Optimization of drug–drug interaction alert rules in a pediatric hospital's electronic health record system using a visual analytics dashboard

RECEIVED 26 November 2013
REVISED 25 July 2014
ACCEPTED 9 August 2014
PUBLISHED ONLINE FIRST 15 October 2014

Allan F Simpao[1], Luis M Ahumada[2], Bimal R Desai[3], Christopher P Bonafide[3], Jorge A Gálvez[1], Mohamed A Rehman[1], Abbas F Jawad[4], Krisha L Palma[5], Eric D Shelov[3]

ABSTRACT

Objective To develop and evaluate an electronic dashboard of hospital-wide electronic health record medication alerts for an alert fatigue reduction quality improvement project.

Methods We used visual analytics software to develop the dashboard. We collaborated with the hospital-wide Clinical Decision Support committee to perform three interventions successively deactivating clinically irrelevant drug–drug interaction (DDI) alert rules. We analyzed the impact of the interventions on care providers' and pharmacists' alert and override rates using an interrupted time series framework with piecewise regression.

Results We evaluated 2 391 880 medication alerts between January 31, 2011 and January 26, 2014. For pharmacists, the median alert rate prior to the first DDI deactivation was 58.74 alerts/100 orders (IQR 54.98–60.48) and 25.11 alerts/100 orders (IQR 23.45–26.57) following the three interventions ($p < 0.001$). For providers, baseline median alert rate prior to the first round of DDI deactivation was 19.73 alerts/100 orders (IQR 18.66–20.24) and 15.11 alerts/100 orders (IQR 14.44–15.49) following the three interventions ($p < 0.001$). In a subgroup analysis, we observed a decrease in pharmacists' override rates for DDI alerts that were not modified in the system from a median of 93.06 overrides/100 alerts (IQR 91.96–94.33) to 85.68 overrides/100 alerts (IQR 84.29–87.15, $p < 0.001$). The medication serious safety event rate decreased during the study period, and there were no serious safety events reported in association with the deactivated alert rules.

Conclusions An alert dashboard facilitated safe rapid-cycle reductions in alert burden that were temporally associated with lower pharmacist override rates in a subgroup of DDIs not directly affected by the interventions; meanwhile, the pharmacists' frequency of selecting the 'cancel' option increased. We hypothesize that reducing the alert burden enabled pharmacists to devote more attention to clinically relevant alerts.

Key words: Electronic health records; drug interactions; medication alert systems; clinical decision support systems; medical order entry systems; visual analytics

RESEARCH AND APPLICATIONS

BACKGROUND AND SIGNIFICANCE

Computerized provider order entry in electronic health record (EHR) systems has been identified as one of the interventions with the greatest potential to reduce medication errors and associated harm in the pediatric inpatient setting.[1] Clinical decision support (CDS) during medication order entry, which is increasingly provided in the EHR, can reduce errors and harm throughout the prescribing process.[2] Interruptive medication alerts (ie, 'pop-up' alerts) are a commonly implemented form of CDS that notify users about potential adverse outcomes from drug allergies, over- and under-doses, and drug–drug interactions (DDIs).[3] Many hospitals elect to have all or a large percentage of available medication alert rules of all severities activated when they implement an EHR, all but guaranteeing a system with high sensitivity and low specificity. As a result, care providers (eg, physicians, nurse practitioners) and pharmacists receive frequent alerts, most of which are not clinically relevant. In response, providers and pharmacists appropriately override most alerts. However, providers and pharmacists who encounter high rates of medication alerts may exhibit alert

Correspondence to Dr Allan F Simpao, Department of Anesthesiology and Critical Care Medicine, Perelman School of Medicine at the University of Pennsylvania and the Children's Hospital of Philadelphia, 34th Street and Civic Center Blvd, Suite 9329, Philadelphia, PA 19104-4399, USA; simpaoa@email.chop.edu

For numbered affiliations see end of article.

fatigue, potentially missing or overlooking important, clinically relevant alerts that can prevent adverse drug events and harm.[4,5]

In response to this problem, we initiated a quality improvement project to reduce the volume of clinically irrelevant alerts. To facilitate alert analysis and evaluate our progress, we recognized the need for frequent 'self-service' comprehensive access to this complex and very large dataset in near real time, without having to continually rely on information services support staff to run static queries that would need to be combined manually. Thus, we devised an automated, user-friendly, data analysis software tool—a visual analytics dashboard—to facilitate both the operational work and subsequent analysis.[6]

MATERIALS AND METHODS

Setting

This cross-sectional study was performed between January 31, 2011 and January 26, 2014 at The Children's Hospital of Philadelphia, an urban, tertiary care children's hospital with 535 beds. The Institutional Review Board at The Children's Hospital of Philadelphia approved the study protocol. Our development team consisted of experts from pediatrics, clinical informatics, data analytics, and pharmacy.

The hospital implemented a new inpatient EHR (Epic Systems, Verona, Wisconsin, USA) in January 2011. Based on recommended interaction severity and documentation level filtration settings, approximately 3549 of the available 4033 DDI alert rules from a third-party vendor (Medi-Span, Indianapolis, Indiana, USA) were implemented for pharmacists along with the new EHR, while a substantially fewer 1351 rules were implemented for inpatient providers. Within each DDI described by the vendor is a set of rules that specify the medications that are included in that particular class as well as the formulations and routes. Many of these rules are dynamic, subject to periodic updates by the database vendor. After the implementation, care providers and especially pharmacists experienced a sharp increase in the number of medication alerts, some categories of which were overridden more than 90% of the time. This high false-positive rate was not surprising, as both the third-party medication alert and the EHR vendors recommended relatively high sensitivities by default (presumably as a safety measure). At our institution there was growing concern, however, that the frequent irrelevant alerts would impair pharmacists' and providers' ability to discern valid, clinically relevant alerts.

In our EHR, all alerts including allergy, maximum dose, duplicate medication, and DDIs for a single ordering session are displayed in one window. A single medication order can generate multiple alerts that are counted separately in this dataset even though they are delivered within the same window. For example, if five alerts in the context of one ordering session were presented in the same window, all five would be counted as individual alerts. Alert actions include 'override', 'cancel', 'remove', or 'view'. Alerts in the window can be overridden with a single button or the clinician can cancel the window to return to their ordering or verification session. 'Viewing' of alerts occurs when clinicians voluntarily elicit alerts that will fire for that patient should they proceed with signing or verifying. In our system override reasons are required for maximum dose and drug allergy alerts, but not for DDI and duplicate medication alerts.

Prescribers have all alert actions available to them. Within the alert window, a prescribing clinician has the option to 'remove' or discontinue medications contributing to the DDI, maximum dose, or duplicate alert. Alerts are also specific to the ordering session, so should the same alert be triggered for the same patient or prescriber at a later date in a separate session, it would fire regardless of what actions were taken previously.

In the pharmacists' workflow, local scope of practice does not allow pharmacists to enter or discontinue orders, so their actions are limited to override, view, and cancel. In addition, if interacting medications are being verified in the same session, a pharmacist will only see a DDI alert *once* for the pair. Pharmacists are also able to see the actions taken by prescribers when the alert appears in the verification process.

We recognized the risk of removing perceived safety measures from our hospital's EHR even if that risk was deemed less than that of widespread alert fatigue. Therefore, we consulted hospital stakeholders prior to establishing our medication alert quality improvement initiative: the CDS Committee, Patient Safety and Quality, Clinical Pharmacy, Medication Safety, Therapeutic Standards, Risk Management/Legal, Clinical Informatics Council, and the Enterprise Analytics and Reporting Group.

Dashboard development

We developed the secure, cross-platform, web-accessible, visual analytics dashboard in five steps. First, we wrote Structured Query Language queries to retrieve the necessary data elements from the EHR relational reporting database (Epic Clarity) within enterprise visual analytics software (QlikView, Radnor, Pennsylvania, USA). This allowed us to access all alerts as long as they occurred at least 24 h earlier. We merged over 50 conventional relational database tables into a single dimensional model.[7] Exploratory data analysis was the second step of the dashboard creation process. Visual summarizations (eg, histograms, box plots) of the data were created to improve our understanding of the data structure, as well as to perform data validation and integrity checks. The third step of dashboard creation was the data modeling process. Data transformation and further validation were performed concurrently during the third step. Counts of alerts and medication orders were validated using static vendor supported reports. The fourth step—information visualization—consisted of constructing a highly interactive user interface that adhered to seminal visualization principles in order to achieve graphical excellence and integrity.[8] This step included selecting the appropriate colors, minimizing 'chart junk' (eliminating distractors within each chart), determining the proper data density, and adhering to an aesthetically pleasing 'golden rectangle' interface design layout.[6] Lastly, the fifth step, usability testing, consisted of

end-users interacting with a 'live' dashboard and then providing feedback that influenced the dashboard's design and function.

Initially the dashboard was validated and used by our inpatient EHR leadership, who reviewed and discussed findings collaboratively with the CDS committee. The committee consisted of clinicians and pharmacy representatives who further explored and refined the approach. Preliminary clinical questions included determining medication alert volumes and override rates for each medication alert type (DDIs, allergy, maximum dose, and duplicate medication alerts), along with particular patient and clinician characteristics to explore relationships within the data. Based on this exploration, we prioritized DDIs, which comprised the highest volume of alerts in the inpatient EHR system, in particular for pharmacists.

Intervention phases

A June 2011 review of the dashboard prototype revealed that providers and pharmacists were experiencing a high prevalence of two DDI alerts for benzodiazepines and trace-ingredient ethanol. There was clear consensus to deactivate these clinically irrelevant alert rules for our first intervention (intervention 1). Going forward, the CDS committee concluded that a rigorous and systematic workflow was needed to further identify and deactivate clinically irrelevant DDIs. The CDS committee prioritized intervening upon the most frequently firing DDI alerts to achieve the greatest impact on alert fatigue. There was general communication about the project before the second intervention sent via institution-wide email announcements, but the specific alert rules were not mentioned and there was no formal education of the staff after any of the interventions.

Using the dashboard to identify the most frequently triggered DDIs, a group of 10 pediatric clinical pharmacists then conducted a comprehensive, literature-based review of select additional DDIs to determine if removal was warranted. The clinical pharmacists were divided into pairs; each pair was assigned a different set of DDIs to review. Each team member independently reviewed the Medi-Span DDI monograph; this included a list of offending agents, the mechanism of interaction, management options to mitigate the interaction, and supporting literature. The clinical pharmacists then determined the clinical relevance of each DDI. Due to contradictions noted between available tertiary references and third-party DDI vendors, each Medi-Span monograph's severity level and proposed management strategy was then compared to Lexi-Comp (Lexi-Comp, Hudson, Ohio) and Micromedex (Thomson Micromedex, Ann Arbor, Michigan, USA).[9] When necessary, primary literature was reviewed to aid in assessing the clinical relevance of specific DDIs. If a DDI alert was deemed to be clinically insignificant, it was recommended to deactivate the alert rule so that the alert no longer fired. Override rates did not factor into removal decisions due to the presumed state of alert fatigue.

In the event that team members disagreed upon the clinical significance of their assigned DDI, the alert rule was presented to the entire clinical pharmacy group for discussion until a consensus agreement was reached. Any controversial alerts were discussed with the physician group that would most commonly encounter the DDI alert to ensure that the most likely end users were in agreement regarding the alerts' significance. For example, the potassium-sparing diuretics and potassium supplements DDI alert rule was discussed among the Divisions of Cardiology and Nephrology and determined to be of clinical significance; this alert remains an active, interruptive alert to encourage providers to monitor serum potassium levels. All recommendations were then presented to the CDS committee for approval.

The CDS committee evaluated two sets of removal recommendations from the Clinical Pharmacy group. The second round of deactivation of select DDIs occurred in December 2011 (intervention 2, see online supplementary appendix) and the third round took place in March 2012 (intervention 3, see online supplementary appendix). During both of the formal reviews, the visual analytics dashboard was referenced extensively and projected on a screen for all of the meeting attendees to review. The dashboard enabled the team to filter and examine the alert data according to patient location and ordering provider type, and to determine which specific medication orders triggered the DDI alerts. Following CDS committee review and subsequent approval, our institution's Therapeutic Standards Committee acted as the final arbiter of the removal of clinically irrelevant alerts. Only DDI alert rules were removed during this process; no other systematic changes were made to any other alert types during this time.

Statistical analysis

We identified two groups of alerts for analysis: all medication alerts and 'untouched' DDI alerts (alerts that remained active and unchanged throughout the entire study period). This second cohort of alerts was of interest because we hypothesized the untouched alert override rate would decrease due to a reduction in alert fatigue from clinically irrelevant alerts, thereby allowing pharmacists and providers to pay more attention to clinically relevant ones.

We evaluated the following longitudinal measures among both pharmacists and providers: (1) alerts per 100 medication orders, (2) 'untouched' DDI alerts per 100 medication orders, (3) overrides per 100 alerts and per 100 untouched alerts, (4) 'cancel' actions per 100 alerts and per 100 untouched alerts, (5) 'remove' actions per 100 alerts and per 100 untouched alerts, and (6) 'view' actions per 100 alerts and per 100 untouched alerts. We compared baseline median rates and IQRs to median rates and IQRs following three intervention rounds of DDI deactivation, and tested for statistical significance using the Wilcoxon rank-sum test. To arrive at the weekly proportions in the figures, we calculated the overall proportions for that week. For the pharmacist overrides per 100 alerts, we took the total number of overrides (as well as cancel and view actions) for the week divided by the total number of alerts for the week and multiplied by 100. For the provider overrides per 100 alerts, we took the total number of overrides (as well as cancel, remove, and view actions) for the week divided by the total

number of alerts for the week and multiplied by 100. We used piecewise linear regression models to estimate slopes by the predefined intervention's phases. Statistical significance was defined at $p < 0.05$.

As a balancing measure, we monitored the hospital incident reporting system that tracks patient safety events related to medications and near-miss rates throughout the intervention period. At our institution, a medication serious safety event (SSE) is defined as class G ('An error occurred that may have contributed to or resulted in permanent patient harm') or greater, in accordance with the National Coordinating Council for Medication Error Reporting and Prevention (NCC MERP) Medication Error Index.[10] Our monthly medication SSE rate is calculated by dividing the rolling 12-month sum of medication SSEs by 10 000 adjusted patient days.

RESULTS

The resulting visual analytics dashboard user interface consists of a central display region with tabular and graphical data representations. A variety of relevant filtering options frame the central display region and are accessible via open item selectors, drop-down menus, and a date range selector slider (figure 1). Medication alert and override rates, alert types, and various patient and care provider characteristics can be displayed, explored, and selected at a particular time point or across a user-defined time interval using the filters and limits. The dashboard also enables the user to explore the data in granular detail via a wide range of tabular and graphical display formats (eg, time series, heat maps, small multiples).

As the first intervention, we deactivated two benzodiazepine/ethanol DDI alert rules (intervention 1). During the second and third rapid-cycle improvement stages (interventions 2 and 3), 42 and 19 DDI alert rules were deactivated, respectively. A table listing the deactivated DDIs is available as online supplemental digital content. These changes primarily impacted pharmacists, since the majority of the deactivated DDIs were not firing for prescribers due to previously discussed severity settings.

We evaluated 2 391 880 medication alerts presented to physicians and pharmacists between January 31, 2011 and January 26, 2014. For pharmacists, the baseline median alert

Figure 1: The visual analytics dashboard user interface consists of a central display region with tabular and graphical data representations, as well as a variety of relevant filtering options framing the central display region.

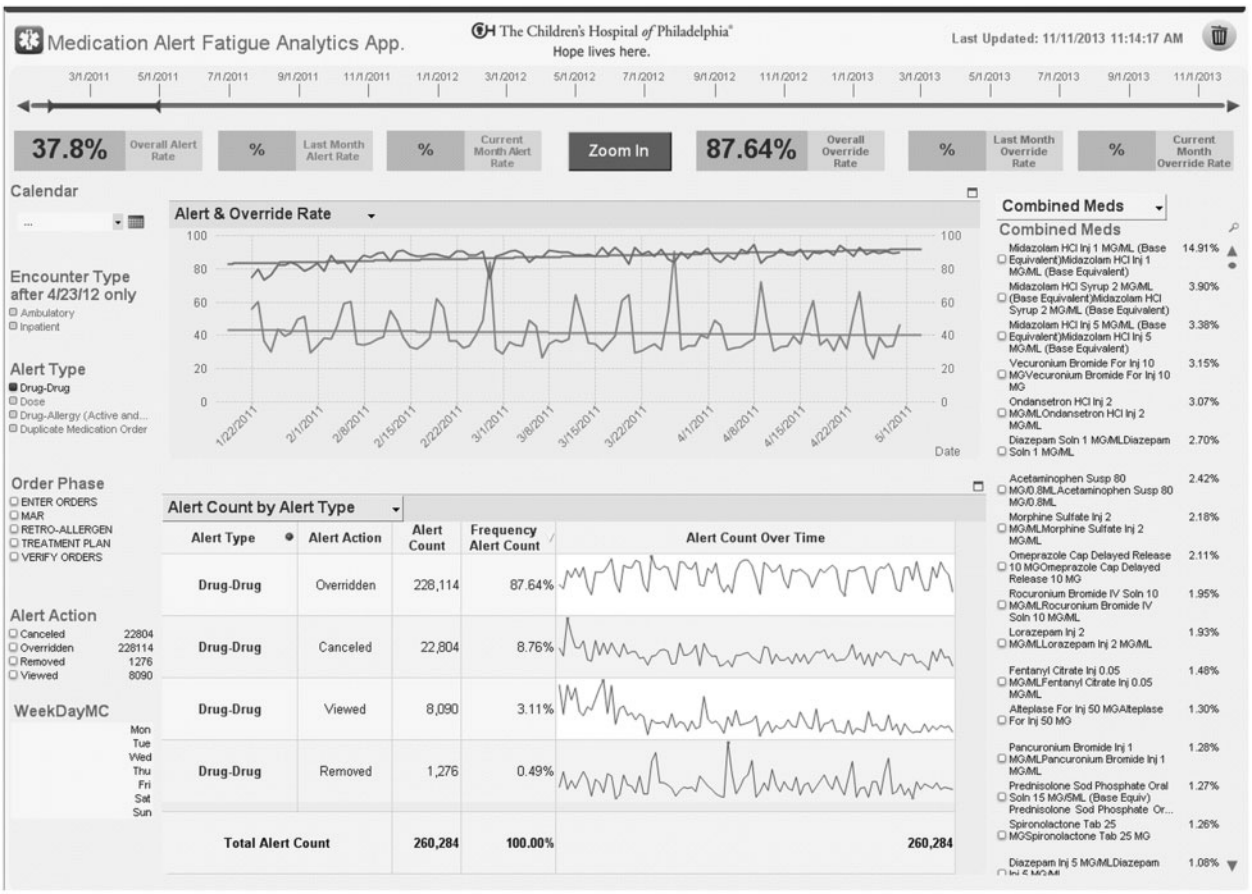

rate prior to the first round of DDI deactivation was 58.74 alerts/100 orders (IQR 54.98–60.48) and 25.11 alerts/100 orders (IQR 23.45–26.57) following the three interventions ($p < 0.001$). For providers, baseline median alert rate prior to the first round of DDI deactivation was 19.73 alerts/100 orders (IQR 18.66–20.24) and 15.11 alerts/100 orders (IQR 14.44–15.49) following the three interventions ($p < 0.001$).

For pharmacists, baseline median override rate prior to the first round of DDI deactivation was 95.14 overrides/100 alerts (IQR 94.16–95.46) and 84.38 overrides/100 alerts (IQR 83.04–85.44) following the three interventions ($p < 0.001$). For providers, baseline median override rate prior to the first round of DDI deactivation was 84.22 overrides/100 alerts (IQR 83.27–85.69) and 84.91 overrides/100 alerts (IQR 83.96–85.21) following the three interventions ($p = 0.16$).

Figures 2 and 3 show the weekly override, cancel, remove, and view rates per 100 orders. In a subgroup analysis, we observed a decrease in pharmacists' untouched override rates for DDI alerts that were not modified in the system from an average rate of 93.56 overrides/100 alerts to 85.71 overrides/100 alerts, a decrease of 7.85% (95% CI 6.63 to 9.07%, $p < 0.001$). In the subgroup, the decrease in pharmacists' override rates coincided with increased rates of cancel actions taken by pharmacists, while view rates remain unchanged; 'remove' actions taken by providers did not change during the study period (figure 3). In week 65, coinciding with a system upgrade that included numerous user interface changes, the number of alerts 'viewed' by providers decreased to zero and the total alert 'view' rate for pharmacists increased. Of note, the view rate for the untouched pharmacy subgroup analysis remained constant (figure 3). We estimated the change in rates (slopes) during each intervention period using piecewise linear regression analysis as shown in table 1.

The hospital's medication SSE rate decreased from 0.18 events per 10 000 adjusted patient days before the study period to 0.08 after the study period. No SSEs were reported for the medications associated with the deactivated DDI alert rules.

DISCUSSION

We developed a visual analytics dashboard that facilitates the continuous, dynamic monitoring of EHR medication alerts and providers' and pharmacists' responses. The dashboard was and remains an integral part of a hospital quality improvement initiative to customize a commercial rule base by deactivating clinically irrelevant alert rules in the interest of improving medication safety and reducing alert fatigue. The dashboard was particularly helpful when we looked at the high frequency alerts in a specialty area versus in a trainee or non-specialty environment. For example, we found while using the dashboard that the DDI alerts for methotrexate and steroids (or Bactrim) were firing with a high frequency for oncology providers. The dashboard helped us to quickly discern that the alert was also appearing for trainees and non-oncology areas, so lacking the ability to filter by specialty we refrained from deactivating that particular alert rule.

Our study produced three key findings. First, the dashboard was an effective tool for informing and monitoring the impacts of the three successive interventions phases that resulted in reductions in pharmacists' alert and override rates. Second, the reductions were not dangerous or harmful; there were no reported medication safety events associated with the 63

Figure 2: Mean weekly provider and pharmacist alerts per 100 orders (includes drug–drug interactions (DDIs), allergy, maximum dose, and duplicate-medication alerts).

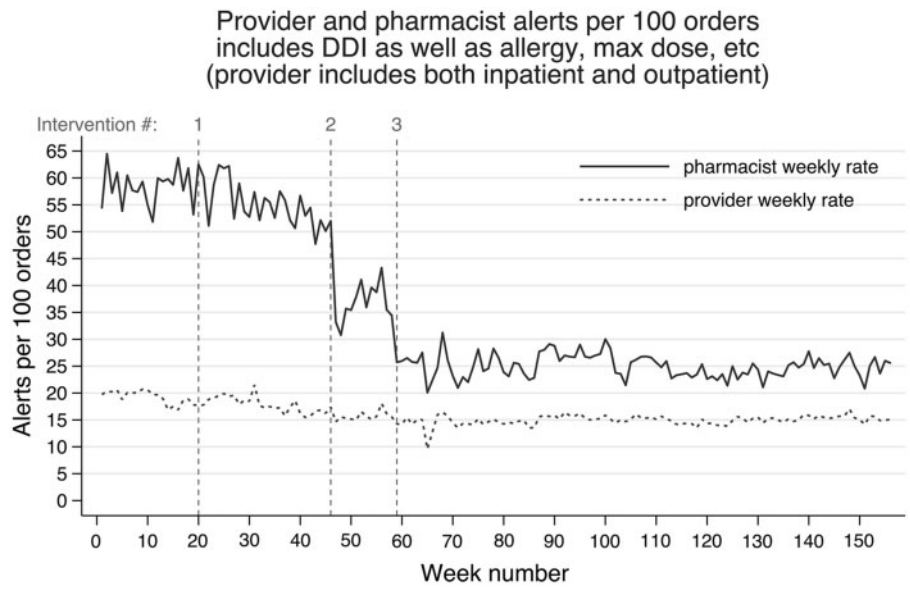

Figure 3: Mean weekly override rates for pharmacists and providers during the study period. The two graphs on the left side show the override rates for all medication alerts for both care providers and pharmacists. The two graphs on the right side show the pharmacists' and providers' override rates for the 'untouched' drug–drug interaction (DDI) alerts that were active and unmodified throughout the study period.

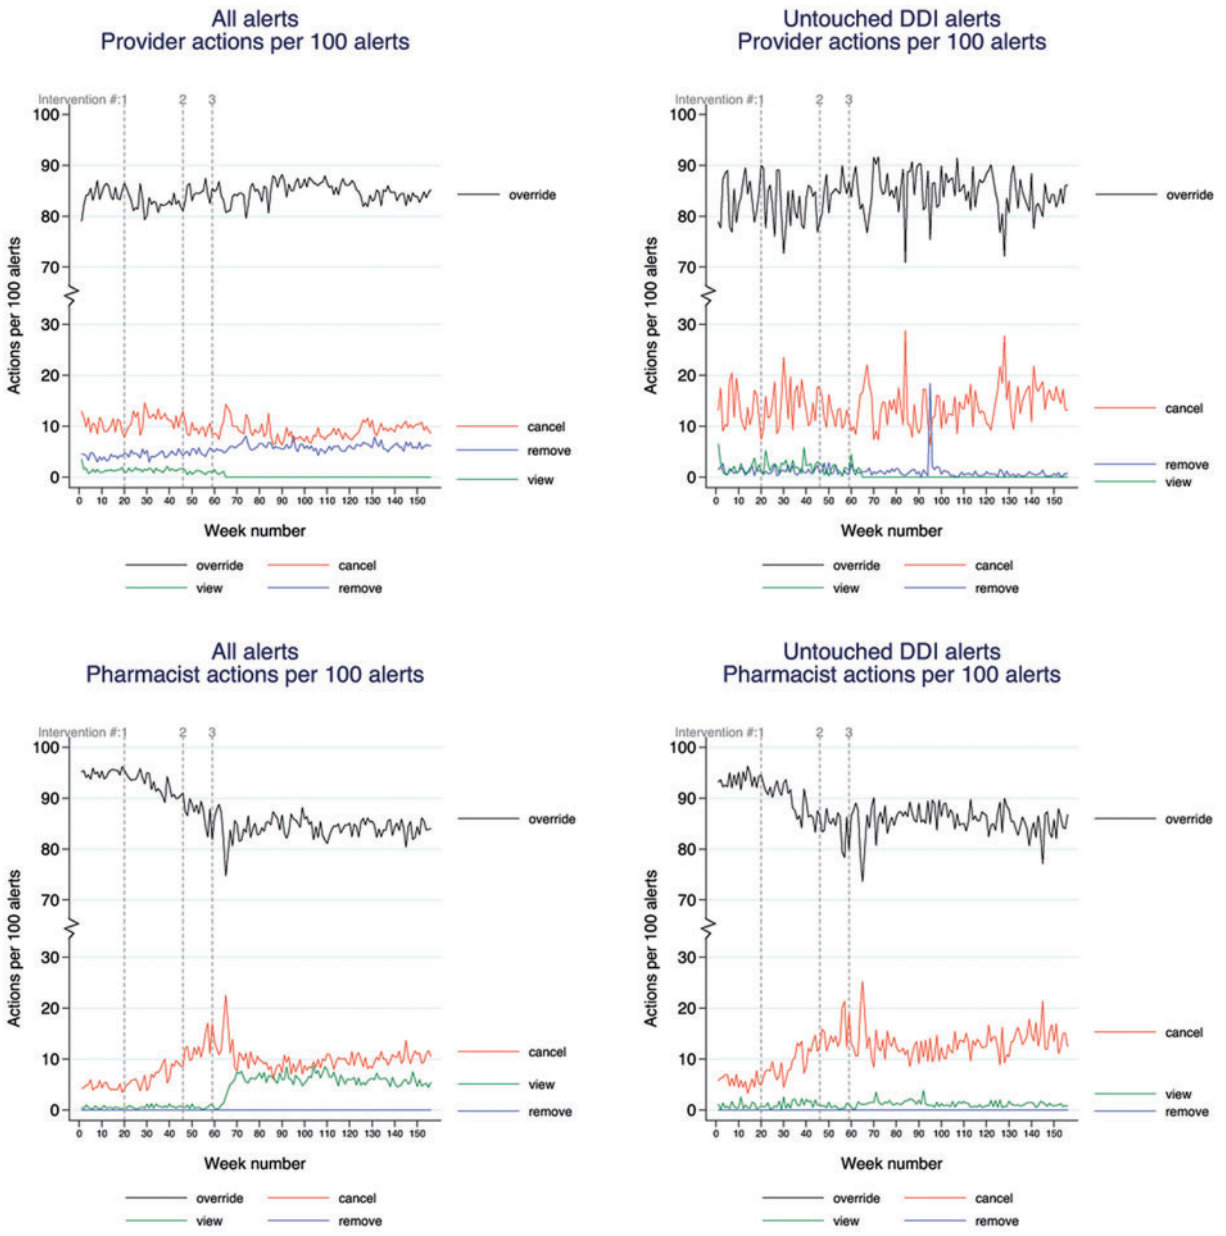

deactivated DDI alert rules. Third, pharmacists' override rates of untouched DDIs decreased as their alert burden was reduced. One possible explanation for this phenomenon is that pharmacists experienced a reduction in alert fatigue due to an improvement in the EHR alert system's signal-to-noise ratio.

Visual analytics is the science of analytical reasoning facilitated by interactive visual interfaces and has been utilized to evaluate large, complex datasets in a variety of fields, including epidemiology, genetics, and immunology.[11–13] The dashboard that we created enabled us to determine and analyze the millions of rows of alert data to answer the CDS committee's questions in minutes, rather than days using the manual database query method. Utilizing the dashboard for visual analysis, we quickly identified decreases in alert and override rates for pharmacists, our intervention target. Neither of these results was surprising, since we knew that the potential impact on alert volume was significant and that inactivated DDI alerts tended to have higher rates of override. Visual analytics was

Table 1: Pharmacists' mean override rate and estimated slope of rate change for untouched DDIs at baseline and during each intervention phase

Intervention phase	Mean override rate	95% CI	Slope (per year)	95% CI	Slope p value
Baseline	93.56	92.95 to 94.17	Increasing 3.08%/year	−7.51% to +13.67%	.57
After intervention 1, before intervention 2	89.58	88.40 to 90.77	Decreasing 15.87%/year	−22.11 to −9.63%	<0.001
After intervention 2, before intervention 3	84.12	82.43 to 85.82	Decreasing 17.58%/year	−36.31 to +1.16%	.066
After intervention 3	85.71	85.17 to 86.25	Decreasing 0.81%/year	−1.74 to +0.12%	.087

A slope p value less than 0.05 denotes a slope that is significantly different from zero (a flat line).
DDI, drug–drug interaction.

also critical in the important exploratory step of analyzing the untouched DDIs, in the pre-intervention period, which would otherwise have required complex database queries using SQL.

High EHR alert burden is a common issue in hospitals. Many use a third-party medication alert database for quick deployment and to avoid the significant resources required to manage alert content locally. However, these products are designed to be as sensitive as possible, often with poor specificity and high false positive rates, especially for specific populations such as children.[14,15] Comparisons of DDI alerts with manual checking by clinical pharmacologists showed that CDS software performs poorly due to over-alerting.[16,17]

Alert fatigue has been implicated as a potential unintended consequence of alerting systems with high sensitivity but poor specificity. In response to the large number of inappropriate 'noise' alerts, clinicians override many alerts and are at risk of alert fatigue—the neglect of 'signal' alerts that are intended to prevent errors and adverse drug events.[4,5,18] Prior work has identified key areas for improvement of DDI alert fatigue, such as optimizing the quality of the alert display, refining the identification of clinically significant alerts, and reducing signal-to-noise ratios.[19,20] Work is ongoing at other institutions to establish standards and best practices for the clinical content of DDI alerts, yet most reports of such efforts in the literature have not included the use of pediatric medication recommendations and guidelines, and thus bear minimal overlap with our customization efforts.[5,21] Standards are beginning to emerge and of note we did not deactivate any of the DDI alerts from the Children's Hospital Association's core list of 19 high value DDI alerts.[22] Until standards become clearer and more widely implemented by vendors, however, institutions must rely on local, manual reclassification of DDI alert rules to develop CDS systems with higher specificity and less frequent over-alerting.[15,23]

We suspected a decreased burden of nuisance DDI alerts (ie, 'noise') would result in a change in the pharmacists' responses to the alerts that remained in the system (ie, 'signal'). We confirmed our hypothesis after our analysis showed a decreased override rate in untouched DDIs after deactivating clinically irrelevant alert rules. The steepest decrease in override rates for unchanged DDIs occurred surprisingly following the removal of two alert rules during intervention 1, and not after interventions 2 and 3 when a far greater number of alert rules were deactivated. The cause of this finding is unclear. The minimal change in prescriber override rates following the first intervention was likely due to prescribers seeing fewer alerts at baseline and presumably experiencing less alert fatigue than the pharmacists. While the observed decrease in override rates for the untouched DDIs was not very large, we speculate that this work might be an incremental step toward identifying empiric alert rates at which clinicians experience alert fatigue.

This study had several limitations. First, the visual analytics dashboard was developed for use with our hospital's EHR. Our results may not be fully generalizable to hospitals using EHR systems that differ from ours. Second, while the visual analytics dashboard enables the user to detect medication alerts with high override rates, the dashboard does not help to detect and add missing relevant alerts. Third, while we did not see an increase in the incidence of medication safety adverse events after deactivating alerts, the risk remains that an adverse event will occur that a deactivated alert would have prevented. Moreover, we relied on reported patient events as our balancing measure, which only represent a fraction of actual medication related safety events.[24] However, NCC MERP has stated explicitly the limited utility of using medication error rates to compare healthcare organizations.[25] Fourth, this study was a quality improvement effort in a live, uncontrolled environment that included numerous EHR upgrades and patches that may have impacted our results without our knowledge. Fifth, our analysis of the DDI override and alert rates did not include subgroup analysis based on DDI severity rankings, as our local review of the DDI rules brought the reliability of these rankings into question. Thus, the severity rankings alone did not factor

into the decision on whether a DDI was a candidate for removal. However, severity was taken into account during pharmacy's analysis. Furthermore, we did not review medical records to determine the actions the providers or pharmacists took after viewing a DDI alert (eg, ordering an electrocardiogram after overriding the alert for co-administration of 2 QT-interval-prolonging agents). If a prescriber overrode the alert but then shortly thereafter discontinued an order in a subsequent activity, then clearly behavior was affected yet our analysis as reported would not capture that possibility. Lastly, while we assert that the dashboard improved the efficiency of our operational work, we neither utilized objective measures of efficiency, nor conducted a parallel control arm with the manual database query method to quantitatively assess the amount of time that was saved by using the dashboard.

CONCLUSIONS

We have shown that it is possible to safely reduce the medication alert burden by systematically deactivating clinically irrelevant alert rules while concurrently monitoring for medication-associated harm. Hospitals may find implementing a dashboard like ours useful when performing quality improvement customization of medication alert rules and developing a prospective strategy for alert management. We have also shown preliminary evidence of a decrease in pharmacist override rates with a concomitant increase in pharmacists selecting the 'cancel' option for alerts that remained constant throughout the study period (the untouched alerts). This could represent a reduction in alert fatigue due to an improvement in the signal-to-noise ratio. Further research will be necessary to investigate this issue.

ACKNOWLEDGEMENTS

The authors appreciate deeply the enthusiasm and support for this project from their colleague Bryan Wolf, MD, PhD at The Children's Hospital of Philadelphia.

CONTRIBUTORS

The following contributed to conception and design, acquisition of data or analysis and interpretation of data, drafting the article or revising it critically for important intellectual content, and final approval of the version to be published: AFS, LMA, BRD, CPB, JAG, MAR, KLP, EDS. AFJ contributed to conception and design, analysis and interpretation of data and revising the article for important intellectual content. AFS and EDS are guarantors for the integrity of the work, from inception to published article.

COMPETING INTERESTS

None.

ETHICS APPROVAL

The Institutional Review Board at The Children's Hospital of Philadelphia approved the study protocol.

PROVENANCE AND PEER REVIEW

Not commissioned; externally peer reviewed.

SUPPLEMENTARY MATERIAL

Supplementary material is available online at http://jamia.oxfordjournals.org/.

REFERENCES

1. Fortescue EB, Kaushal R, Landrigan CP, *et al*. Prioritizing strategies for preventing medication errors and adverse drug events in pediatric inpatients. *Pediatrics* 2003;111:722–729.
2. Abramson EL, Kaushal R. Computerized provider order entry and patient safety. *Pediatr Clin North Am* 2012;59:1247–1255.
3. Jha AK, Laguette J, Seger A, *et al*. Can surveillance systems identify and avert adverse drug events? A prospective evaluation of a commercial application. *J Am Med Inform Assoc* 2008;15:647–653.
4. Weingart SN, Toth M, Sands DZ, *et al*. Physicians' decisions to override computerized drug alerts in primary care. *Arch Intern Med* 2003;163:2625–2631.
5. 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;20:489–493.
6. Thomas JJ, Cook KA. A visual analytics agenda. *IEEE Comput Graph Appl* 2006;26:10–13.
7. Kimball R, Ross M, Thornthwaite W, Mundy J, Becker B. eds. *The data warehouse lifecycle toolkit*. 2nd edn. Hoboken, NJ: Wiley, 2008.
8. Tufte ER. *The visual display of quantitative information*. 2nd edn. Cheshire, CT: Graphics Press, 2001.
9. Shah VS, Weber RJ, Nahata MC. Contradictions in contraindications for drug-drug interactions. *Ann Pharmacother* 2011;45:409–411.
10. Medication Error Index. NCC MERP web site. http://www.nccmerp.org/medErrorCatIndex.html. Accessed July 20, 2014.
11. Chui KK, Wenger JB, Cohen SA, *et al*. Visual analytics for epidemiologists: understanding the interactions between age, time, and disease with multi-panel graphs. *PLoS ONE* 2011;6:e14683.
12. Kumasaka N, Nakamura Y, Kamatani N. The textile plot: a new linkage disequilibrium display of multiple-single nucleotide polymorphism genotype data. *PLoS ONE* 2010;5:e10207.
13. Naumova EN. Visual analytics for immunologists: data compression and fractal distributions. *Self Nonself* 2010;1:241–249.
14. Resetar E, Reichley RM, Noirot LA, *et al*. Customizing a commercial rule base for detecting drug-drug interactions. *AMIA Annu Symp Proc* 2005:1094.
15. Kuperman GJ, Reichley RM, Bailey TC. Using commercial knowledge bases for clinical decision support: opportunities, hurdles, and recommendations. *J Am Med Inform Assoc* 2006;13:369–371.
16. Smithburger PL, Kane-Gill SL, Benedict NJ, *et al*. Grading the severity of drug-drug interactions in the intensive care unit: a comparison between clinician assessment and

proprietary database severity rankings. *Ann Pharmacother* 2010;44:1718–1724.
17. Taegtmeyer AB, Kullak-Ublick GA, Widmer N, *et al*. Clinical usefulness of electronic drug-drug interaction checking in the care of cardiovascular surgery inpatients. *Cardiology* 2012;123:219–222.
18. Ash JS, Berg M, Coiera E. Some unintended consequences of information technology in health care: the nature of patient care information system-related errors. *J Am Med Inform Assoc* 2004;11:104–112.
19. Smithburger PL, Buckley MS, Bejian S, *et al*. A critical evaluation of clinical decision support for the detection of drug-drug interactions. *Expert Opin Drug Saf* 2011;10:871–882.
20. Seidling HM, Phansalkar S, Seger DL, *et al*. Factors influencing alert acceptance: a novel approach for predicting the success of clinical decision support. *J Am Med Inform Assoc* 2011;18:479–484.
21. 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;19:735–743.
22. Harper MB, Longhurst CA, McGuire TL, *et al*. Children's Hospital Association CDS working group. Core drug-drug interaction alerts for inclusion in pediatric electronic health records with computerized prescriber order entry. *J Patient Saf* 2014;10:59–63.
23. Guzek M, Zorina OI, Semmler A, *et al*. Evaluation of drug interactions and dosing in 484 neurological inpatients using clinical decision support software and an extended operational interaction classification system (Zurich Interaction System). *Pharmacoepidemiol Drug Saf* 2011;20:930–938.
24. Cullen DJ, Bates DW, Small SD, *et al*. The incident reporting system does not detect adverse drug events: a problem for quality improvement. *Jt Comm J Qual Improv* 1995;21:541–548.
25. Statement from NCC MERP: Use of Medication Error Rates to Compare Health Care Organizations is of No Value. NCC MERP web site. http://www.nccmerp.org/council/council2002-06-11.html. Accessed July 20, 2014.

AUTHOR AFFILIATIONS

[1]Department of Anesthesiology and Critical Care Medicine, Perelman School of Medicine at the University of Pennsylvania and the Children's Hospital of Philadelphia, Philadelphia, Pennsylvania, USA

[2]Department of Enterprise Analytics and Reporting, The Children's Hospital of Philadelphia, Philadelphia, Pennsylvania, USA

[3]Department of Pediatrics, Perelman School of Medicine at the University of Pennsylvania and the Children's Hospital of Philadelphia, Philadelphia, Pennsylvania, USA

[4]Department of Biostatistics in Pediatrics, Perelman School of Medicine at the University of Pennsylvania and the Children's Hospital of Philadelphia, Philadelphia, Pennsylvania, USA

[5]Department of Pharmacy Services, The Children's Hospital of Philadelphia, Philadelphia, Pennsylvania, USA