alerting can quickly lead to alert fatigue and decreased effectiveness and ignoring of alerts. In this review, we discuss the evidence for effective alert stewardship as well as practices and methods we have found useful to assess interruptive alert burden, reduce excessive firings, optimize alert effectiveness, and establish quality governance at our institutions. We also discuss the importance of a holistic view of the alerting ecosystem beyond the electronic health record.

Background and Significance

Clinical decision support (CDS) systems provide "a process for enhancing health-related decisions and actions with pertinent, organized clinical knowledge and patient information to improve health and healthcare delivery."[1] Electronic CDS has been used successfully in a variety of clinical areas in the pediatric and adult literature.[2,3] Many CDS systems provide passive guidance in the form of order sets, templated forms or notes, or reports displaying relevant information, but interruptive or pop-up alerts are the most visible form of CDS. Meanwhile, electronic health record (EHR) interruptive alerts represent only one element of the full interruptive alert ecosystem (►Fig. 1).

Interruptive alerts have the perceived benefit of forcing a user to notice and respond to the prompt. However, this interruption comes with significant costs: the immediate cost of increasing cognitive burden and interrupting tasks, and the longer-term cost of alert fatigue and decreased provider receptiveness to future alerts both in the EHR and monitor alarms.[4] Fundamentally, we risk increasing provider burnout with the cumulative effect of interruptive alerts[5]; thus, interruptive alerts should be used only when other less intrusive options have been thoroughly considered.[6]

In this review, we discuss our institutional experiences and best practices to preserve the value of the interruptive alert, whether via EHR, physiologic monitor, or mobile device. We often perform tasks analogous to antimicrobial stewards; we educate clinicians about other tools that may work better with narrower coverage, optimize CDS strategies to reduce side effects, and, when necessary, discourage or even prevent the implementation of interruptive alerts when they may cause more harm than good. We also describe the evaluation and monitoring necessary for a proper alert stewardship program, including timely and accurate tracking of CDS tool use and effectiveness as well as CDS governance. We anticipate others may use these experiences to develop their own comprehensive CDS program.

The Case for Interruptive Alert Stewardship

Medical providers have long had a conflicted relationship with interruptive CDS. In one of the first implementations of computerized physician order entry (CPOE), the response to a medication interaction alert was described thus: "*When this feed-back first occurs, the user is very impressed; however, by the tenth time it occurs he is annoyed, and when it appears for the twentieth time, he is insulted and frustrated at the computer's insensitivity to the fact that, by this time, the operator is aware of this bit of medical knowledge.*"[7]

While a novel experience back in 1968, the modern provider is inundated with alerts as more elements of patient care occur via electronic platforms. Physiologic alarms in one adult intensive care unit generated 187 audible alerts/bed/day, many of which were erroneous or unactionable.[8] The EHR is also the source of an increasing number of alerts.[9]

Interruptive alert frequency is a risk factor for medical errors. Workflow interruptions are correlated with increased number of errors as well as failure to return to the original task.[10–13] Additionally, interruptive alerts have low rates of practitioner acceptance (4–11%).[14–16] Although much of these data have come from vendor-provided CPOE alerts, override rates of custom-built interruptive alerts are similarly high.[17]

Low acceptance rates may result from poor targeting of recipients, incorrect information, un-actionable guidance, or misaligned workflow and contribute to alert distrust and alert fatigue. Humans are adept at probability matching and identify unreliable alerts after only a small number of exposures.[18] High cognitive load, common in most clinical practice, exacerbates distrust of unreliable alerts.[19] Overriding alerts also leads to habitual alert override behavior. An alert is routinely shown in a stable context and frequently overridden. Because in most cases there is no immediate negative consequence, the override behavior is positively reinforced.[20] Aside from these behavioral adaptations to alerts, more concerning is the "boy who cried wolf" phenomenon whereby the cacophony masks the one alert with the potential to save life or limb. Evidence for this already exists: in one quality improvement effort, criteria for a series of medication alerts were made more strict to reduce total firings, while other medication alerts in the system remained untouched. The acceptance rates improved, not only for the edited alerts as expected but also for alerts where the design did not change at all.[21] Thus, the presence of poor alerts in the system can reduce the effectiveness of all CDS.

In addition to the challenges of alert fatigue and burnout, each new alert requires additional upkeep and monitoring for unintended outcomes and performance after initial implementation. Alert prioritization and maintenance depend on other health system factors such as knowledge management and sharing capability, health system priorities, quality

Fig. 1 Interruptive alert ecosystem in health care. Interruptive alerts can come from electronic health records as well as patient monitors, phones, pagers, and other channels that may each lead to frustration and potentiate alert fatigue.

improvement impacts, and feedback generated from evaluation of CDS.[22] Over time, these interactions and outcomes can change and produce unintended consequences. Clinical guidelines may change requiring updates to CDS interventions or technical changes may lead to CDS malfunctions such as inadvertent inactivation.[23]

The decision to introduce new alerts and their interruptions into clinical work must be a thoughtful one, akin to the tasks of antimicrobial stewards. We must educate clinicians about tools that may work with narrower coverage to limit overexposure, revise CDS choices to minimize side effects, and, when necessary, nudge CDS requestors away from interruptive alerts so they will be available for use in the future. However, to identify those alerts worthy of interrupting workflows, we must understand the current state of interruptive alerts and how we can measure alert success or failure.

Measuring Alert Burden and Effectiveness

Currently, no standard metrics exist to easily compare the burden of alerts on EHR users or their effectiveness at improving outcomes.[24] Commonly used alert burden metrics focus on alert frequency (e.g., alerts per 100 orders), override rates, or time burden (e.g., think time or dwell time).[25–27] However, differences in denominators for alert frequency (e.g., alerts per 100 orders versus alerts per inpatient day or clinician time in the EHR) may affect which alerts or care settings an organization categorizes as high burden, leading to biases in prioritization efforts. Alert effectiveness metrics may focus on *proximal* outcomes[28–30] (e.g., appropriateness or use of the intended action) or *distal* outcomes specific to the goal of each alert (e.g., whether a vaccine was administered or care gap closed).[24] Proximal measures can be calculated and compared more easily across alerts but may not reflect how well the alert is achieving its intended purpose. Standardizing alert metrics for both burden and effectiveness could enable benchmarking across institutions and development of translational dashboards directly comparing CDS strategies for the same use case across institutions.

Alert Burden

In a cross-sectional analysis of six academic pediatric health systems, we compared interruptive alert burden from September 2016 to September 2019 using four metrics including two patient-focused denominators (alerts per inpatient-day and alerts per encounter) and two clinician-focused denominators (alerts per 100 orders and alerts per clinician day, defined as the number of unique clinicians on each calendar day with at least one EHR log in the system).[31] We found wide variation of alert burden, with alerts per clinician day at the highest burden site a staggering 43.8 times higher than the lowest burden site. The rank order of institutions by alert burden did not substantially vary across the four different alert metrics. Custom alerts accounted for a higher proportion of alert burden compared with drug–drug interaction alerts or medication administration alerts across all sites and metrics. By contrast, when we examined intrainstitutional variation, we found that areas of highest burden varied by

metric chosen. For example, nurses had the highest alert burden across all sites when looking at alerts per 100 orders, whereas pharmacists experienced 3.1 times higher alert burden than all other provider types when using the metric alerts per clinician day.

These findings demonstrate that while interinstitutional comparison is possible with existing metrics, understanding intrainstitutional variation in current state requires using multiple metrics, each providing a different lens to guide burden reduction strategies. Future studies establishing the predictive validity of specific burden metrics for alert fatigue behaviors—for example, determining the association between alerts per hour in the EHR and deleterious effects on patient care or how a user responds to subsequent alerts—would advance our ability to reduce burnout and patient safety concerns from excess alerts.[32–34]

Alert Effectiveness

Proximal measures of alert effectiveness focus on the user's response in the moment the alert fires. Simple examples include metrics where the alert firing is the unit of analysis (e.g., override rates or acceptance rates). These metrics are easy to calculate and quickly compare, but mask nuances such as justifiable overrides (i.e., the alert was inappropriate and the user correctly ignored it) and unintended consequences (i.e., the alert was inappropriate, but the user "accepted" it, leading to an error). Alternatively, alert outcomes can be classified as a success, an appropriate override, provider nonadherence, or an unintended consequence. However, while this classification yields more useful information, it generally requires manual chart review from clinicians with sufficient context to judge the user response.[29] Machine learning approaches to reduce the resources required for this classification task have shown promise for some use cases but are yet to achieve widespread uptake.[28]

Distal measures of alert effectiveness focus on the degree to which the alert addresses the targeted quality or safety problem. For example, if an alert was developed to notify providers of a patient's eligibility for COVID vaccine during an office visit, the appropriate distal outcome measure might be the proportion of eligible visits with a vaccine administered or the proportion of the total population cared for by the clinic that is vaccinated. These metrics are of greater utility but (1) require more resources to develop, (2) cannot be easily compared across alerts with different intended purposes, and (3) suffer from study design challenges to determine if the alert itself was responsible for a change in an outcome metric compared with other interventions or secular trends.

In the absence of an easily scalable, valid way to measure alert effectiveness, combining process and outcome measures in a quality improvement or process evaluation framework may be the most useful approach for assessing alert utility and comparing alert effectiveness. Traditional quality improvement studies tend to use the Donabedian framework, which separates measures into structure (e.g., how many alerts, order sets, or other CDS artifacts are available for a care process), process (e.g., alert acceptance rates), and outcomes.[35] This approach allows quality improvement
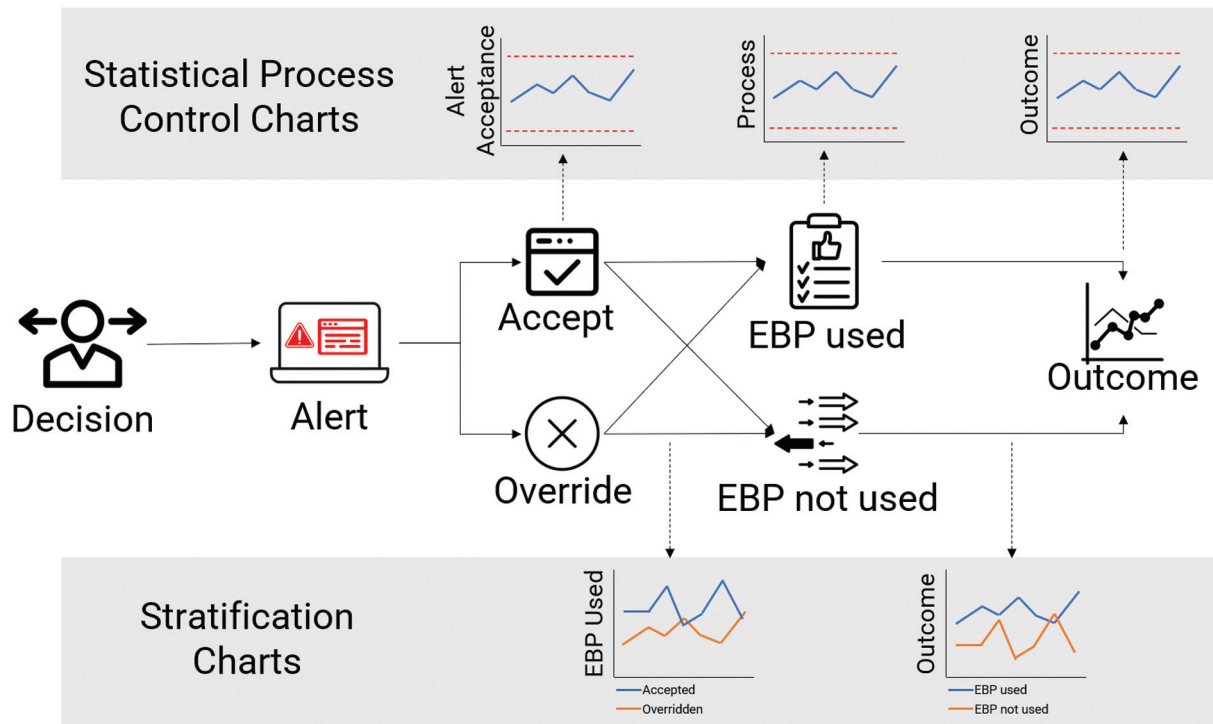
Fig. 2 Alert evaluation framework connecting proximal alert measures (e.g., alert acceptance) with distal measures such as care processes and outcomes. EBP, evidence-based practice.

advocates to start by evaluating the outcome and then examine process adherence to drive plan-do-study-act cycles. The Medical Research Council's Process Evaluation Framework for Complex Interventions emphasizes studying implementation effectiveness, the mechanism of impact, and outcomes stratified by contextual factors.[36] In this approach, proximal alert effectiveness measures can be used as a proxy for implementation measures—for example, the number of firings can help estimate the reach of the intervention, while alert acceptance rates can assess fidelity.[37] Changes in care processes with and without the alert can help assess the mechanism of impact or change theory. One scalable approach evaluates for the intended action within 1 hour of an alert firing to determine if the alert is leading to the intended behavior change. Finally, the care process is attached to the final patient outcome of interest when available. This framework can help quality improvement advocates examine their theory for how an alert should help achieve quality goals (►**Fig. 2**). By measuring outcomes of interest, adherence to guideline-recommended care, and proximal measures of alert response, organizations can frame CDS performance alongside outcomes of interest and use this insight to target interventions.[38]

Optimizing Alert Effectiveness

While measuring the success of a custom alert is a difficult and often very individualized process, there are systematic approaches to optimize the usefulness of interruptive alerts. This involves the detection of alerts that are malfunctioning and not working due to technical issues, and the more challenging task of identifying alerts that are functioning but were not built following the five rights of CDS. We describe our experiences here, but we recommend reviewing the excellent best practices developed by Wright et al.[39]

Identifying malfunctioning alerts primarily occurs through either direct reporting of the issue or anomaly detection using statistical methods. In the past, direct reporting occurred either as an override comment within the alert, as a help desk ticket, or via other informal communication to the EHR team. Because the latter two of these require additional effort, they often do not occur unless the error is particularly frequent or egregious.

However, some institutions as well as EHR vendors now embed feedback mechanisms within interruptive alerts, allowing immediate qualitative feedback as well as providing user and patient information for troubleshooting. At one of our member's institutions, Nationwide Children's Hospital (NCH), both positive and negative feedback links were gradually implemented to most interruptive alerts. End users submitted feedback 806 times in a 30-month period, and surprisingly, 53% of feedback was positive. Through the critical feedback received, 21 unique alerts were fixed/improved. Additionally, feedback surveys often led to email outreach from the informatics team to the end user either to request more detail or to provide updates on changes. This direct communication may help put a human face to CDS and allow users to feel heard. Taking this one step further, Wright and coworkers used a cranky comments heuristic to identify feedback that could potentially indicate a broken alert.[40] This heuristic involved searching for a limited number of words that indicated user frustration or annoyance.

In addition to user feedback, anomaly detection can successfully identify alert malfunctions.[41–43] Although not all anomalies are malfunctions, finding abnormal patterns of firing can act as a screening tool to prompt review for malfunctions. Malfunctions that lend themselves to discovery by anomaly detection include changes elsewhere in the system (e.g., diagnostic code sets, drug reference databases, record names, or workflow) that unintentionally "break" alerts, causing excessive firing or absence of firing. Alert users may bring the former to attention after receiving extra alerts, but the latter rarely elicits complaints.

Identifying correctly functioning but poorly designed alerts is more difficult. In a previous publication, Chaparro et al described a systematic approach to reduce inappropriate alerting and improve the quality of interruptive alerts.[9] Initial efforts using quality improvement methodology targeted high-volume alerts, as reducing inappropriate firing of these alerts would provide more return on effort spent revising the alerts. After these initial gains, it was clear that different approaches would be needed to achieve further improvements. To this end, further alert review/revision was prioritized based on two other dimensions, patients on whom a disproportionate number of alerts fired and providers for whom certain alerts fired at much higher rates than other providers. The former often indicated edge cases where a patient had certain unusual characteristics in their digital profile not accounted for in build or testing, while the latter was frequently a symptom of a misaligned workflow. One of the takeaways is that while considering Osheroff's 5 Rights of CDS is important in the design and implementation of new alerts, identifying existing alerts that were not built with these in mind requires viewing alert firing data using different dimensions such as high-firing alerts, providers, or patients.

In addition to ensuring the technical build is correct, human factors principles are a key component to increasing acceptance rates of alerts.[44] In the work at NCH, Nielsen's 10 Usability Heuristics were adapted for the redesign of their alerts. A standard template was used clearly explaining why the alert displayed (including relevant clinical data), with directions as to what the user was expected to do. Effort was made to keep alert text as concise as possible. Lastly, they worked to make alert acknowledgment reasons clearer as to their behavior, whether it was a temporary or permanent silencing of the alert, and whether this action affected only the current user, or potentially suppressing for other users as well. Effective governance to establish these standards and style guides was essential to this endeavor.

Governance of Alerting

By regulation, most institutions have at least a rudimentary process for regular review of order sets in their EHR. However, the intake/evaluation process for CDS implementation is an important "gatekeeper function" and, more importantly, educates requestors and encourages best practices. It is also an opportunity to build in accountability for all parties, ensuring requestor engagement in the CDS implementation and change management processes. Otherwise, CDS requests may be used to attempt to implement people and process changes in isolation. Some institutions have robust governance structures for CDS, with representation from clinical stakeholders, quality leaders, informatics, and IT professionals. However, unlike antibiotic stewardship where Centers for Disease Control and Prevention defined a set of core elements for antibiotic stewardship programs, such guidance does not yet exist for CDS governance. With alerting, there are at least three elements essential to any governance approach in addition to stakeholder representation and accountability for outcomes.[45]

First, an overall model for approval, maintenance, and review should be established. Depending on how CDS is developed and implemented at an institution, both centralized and federated approaches may be appropriate. A centralized CDS committee chartered to manage governance can ensure a consistent approach to alerting. However, a federated approach with more local governance of alerts but with shared standards for design and implementation may be more capable of rapidly responding to changes in clinical knowledge and allow for greater institutional capacity to build and maintain CDS. Similar to structure, accountability may also vary across health systems, with overall accountability potentially including information technology leadership, clinical leadership, or quality and safety leadership. However, effective governance requires representation of all stakeholders that include clinical roles (e.g., physicians, RNs, RTs), informatics, information technology teams, and administration including quality and safety as well as regulatory roles. Nevertheless, this ideal state should not preclude pragmatic tradeoffs that will allow governance to proceed, especially in the context of a pandemic that has left many clinicians burned out and without the bandwidth to participate in governance efforts. While a detailed review of CDS governance is beyond the scope of this work, we encourage readers to review the discussion of governance structures found in McGreevey et al's summary of a previous AMIA panel discussion on EHR alerts.[24]

Second, standards should exist to determine when the interruption created by an alert is justified, as every interruption has the capacity to increase cognitive burden and the likelihood of errors in decision-making.[24] One approach may be to adapt risk prioritization frameworks such as the Healthcare Failure Modes and Effects Analysis.[46] This approach can provide structured questions such as: "In the absence of this knowledge, how severe a harm could occur? How probable is the harm? What is the likelihood that the user would already have this knowledge or access to this knowledge?" Predetermined thresholds for the answers to these questions can be used to determine when an interruption is warranted. Additionally, the decision to implement an interruptive alert, in particular, should include a representative of the user groups who will receive the alert. Otherwise, generalist providers such as hospitalists may find themselves on the receiving end of interruptive alerts from multiple specialties.

Lastly, standards should be developed and enforced around alert design and construction to maximize the

effectiveness of alerts and ensure predictable presentation of information. The Electronic Health Records Association provides basic design patterns to be adapted by informatics and human factors professionals to establish local standards. Specifically, four basic components are recommended: (1) a consistent signal to indicate the seriousness of the alert, (2) information about the hazard, (3) instructions or actions to mitigate the hazard, and (4) specific clinical consequences that may ensue if the hazard is not averted.[47] Another best practice is alert verification and validation prior to implementation, for example, using approaches such as retrospective analysis of alert criteria and/or running the alert in the "background" without being visible to any users.[24]

As noted earlier, these three elements of alert governance are largely suggested based on institutional experience. Alert stewardship and CDS governance need robust evidence to guide implementation across institutions and may be incentivized by policy.

Beyond the EHR: A Holistic View of the Alert Ecosystem

While previous sections have focused on EHR-based alerting mechanisms, a holistic view requires attention to the expanding contribution of interruptive alerting mechanisms outside the EHR. Without this perspective, we risk losing sight of the broader alerting ecosystem and the important ways alerts outside the medical record contribute to a provider's ability to maintain attention, complete critical tasks, and respond appropriately to alerts of any kind.

The contribution of physiologic monitor and device-based alarms is well described.[48,49] These alerts have long comprised the aural tapestry of many inpatient wards and intensive care units and have been a source of concern for more than a decade, including contribution to alarm fatigue and workflow interruptions.[4,50] However, health care systems have proceeded apace with additional vehicles for delivering interruptive alerts. Rarely, however, are these implementations undertaken with consideration of the impact within the broader alerting ecosystem.[51]

Many of these implementations start with the best of intentions and are often successful at achieving their aims within a focused scope. Consider the following three example projects to (1) reduce in-room monitor-based alarms for patients and families, (2) improve antibiotic prescribing based on culture resistance profiles, and (3) drive proper hand-hygiene practice.[52–54] All three projects target key areas for improvement in patient experience, care delivery, and safety. Quality improvement teams may identify novel and compelling methods to deliver notifications considering their project aims and pay close attention to the burden of these notifications *within the scope of their work*. However, there is often little attention paid to the broader alerting ecosystem and how the introduction of these additional alerts might impact users at the receiving end of many other alerting systems.

Unfortunately, few mechanisms exist to federate alerting data from different clinical information systems and track the impact of one project on the whole alerting ecosystem.

Health care systems need to develop proactive and deliberate approaches to address this cumulative burden; this means processes and tools to understand alerting trends at the macrosystem down to the microsystem. Health care systems should develop and deploy tools for leaders to understand what might amount to a "global interruption index" for individual providers or work contexts within a hospital. Such tools would help identify hotspots at risk of task disruption and associated errors as well as provide insight for more deliberate system design and improvement.

Taking a lesson from another industry, the car insurance industry recognizes the ample evidence that performance in menial tasks, like safely driving an automobile, suffers when an individual is interrupted by device-based distractions.[55] Recent years have seen a proliferation of "drive monitoring apps," which surveil a driver's proclivity to using a device while driving. Similar evidence exists in health care–specific literature: interruptions lead to errors.[13,34] Health care systems need tools to understand and mitigate interruptions for frontline workers and, in doing so, allow members of the health care team to leverage the advantages of emerging approaches without suffering needlessly at the receiving end of ever-proliferating interruptions.

Future Directions

Here, we expand upon earlier themes, looking to what the future might hold for CDS efforts and research. There is much progress to be made in CDS design and delivery, but there will inevitably be disagreement on the focus of future research and operational efforts. The following topics are intended to guide these discussions.

Governance

Users of EHRs are increasingly asking for specific CDS interventions as they become comfortable with available CDS tools. Similarly, regulatory bodies, as well as other external entities, continue to increase requirements often requiring EHR implementation. These pressures make governance critical to appropriately implement requirements and to minimize unintended consequences. Historically, governance groups were largely technical. Informed, engaged, multidisciplinary governance including expertise in data science and human factors will need to grow throughout the full lifecycle of CDS.[56,57] Governance standards should become more evidence-based, allowing regulatory guidance to go beyond simple establishment of review processes to true best practices for CDS use in clinical care.

Design and Implementation (and De-implementation)

Once a new request for CDS passes the governance intake process, the next critical phase is usually design and implementation. There are many well-known design principles, heuristics, and frameworks to adhere to, including those in human-centered design, human–computer interaction, and human factors disciplines.[58–61] As with governance, we anticipate more regular and reliable infusion of these principles and frameworks into the design of CDS. Vendors are

already taking a more rigorous approach to the basic designs of their software and hardware, but the application of design and usability techniques must also occur at the local configuration level. While vendors must supply appropriate flexible tools to permit good design, local customization dictates the end-user experience. Implementation must be equally rigorous, locally customized, and well thought-out. At the other end of the lifecycle, poorly performing and/or out-of-date CDS must be deprecated to avoid negative impacts on clinical care and end-user experience. To de-implement these CDS artifacts, however, one must first systematically identify them. We anticipate a formalization of design and standardization of implementation both within and between organizations, with growing interest and effort in de-implementation.

Measuring Performance Intelligently

To understand how different forms and instances of CDS are performing, one must be able to measure and track performance in a meaningful and consistent way. Many organizations under-allocate resources to this task, deeming it lower priority than other "operational" tasks. Allocation of resources must also consider the sustainability of measurement and monitoring programs. These programs must be based on accurate and trusted data, be standardized and repeatable, and be contextually specific.[31] An increasing number of monitoring efforts will focus on the holistic impact of multiple CDS and operational systems in concert, in lieu of the current state of measurements being siloed to one work task or form of CDS. As CDS evolves, clinical outcomes-based measures will become expected measures of performance, as opposed to purely process (proximal) measures.

Clinical Decision Support at Scale

Lastly, we expect that there will be increasing focus and a shift to centralized, scalable CDS as technical interoperability increases. The 21st Century Cures Act and other legislative and regulatory mandates are pushing vendors to incorporate more interoperability standards.[62–65] The instance-specific creation of CDS is notoriously costly and often the primary barrier to adoption of third-party tools. The desire to be resource efficient and standardized, along with the factors above, will encourage the adoption of centralized CDS systems. The rise of learning networks, learning health systems, and resource sharing collaboratives will only further push us in this direction.[66–70] Federated learning techniques are one potential pathway for adoption of advanced CDS, including artificial intelligence-based systems. The rise of machine learning and natural language processing in medicine is highlighted by exponentially increasing trends in publications. Becoming more permissible to the adoption of third-party CDS mechanisms will potentially confer many benefits, including standardization of implementation and data reporting, practical implementation of more robust testing methods (such as A/B testing by site), synchronization of CDS implementations for studies, and responsible resource utilization, among other benefits.

Clinical Relevance Statement

Although interruptive alerts have their role in these systems, excessive reliance on or poorly built interruptive alerts can lead to alert fatigue and other downstream effects. Further refinement of alert burden metrics is needed as current metrics do not adequately represent the impact on end users when viewed through different dimensions. The best practices we describe here will allow institutions to establish monitoring and optimization programs to reduce alert burden.

Multiple Choice Questions

1. A new asthma alert is implemented that fires for any patient who has risk of an emergency department visit due to asthma. Unfortunately, the alert is firing for more patients than it should due to incorrectly written criteria. What is the expected impact of this alert on user behavior?
   a. Users will ignore this alert but still respond to other alerts as previously.
   b. Users will ignore this alert and may decrease their responsiveness to other alerts.
   c. Users will accept this alert and also respond to other alerts as previously.
   d. Users will ignore this alert but increase their responsiveness to other alerts.

   Correct Answer: Option b is the correct answer. Alerts do not exist in a vacuum but rather in a dynamic ecosystem of other decision support systems that affect the overall alert burden and the user experience. Previous studies have shown that excessive incorrect alerts may affect acceptance rates of unrelated alerts. By adding a poorly built alert to the system, we may increase rates of alert fatigue and affect unrelated alerts.

2. Proximal measures of alert effectiveness reflect which of the following:
   a. The frequency of alert displays for a given provider action.
   b. The action of the user at the time the alert is displayed.
   c. Clinical outcomes for the process attempting to be changed.
   d. The action of the provider at any time after the alert is displayed reflecting the desired action of the alert.

   Correct Answer: Option b is the correct answer. Proximal measures are best defined as reflecting the action taken immediately upon display of the alert. While they are often the easiest to measure, they often do not reflect the full process and may miss actions taken after the alert is addressed (e.g., orders placed after further reflection or actions reversed after initial acceptance in the alert). Clinical outcomes (c) reflect distal outcomes of alerts and are much more difficult to measure and to standardize across alerts. Similarly, tracking actions

that are suggested within an alert but occur at a later time (d) would also reflect more distal outcomes and are often more difficult to directly attribute as a result of the alert.

Protection of Human and Animal Subjects

There were no human subjects involved in the project.

Funding

None.

Conflict of Interest

None declared.

References

1 Osheroff JA, Teich J, Levick D, et al. Improving Outcomes with Clinical Decision Support: An Implementer's Guide. 2nd ed. HIMSS; 2012

2 Kwan JL, Lo L, Ferguson J, et al. Computerised clinical decision support systems and absolute improvements in care: meta-analysis of controlled clinical trials. BMJ 2020;370:m3216

3 Varghese J, Kleine M, Gessner SI, Sandmann S, Dugas M. Effects of computerized decision support system implementations on patient outcomes in inpatient care: a systematic review. J Am Med Inform Assoc 2018;25(05):593–602

4 Bonafide CP, Localio AR, Holmes JH, et al. Video analysis of factors associated with response time to physiologic monitor alarms in a children's hospital. JAMA Pediatr 2017;171(06):524–531

5 ECRI. Special report: top 10 health technology hazards for 2020. Accessed January 11, 2022 at: https://assets.ecri.org/PDF/White-Papers-and-Reports/ECRI-Top-10-Technology-Hazards-2020-v2.pdf

6 Escovedo C, Bell D, Cheng E, et al. Noninterruptive clinical decision support decreases ordering of respiratory viral panels during influenza season. Appl Clin Inform 2020;11(02):315–322

7 Gouveia WA, Diamantis C, Barnett GO. Computer applications in the hospital medication system. Am J Hosp Pharm 1969;26(03):141–150

8 Drew BJ, Harris P, Zègre-Hemsey JK, et al. Insights into the problem of alarm fatigue with physiologic monitor devices: a comprehensive observational study of consecutive intensive care unit patients. PLoS One 2014;9(10):e110274

9 Chaparro JD, Hussain C, Lee JA, Hehmeyer J, Nguyen M, Hoffman J. Reducing interruptive alert burden using quality improvement methodology. Appl Clin Inform 2020;11(01):46–58

10 Westbrook JI, Coiera E, Dunsmuir WT, et al. The impact of interruptions on clinical task completion. Qual Saf Health Care 2010;19(04):284–289

11 Grundgeiger T, Sanderson P. Interruptions in healthcare: theoretical views. Int J Med Inform 2009;78(05):293–307

12 Ashcroft DM, Quinlan P, Blenkinsopp A. Prospective study of the incidence, nature and causes of dispensing errors in community pharmacies. Pharmacoepidemiol Drug Saf 2005;14(05):327–332

13 Westbrook JI, Woods A, Rob MI, Dunsmuir WT, Day RO. Association of interruptions with an increased risk and severity of medication administration errors. Arch Intern Med 2010;170(08):683–690

14 Ariosto D. Factors contributing to CPOE opiate allergy alert overrides. AMIA Annu Symp Proc 2014;2014:256–265

15 Humphrey K, Jorina M, Harper M, Dodson B, Kim SY, Ozonoff A. An investigation of drug-drug interaction alert overrides at a pediatric hospital. Hosp Pediatr 2018;8(05):293–299

16 van der Sijs H, Aarts J, Vulto A, Berg M. Overriding of drug safety alerts in computerized physician order entry. J Am Med Inform Assoc 2006;13(02):138–147

17 Muthu N, Shelov E, Tobias MC, Karavite DJ, Orenstein EW, Grundmeier RW. Variability in user response to custom alerts in the electronic health record: an observational study. Paper presented at: the AMIA 2019 National Symposium Washington DC; November 16–20, 2019

18 Bliss JP, Gilson RD, Deaton JE. Human probability matching behaviour in response to alarms of varying reliability. Ergonomics 1995;38(11):2300–2312

19 Bliss JP, Dunn MC. Behavioural implications of alarm mistrust as a function of task workload. Ergonomics 2000;43(09):1283–1300

20 Baysari MT, Tariq A, Day RO, Westbrook JI. Alert override as a habitual behavior - a new perspective on a persistent problem. J Am Med Inform Assoc 2017;24(02):409–412

21 Simpao AF, Ahumada LM, Desai BR, et al. Optimization of drug-drug interaction alert rules in a pediatric hospital's electronic health record system using a visual analytics dashboard. J Am Med Inform Assoc 2015;22(02):361–369

22 Greenes RA, Bates DW, Kawamoto K, Middleton B, Osheroff J, Shahar Y. Clinical decision support models and frameworks: seeking to address research issues underlying implementation successes and failures. J Biomed Inform 2018;78:134–143

23 Wright A, Hickman TT, McEvoy D, et al. Analysis of clinical decision support system malfunctions: a case series and survey. J Am Med Inform Assoc 2016;23(06):1068–1076

24 McGreevey JD III, Mallozzi CP, Perkins RM, Shelov E, Schreiber R. Reducing alert burden in electronic health records: state of the art recommendations from four health systems. Appl Clin Inform 2020;11(01):1–12

25 Schreiber R, Gregoire JA, Shaha JE, Shaha SH. Think time: a novel approach to analysis of clinicians' behavior after reduction of drug-drug interaction alerts. Int J Med Inform 2017;97:59–67

26 McDaniel RB, Burlison JD, Baker DK, et al. Alert dwell time: introduction of a measure to evaluate interruptive clinical decision support alerts. J Am Med Inform Assoc 2016;23(e1):e138–e141

27 Elias P, Peterson E, Wachter B, Ward C, Poon E, Navar AM. Evaluating the impact of interruptive alerts within a health system: use, response time, and cumulative time burden. Appl Clin Inform 2019;10(05):909–917

28 McCoy AB, Thomas EJ, Krousel-Wood M, Sittig DF. Clinical decision support alert appropriateness: a review and proposal for improvement. Ochsner J 2014;14(02):195–202

29 McCoy AB, Waitman LR, Lewis JB, et al. A framework for evaluating the appropriateness of clinical decision support alerts and responses. J Am Med Inform Assoc 2012;19(03):346–352

30 Strategies to the five rights of clinical decision support. January 23, 2019. Accessed January 11, 2022 at: https://userweb.epic.com/Thread/82831/Strategies-to-the-Five-Rights-of-Clincial-Decision-Support/?reply=385118

31 Orenstein EW, Kandaswamy S, Muthu N, et al. Alert burden in pediatric hospitals: a cross-sectional analysis of six academic pediatric health systems using novel metrics. J Am Med Inform Assoc 2021;28(12):2654–2660

32 Sinha A, Stevens LA, Su F, Pageler NM, Tawfik DS. Measuring electronic health record use in the pediatric ICU using audit-logs and screen recordings. Appl Clin Inform 2021;12(04):737–744

33 Bonafide CP, Lin R, Zander M, et al. Association between exposure to nonactionable physiologic monitor alarms and response time in a children's hospital. J Hosp Med 2015;10(06):345–351

34 Bonafide CP, Miller JM, Localio AR, et al. Association between mobile telephone interruptions and medication administration errors in a pediatric intensive care unit. JAMA Pediatr 2020;174(02):162–169

35 Donabedian A. Evaluating the quality of medical care. 1966. Milbank Q 2005;83(04):691–729

36 Moore GF, Audrey S, Barker M, et al. Process evaluation of complex interventions: Medical Research Council guidance. BMJ 2015;350:h1258

37 Proctor E, Silmere H, Raghavan R, et al. Outcomes for implementation research: conceptual distinctions, measurement challenges, and research agenda. Adm Policy Ment Health 2011;38 (02):65–76

38 Kandaswamy S, Karavite DJ, Muthu N, et al. User and task analysis for evaluation of clinical decision support for quality improvement. Proc Hum Factors Ergon Soc Annu Meet 2020;64(01):750–754

39 Wright A, Ash JS, Aaron S, et al. Best practices for preventing malfunctions in rule-based clinical decision support alerts and reminders: results of a Delphi study. Int J Med Inform 2018; 118:78–85

40 Aaron S, McEvoy DS, Ray S, Hickman TT, Wright A. Cranky comments: detecting clinical decision support malfunctions through free-text override reasons. J Am Med Inform Assoc 2019;26(01): 37–43

41 Ray S, McEvoy DS, Aaron S, Hickman TT, Wright A. Using statistical anomaly detection models to find clinical decision support malfunctions. J Am Med Inform Assoc 2018;25(07):862–871

42 Kassakian SZ, Yackel TR, Gorman PN, Dorr DA. Clinical decisions support malfunctions in a commercial electronic health record. Appl Clin Inform 2017;8(03):910–923

43 Yoshida E, Fei S, Bavuso K, Lagor C, Maviglia S. The value of monitoring clinical decision support interventions. Appl Clin Inform 2018;9(01):163–173

44 Seidling HM, Phansalkar S, Seger DL, et al. Factors influencing alert acceptance: a novel approach for predicting the success of clinical decision support. J Am Med Inform Assoc 2011;18(04):479–484

45 Kawamoto K, Flynn MC, Kukhareva P, et al. A pragmatic guide to establishing clinical decision support governance and addressing decision support fatigue: a case study. AMIA Annu Symp Proc 2018;2018:624–633

46 DeRosier J, Stalhandske E, Bagian JP, Nudell T. Using health care failure mode and effect analysis: the VA National Center for Patient Safety's prospective risk analysis system. Jt Comm J Qual Improv 2002;28(05):248–267

47 Electronic Health Record Association. Electronic health record design patterns for patient safety. 2017. Accessed January 11, 2022 at: https://www.ehra.org/sites/ehra.org/files/docs/ehra-design-patterns-for-safety.pdf

48 Gross B, Dahl D, Nielsen L. Physiologic monitoring alarm load on medical/surgical floors of a community hospital. Biomed Instrum Technol 2011(Suppl):29–36

49 Yu D, Obuseh M, DeLaurentis P. Quantifying the impact of infusion alerts and alarms on nursing workflows: a retrospective analysis. Appl Clin Inform 2021;12(03):528–538

50 Ruskin KJ, Hueske-Kraus D. Alarm fatigue: impacts on patient safety. Curr Opin Anaesthesiol 2015;28(06):685–690

51 Hagedorn PA, Singh A, Luo B, Bonafide CP, Simmons JM. Secure text messaging in healthcare: latent threats and opportunities to improve patient safety. J Hosp Med 2020;15(06):378–380

52 Pater CM, Sosa TK, Boyer J, et al. Time series evaluation of improvement interventions to reduce alarm notifications in a paediatric hospital. BMJ Qual Saf 2020;29(09):717–726

53 Tchou MJ, Andersen H, Robinette E, et al. Accelerating initiation of adequate antimicrobial therapy using real-time decision support and microarray testing. Pediatr Qual Saf 2019;4(04):e191

54 Singh A, Haque A, Alahi A, et al. Automatic detection of hand hygiene using computer vision technology. J Am Med Inform Assoc 2020;27(08):1316–1320

55 Redelmeier DA, Tibshirani RJ. Association between cellular-telephone calls and motor vehicle collisions. N Engl J Med 1997;336 (07):453–458

56 Wright A, Sittig DF, Ash JS, et al. Governance for clinical decision support: case studies and recommended practices from leading institutions. J Am Med Inform Assoc 2011;18(02):187–194

57 Orenstein EW, Muthu N, Weitkamp AO, et al. Towards a maturity model for clinical decision support operations. Appl Clin Inform 2019;10(05):810–819

58 Horsky J, Schiff GD, Johnston D, Mercincavage L, Bell D, Middleton B. Interface design principles for usable decision support: a targeted review of best practices for clinical prescribing interventions. J Biomed Inform 2012;45(06):1202–1216

59 Zhang J, Walji M. Better EHR: Usability, Workflow & Cognitive Support in Electronic Health Records. 1st ed. National Center for Cognitive Informatics & Decision Making in Healthcare, University of Texas Health Science Center at Houston & School of Biomedical Informatics;2014

60 Miller K, Capan M, Weldon D, et al. The design of decisions: matching clinical decision support recommendations to Nielsen's design heuristics. Int J Med Inform 2018;117:19–25

61 Kannampallil TG, Kaufman DR, Patel VL. Cognitive Informatics for Biomedicine: Human Computer Interaction in Healthcare. 1st ed. Springer International Publishing; 2015

62 Rodriguez JA, Clark CR, Bates DW. Digital health equity as a necessity in the 21st century Cures Act era. JAMA 2020;323 (23):2381–2382

63 Majumder MA, Guerrini CJ, Bollinger JM, Cook-Deegan R, McGuire AL. Sharing data under the 21st Century Cures Act. Genet Med 2017;19(12):1289–1294

64 Pageler NM, Webber EC, Lund DP. Implications of the 21st Century Cures Act in pediatrics. Pediatrics 2021;147(03):e2020034199

65 Gordon WJ, Mandl KD. The 21st Century Cures Act: a competitive apps market and the risk of innovation blocking. J Med Internet Res 2020;22(12):e24824

66 Dzau VJ, Cho A, Ellaissi W, et al. Transforming academic health centers for an uncertain future. N Engl J Med 2013;369(11): 991–993

67 Horwitz LI, Kuznetsova M, Jones SA. Creating a learning health system through rapid-cycle, randomized testing. N Engl J Med 2019;381(12):1175–1179

68 Friedman C, Rubin J, Brown J, et al. Toward a science of learning systems: a research agenda for the high-functioning Learning Health System. J Am Med Inform Assoc 2015;22(01): 43–50

69 Crandall W, Kappelman MD, Colletti RB, et al. ImproveCareNow: The development of a pediatric inflammatory bowel disease improvement network. Inflamm Bowel Dis 2011;17(01): 450–457

70 Marsolo K, Margolis PA, Forrest CB, Colletti RB, Hutton JJ. A digital architecture for a network-based learning health system: integrating chronic care management, quality improvement, and research. EGEMS (Wash DC) 2015;3(01):1168