



Technology and Health Care Improvement

8

Jessica Herstek and Eric Shelov

Health information technology (IT) encompasses multiple overlapping electronic hardware and software systems that enable clinical services across health care settings today. Federal inventive programs propelled the adoption of electronic health records (EHRs) and other medical software systems in a majority of hospitals and clinics over the past decade. Much focus during this time was on implementation, getting up and running on electronic equivalents of preexisting paper processes. While many types of errors and inefficiencies have been eliminated in this phase of expanding use, new risks and unintended consequences have emerged. Clinicians, informaticists, and IT developers, engineers, and companies are now partnering as health IT evolves to maximize its potential to support safer and high-quality health care for patients and decreased burden and frustration for clinicians.

J. Herstek (✉)

Department of Pediatrics, The George Washington University School of Medicine and Health Sciences, Children's National Hospital, Washington, DC, USA

e-mail: jherstek@childrensnational.org

E. Shelov

Department of Biomedical & Health Informatics, Division of General Pediatrics, Perelman School of Medicine, University of Pennsylvania, Children's Hospital of Philadelphia, Philadelphia, PA, USA

e-mail: shelov@chop.edu

The Next Phase

Gone are the days when IT was just another department in the fluorescent-lit basement of a hospital or office building. Health IT systems now underpin every clinical care activity, from the way doctors and nurses communicate and document care to the automated tools that are intended to add layers of safety to inherently risky and complex processes and procedures. EHRs are the foundational systems that capture the bulk of clinical data created by the minute and hour in the course of care, though health IT today comprises a multitude of parallel interfaced devices and systems.

Even the most basic EHRs were implemented in only around 10% of hospitals in 2009, the year that the federal HITECH Act was passed to incentivize adoption, but was approaching 100% a decade later [1]. Given the transformational potential of this trend, this has been a phenomenal accomplishment. We understand intuitively now the safety of legible prescriptions, the convenience of electronic communication, and the value of instant access to information for both patients and health care professionals. But what was not broadly appreciated until more recently was the double-edged sword that digitization and automation represent.

EHRs are plagued by a litany of complaints from health care workers and patients alike. Poor usability, cognitive overload, alert fatigue, and automation complacency are examples of health hazards that technology has introduced. Counterintuitive user interfaces stem from a lack of user-centered design, where clinicians have input on engineering decisions early in the project to make the system work best for them. Propagation of information overload leads to dangerous phenomena such as alert fatigue, where clinicians become conditioned to ignore meaningless yet interruptive signals from the system, at the risk of missing actual critical signals. Even when systems are perceived as highly reliable, an overreliance on technology can result in automation bias and operator complacency, an inappropriate degree of trust in the system over human input or common sense. Today's EHRs and

related health IT products offer innumerable advantages over legacy systems and have been an integral part of improving patient safety. Yet even as they have addressed many inherently unsafe processes by digitizing handwritten medical records, electronic systems introduce these and other new safety risks into health care.

Health IT systems can and should be leveraged as the robust safety and quality improvement tools they are. But no matter how much time, effort, and money are spent on implementation, failure to recognize that no electronic or automated system is intrinsically, infallibly safe can have startling consequences.

In Error

A teenaged boy holds a cup of pills in his hand. He places a pill in his mouth, swallows, and takes another. Then another and another. Then a handful at a time. He doesn't stop until he has swallowed 39 pills.

He is not alone. He is with a young woman who watches him closely. She doesn't leave until she sees that he took all 39 pills. Doctor's orders.

The woman is his nurse. The boy is in the hospital. His mother is nearby, in another room with his younger brother, who is also sick in the hospital on this particular night. The pill is an antibiotic the boy takes every day at home to prevent infections. He has a genetic condition that affects his immune system, and the pill usually helps keep him out of the hospital. At home he takes a single pill twice a day, every day. He was not admitted to the hospital for an infection this time – the pill had been doing its job. He was there for a routine procedure.

Later that night, the boy notices numbness and tingling all over his body. Then, after texting with a friend, he suddenly screams for his mother. His body goes stiff, limbs shaking, jaw clenched, and back arched against the hospital bed. His breathing stops. A “code blue” is called. He survives the seizure but is transferred to the intensive care unit, where he will have to remain in the hospital much longer than planned.

The 39 pills were an unintentional overdose ordered by the resident physician, approved by the pediatric pharmacist, handed over by the bedside nurse, and dutifully taken by the boy. Between each of these – by all accounts – competent and caring health care professionals and the patient lay an elaborate electronic safety system, meant to remove all the points at which such errors could transpire. This case occurred at a prestigious academic medical center, within reach of Silicon Valley, that had invested heavily in technology to further patient safety [2]. Hundreds of millions of dollars in equipment and staff time spent training and typing and clicking and scanning conspired to create a nonsensical medication error.

Necessary But Not Sufficient

This type of massive overdose would almost certainly have not occurred in the era before widespread use of health IT systems. The particular circumstances that allowed the error to occur stemmed directly from the multiple layers of technology in place, as detailed by Dr. Robert Wachter in his examination of this case in his book, “The Digital Doctor: Hope, Hype, and Harm at the Dawn of Medicine’s Computer Age” [2]. The advent of technology in health care was a necessary surge forward in patient safety, but clearly, the use of high-tech tools is not sufficient to prevent harm to patients nor adequate to support clinicians in their work.

Many doctors practicing today trained in an era of handwritten notes, hand-signed orders, and scribbled prescriptions. In this world, a doctor 24 hours in to a 30-plus-hour shift could hastily scrawl an incomplete order for a medication on a piece of paper, which was then faxed to the pharmacy, yielding faint scratchings that translated to a drug or dosing error. Pharmacy staff could grab the wrong bottle off the shelf before counting out pills and sending them to the unit. A nurse could give a dose intended for an obese adult patient to the small child down the hall instead. These flagrant errors all could have happened and did with frightening regularity. We can now prevent these in part with the installation of complex and expensive technology meant to keep patients safe,

yet an unprecedented overdose still reached a patient in the current era. These types of errors continue to occur with regularity across health care settings. Implementation and optimization of EHRs and related systems are a necessary but not sufficient component of any patient safety program.

A joint investigation conducted by Kaiser Health News and Fortune magazine reviewed the federal government's decade-long, 36-billion-dollar incentive program to convert medical records from paper to electronic. This conversion promised freedom from all the limitations, inefficiencies, and well-documented safety issues in previous decades that could be traced back to reams of paper. In the scathing review, the reporters cataloged medicolegal cases implicating EHRs. Cases are detailed of missed or delayed diagnoses due to EHR errors with catastrophic consequences. The EHR vendors in these cases blamed clients for user error, inadequate training, or improper setup of their systems. Hospitals blamed EHR companies for poor visual layout and unintuitive user design. "It can be hard to tell where human error begins and the technological shortcomings end" [3]. Physicians and nurses in general don't have any better views of EHRs. Surgeon and writer Atul Gawande penned an article about EHRs in 2018 in *The New Yorker* called "Why Doctors Hate Their Computers," in which he notes the frustrations, limitations, and information overload many clinicians experience working with current systems [4].

There is no question that electronic systems save lives, but thoughtful design, careful implementation, and continuous quality improvement are required to prevent or mitigate unintended consequences and novel error types, too.

Closing Loops

The process of ordering medications in any hospital is fraught with potential for harm. Whether paper or electronic, it involves dozens of steps. Each step represents an opportunity for error, but also an opportunity for assistance from technology. Medication processes that can be facilitated by EHRs and related systems, and

are recommended by the federal government, include computerized provider order entry, clinical decision support, and bar-coded medication administration systems [5, 6]. Primarily over the last decade, hospitals spent tens or even hundreds of millions of dollars to purchase and implement computerized provider order entry systems. These systems finally eliminated the risk of misinterpretation of a doctor's handwriting, loss of paper order sheets, or poor fax transmission quality to the pharmacy leading to wrong drug or dose errors, among a myriad of other well-documented issues. Additionally, rather than starting with a blank sheet of paper, and consulting a pocket reference or other external dosing guide, an ordering provider can be presented with appropriately calibrated choices with built-in guidance specific to that patient and clinical scenario in order to support their decision-making. This system of computerized knowledge is referred to as clinical decision support. A third component towards safer, "closed-loop" medication processes entails bar-coded medication administration tools. With these tools, nurses at the patient bedside can scan bar codes or other machine-readable tags imprinted on medication packaging and an identification bracelet on the patient, prior to administering the drug. The system cross-checks the patient and drug with the original medication ordered by the physician in the EHR. The final step confirms that the correct drug was delivered and is administered to the intended patient.

In addition to a robust EHR with computerized provider order entry, clinical decision support, and bar-coded medication administration modules, the hospital where this massive overdose occurred had also spent millions of dollars on a pharmacy robot to select, dispense, and deliver routine medications. This investment freed the pharmacy staff to focus on more nuanced work that requires human attention. The combination of these systems embedded in the EHR, the pharmacy robot, and the bar-coded medication administration process were meant to prevent errors at each of the major inflection points in medication ordering, dispensing, and administration. In other words, this hospital had put in place all the latest recommended technology systems to provide closed-loop medication administration, starting from correct order entry and ending with giving the right medication to the right patient.

Despite all these high-tech safety systems, errors are still widely reported. An identification bracelet can be lost or placed on the wrong patient, a pharmacist or manufacturer can mislabel a medication bag or vial. A nurse can – and sometimes must – find a workaround to administer a fluid or medication to the patient without delay when the bar code scanner is not connecting to the system via Bluetooth. All the issues with electronic systems we all run into in our own daily lives happen with health IT, too. Wires get unplugged, batteries run out, mobile units go missing, systems freeze, and a wireless pairing fails. Yet the delivery of health care must proceed, with or without these safety supports in place.

So what went wrong in this case? There was in fact no spectacular technologic failure, but a combination of issues with people, process, and technology that culminated in the dozens of pills reaching the patient. The introduction of each new layer of safety comes with both benefits and risks. In fact, technology can create novel and sometimes unanticipated sources of error. In this case, poor usability, alert fatigue, and automation bias were complicit. These phenomena are essential for quality and safety teams to understand, as they must be considered when reviewing safety incidents involving health IT systems, as well as when leveraging health IT to address patient safety issues or quality improvement efforts. The introduction of EHR systems designed, historically, by non-clinicians has led to disastrous examples of the limits of technology to eliminate errors altogether. As with any complex, high-risk, and evolving system, a continuous quality improvement approach should be taken to implementing and maintaining health IT systems, as the hardware and software as well as user training and cultural norms will inevitably change over time.

Calibrating Trust

The doctor who ordered the overdose had actually entered the order correctly in the computer system earlier that day. But because the dose calculator was overly precise, there was a rounding issue. And because there was a rounding issue, the reviewing

pharmacist, following hospital policy, called the doctor back to reenter the correct, rounded dose. This time, the doctor entered the order a slightly different way, with little visual indication of a mistaken unit of measure. The ease with which this doctor committed an extreme prescribing error on a commonly used medication in the EHR can be blamed on poor usability of the order interface. A well-designed user interface can prevent or at least deter simple mistakes in order entry. After signing the incorrect order, she did, however, receive a pop-up alert warning of the overdose, as did the pharmacist. They both clicked habitually through these warnings, due to the frequency of such alerts that were often meaningless. This is attributable to alert or notification fatigue – the concept that the quantity, design, and calibration of alerts in any system can dramatically alter cognitive processing and response to the information being presented. Based on their past experiences with unhelpful alerts, both the doctor and the pharmacist were unintentionally trained to mistrust the EHR’s alerts, even though in this case it was giving them critical feedback on their actions.

Clearly, all the built-in knowledge embedded in the EHR can backfire, manifesting as alert fatigue among clinicians, with real consequences on patient safety and clinician satisfaction. Articles in the medical literature describe techniques to address alert fatigue and poor usability of EHRs. In one paper, researchers from Harvard Medical School created an algorithm to use “cranky comments” that doctors typed in responses to pop-up alerts in the EHR to detect programming errors [7]. The prestigious *New England Journal of Medicine* published a hospital system’s popular EHR improvement campaign titled “Getting Rid of Stupid Stuff,” detailing their use of clinician input to address serious problems such as alert fatigue [8]. The phenomenon of alert fatigue is now so well-recognized that it is called out in the non-profit ECRI Institute’s annual report of Top 10 Health Technology Hazards [9].

To mitigate this risk, the process of designing alerts and other electronic tools can be improved by focusing on usability. Usability refers to the design of the user interface of a system,

whether a button on a car dashboard or in a hospital's EHR. Products or systems with good usability are intuitive, not requiring hours or days of training, to start using and use correctly. They are efficient, not engendering frustration by requiring multiple clicks or interruptions that are of low value or yield to the user. For EHRs, usability addresses how well the system helps health care workers complete their tasks and how well its user interface balances efficiency and safety by minimizing human error. This relates to both broad functionality and the detailed design of visual layout. The risk of unintentional selection of the wrong item, or of missing an important system prompt, is directly affected by design choices. Errors can stem from displaying too much information, requiring extensive scrolling, grouping items too close together, or using too small a font. One factor known to lead to poor usability is a lack of user-centered design or appreciation for the criticality of the human-computer interface.

Improving the usability of EHRs is not just an imperative to mitigate clinician frustration but also for patient safety. Many EHR vendors now employ human factors engineers and usability experts and leverage user-centered design methods that were often absent from earlier iterations of current EHR systems. Clinicians, too, have responded to the clear need to filter design and decision-making through the lens of usability. In fact, the need is so great that an entirely new specialty of medicine was created to train and certify physicians to work in this area. Clinical informatics became an official subspecialty of the American Board of Preventive Medicine in 2013 – the same year the boy described above was given an extreme overdose both despite, and because of, health IT systems in place. Physicians, nurses, and other health care professionals who work in informatics act as crucial partners with frontline clinicians, software engineers who design the systems, and hospital IT staff who are tasked with configuring and customizing these systems. Inclusion of informatics-trained clinicians in improvement projects can help the project team select and design appropriate EHR interventions, as well as take advantage of electronic systems to measure change.

Making the Right Thing Easy

Most fundamentally, the goal of clinical decision support is to present information to a clinician in real time and expect that they take action of some kind. At first pass this may seem straightforward, but there are many elements to be considered when developing decision support, and the consequences of not being thoughtful in this development can be significant. Medicine is hardly the first industry with technology attempting to assist in its productivity, and over the years clinical informaticists have learned valuable lessons and started to define and refine best practices. Perhaps the most important lesson learned in the health IT field is that out-of-the-box technology by itself is unlikely to solve a clinical problem. Rather, for effective solutions to be deployed, significant analysis of the problem at hand and the workflow involved are critical prerequisites to determine how best to incorporate technology tools like clinical decision support.

Relatively early in the field of clinical informatics, a group of researchers defined the “Ten Commandments” for successful clinical decision support [10]. These best practices include redirecting as opposed to stopping the user, recognizing the importance of speed in clinical work, and paying careful attention to workflow. Although this workflow analysis requires an investment of resources, it is nearly always worth the effort when the clinical decision support tool is built. Of note, there are excellent quality improvement tools that can play a vital role in this kind of analysis. Techniques such as driver diagrams and swimlane analysis can help represent the details of complex workflows that are needed before beginning the process of proper decision support development.

An additional framework that has evolved in the informatics literature is the concept of the “Five Rights” of clinical decision support [11]. The rights include the following:

- *Right information:* Evidence- or consensus-based, suitable to guide action.
- *Right person:* Including all members of the care team. Increasingly with electronic patient portals, this may also include patients and families.
- *Right time:* At the time of decision-making and desired action.

- *Right channel:* These may include both digital channels, such as the EHR, and non-digital such as signs at the bedside.
- *Right format:* Once a channel is determined, the right tool must be used, for example, an order set versus an alert.

These principles are helpful when considering a specific decision support tool, but during the design phase, it may be more practical to consider a question-oriented format. In the ideal scenario, standard tools are first used to analyze a problem, complete a workflow analysis, and identify necessary behavior changes. At that point, a series of questions can help to determine the right decision support approach. Who is the person or role that needs to change their behavior? What information do they need? When do they need it? What action do we expect them to take? These questions are the five rights of clinical decision support framed in a question-oriented format.

By answering these questions before committing to a particular approach in the EHR, quality improvement teams have a much better chance of choosing the right tool as opposed to one that is familiar but is poorly aligned with workflow. Unfortunately, this kind of analysis has historically not taken place when clinical teams design decision support. Instead, they are often developed quickly, in a reactive manner, jumping to common but blunt tools like pop-up alerts. Alerts developed in this manner tend to be ineffective in the long run, as the target audience quickly becomes accustomed to the alert's presence if it is not properly calibrated. Worse still, they propagate the sense that the EHR system is poorly designed, inefficient, and frustrating for busy frontline clinicians and, as we have seen, can lead to patient harm rather than prevent it as intended.

Clinical decision support systems can both save physicians clicks and time and steer them toward desirable choices or away from risky or otherwise undesirable prescribing behaviors. Their actions can even be forced, depending on how restrictive the design of the clinical decision support tool is. Linking in access to vast databases can support clinical decision-making by having the computer tap into knowledge resources and providing, for example, cross-checks on standard dose ranges, drug interactions with other medications, or allergies that may otherwise go unnoticed. More advanced displays provide the physician with access to the

latest clinical guidelines for their patient's condition or, better yet, implement it as the default pathway so that it requires extra effort and attention on the part of the physician to deviate from the standard of care. This is sometimes appropriate, and physicians may require and desire flexibility to customize patient care at their discretion. Often, however, patient safety and quality of care benefit from standardization where such approaches exist. In many cases, this can be readily accomplished with built-in templates of orders, termed order sets, which can help standardize care and promote education of providers at the same time. Order sets are perhaps the most frequently used decision support tool employed to support quality improvement efforts, with every modern EHR supporting this functionality in one way or another.

Order sets can most fundamentally be thought of as a collection of provider orders commonly grouped together to facilitate entry. On the less sophisticated side, an order set may simply be a reference for commonly used tests or medications for a particular scenario or clinical condition, such as treating an asthma attack or managing postoperative pain. In more sophisticated forms, order sets can represent a clinical pathway, algorithm, or decision tree with highly prescriptive guidance through each anticipated phase of care. Additional features provided by many EHRs include default selection of particular orders to encourage or even mandate their inclusion, nested groups of orders where a single selection leads to a cascade of additional selections, and the ability to hide or show orders based on available data about that particular patient, such as age and gender. Logic can be built in, for example, to automatically add a pregnancy test when indicated for female patients in an appropriate age range, but not display at all for male patients. This ensures extra protections for pregnant patients while decreasing the chances that a provider inadvertently orders the test when not applicable.

In the overdose case reviewed above, the ordering physician had to type in a specific numeric dose for the patient's home medication, leaving vulnerable several variables in each order, including unit of measure, intended strength, and formulation. Because this particular drug came in tablet form, with very standardized

dosing across almost all patients, forcing the doctor to enter these details anew every time adds unnecessary steps to the ordering process and increases opportunity for error. If the order fields for that medication had been pre-filled out to include only the numeric dose available per pill, she would have had less clicks to enter and less chance of doing it incorrectly. Better still, she could have selected the correct choice for that patient automatically from within an order set dedicated to, for example, patients with compromised immune systems who often take this medication every day. Moreover, modern EHRs can filter options being displayed, so in this case the system “knew” the patient was an adolescent and therefore could have predicted that he would take the standard adult dose of that medication. The default choice displayed to the doctor, such as a pre-checked dose, usually represents the path of least resistance. In other words, the easiest or most obvious available action on any view presented to the provider should contain the most common or correct option, raising the likelihood of a busy, stressed, distracted, or inexperienced provider entering a safe order. To maximize the benefits of order sets, careful consideration must be given to align the order set with the clinical decision-making process in an algorithm or pathway. The ease and effectiveness with which a clinical guideline can be translated into meaningful decision support will depend on its design and language. If a guideline or clinical pathway is full of ambiguity and points of indecision, it will be challenging to create an order set that clinicians find useful and are therefore willing to use [12].

There are many success stories of order sets that resulted in significant improvements in clinical care when built on the foundation of a thoughtful clinical pathway. This is facilitated by considering a pathway as a clinical decision tree, with the information presented at moments of decision to guide the clinician in the right direction for the patient. Order sets can be created to reflect this very same structure, with key data from the patient and pertinent reference information included right at the point of decision-making in the workflow. The provider then has all the information needed to make the correct or best decision for that patient. This approach has been taken to decrease the use of popular but overly

expensive drugs or tests, such as an intravenous form of the common fever medicine acetaminophen in children who can just as easily take it by mouth, or to make sure physicians order a blood test for certain patients prior to starting antibiotics. That test being ordered at the proper time provides crucial information to the medical team that cannot be obtained once antibiotics have been started. The order set prompts the physician to include it every time, instead of relying on them to remember to add it manually and risk missing a diagnosis that can help tailor the subsequent treatment plan.

On the other hand, there are also many cautionary tales where order sets do not align with a particular clinical pathway and result in unintended consequences. For example, a team designing a pathway may want to discourage the use of a particular test for a certain condition, due to expense, harm, or other downstream effects. In children who develop a common and uncomplicated respiratory infection, ordering a chest X-ray is often unnecessary to safely diagnosis and treat the condition. The X-ray not only adds cost to the parents' health care expenses but exposes the child to radiation without direct benefit and, further, could lead to overdiagnosis of otherwise benign conditions with a cascade of subsequent added costs, risks, and worries. The obvious approach may be to simply not include that X-ray order in an order set used for this type of patient. While this may certainly achieve some gains, it is likely some providers will go outside the order set to find the X-ray order and order it anyway. These providers would then deem that order set not useful since it does not contain all the orders they are accustomed to using. The next time they see a similar patient, they are less likely to use the order set and thereby miss out on the other benefits of streamlined and standardized care driven by evidence-based electronic pathways. One possible approach to prevent this from happening could be to actually include the X-ray order in the order set but display specific instructions or links to evidence-based resources supporting the rationale behind limiting its use. This serves not only to discourage inappropriate use of the study, but it also provides relevant education and real-time feedback to the provider who has been in the habit of using such a test indiscriminately.

The Signal and the Noise

In addition to well-designed orders and order sets, one of the other most common interventions that can support quality improvement work is an electronic alert. Alerts can be broadly defined as an automated flag or indicator meant to get the user's attention. They come in countless forms and representations such as icons, banners, or pop-up windows with text, figures, and buttons that may need to be clicked to bypass the alert and return to the original workflow. Alerts are a tool fraught with challenges in the EHR. While they hold much promise to change behavior, their overuse has led to significant frustration and unintended, potentially fatal, consequences in the EHR [13]. Health care is not alone in struggling with how to best deploy alerts. Aerospace and nuclear power industries have also learned hard lessons on the benefits and unintended consequences of alert design, particularly those that interrupt the thoughts and actions of a doctor, pilot, or other professional performing high-risk work. These streams constitute workflow, and its analysis is integral to designing useful alerts.

A well-designed alert constitutes a critical signal from the system to the clinician of a scenario or action with potentially dire consequences. When that signal is sent to the wrong person, at the wrong time, or through the wrong channel, it can be lost in the noise of all the other, less important alerts. Ideally, the EHR should serve as a trusted advisor to clinicians, delivering timely guidance and relevant suggestions integrated smoothly within their workflow. Electronic alerts have historically taken a more adversarial tone with clinicians, but today clinical informatics professionals and EHR vendors are actively trying to course-correct, to everyone's benefit.

Alerts have several characteristics to consider before use. They can be proactive, guiding clinicians towards the right thing, or reactive, stopping them when they have potentially done the wrong thing. They can be made to be intentionally interruptive, commonly in the form of a separate window popping up in a workflow, as opposed to non-interruptive, such as a banner or flag

that can be acknowledged at their convenience without interruption. Whenever possible, decision support should be proactive and non-interruptive, striving to make the right thing easy rather than penalizing the clinician for having done the wrong thing.

The consequences of poor alert design and implementation are important to recognize. Alert fatigue emerges when clinicians must bypass numerous insignificant or irrelevant alerts in the course of their usual work and then miss or unintentionally override those few alerts that may be most important. Alerts that contain incorrect or misleading information also cause harm when clinicians, accustomed to relying on alerts to catch errors, act on false information without independent verification. This is an example of automation bias, which increases the more accurate any clinical decision support system is. Well-designed alerts and other clinical decision support can decrease harm overall to patients, but clinicians should be educated on automation bias. Clinicians must remain vigilant in environments with high automation and should be trained to maintain a culture of safety and use clinical judgment in conjunction with clinical decision support systems. In the case of ordering medications in a computerized prescribing system, for example, prescribers should consider the decision support system a secondary, independent check on the dose or indication for a medication [14].

Before selecting an EHR alert as an intervention for a quality improvement project, consider the test characteristics, such as how often it is right and how often it may fire inappropriately. Consider the design, from use of color to font size to wording, and consider the workflow, where in the many cognitive and physical steps the optimal timing is for the alert to fire. The severity of the clinical scenario should dictate how forceful the alert is, distinguishing life-threatening situations from best practice or cost-associated concerns, and the action(s) required of the clinician to either accept, override, or bypass the warning. The text of the alert should be designed in collaboration with the frontline clinicians who are the intended recipient of that contextual information, alongside informatics-trained clinicians when available, rather than by the quality improvement or IT teams independently – this is an example of the application of user-centered design. This

approach helps achieve the desired goals of high usability and avoidance of unintended consequences. Even though changes can and should be made after go-live, the risks of inadequate initial design are high. In addition to immediate harm from use in the live patient environment, busy clinicians may get accustomed to the first impression made by a poorly performing alert, and their minds are not easily changed even if it is subsequently improved.

Even using good design principles, following up on the utility of the alert in the live environment remains as important as the initial design. The evidence is growing that clinical decision support testing can and should be performed “silently,” or more accurately, invisibly, in live EHR systems with real patient data and fully functional interfaces, instead of in an isolated test environment as has been done traditionally. The results of this testing can thereby better inform the alert criteria and design. The criteria can be refined and improved based on actual patient data rather than scripted testing scenarios. Once alert criteria have been optimized, the final test characteristics should inform elements including the degree of interruption and the language displayed to the clinician. Methods of quantifying alert performance have improved and can facilitate ongoing improvement cycles.

In the previous overdose example, both the physician and the pharmacist received a number of pop-up alerts in the process of ordering and approving the medication. The alerts were reactive, occurring after the erroneous order had already been signed, and interruptive, triggering the instinctive behavior to click through the screens as quickly as possible to resume patient care activities. Furthermore, there was little visual distinction between these critical – and correct – massive overdose alerts and innumerable other trivial alerts all staff received routinely through the usual course of care. This is evidence of lack of user-centered design and is the setting that gives rise to alert fatigue.

On the other hand, the nurse in this case had come to rely on the accuracy of the bar code scanning system. It emitted a reassuring audible and visual signal every time a medication was scanned successfully showing a match between patient, drug, and order, and no mismatch was detected in this case. That system had worked so well, in fact, that the nurse ignored common sense and

her own gut feeling that the dose was off. This natural tendency to trust automated systems despite evidence to the contrary is a manifestation of automation bias. In the everyday use of health IT systems, it is important to instill in all clinicians a healthy skepticism, especially during high-risk activities such as medication processing. Regardless of how automated or reliable an electronic system is, human behavior matters. Growing overreliance on technology, or automation complacency, is human nature. Maintaining a culture of safety becomes even more important with advanced technologies and automation.

Making It Count

Any quality improvement effort requires measurement, and informatics-based interventions are no exception. Informatics interventions can be assessed by the impact on structure, process, and/or outcome, with the ultimate goal of improving clinical outcomes. Not every informatics project can or should measure clinical outcomes, for example, if a specific process has clearly demonstrated tight linkage to a clinical outcome in previous research. This is particularly relevant for very rare events that may not occur with enough frequency at an individual institution to effectively demonstrate change. Process measures in the field of informatics have some unique challenges, however. The determination of what to measure is not always straightforward. As with any improvement effort, a discussion of what will be measured must be part of the planning for the intervention. As described above, the goal of many informatics interventions in the EHR are to change provider behavior in some way, so choices of what to measure boil down to the critical element of the workflow. In the current state of overburdened EHR users, informaticists strive to measure passively, leveraging actions clinicians take in the routine care of patients without introducing additional clicks or steps solely for the purpose of tracking. Some of these measures can be relatively straightforward, while others require greater sophistication for quality improvement project teams to assemble and interpret.

Orders are perhaps the easiest discrete unit to measure from the EHR and are a very reliable indicator of changes to patient care. Implementation of a pneumonia pathway, for example, could measure specific antibiotic usage in a defined cohort and determine whether or not the pathway achieved a particular goal in standardizing care. Nonmedication orders, such as those for specific nursing or supportive care, can also serve as useful metrics, but one must keep in mind that unlike medications, these orders may not always precisely correlate with the intended action. One type of order is nurse communication orders, which serve as standing free text instructions to nursing staff. An order that simply instructs a bedside nurse to “Apply vascular access care bundle,” for example, may not be the best process metric to determine whether or not the bundle was applied. Fortunately, there are other discrete data elements that can serve as process metrics. Flowsheet documentation, commonly used by nursing staff, is a reportable discrete data source. These charts indicating completion of bundle elements such as central line inspection, dressing changes, and flushes would be more clinically meaningful metrics and just as readily retrieved from the EHR database.

Other useful data, although somewhat more difficult to extract and interpret, to serve as process metrics include metadata (data about data), alert interaction, and documentation data. An example of metadata would be attributes of particular orders, such as time of day ordered (e.g., during morning rounds), or whether or not the user ordered it as a standalone order or as part of a specific order set. A team working to improve sepsis care may want to measure whether orders for antibiotics, blood tests, and intravenous fluids derived from a specific order set they designed to standardize sepsis care. Alert data may initially seem straightforward, simply looking at how often an alert appears and is either bypassed by the user, if this is permitted, or leads the user to change what they were doing, such as changing a drug dose after seeing an overdose alert. Considering the numerous ways non-interruptive alerts may appear, however, and the variable actions or responses a clinician can take subsequently, it becomes more challenging. Sophisticated alerts with complex criteria and multiple action options may be challenging to compare to one

another. Nonetheless, an informatics team may still find value in tracking metrics of a specific alert over time. Even with relatively straightforward interruptive alerts such as allergy or dose alerts, override rates are not always an accurate measure of alert performance.

Clinical documentation is another potential source of metric data, but it is one that presents unique challenges. Historically, clinical documentation was non-discrete and highly narrative. With the introduction of the EHR, there has been a shift towards discrete documentation with elements like mandatory fields and checklists. While this has been helpful for data capture, many in the medical field have lamented the loss of the narrative, arguably the most important part of medical care, cognitive processes, and learning. As it stands currently, documentation is largely a mixture of discrete and non-discrete elements. While the ability of natural language processing tools to extract reportable data from narrative text has advanced significantly in recent years, it remains a tool largely beyond the reach of most medical systems and clinical quality improvement teams. Thus, when physician notes are the only source of truth in the EHR for a particular question, such as a patient-reported symptom or the physician's thought process and decision-making, a manual review of notes remains the only option to leverage this kind of data.

More recently, techniques have emerged to create clinical documentation templates that provide data on the author's thought process. For example, a note template for pneumonia can be designed to offer select choices of text, rather than prompting free text entry, depending on clinical considerations such as severity of presenting symptoms, choosing from a drop-down menu of "moderate" or "severe," or their management plan referencing consideration of pneumonia. In many cases, the choices clinicians make as they complete the template can be recorded and used as a project metric. Embedded note data elements could help a quality team track how often physicians considered a diagnosis of pneumonia when patients present with severe respiratory symptoms.

Finally, to support continuous quality improvement efforts regarding the safety and usability of the EHR itself, informatics and quality improvement teams can now measure how clinicians are interacting with the EHR software more directly. Built-in software can monitor the time spent in a particular section of the chart, determine whether or not specific data was reviewed, or track the number of clicks spent on particular tasks. Unfortunately, this data can be quite challenging to work with. Not all EHR vendors make this data available, and when they do, it can be difficult to interpret without a deep understanding of EHR database structures. Increasingly this kind of data is being studied by informatics researchers to address systemic safety concerns that include documentation burden and alert fatigue.

Onward

The past decade of experience has illustrated the power and pitfalls of health information technology to readily implement clinical decision support interventions, including well-designed orders, order sets, alerts, and other tools to prompt or prevent targeted provider actions. Quality improvement teams can implement and enforce change quickly and broadly. EHRs support the measurement of successes and failures to support rapid-cycle change and iterative progress. With the majority of health care systems, from small private practices to large hospital networks, now using EHRs and complementary technologies, we are entering a new phase of more advanced thinking to design, use, and improve these all-encompassing systems proactively. This shift to EHRs and accompanying technologies impacts almost every facet of patient safety and quality of care. We all must bear in mind the principles of usability, mitigation of alert fatigue, and education and training to counterbalance the tendency toward automation bias and automation complacency. Technology solutions will continue to mature. Our approach to building and interacting with these systems must evolve, too, to address errors old and new.

References

1. The Office of the National Coordinator for Health Information Technology. Health IT dashboard. <https://dashboard.healthit.gov/quick-stats/pages/FIG-Hospital-EHR-Adoption.php>
2. Wachter R. The digital doctor: hope, hype, and harm at the dawn of medicine's computer age. New York: McGraw Hill; 2015.
3. Shulte F, Fortune FE. Death by 1,000 clicks: where electronic health records went wrong. Mar 18, 2019. <https://fortune.com/longform/medical-records/>
4. Gawande A. Why doctors hate their computers. The New Yorker Magazine. Nov 5, 2018. <https://www.newyorker.com/magazine/2018/11/12/why-doctors-hate-their-computers>
5. The Office of the National Coordinator for Health Information Technology. Health IT playbook. <http://healthit.gov/playbook/quality-and-patient-safety/>
6. Institute for Healthcare Improvement. Improve core processes for administering medications. <http://www.ihl.org/resources/Pages/Changes/ImproveCoreProcessesforAdministeringMedications.aspx>
7. Aaron S, McEvoy D, 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(1):37–43. <https://doi.org/10.1093/jamia/ocy139>.
8. Ashton M. Getting rid of stupid stuff. *N Engl J Med.* 2018;379(19):1789–91. <https://doi.org/10.1056/NEJMp1809698>.
9. ECRI Institute. Top 10 patient safety concerns 2020. <https://www.ecri.org/landing-top-10-patient-safety-concerns-2020>
10. Bates D, Kuperman G, Wang S, et al. Ten commandments for effective clinical decision support: making the practice of evidence-based medicine a reality. *J Am Med Inform Assoc.* 2003;10(6):523–30. <https://doi.org/10.1197/jamia.M1370>.
11. Agency for Healthcare Research and Quality. Improving medication use and outcomes with clinical decision support: a step-by step-guide. <https://digital.ahrq.gov/ahrq-funded-projects/current-health-it-priorities/clinical-decision-support-cds/chapter-1-approaching-clinical-decision/section-2-overview-cds-five-rights>
12. Li RC, Wang JK, Sharp C, et al. When order sets do not align with clinician workflow: assessing practice patterns in the electronic health record. *BMJ Qual Saf.* 2019;28:987–96. <https://doi.org/10.1136/BMJQS-2018-008968>.

13. Kolwalczyk L. Brockton hospital staff overlooked warning about a deadly allergy. 'Alert fatigue' may be to blame. Boston Globe. June 6, 2018. <https://www.bostonglobe.com/metro/2018/06/06/brockton-hospital-staff-overlooked-warning-about-deadly-allergy-alert-fatigue-may-blame/TAacermEyROxRdcdMvyCvM/story.html>
14. Lyell D, Magrabi F, Raban M, et al. Automation bias in electronic prescribing. *BMC Med Inform Decis Mak.* 2017;17:28. <https://doi.org/10.1186/s12911-017-0425-5>.