Reducing Alert Burden in Electronic Health Records: State of the Art Recommendations from Four Health Systems

John D. McGreevey III[1,2] Colleen P. Mallozzi[1] Randa M. Perkins[3] Eric Shelov[4] Richard Schreiber[5]

[1] Office of the CMIO, University of Pennsylvania Health System, Philadelphia, Pennsylvania, United States
[2] Section of Hospital Medicine, Division of General Internal Medicine, Perelman School of Medicine at the University of Pennsylvania, Philadelphia, Pennsylvania, United States
[3] H. Lee Moffitt Cancer Center and Research Institute, Tampa, Florida, United States
[4] Division of General Pediatrics, Department of Biomedical and Health Informatics, Children's Hospital of Philadelphia, Philadelphia, Pennsylvania, United States
[5] Physician Informatics and Department of Medicine, Geisinger Health System, Geisinger Holy Spirit, Camp Hill, Pennsylvania, United States

Address for correspondence John D. McGreevey III, MD, 3400 Spruce Street, Philadelphia, PA 19104, United States (e-mail: John.mcgreevey@pennmedicine.upenn.edu).

Appl Clin Inform 2020;11:1–12.

Abstract

Background Electronic health record (EHR) alert fatigue, while widely recognized as a concern nationally, lacks a corresponding comprehensive mitigation plan.

Objectives The goal of this manuscript is to provide practical guidance to clinical informaticists and other health care leaders who are considering creating a program to manage EHR alerts.

Methods This manuscript synthesizes several approaches and recommendations for better alert management derived from four U.S. health care institutions that presented their experiences and recommendations at the American Medical Informatics Association 2019 Clinical Informatics Conference in Atlanta, Georgia, United States. The assembled health care institution leaders represent academic, pediatric, community, and specialized care domains. We describe governance and management, structural concepts and components, and human–computer interactions with alerts, and make recommendations regarding these domains based on our experience supplemented with literature review. This paper focuses on alerts that impact bedside clinicians.

Results The manuscript addresses the range of considerations relevant to alert management including a summary of the background literature about alerts, alert governance, alert metrics, starting an alert management program, approaches to evaluating alerts prior to deployment, and optimization of existing alerts. The manuscript includes examples of alert optimization successes at two of the represented institutions. In addition, we review limitations on the ability to evaluate alerts in the current state and identify opportunities for further scholarship.

Conclusion Ultimately, alert management programs must strive to meet common goals of improving patient care, while at the same time decreasing the alert burden on clinicians. In so doing, organizations have an opportunity to promote the wellness of patients, clinicians, and EHRs themselves.

Keywords

- clinical decision support
- alerting
- alert fatigue
- safety
- governance

received
August 19, 2019
accepted after revision
November 12, 2019

DOI https://doi.org/10.1055/s-0039-3402715.
ISSN 1869-0327.

Background and Significance

Electronic health record (EHR) alert fatigue, while widely recognized as a concern nationally, lacks a corresponding action plan for management.[1–4] This manuscript synthesizes several approaches and recommendations for better alert management, derived from four U.S. health care institutions that presented their experiences and recommendations at the American Medical Informatics Association (AMIA) 2019 Clinical Informatics Conference (CIC) in Atlanta, Georgia, United States. The assembled health care institution leaders represent academic, pediatric, community, and specialized care domains.

There is increasing national attention on the impact of EHRs on clinician wellness, with potential threats including usability limitations, imposition of functions and processes that do not correspond to actual clinical work, onerous documentation and data entry requirements, and a perception of excessive, burdensome alerting experienced by frontline clinicians, among others.[5–12] Nearly one in four medication orders generate an alert.[13] While alerts can change clinician behavior and improve care processes, clinicians largely dismiss them, an action known as overriding.[14] Overrides may be clinically appropriate, such as when a clinician deems the likely benefit of administering a medication to far exceed the potential medication risks. In other cases, overrides may represent not carefully considered clinical decisions, but reflexive dismissals by clinicians who have become inured to the large number of EHR alerts. EHR alert override rates are as high as 96%[15] and the override rates of drug-allergy alerts have increased over time.[11]

Several authors have considered the reasons behind these high override rates. More than two-thirds of all drug-allergy alerts presented to clinicians were for non-life-threatening allergic reactions,[11] raising the question of whether these alerts were important enough to have shown to clinicians in the first place. How much of a threat to a patient's health does a medication or other therapy need to pose to warrant interrupting a clinician's workflow with an alert? Only immediately life-threatening reactions? A low chance of temporary reversible harm? Perhaps something in between? In addition, almost one-third of medication alerts shown to primary care physicians (PCPs) in one study were repeats of alerts that had fired for the same patient in the last year.[16] These findings illustrate (1) the persistent and repeated decisions (appropriate or not) by clinicians to ignore alert guidance, suggesting the guidance is not clinically helpful nor likely to alter clinical management, and (2) the inability of current-state EHR systems to adapt alert-firing based on prior end-user actions.

The more alerts one experiences, the more likely one is to ignore them.[17] Clinicians overrode drug allergy alerts appearing two or more times more frequently than drug allergy alerts that appeared only once.[11] Reminder alerts to perform certain tasks were also less likely to be heeded with increasing numbers of alerts presented[16] as were responses to clinical trial recruitment alerts.[18]

Alerts can have significant costs. Notably, there is an opportunity cost to the time it takes clinicians to process alerts. McDaniel studied nearly 26,000 drug-drug interaction (DDI) alerts and found that the median time spent processing those alerts was about 8 seconds.[19] Schreiber and colleagues estimated the hourly cost of a physician at between $108 (U.S. dollars [USD])) and $234 per hour.[12] Using the lower value yields a physician time cost of about $0.03 per second; multiplying by 8 seconds gives an approximate cost of $0.24 per physician per DDI alert. This cost increases to $0.52 per DDI alert using the higher value for physician time. With alert override rates typically above 90%, it is reasonable to believe that much of the time physicians spend engaging with alerts represents lost productivity. Given that the above calculations are derived from DDI alerts, it is not possible to extrapolate the lost productivity cost related to other types of alerts. However, the cumulative cost to health care organizations across many thousands of alerts is still likely to be substantial.

Clinicians perceive alert fatigue. Peterson and Bates define alert fatigue as a "condition in which too many alerts consume time and mental energy to the point that both important warnings and clinically unimportant ones can be ignored."[20] In a survey of 2,590 PCPs, 86.9% reported the alert burden was excessive and more than two-thirds indicated the number of alerts they received was more than they could manage effectively.[21] PCPs report experiencing a median of 63 alerts per day.[21] The perception of, and not the actual alert burden, was positively associated with burnout among PCPs.[22] Improving EHRs will require far more than measuring alert performance if we are to address alert effectiveness and burden.[23]

What has been done to manage alerts better, to increase their effectiveness, and reduce the burden that alerts pose to clinicians? Targeted deactivation of alerts deemed to be of low quality or low effectiveness is one approach.[24] In another study, 18 DDI pairs were downgraded to be no longer visible to clinicians (e.g., warfarin and enoxaparin).[12] This reduced the DDI alert rate by 10.3%, yet the alert override rate remained 96.7%. Another approach is for organizations to implement alerts according to severity settings that their drug knowledge vendors provide. However, such an approach can lead to variation in alert implementation, as drug knowledge vendors define alert severity, classify drugs, and categorize DDI differently.[25] Other authors have advocated running alerts silently as a starting point, so that the alerts are not visible to clinicians.[13] In doing so, organizations can gather data on alert performance to inform implementation decisions. Still others have promoted the notion of adaptive clinical decision support (CDS) that learns user behavior over time and has the capability to filter alerts to which a clinician has previously responded.[26] While current-state EHRs may allow individual users some discretion over the alerts that they wish to continue receiving, it is not clear how many EHRs, if any, may have this built-in adaptive alert filtering capability at present.

Another approach to effective alert management is to follow expert guidance about alert configuration and settings. Payne and colleagues advocated for the consistent use of seven core elements in DDI-alert presentations.[27] Despite the publication of high- and low-priority DDI alerts that should be interruptive and that should be demoted to noninterruptive status, respectively,[10,28] a later study revealed that even within the same vendor product no two EHR instances had the same alert settings.[25] There remains wide variability in alert settings and no prevailing standard of care about how to implement alerts.[25]

Why is managing alerts effectively so difficult? A variety of factors may contribute including regulatory mandates, public reporting initiatives, liability concerns, and other external pressures that may incline institutions to advocate for more rather than fewer alerts to avoid preventable harm.[29] In addition, variations in clinicians' drug knowledge and experience can make alerts that are appropriate for one clinician, perhaps a July intern, inappropriate for another, such as an experienced attending physician. Lastly, there is system inertia where alerts, once created and deployed, may accumulate over time and compete for ongoing attention and resources with new, higher health system priorities.

Several unanswered questions deserve researchers' attention. Is there an optimal override rate and if so does it vary depending on the type of alert and other contextual factors? Is there an absolute or relative reduction in the number of alerts that fire that can reduce the alert override rate? What is the best metric for evaluating the effectiveness and appropriateness of alerts? What guiding principles should determine whether alerts should be interruptive versus noninterruptive, hard stops versus soft stops? Are there CDS alternatives to alerts that could be widely scalable to achieve safety goals? What is the best way to lighten alert burden without jeopardizing patient safety and ideally in a way that demonstrably improves patient safety?[30] This paper explores the approaches found in the literature and that the authors, all of whom are clinicians, have found of benefit to answer these questions whenever possible, although much more research is needed.

Alert Governance

There is no prescriptive form of CDS governance that assuredly works for every institution.[4,31] What works for one may fail gloriously for another. An institution's organizational culture generally predicts whether a given governance structure works for it. There are numerous governance structures, worthy of a focused literature review. As an overview, most large, integrated delivery health care systems govern in a top-down approach,[31] often called hierarchical governance. Smaller institutions may use this approach but often use a consensus-driven model, usually termed either spoke and wheel governance or star network, where anyone can say "no," but no one person can say "yes."[32]

Regardless of the approach to governance, we recommend that alert governance include regularly scheduled reviews of all alerts that breach the organization's established thresholds, discussed in more detail below. Alerts should not be removed automatically but presented to a committee for review and consideration for removal or revision (if possible). Regular review will gradually prune the existing alerts with the threshold then adjusted as needed for more specific fine-tuning. In cases where alerts are turned off, vigilance for unintended consequences is an important goal, recognizing that identifying any safety event, let alone a safety event clearly attributable to deactivation of a single alert, is known to be extremely challenging.[33]

There is no widely accepted definition of successful alert governance. In the absence of one, it may be reasonable to define governance success as a process for evaluation and decision making that enables an organization to achieve predefined goals for its clinical decision support systems.

Governance Structure and Foundational Processes

On what do information technology (IT) departments and end-users agree? There is little argument about the need for the individuals and groups shown in ►Table 1 to participate in alert governance and to be part of alert decision making from the beginning.

Despite the necessity of these governance participants, they are accompanied by sociotechnical domains which may impede effective governance as shown in ►Table 2.

Governance structures evolve over time, sometimes as the result of changing organizational needs or leadership. Health care integration, too, can impact governance structures. For example, as community hospitals become members of larger health systems, the culture of governance often changes from local- to system-level control. A previously consensus-driven structure often yields to top-down governance, compelling institutions to reflect on and refine governance structures. Moreover, resources may expand with mergers in such a way

Table 1 Key participants in alert governance[a]

Subject matter expert groups	Role as participant in governance
Clinicians Advanced practitioners Nurses Pharmacists Physicians	Define clinical goal of alerts May overlap with informatics specialists, if available
Informatics Nursing Pharmacy Physician	Expertise regarding knowledge management Arbiters between clinicians, technical staff, and administrators Broad cross-sectional knowledge of clinical and technical domains
Information Technology Analysts, builders Data scientists, researchers Education leaders Human factors engineers Optimization staff	Inform decision makers what is possible from a technical capability standpoint Subject matter experts for issues as they arise
Administration Legal staff Regulatory Safety officers	Define personnel, budgetary, time resources

[a] Not all institutions have the breadth or depth of all participants. If available, each has a unique role.

Table 2 Sociotechnical domains which may impede effective governance

Domains	Issues	References
Physicians Other clinicians	Limited availability to participate Lose interest rapidly May not receive compensation	Authors' experience
EHR analysts	May not know system limitations early in project May be few knowledgeable analysts available in health IT market for some applications Failure to recognize alert errors and anomalies	Authors' experience[34,35]
Budget	Total cost of ownership not always clear Allotments may not be realistic	[36,37]
Administration	Impaired or interrupted institutional knowledge due to: Frequent hospital leadership turnover (avg.: 5.5 years) Frequent medical executive leadership turnover (often 1 year) Organizational priorities unlikely focused on system optimization Prioritization of other projects Risk intolerance variability	[38]
Regulatory	Rapid cycle changes (e.g., meaningful use, ICD-10)	Authors' experience

Abbreviations: avg., average; EHR, electronic health record; ICD, international classification of diseases; IT, information technology.

that facilitates better governance. More informatics resources, however, may become available to community hospitals after health system integration.

A strong governance process adheres to a defined strategy and sets metrics and specific goals. Decision making accelerates when explicit rules exist regarding prioritization, metrics, and accountability.[39] In addition to making decisions about creation and removal of alerts, each operational governance body will have to weigh the value of alert optimization against other optimization efforts that may compete for the same resources. In the interest of efficiency, each governance committee should have the authority to make decisions about the topics presented to them, rather than serving in a purely advisory role to other decision makers.

When attempting to evaluate existing alerts, there may be staff available with the institutional knowledge of who created alerts, and why and when, who may be able to shed light on prior processes. Alert assessments should be performed in a nonjudgmental manner, recognizing that prior parties created the current alerts with good intentions and felt doing so was important. The original impetus for alert creation must be addressed, especially when removing an alert.

Governance includes setting up a categorization of alerts into standard groups, such as medication alerts or alert recipients which facilitates alert build processes, evaluation, and maintenance.[40,41] Prioritization models used by clinical informaticists, focusing on patient safety, revenue, and clinical workflow optimization, have been published.[42]

It is not uncommon for groups who will not receive an alert to request it on another group's behalf. Regardless of the requestor and respecting the intent of the request, a facilitator should clarify the goal and try to determine the original problem that the requestor is trying to resolve. Ideally both the requestor and the alert receiving party (if different) should be involved in a collaborative discussion to seek resolution of the original problem, with or without an alert.

New alert requests require justification. First, prior to creating an alert, there should be a system search for any alert or other CDS tool which already meets the intent of the proposed alert. Second, alert requests should be vetted by an informaticist (or by whatever means an organization determines) to determine whether an alert or reminder is the appropriate tool with which to achieve the desired objective. Once these two prerequisites are met, the request may come before an appropriate governance group. Consider use of an intake form which asks the submitter to delineate information consistent with the five rights of CDS.[43]

If an alert is the best solution to the issue and meets build criteria, we recommend following the guidance in ►**Table 3** before building the alert. In addition, ideal alerts meet our proposed seven CREATOR rules for alerts, also shown in ►**Table 3**.

After alerts have been implemented, health care organizations must recognize that a clinician who overrides an alert is making an individual patient level decision, while at the same time creating an important signal for the organization. Organizations should appreciate the value of these messages and seek to learn from them.[44] While we do not know of specific override rate goals, each organization should determine what alert override rates should prompt further alert revision, changes in alert actions, or further clinical education.

Following the guidance to "start with the end in mind,"[45] ideally governance design should welcome feedback about alerts from end-users. There should be a process to understand how alerts function in the production environment (see sections "Alert Metrics" and "Thinking of Alerts like Diagnostic Tests"), whether the result is as expected (see section "Designing Alerts"), and whether users find the alerts helpful, discussed further in section Case Examples of Alert Maintenance and Reduction "Geisinger Health System."

Description of a Governance Structure at Geisinger Health System

At Geisinger Health System, there is a system Chief Medical Informatics Officer (CMIO) and three associate CMIOs who report to the system CMIO. Separate from the CMIO team,

Table 3 Author-recommended checklist to justify alert appropriateness

Alert justification criteria	Questions to determine appropriateness
Considerations prior to build	
Problem identification	What problem will the alert solve?
Beneficial	Is there a clearly defined return on investment (e.g., increased screening referrals)? Will alert reduce potential adverse events?
Appropriateness	Is the alert consistent with clinician workflow?
No better alternative	Is there an existing clinical decision support that accomplishes the same thing? Is there a less intrusive mechanism that may succeed?
Not so complex as to inhibit system performance	Does testing reveal system malfunction or slowing?
Metrics defined	How will alert success be measured?
Scheduled review of alert	When are the first and subsequent dates of review?
Consistent with organizational strategy and principles	Is the alert compatible with institutional policy, financial goals, and clinician workflows?
Considerations for alert build	
CREATOR: 7 new rules for ideal alerts:	
Consistent with organizational strategy and principles	Addresses high priority goals, adheres to established alert guidelines
Relevant and timely	Appropriate for clinical workflow impacted
Evaluable	Predefined metrics
Actionable	Allow delete or modify of triggering orders from within alert
Transparent	Rationale of alert is clear, highlighting patient-specific data which triggered the alert
Overridable	Clinical workflows may not be predictable by alert designers; clinician may be presented with a scenario which exceeds the alert designers' foresight
Referenced	Citing literature as appropriate, supporting intent of alert

there are unique structures for nursing and pharmacy informatics. These groups collaborate regularly at conjoint informatics huddles, optimization meetings, and collaborative informational and educational meetings. Numerous ad hoc committees are created as different issues arise, such as order set optimization for enhanced postoperative recovery.

Geisinger has revamped its process for managing alert requests. Requestors must first obtain managerial approval. A team of nursing and physician informaticists and analyst staff then triages requests for new or altered alerts for assignment to one of seven major governance committees: alerts and reminders, documentation, orders and order sets, interface, optimization, usability, and education. Importantly, each of these groups is authorized to make decisions regarding rejection, acceptance, or altering of the requests. For more complex issues involving more than one constituency, or when there is disagreement, the request rises to the appropriate optimization group.

While this structure appears to function well in this large, integrated, dispersed, complex medical system, those creating governance in other settings must be cognizant of local culture, existing reporting structures, and mindful of prior governance that succeeded or failed. At Geisinger, this governance has been more successful than prior governance versions by decreasing the number of pending requests, accelerating the time to request completion, and increasing the ability to create system-wide consistency in alerting.

Governance is challenging yet essential for operational efficacy and efficiency. Informatics teams help, but are not always available, especially in smaller community hospitals. Participation by noninformatics staff is difficult, time-consuming, and often unsustained. At all institutions, governance structures evolve over time.[4]

Alert Metrics

Quantitative alert assessment is difficult in that there is no agreed upon measure to assess alert effectiveness and burden. Consider metrics such as the total number of alerts or an alerts per orders ratio, number of interruptive versus noninterruptive alerts, alert override rate, or time required to act on an alert.[12,19] Whether these metrics or some combination are the appropriate indicators of alert effectiveness and burden is unclear. Ideally, outcome metrics for each alert would establish the alert's clinical impact but measuring clinical outcomes directly attributable to alerts is challenging.[46] In the absence of outcome measures, process measures may still be beneficial. ►**Table 4** is a compendium of alert metrics which the authors have found in the alert literature.

One intriguing metric distinguishes alert effectiveness and efficiency. The former is the number of patients for whom the alert's intended action was taken divided by the total number of patients on whom the alert fired (what proportion of patients on whom the alert fired did it prompt the desired action?), while the latter is the number of alerts on which the desired action was taken divided by the total alerts fired (how many times did the alert fire on all patients to achieve its purpose?).[47] Values for effectiveness and efficiency could be very similar but could diverge when an alert fires multiple times for one clinician for a single patient or when the alert fires to multiple clinicians for a single patient.

Another complication with alert metrics is that there is not a uniform definition for what counts as an alert override. There is considerable commentary in the CDS literature that high override rates are unacceptable, yet there is no agreement about a proper rate and it is unclear if the override rate is even a valid measure of alert effectiveness. The authors' experience in the case of DDI alerts[25] showed that there were several options

Table 4 Alert metrics with definitions, advantages, and disadvantages

Parameter	Definition	Advantages	Disadvantages
Override rate	% of alerts dismissed[a]	Easy to calculate	No clear baseline or desirable target goal
Acceptance rate	% of alerts where user selected suggested action	Easy to calculate	No clear baseline or desirable target goal
Volume of alerts	Total number of alerts fired	Easy to calculate	Crude metric
Alerts/patient	Ratio of total alerts divided by number of patients on whom alerts fired	Easy to calculate	Crude metric that depends on illness severity, types of medications and orders
Alerts/clinician	% of alerts fired for an individual clinician	Easy to calculate; may offer ability to compare clinicians with similar types of patients	Comparing clinicians with different patient populations may not be meaningful
Alerts/patient/day	Ratio of total alerts divided by number of patients on whom alerts fired per day	An improved assessment of daily workload	Comparing different patient populations may not be meaningful
Alerts/clinician/day	% of alerts fired for an individual clinician	An improved assessment of daily workload	Comparing clinicians with different patient populations may not be meaningful
Alerts/orders entered	Proportion of total orders entered on which an alert fired	Another analysis of workload burden	Comparing different individual's order burden may not be meaningful
Alerts/order session	Average number of alerts fired per order session	May offer an alternative view of workload burden	Comparing different individual's order burden may not be meaningful
Alerts/specific order item	Alerts that fire on a given order	Helpful for analysis of specific orders where clinicians see frequent alerts	No clear baseline or desirable target goal
Harm occurred when clinician overrode true positive alert	Number of alerts correlated with adverse events reports	Perhaps the most pertinent in terms of PPV	Very difficult data to obtain
Dwell time[19]	Time that an alert is on screen	Easy to obtain metadata	May not reflect total cognitive load of alert
Think time[12]	Total time between firing of alert and resolving it	Easy to obtain metadata, only analyzes time of dealing with alert	May be an overestimate if user unfamiliar with screens, or underestimate for those who dismiss alerts rapidly
Effectiveness[47]	Proportion of patients on whom an alert fired where the clinician chose the alert's intended action	Measures success of alert on a per patient basis at achieving intended action	May be difficult to quantify if alert not actionable or several acceptable choices are possible; May require chart review
Efficiency[47]	Proportion of alerts for which the clinician chose the intended action	Measures burden of alerts required to fire before intended action occurs	May be difficult to quantify if alert not actionable or several acceptable choices are possible; May require chart review
Number needed to alert or prompt	Number of times alert needs to fire to elicit the intended action	A variation on effectiveness and efficiency, but measures total alert firing rate, rather than per patient	As a number, and not a ratio, may suffer from differing denominators
Outcome	Many possible definitions	Optimal metric for effectiveness	Proving cause (the alert) and effect (outcome) is challenging; concurrent and evolving changes create bias and confounding

Abbreviation: PPV, positive predictive value.
[a] Definitions of override vary.

that were equivalent to an override action. Clinicians could click on a button called "override" to reject the alert guidance, with or without indicating a reason for doing so. Alternatively, clicking on the "X" dismisses the alert as does clicking "continue," but in neither case is there an indication of the clinician's intention other than to skip the alert. For data analysis purposes, it is worth noting that some available override reasons may not even be relevant to the alert displayed.[48] In addition, dismissal of an alert does not always mean that the user ignored the advice; rather, the user may sometimes enter an appropriate order later. These examples illustrate the complexity of interpreting alert metrics.

Thinking of Alerts like Diagnostic Tests

Whenever possible, informaticists and others designing alerts and CDS tools need to think of them as they would a diagnostic test with true and false positives and negatives. Decision support relies on positive and negative criteria and whether the patient truly meets the condition in question (e.g., sepsis). These data are necessary to calculate the familiar test characteristics of sensitivity, specificity, and positive/negative predictive value (PPV/NPV; ►**Table 5**). Importantly, these test characteristics apply to the alert performance itself (whether it fired appropriately or not), rather than to the clinician's response to the alert (such as to accept, override, or dismiss).

The consensus is that current alert configurations are overly sensitive, with rampant false positives and subsequent pervasive alert fatigue.[16] Where the field has struggled, and perhaps introduced the greatest potential harm, is the failure to incorporate PPV into the design of CDS and to consider the level of control (LOC). Level of control is the degree to which the alert is attempting to alter clinical decision making. For example, interruptive alerts which require entry of an override reason are much more controlling than those that are dismissible with one click. While a higher PPV is always a goal, this can be challenging, especially in the case of conditions with low prevalence or where the severity of outcomes, such as sepsis or cardiac arrest in the pediatric population, warrants a higher alerting LOC.

The key principle is that the PPV should align with the LOC. Without this alignment, which is typically low PPV and high LOC, there are two prominent risks: first, an often incorrect alert with onerous requirements contributes to alert fatigue, and second, a clinician may heed a false positive alert with high LOC (i.e., strongly recommending something) and take the wrong action for the patient, again potentially leading to harm. Over time, informaticists have concluded that alerts will always be more successful when the design and format, specifically the level of control, align with its test characteristics, perhaps most importantly the PPV.[24]

Designing Alerts

There are several types of interruptive alerts. A complete hard stop prevents the user from proceeding (e.g., trying to prescribe isotretinoin for a pregnant woman). Partial hard stops require that one cannot proceed without supplying required elements, while soft stops require the user to pause, even if data entry is not required. Soft stops are less controlling but may still contribute to end-user perceptions that all alerts can be a nuisance.

Noninterruptive or advisory alerts do not interfere with the user's workflow but may not be seen and clinicians may ignore them more readily. Language of an alert can also play a role in LOC, depending on the strength of the verbiage regarding the recommendation in the alert.

There is growing interest in novel designs of alert appearance which may improve user-computer interaction when applied to alerts.[49–51] Although this is a promising area of research, a full discussion of these topics is beyond the scope of this paper.

Alert Testing

There are limitations in the ability to test alerts, most notably that scenarios in testing environments typically represent a small fraction of the potential clinical variations that can trigger an alert in an environment with actual patient data. Moreover, scenarios for testing almost always are limited to anticipated true positive and negative behavior. This risks underestimation of false positive and negative alerts. Although

Table 5 Calculation of an alert performance measure

		Patient clinical condition		Performance measures
		Condition present	Condition absent	
Alert behavior	Condition criteria triggered (alert fired)	(A) True positive: condition identified	(B) False positive: condition incorrectly identified	Positive predictive value: A/(A + B)[b,c]
	Condition criteria not triggered (no alert)	C) False negative: condition missed	(D) True negative: condition truly absent and alert correctly did not fire	Negative predictive value: D/(C + D)[b]
	Performance measures	Sensitivity: A/(A + C)[a]	Specificity: D/(B + D)[a]	

[a] Not affected by condition prevalence.

[b] Affected by condition prevalence.

[c] Easiest to calculate since both numerator and denominator are based on easily retrieved alert firing data. All other performance measures have either true or false negatives, both of which require methods like manual chart review or retrospective query of validated cohort criteria to calculate.

feedback, after the alert is live, can lead to adjustments that improve an alert's performance, this is only after the potentially poorly performing alert has been live in the system with the associated risks that it could cause unintended patient harm and negative impressions on clinicians.[17] Finally, the task of correcting a problematic alert postdeployment can be time sensitive and is less likely to enable a thoughtful work environment for informaticists relative to the more planned and deliberative conditions predeployment.

Is there a better way to test alerts? The evidence is growing that alert testing can and should be performed with real, dynamic patient data, and fully functional interfaces.[52] Actual patient data rather than scripted testing scenarios will help to refine and improve the alert criteria and design. Once alert criteria are optimized, final test characteristics, specifically the PPV, should update alert formatting elements, such as LOC and language. There are two possible approaches that alert designers can take to achieve more rigorous level of testing and alert performance evaluation: retrospective analysis and background deployment.

Retrospective Analysis

Retrospective alert analysis is a relatively high-resource requiring method. It involves application of potential alert criteria using retrospective data. This approach offers several clear advantages. First, it can apply the logic of an alert to a very large cohort. Second, one can derive all four test characteristics (sensitivity, specificity, PPV, and NPV) for the alert, if there is a validated cohort group of patients who truly meet the criteria of interest (►Table 5). In retrospective analysis, determining a specific cohort of patients defining a "true condition" may be a distinct analytics project, since the cohort definition data may not necessarily be the same as the criteria in an alert. For example, the criteria used for sepsis screening of a general population (vital signs, certain laboratory tests and assessment documentation) may differ from those describing a true sepsis cohort using data that may not be present at the time of screening (end-organ damage and positive culture data), but is present in the retrospective analysis, which will impact the PPV and NPV.

An additional programming challenge is that the degree of data manipulation necessary in retrospective analysis is greater than the manipulation and tools available in commercial EHRs. As such, informatics teams must guard against building a retrospective alert that is impossible to build and deploy for use with actual patients.

Background Deployment

"Background deployment" or "silent mode" means activating an alert in the live environment without making it visible to clinicians while recording all potential firings. This "lower tech" approach offers many of the same advantages of retrospective analysis while avoiding some of the limitations. It allows testing of the alert in a setting with similar limitations to those encountered by deployed, clinician-visible alerts, for example, erroneous data entries and lack of final cultures or final billing codes.

The primary limitation of the background approach, especially when compared with retrospective analysis, is the time it may take to gather enough data to perform an adequate performance analysis. While a retrospective query will have a large number of patient visits immediately available, a background alert must be left to run for some period of time for initial analysis and any subsequent criteria refinement steps. For high-volume alerts, this may be of little practical consequence, whereas for rare conditions, it could be a significant impediment to analysis. Fortunately for the purposes of minimizing alert fatigue and false positives, typically even a limited period of background analysis yields significant insights that can significantly reduce the potential alert burden.[53]

Either of these approaches to "going live before go-live," in which teams can include in the alert planning phase accurate performance data, is strongly recommended, in particular for alerts where complex logic is involved. Both can then provide real-time prospective data which when combined with other variables, such as clinical severity of the targeted condition, can enable the configuration of an alert that is more likely to achieve its goals and avoid unintended consequences.

Case Examples of Alert Maintenance and Reduction

Geisinger Health System

Geisinger Health System has over 1,500 alerts and reminders in its Epic EHR (Verona, Wisconsin, United States) which it installed in 1996 (ambulatory) and 2006 (inpatient). More recently Geisinger contracted with a third party CDS software company (Stanson Health, Sherman Oaks, California, United States) which supplies its own CDS, as well as analyzes currently installed alerts, including those in silent mode. Stanson supplies and Geisinger reviews alert statistics regarding firing rates, override percentages, acceptance rates, alert comments, and other vital alert data.

Armed with this data, Geisinger turned off alerts that users always override or ignore, or which violate one or more of the five rights.[43,54] Comparing monthly alerting rates between January 2018 and January 2019 Geisinger reduced the absolute number of active alerts fired to nurses from 1,674,429 to 763,132 (54%) and for physicians, from 630,690 to 511,705 (19%). Alerts that were turned off included those directed to incorrect users or that were poorly timed (e.g., reminding nurses to obtain a patient's smoking history before the nurse had a chance to complete the nursing intake interview) or that included guidance that was inconsistent with current workflow.

Penn Medicine

Penn Medicine is a large academic and community-based institution that has undertaken efforts to optimize alerts. Penn Medicine also uses the Epic EHR and calls this optimization initiative EHR wellness, the practice of continuously analyzing the performance and efficacy of clinical decision support and other tools to assure that they are functioning appropriately and supporting clinician workflow as intended. The goal of this iterative process is to eliminate noisy and burdensome alerts that cause cognitive load on ordering providers, nurses, and pharmacists while optimizing the important and necessary alerts, all while continuing to support patient care. The EHR wellness program addresses interruptive and noninterruptive care-guidance alerts, medication alerts, and order sets. Only 14% of physicians find that they have the time they need to provide the highest standard of care,[55] which served as a driving force for the EHR wellness campaign. The campaign aims to proactively guide providers and make it easier to do the right thing at the right time in the EHR.

The Penn Medicine EHR wellness team includes operational and clinical leaders, technical analysts, and informaticists from the CMIO team. The team uses a third-party platform (Phrase Health Inc., Philadelphia, Pennsylvania, United States)[56] that provides detailed performance data on EHR alerts, which complements the EHR vendor-supplied reporting tools to pinpoint the largest areas of opportunity. ►**Table 6** lists the goals of the team.

The initial focus of the optimization efforts targeted 17 of the most "burdensome" alerts that accounted for nearly 1.7 million alerts/month across the health system. At Penn Medicine, an alert is considered potentially burdensome if it fires >100,000 time per month, if the alert has an elevated average number of interruptive alert firings per day for the population that is exposed to it relative to other alerts,[56] or if the alert has inconsistent build/design according to institutional standards. Once troublesome alerts were identified, a detailed review of individual alert settings was performed to assess for elements, such as acknowledgment option consistency compared with other alerts, ease of jumping to the intended action from the alert itself, and alert triggers. A fundamental question guiding the process was, "is this alert even necessary?" Three months of analysis and alert optimization resulted in complete removal of three of the most burdensome alerts and editing of the remaining 14 alerts. These changes resulted in the reduction of interruptive alerts by 67,863 alerts/month (45%), and overall alerts by 251,505 alerts/month (15%).

In August 2017, Penn Medicine initiated a secondary effort to evaluate EHR medication alert performance specifically. Baseline alert data revealed 675,613 alerts per month that users overrode 94.6% of the time. The goal was to safely reduce unnecessary medication alerts by 3 to 5% and provide more effective guidance to ordering providers, pharmacists, and nurses. Pharmacy residents, under the supervision of the Director of Pharmacy, conducted a literature review, analyzed the evidence available to support medication alerts, such as drug–drug and dose-range alerts, and made recommendations for which of the 20 alerts that fired most often should be continued, edited, or retired. These recommendations gained approval from key medical, nursing, and pharmacy leaders. The results of these combined efforts exceeded original expectations. Overall medication alerts decreased by 23%. The number of alerts per 100 orders dropped by nearly 34% as shown in ►**Table 7**.

Of note, these efforts only minimally impacted the override rate for medication alerts, a conundrum previously noted.[12]

Discussion

Managing alerts within an EHR is a complex undertaking, with notable considerations being organizational history, expectations, and governance related to alerts. Complicating management is the lack of a widely accepted metric to judge the effectiveness and burden of alerts. New, scalable methods of evaluating alerts and responses to them are necessary[57] but this requires further research. High-alert volume can lead to alert fatigue which contributes to increased mental workload, potential patient harm via workarounds, and mistakes in ordering and treatment.[17,57,58]

There are no nationally-developed, endorsed standards for which alerts are appropriate and which are not, which ones should be interruptive and which ones should be passive. Some proposals are gaining acceptance, such as the list of DDIs that should be interruptive.[28] Standardization of alert nomenclature would enhance understanding and promote better research. This may also assist vendors to coordinate their alert design to enhance crossvendor comparisons. The authors believe that new standards, such as Fast Healthcare Interoperability Resources (FHIR), will enable enhanced CDS in general and alerting in particular by sharing successful processes across institutions. Machine learning will also offer new approaches to improve alerting through analysis of huge datasets.

Notwithstanding these limitations organizations have taken proactive steps toward optimizing alerts with some early success. Organizations can begin an alert optimization program by evaluating alerts with high firing or override rates, or those assessed to be burdensome on clinicians. Doing so will likely uncover alerts that are relatively less valuable and that may be optimized for better effectiveness or alternatively, deactivated. More work is needed to understand at what level of alert reduction clinicians respond more appropriately to the guidance of remaining alerts. When considering deployment of new alerts, analysis of alert performance prior to go-live can

Table 6 Guiding principles of electronic health record wellness

1. Correct design inconsistencies and tailor alerts to meet the needs of the target population
2. Engage directly with impacted clinicians to redesign workflows (user-centered design/optimization)
3. Make all alerts actionable: assure the ability to jump directly to the intended action
4. Set standards for inclusion/exclusion logic across all care settings so that alerts do not impact unintended areas or users
5. Review trigger actions, align acknowledgment reasons, streamline verbiage
6. Standardize analyst capture of metadata when alerts are changed, to assure a reliable record of alert adjustments and the reasons for them, as well as the routine review of the alerts during change control

Table 7 Results of medication alert reduction

	Medication alerts (per mo)	Number of alerts (per 100 orders placed)	Number of overridden alerts (per mo)	Override rate[b]
July 2017	675,613	55.4	581,958	94.6%
August 2018	521,005[a]	36.8	445,088	92.9%
Difference	−154,608	−18.6	−136,870	−1.7%
Overall change	23% reduction in all medication alerts	34% fewer alerts per 100 orders	24% reduction in overridden alerts	~2% reduction in overall override rate

[a] An additional approximately 215,000 orders placed (per month) attributed to an additional hospital going live on the EHR.
[b] Override rate only includes unfiltered alerts.

improve PPV and design before any clinician ever experiences the alert.

Conclusions and Recommendations

Alert management programs must strive to meet common goals of improving patient care while at the same time decreasing the alert burden on clinicians. In doing so, organizations have an opportunity to promote the wellness of patients, clinicians, and EHRs themselves. There are multiple components to ensure a successful alert management program:

- Governance is complex but essential infrastructure for effective alert management.
- Organizations should conduct ongoing analysis and review of alerts.
- Absent an agreed-upon optimal metric for analyzing alert performance, each organization must select metrics appropriate for itself.
- For guidance regarding implementation of an alert management program, look to organizations that have been successful and have reported their experiences.
- New design paradigms, data and alert visualization displays, and emerging technologies also offer promise for improved alerting.

Clinical Relevance Statement

Electronic health record alerting provides tools for clinical decision support and can help clinicians to provide improved care, while also preventing medical errors. Yet there is widespread agreement that over-alerting leads to alert fatigue, with the subsequent risks of potential patient harm and clinician burnout. This paper presents the analysis and recommendations for mitigation of this problem from informatics leaders from four major health care organizations which may provide useful guidance for small and large as well as community and academic institutions.

Multiple Choice Questions

1. Recognizing that an alert is a form of one-way communication (system to user), what is an objective way to measure its effectiveness?
   a. Committee review and discussion to modify or delete alerts.
   b. Reports showing the rate and changes of alert firing over time.
   c. Total percentage of alerts which elicit the intended action.
   d. User narrative feedback to modify, delete, or add alerts.

   Correct Answer: Option c is the most objective and direct measure. The other answers have a higher potential for bias or false attribution of effect.

2. Institutions with a long history of clinical decision support (CDS) and those newer to creating governance for CDS, including alerts, share the same struggles regarding the scope and complexity of the task. The most critical factor to assure success in alert governance is:
   a. Establish best infrastructure for alert installation.
   b. Manage and lead change transformation processes.
   c. Purchase of the newest technology from an outside vendor.
   d. Reward staff for meeting or exceeding goals.

   Correct Answer: The best option is b because use of CDS depends on supporting the workflow requirements of users. Although solid infrastructure is necessary, it is not sufficient. The latest and greatest technology is only as good as the ability of personnel to use it. Rewarding staff has merit but is also not sufficient.

3. Interruptive alerts should be used for critical decision-making processes. Which statistical measure is the most helpful in determining whether an alert should be interruptive?
   a. Negative predictive value.
   b. Positive predictive value.
   c. Sensitivity.
   d. Specificity.

   Correct Answer: The best option is b. Several factors can determine the format of a CDS intervention, but of the statistical measures positive predictive value (true positives/all positives) is most useful in determining whether CDS should be interruptive and prescriptive.

Protection of Human and Animal Subjects

No human subjects were involved in this project and Institutional Review Board approval was not required.

Conflict of Interest

None declared.

References

1 Gawande A. Why doctors hate their computers. The New Yorker. Available at: https://www.newyorker.com/magazine/2018/11/12/why-doctors-hate-their-computers. Accessed October 18, 2019

2 Jha AK, Iliff AR, Chaoui AA, et al. A crisis in health care: A call to action on physician burnout. Available at: https://cdn1.sph.harvard.edu/wp-content/uploads/sites/21/2019/01/Physician-BurnoutReport2018FINAL.pdf. Accessed October 18, 2019

3 Schulte F, Fry E. Death by 1,000 clicks: Where electronic health records went wrong. Kaiser Health News. Available at: https://khn.org/news/death-by-a-thousand-clicks/. Accessed October 18, 2019

4 Kawamanto 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

5 Monica K. 5 ways to prevent physician burnout in the age of the EHR System. Available at: https://ehrintelligence.com/news/5-ways-to-prevent-physician-burnout-in-the-age-of-the-ehr-system. Accessed October 18, 2019

6 Howe JL, Adams KT, Hettinger AZ, Ratwani RM. Electronic health record usability issues and potential contribution to patient harm. JAMA 2018;319(12):1276–1278

7 Wears RL, Berg M. Computer technology and clinical work: still waiting for Godot. JAMA 2005;293(10):1261–1263

8 Downing NL, Bates DW, Longhurst CA. Physician burnout in the electronic health record era: Are we ignoring the real cause? Ann Intern Med 2018;169(01):50–51

9 Verghese A. How tech can turn doctors into clerical workers. The New York Times Magazine. Available at: https://www.nytimes.com/interactive/2018/05/16/magazine/health-issue-what-we-lose-with-data-driven-medicine.html. Accessed October 18, 2019

10 Phansalkar S, van der Sijs H, Tucker AD, et al. Drug-drug interactions that should be non-interruptive in order to reduce alert fatigue in electronic health records. J Am Med Inform Assoc 2013;20(03):489–493

11 Topaz M, Seger DL, Slight SP, et al. Rising drug allergy alert overrides in electronic health records: an observational retrospective study of a decade of experience. J Am Med Inform Assoc 2016;23(03):601–608

12 Schreiber R, Gregoire JA, Shaha JE, Shaha SH. Think time: A novel approach to analysis of clinician's behavior after reduction of DDI alerts. Int J Med Inform 2017;97:59–67

13 Saiyed SM, Greco PJ, Fernandes G, Kaelber DC. Optimizing drug-dose alerts using commercial software throughout an integrated health care system. J Am Med Inform Assoc 2017;24(06):1149–1154

14 Silbernagel G, Spirk D, Hager A, Baumgartner I, Kucher N. Electronic alert system for improving stroke prevention among hospitalized oral-anticoagulation-naïve patients with atrial fibrillation: A randomized trial. J Am Heart Assoc 2016;5(07):e003776

15 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

16 Ancker JS, Edwards A, Nosal S, Hauser D, Mauer E, Kaushal R; with the HITEC Investigators. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. BMC Med Inform Decis Mak 2017;17(01):36–44

17 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

18 Embi PJ, Leonard AC. Evaluating alert fatigue over time to EHR-based clinical trial alerts: findings from a randomized controlled study. J Am Med Inform Assoc 2012;19(e1):e145–e148

19 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

20 Peterson JF, Bates DW. Preventable medication errors: identifying and eliminating serious drug interactions. J Am Pharm Assoc (Wash) 2001;41(02):159–160

21 Singh H, Spitzmueller C, Petersen NJ, Sawhney MK, Sittig DF. Information overload and missed test results in electronic health record-based settings. JAMA Intern Med 2013;173(08):702–704

22 Gregory ME, Russo E, Singh H. Electronic health record alert-related workload as a predictor of burnout in primary care providers. Appl Clin Inform 2017;8(03):686–697

23 Payne TH. EHR-related alert fatigue: minimal progress to date, but much more can be done. BMJ Qual Saf 2019;28(01):1–2

24 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

25 McEvoy DS, Sittig DF, Hickman T-T, et al. Variation in high-priority drug-drug interaction alerts across institutions and electronic health records. J Am Med Inform Assoc 2017;24(02):331–338

26 Lee EK, Wu TL, Senior T, Jose J. Medical alert management: a real-time adaptive decision support tool to reduce alert fatigue. AMIA Annu Symp Proc 2014;2014:845–854

27 Payne TH, Hines LE, Chan RC, et al. Recommendations to improve the usability of drug-drug interaction clinical decision support alerts. J Am Med Inform Assoc 2015;22(06):1243–1250

28 Phansalkar S, Desai AA, Bell D, et al. High-priority drug-drug interactions for use in electronic health records. J Am Med Inform Assoc 2012;19(05):735–743

29 Carspecken CW, Sharek PJ, Longhurst C, Pageler NM. A clinical case of electronic health record drug alert fatigue: consequences for patient outcome. Pediatrics 2013;131(06):e1970–e1973

30 Grissinger M. Medication errors involving overrides of healthcare technology. Pennsylvania Patient Safety Advisory 2015;12(04):141–148

31 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

32 Alexander B. Schreiber R. WakeMed Health & Hospitals, Raleigh NC. Personal Communication

33 Classen DC, Resar R, Griffin F, et al. 'Global trigger tool' shows that adverse events in hospitals may be ten times greater than previously measured. Health Aff (Millwood) 2011;30(04):581–589

34 Wright A, Ai A, Ash J, et al. Clinical decision support alert malfunctions: analysis and empirically derived taxonomy. J Am Med Inform Assoc 2018;25(05):496–506

35 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

36 Koppel R. Is healthcare information technology based on evidence? Yearb Med Inform 2013;8:7–12

37 Koppel R, Lehmann CU. Implications of an emerging EHR monoculture for hospitals and healthcare systems. J Am Med Inform Assoc 2015;22(02):465–471

38 Khaliq AA, Thompson DM, Walston SL. Perceptions of hospital CEOs about the effects of CEO turnover. Hosp Top 2006;84(04):21–27

39 Blenko MW, Mankins M, Rogers P. The decision-driven organization. Available at: https://hbr.org/2010/06/the-decision-driven-organization. Accessed October 18, 2019

40 Kuperman GJ, Bobb A, Payne TH, et al. Medication-related clinical decision support in computerized provider order entry systems: a review. J Am Med Inform Assoc 2007;14(01):29–40

41 Murphy DR, Reis B, Sittig DF, Singh H. Notifications received by primary care practitioners in electronic health records: a taxonomy and time analysis. Am J Med 2012;125(02):209.e1–209.e7

42 Longhurst C, Sharp C. Implementation and transition to operations. In: Payne TH, ed. Practical Guide to Clinical Computing Systems:

Design, Operations, and Infrastructure. 2nd ed. Amsterdam: Elsevier; 2015:99–110
43 Schreiber R, Knapp J. Premature condemnation of clinical decision support as a useful tool for patient safety in computerized provider order entry. J Am Geriatr Soc 2009;57(10):1941–1942
44 Aaron S, McEvoy DS, Ray S, Hickman T-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
45 Covey SR. The 7 Habits of Highly Effective People. New York, NY: Simon & Schuster; 2013
46 Middleton B, Sittig DF, Wright A. Clinical decision support: a 25 year retrospective and a 25 year vision. Yearb Med Inform 2016 (Suppl 1):S103–S116
47 Truong Q. Strategies to the five rights of clinical decision support. Epic User Web 82831. Verona, WI: Epic Systems Corporation
48 Wright A, McEvoy DS, Aaron S, et al. Structured override reasons for drug-drug interaction alerts in electronic health records. J Am Med Inform Assoc 2019;26(10):934–942
49 Horsky J, Phansalkar S, Desai A, Bell D, Middleton B. Design of decision support interventions for medication prescribing. Int J Med Inform 2013;82(06):492–503
50 Feblowitz J, Henkin S, Pang J, et al. Provider use of and attitudes towards an active clinical alert: a case study in decision support. Appl Clin Inform 2013;4(01):144–152
51 Dowding D, Merrill JA. The development of heuristics for evaluation of dashboard visualization. Appl Clin Inform 2018;9(03):511–518
52 Wright A, Aaron S, Sittig DF. Testing electronic health records in the "production" environment: an essential step in the journey to a safe and effective health care system. J Am Med Inform Assoc 2017;24(01):188–192
53 Ashton M. Getting rid of stupid stuff. N Engl J Med 2018;379(19): 1789–1791
54 Hussain MI, Reynolds TL, Zheng K. Medication safety alert fatigue may be reduced via interaction design and clinical role tailoring: a systematic review. J Am Med Inform Assoc 2019;26(10):1141–1149
55 Hasan H. Combating physician burnout: Five insights to help restore the balance. Available at: https://www.advisory.com/research/medical-group-strategy-council/white-papers/2016/combating-physician-burnout?WT.mc_id=Web|Vanity|LB|PhysicianIssues|2017Jan26|Burnout|. Accessed August 10, 2019
56 Phrase Health; analytics and governance for decision support. Available at: https://www.phrasehealth.com. Accessed October 18, 2019
57 Kane-Gill SL, O'Connor MF, Rothschild JM, et al. Technological distractions (part 1): summary of approaches to manage alert quantity with intent to reduce alert fatigue and suggestions for alert fatigue metrics. Crit Care Med 2017;45(09): 1481–1488
58 Health Research & Educational Trust; U.S. Department of Health and Human Services. Implementation guide to reducing harm from high-alert medications. Available at:http://web.mhanet.com/ade_changepackage_508.pdf. Accessed October 18, 2019