

Adherence to drug–drug interaction alerts in high-risk patients: a trial of context-enhanced alerting

Jon D Duke,^{1,2} Xiaochun Li,³ Paul Dexter^{1,2}

¹Department of Medical Informatics, Regenstrief Institute, Indianapolis, Indiana, USA

²Department of Medicine, Indiana University School of Medicine, Indianapolis, Indiana, USA

³Division of Biostatistics, Indiana University School of Medicine, Indianapolis, Indiana, USA

Correspondence to

Dr Jon D Duke, Department of Medical Informatics, Regenstrief Institute, 410 W. 10th Street, Indianapolis, IN 46202, USA; jduke@regenstrief.org

Received 30 April 2012

Accepted 26 October 2012

Published Online First

17 November 2012

ABSTRACT

Objective Drug–drug interaction (DDI) alerting is an important form of clinical decision support, yet physicians often fail to attend to critical DDI warnings due to alert fatigue. We previously described a model for highlighting patients at high risk of a DDI by enhancing alerts with relevant laboratory data. We sought to evaluate the effect of this model on alert adherence in high-risk patients.

Methods A 6-month randomized controlled trial involving 1029 outpatient physicians was performed. The target interactions were all DDIs known to cause hyperkalemia. Alerts in the intervention group were enhanced with the patient's most recent potassium and creatinine levels. The control group received unmodified alerts. High-risk patients were those with baseline potassium >5.0 mEq/l and/or creatinine ≥ 1.5 mg/dl ($132 \mu\text{mol/l}$).

Results We found no significant difference in alert adherence in high-risk patients between the intervention group (15.3%) and the control group (16.8%) ($p=0.71$). Adherence in normal risk patients was significantly lower in the intervention group (14.6%) than in the control group (18.6%) ($p<0.01$). In neither group did physicians increase adherence in patients at high risk.

Conclusions Physicians adhere poorly to hyperkalemia-associated DDI alerts even in patients with risk factors for a clinically significant interaction, and the display of relevant laboratory data in these alerts did not improve adherence levels in the outpatient setting. Further research is necessary to determine optimal strategies for conveying patient-specific DDI risk.

INTRODUCTION

Drug–drug interaction (DDI) alerts are commonly employed by computerized physician order entry (CPOE) systems and are considered a basic form of clinical decision support.^{1–4} The potential benefits of DDI alerting have not been fully realized, however, due in part to high physician override rates, with up to 96% of such warnings being overridden.^{5,6} A well-recognized cause of this poor adherence is ‘alert fatigue,’ a state in which physicians become desensitized in the setting of frequent, low-specificity alerts.^{7–10} A major danger of alert fatigue is that physicians will fail to distinguish between high-risk and low-risk warnings, leading to failure to address critical safety concerns. We have recently proposed a strategy for reducing alert fatigue by highlighting alerts in patients at high risk of a clinically significant interaction.¹¹ This approach, known as context-aware drug–drug interaction (CADDI) alerting, enhances alerts with key laboratory data relevant to assessing an individual patient's DDI risk level. In this paper, we present

the results of a study evaluating the impact of the CADDI model on physician adherence to DDI alerts in high-risk patients. We also explore the clinical impact of alert non-adherence in a high-risk population.

To generate patient-specific alerts, the CADDI model combines static information found in traditional DDI alerts (eg, mechanism of action) with patient laboratory data relevant to the clinical outcome of the interaction. For example, CADDI alerts for DDIs causing increased risk of bleeding (eg, warfarin+azithromycin) would display the patient's latest prothrombin time, platelet count, and hematocrit. Alternatively, an alert for an interaction causing prolonged QT duration (eg, levaquin and amitriptyline) would show results such as potassium, calcium, and digoxin level. The underlying theory for CADDI is that displaying relevant data will help physicians quickly separate high-risk from low-risk patients and thus improve the ‘signal-to-noise ratio’ in clinical alerting.

For this pilot study, we focused on DDIs associated with hyperkalemia. Hyperkalemia-inducing interactions are extremely common, as noted in a recent study in which over 20% of all DDIs were due to just two class interactions, both of which were associated with hyperkalemia (ACE inhibitors+potassium-sparing diuretics and ACE inhibitors/angiotensin receptor blockers+potassium supplements).¹⁰ DDI-induced hyperkalemia is also a common reason for hospital admission in the elderly and a frequent hospital complication.^{12,13} In this study, we sought to determine whether enhancing these DDI alerts with laboratory data relevant to hyperkalemia would improve alert adherence in high-risk patients.

METHODS

This study was a 6-month randomized controlled trial that was started on February 22, 2011 and concluded on August 30, 2011.

Setting and participants

The study was conducted at Wishard Health Services in Indianapolis, Indiana. We obtained approval for this trial from the institutional review board at the Indiana University Medical Center. Participants were 1029 physicians involved in outpatient care (as defined by having written at least one outpatient order in the preceding 12 months). The group included 671 residents and 358 staff physicians.

The unit of randomization in this study was the provider. We listed the names of the eligible providers on a spreadsheet and assigned each a random number using an online random number

To cite: Duke JD, Li X, Dexter P. *J Am Med Inform Assoc* 2013;**20**:494–498.

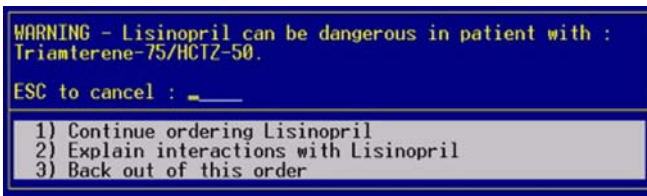


Figure 1 Standard drug–drug interaction alert in the Gopher order entry system. This figure is only reproduced in colour in the online version.

generator.¹⁴ We then sorted the list by the random number and assigned the first half to the intervention group and the second half to the control group. There was no significant difference in the percentage of residents in the intervention arm (65%) and the control arm (63%).

Intervention

The intervention was the integration of context-specific patient laboratory data into the standard DDI alerts displayed by our CPOE system. Figure 1 shows a standard DDI alert. Figure 2 shows an intervention alert with integrated relevant laboratory data.

To identify those interactions associated with hyperkalemia in our CPOE, we used regular expressions to search the clinical effects of all DDIs in our local knowledge base for those containing either ‘hyperkalemia’ or ‘increased potassium.’ This search revealed six class interactions (combinatorial of ACE inhibitors, angiotensin receptor blockers, potassium-sparing diuretics, and potassium supplements). Using the CADDI model, each of these DDIs was tagged as being associated with hyperkalemia. In the CADDI database, hyperkalemia was mapped using the Logical Observation Identifiers Names and Codes (LOINC) terminology to the concepts serum potassium (2823-3) and serum creatinine (2160-0). Thus upon activation of the study, when any of these interactions were triggered by an intervention physician, the most recent values for potassium and creatinine within the past 12 months were included in the DDI alert (figure 2). As shown, the system displayed the result value as well as the date and a textual descriptor if abnormal (eg, high).

Hypothesis

We hypothesized that the integration of patient-specific laboratory data into DDI alerts would improve adherence in patients at high risk for DDI-induced hyperkalemia. We defined high risk as having a baseline potassium >5.0 mEq/l and/or creatinine ≥1.5 mg/dl (132 μmol/l).

Data captured

Each time a DDI was triggered, we captured the relevant patient laboratory data and whether or not the physician adhered to the




Figure 2 Context-aware drug–drug interaction alert enhanced with relevant patient laboratory data. This figure is only reproduced in colour in the online version.

alert. Although the control alerts did not display relevant laboratory results, we captured the same data (ie, most recent potassium and creatinine) for patients in both groups. Adherence was based on the final orders for the session in which the alert appeared, and was considered positive if the physician either (1) did not order the triggering medication, or (2) ordered the triggering medication but discontinued the interacting drug(s) cited by the alert.

Data analysis

We used generalized estimating equations (GEE) to determine whether there was a significant difference in alert adherence between the intervention and control groups. GEE is a method for estimating regression model parameters when dealing with correlated data.¹⁵ This approach helps account for the fact that individual physicians may be more or less likely to override alerts based on their own clinical habits, independent of the study group. Physician responses were modeled as binary outcomes and the exchangeable correlation structure was assumed within physician. Assuming a conservative baseline alert adherence rate of 10%, we need a sample size of 75 physicians in each arm to trigger hyperkalemia-associated DDI alerts in order to detect at least a 10% increase in intervention group adherence, at the two-sided type I error rate 0.05 with power of at least 90%. Analyses were performed in the open-source statistical computing environment R using the ‘geepack’ package.¹⁶

RESULTS

During the 6-month period, 101 intervention physicians and 102 control physicians triggered a total of 2140 alerts involving the target DDIs. More alerts were triggered by the intervention group (n=1174) than by the control group (n=966).

Table 1 shows the alert adherence rates overall as well as stratified by high-risk and normal risk patients. Overall adherence rates were low (<20%), consistent with other studies of DDI alerting.^{5 6 17} As shown, no significant difference was seen between the intervention and control groups in terms of adherence to alerts in high-risk patients. In normal risk patients, adherence was significantly lower in the intervention group than in the control group.

Table 2 shows adherence rates stratified by baseline potassium levels. At lower potassium values (<3.9 mEq/l), adherence was markedly lower in the intervention group than in the control group. While not part of our initial hypothesis, this pattern is not unexpected; it suggests that physicians may have been reassured by lower potassium values as they imply a decreased risk of hyperkalemia.

Table 1 Adherence rates to drug interaction alerts associated with hyperkalemia

	Intervention group		Control group		p Value
	Alerts, n	Adhered to, n (%)	Alerts, n	Adhered to, n (%)	
In high-risk patients	163	25 (15.3%)	167	28 (16.8%)	p=0.71
In normal risk patients	1011	146 (14.4%)	799	152 (19%)	p<0.01
Overall	1174	171 (14.6%)	966	180 (18.6%)	p<0.01

High-risk patients were defined as those with a baseline potassium >5.0 mEq/l or creatinine >1.5 mg/dl.

Table 2 Adherence rates stratified by baseline potassium levels

Baseline potassium	Intervention group		Control group		p Value
	Total alerts, n	Adhered to, n (%)	Total alerts, n	Adhered to, n (%)	
<3 mEq/l	31	9.7	19	26.3	p=0.18
3–3.9 mEq/l	490	14.9	365	22.2	p<0.01
4–4.9 mEq/l	559	14.7	502	16.3	p=0.43
≥5 mEq/l	54	11.1	57	15.8	p=0.50

At higher potassium levels, we encountered two phenomena that ran counter to expectations. First, the display of elevated potassium levels—a clear indicator of risk for more severe hyperkalemia—had no impact on alert adherence. Second, we observed a step-wise worsening of adherence rates in the control group as the baseline potassium increased. In neither group did we see evidence that an elevated potassium level prompted increased concern sufficient to improve adherence.

To explore these findings further and to gain insight into the relative impact of alert overrides in patients with differing baseline risk, we performed manual chart review for adverse events in those patients with the highest and lowest potassium levels in our study. We defined an adverse event as either of the following events occurring within 90 days of the alert: (1) a potassium level >5.5 mEq/l or (2) an emergency department visit or hospital admission with a diagnosis of hyperkalemia. We selected a total of 20 charts for review, representing the patients with the 10 highest and 10 lowest potassium values displayed to and overridden by intervention physicians.

As shown in tables 3 and 4, 10 patients in the high potassium group developed significant hyperkalemia with two requiring hospital admission for further management. No adverse events were seen in the low potassium group (table 4).

DISCUSSION

The most concerning finding in our study was the tendency of physicians to override alerts in even the highest risk patients. In both the control and intervention groups, adherence rates were actually worse in patients with potassium values >5 mEq/l than in patients with lower potassium values. Given the very straightforward effect of the drug interactions studied (ie, they increase

Table 3 Adverse events occurring within 90 days of interacting drug prescriptions in the 10 intervention patients with the highest baseline potassium values in our study

Patient ID	Baseline potassium (mEq/l)	Adverse event	
		Hyperkalemia	Hospital admission
1	5.6	No	No
2	5.5	Yes (5.7 mEq/l)	Yes
3	5.4	No	No
4	5.4	No	No
5	5.4	Yes (6.9 mEq/l)	Yes
6	5.4	No	No
7	5.4	Yes (5.7 mEq/l)	No
8	5.3	Yes (5.7 mEq/l)	No
9	5.2	No	No
10	5.2	No	No

Table 4 Adverse events occurring within 90 days of interacting drug prescriptions in the 10 intervention patients with the lowest baseline potassium values in our study

Patient ID	Baseline potassium (mEq/l)	Adverse event	
		Hyperkalemia	Hospital admission
1	2.5	No	No
2	2.7	No	No
3	2.7	No	No
4	2.7	No	No
5	2.8	No	No
6	2.8	No	No
7	2.8	No	No
8	2.9	No	No
9	2.9	No	No
10	2.9	No	No

potassium), we would have expected more hesitation from physicians in overriding alerts in patients with elevated potassium levels. Such hesitation appears warranted based on our manual chart review, which showed that four of 10 individuals with high baseline potassium developed severe hyperkalemia within 90 days of the overridden alert. The adverse event rate for these patients, while drawn from a small sample, is considerably higher than the expected rate of 2–6% based on previous studies of alert overrides.^{6 18} It is also considerably higher than in our sample of patients with very low baseline potassium levels, in whom we saw no adverse events. Collectively, these findings suggest that the clinical importance of a given DDI alert should not be considered uniform across all patients but is highly dependent on individual risk. This underscores the potential value of DDI alerts that effectively communicate patient-specific risk factors.

Unfortunately in terms of our intervention, we found that the CADDI model of adding relevant laboratory data to outpatient DDI alerts did not improve adherence in this setting. Even in patients with laboratory values suggestive of a heightened risk of clinically significant hyperkalemia, we saw no evidence of increased adherence in physicians who were shown these abnormal values. Interestingly, while our intervention was unsuccessful in raising physician concern about high-risk patients, it appeared highly effective in ‘reassuring’ providers when laboratory values suggested a low risk of adverse events. Indeed, in those patients with baseline results conveying low risk for hyperkalemia (eg, potassium ≤3.9 mEq/l), substantially more alerts were overridden when physicians were shown these values. This finding suggests that such markers of low risk may ultimately be used to suppress alerts when the likelihood of an adverse event is low.

In considering why our intervention did not improve adherence in high-risk patients, we offer three possibilities. The first is that physicians simply ignored the alerts entirely. While possible, the statistically significant impact on adherence (although negative in direction) seen in low-risk patients suggests this is not the case. A second possibility is that patients in the high-risk category clinically differed from lower risk patients in ways that might actually increase the likelihood of alert non-adherence. For example, high-risk patients may be seen more often in clinic, resulting in more frequent monitoring and reduced physician concern about missing an interaction. Similarly, prescriptions for chronic patients may be more likely to be renewals than new treatments, lowering physician concern for an adverse

event. Finally, many patients with abnormal potassium and creatinine levels may already be on hemodialysis, an effective therapy for hyperkalemia and potentially reassuring to physicians that any resultant hyperkalemia would only be transient. Thus, unmeasured patient factors may have contributed to our study's negative findings.

A third potential explanation for our findings is that physicians, accustomed to overriding alerts the vast majority of the time and unaware of any resultant harm, simply do not gauge patient risk accurately in the context of clinical alerts. In other settings (eg, board exams), most clinicians would likely agree that prescribing spironolactone to a patient with a potassium of 5.2 mEq/l is ill-advised. Yet in the moment of a DDI alert, standard clinical logic may be overruled by habit. Further research, including prescriber interviews, will be necessary to determine whether worrisome laboratory data were misinterpreted, correctly interpreted but overruled, or simply ignored. A follow-up study in the inpatient setting, where patient acuity may be higher and laboratory results more recent, is also necessary before drawing firm conclusions on the potential impact of contextual data on DDI alert adherence.

User interface considerations

Our results suggest that we could have done more to intensify the signal when a patient was at high risk of an event. Rather than simply showing the relevant laboratory values and implicitly assuming physicians would make the necessary connections, we might have been better served by explicitly stating the adverse effect of the interaction (eg, hyperkalemia and possible arrhythmia) as well as the risk level of each patient. This statement of elevated risk could then be accompanied by distinct visual cues in the alert display (eg, colors, icons, location) to match the heightened warning.^{19, 20} Such tiering of alert display has previously been shown to improve adherence.²¹ Optimally, physician factors such as specialty, training level, and prior alert history would also be used in defining the context of alert delivery. While the CPOE used in this study did not support such capabilities, our institution is currently deploying a new order entry system that will allow a much broader range of alert delivery styles for clinical decision support research.

Study limitations

Our study considered alert adherence only in the context of whether or not a drug was ordered, but did not consider the impact of alerts on laboratory monitoring or changes in medication dosing. It is possible that DDI alerts could have led to increased monitoring or lower doses for intervention patients. In future studies, we will incorporate drug monitoring and dosing reductions as secondary measures of adherence. Also, our study was carried out in the outpatient setting, where baseline potassium values may have been months old in some cases and perceived as not relevant; had we performed an inpatient study using potassium values from the current admission, the intervention effect may have been greater. Another limitation of our study was that the definition of 'high risk' was based on empirical clinical judgment and reflected thresholds we perceived as likely to elevate physician concern for a potential adverse event. However, these were not validated predictive models of hyperkalemia risk. For future studies we hope to develop such models in a data-driven fashion in order to accurately calculate DDI risk for individual patients. Finally, our manual review for adverse events was limited in size and focused only on comparing those with the highest and lowest potassium values. A larger analysis using a validated predictive

model of hyperkalemia is necessary to draw firm conclusions regarding the clinical consequence of alert overrides based on patient risk.

CONCLUSION

Physicians adhere poorly to hyperkalemia-associated DDI alerts even in patients with risk factors for a clinically significant interaction, and the display of relevant laboratory data does not appear to improve adherence in the outpatient setting. Additional alert enhancements, including an explicit statement of patient risk and tiered alert designs that better distinguish between high- and low-risk patients may help convey DDI risk more effectively and require study in this context. Further research is also warranted to determine physician reasons for DDI alert overrides in patients with clinical evidence suggestive of increased risk for an adverse event.

Acknowledgements We would like to thank Jill Warvel for her assistance with this project.

Contributors JDD designed the study, wrote the analysis plan, analyzed the data, and drafted and revised the paper. He is guarantor. XL helped design the study, analyzed the data, and revised the paper. PD helped analyze the data and revised the paper.

Funding This work was supported by the Regenstrief Institute and by Indiana CTIS Young Investigator Award KL2 RR025760 (A Shekhar, PI).

Competing interests None.

Ethics approval Indiana University IRB approved this study.

Provenance and peer review Not commissioned; externally peer reviewed.

REFERENCES

- 1 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;14:29–40.
- 2 Kaushal R, Shojania KG, Bates DW. Effects of computerized physician order entry and clinical decision support systems on medication safety: a systematic review. *Arch Intern Med* 2003;163:1409–16.
- 3 Glassman PA, Simon B, Belperio P, *et al.* Improving recognition of drug interactions: benefits and barriers to using automated drug alerts. *Med Care* 2002;40:1161–71.
- 4 Hunt DL, Haynes RB, Hanna SE, *et al.* Effects of computer-based clinical decision support systems on physician performance and patient outcomes: a systematic review. *JAMA* 1998;280:1339–46.
- 5 van der Sijs H, Aarts J, Vulto A, *et al.* Overriding of drug safety alerts in computerized physician order entry. *J Am Med Inform Assoc* 2006;13:138–47.
- 6 Weingart SN, Toth M, Sands DZ, *et al.* Physicians' decisions to override computerized drug alerts in primary care. *Arch Intern Med* 2003;163:2625–31.
- 7 Ash JS, Sittig DF, Campbell EM, *et al.* Some unintended consequences of clinical decision support systems. *AMIA Annu Symp Proc* 2007;6–30.
- 8 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;10:871–82.
- 9 van der Sijs H, Mulder A, van Gelder T, *et al.* Drug safety alert generation and overriding in a large Dutch university medical centre. *Pharmacoepidemiol Drug Saf* 2009;18:941–7.
- 10 Zwart-van Rijkom JEF, Uijtendaal EV, ten Berg MJ, *et al.* Frequency and nature of drug-drug interactions in a Dutch university hospital. *Br J Clin Pharmacol* 2009;68:187–93.
- 11 Duke JD, Bolchini D. A successful model and visual design for creating context-aware drug-drug interaction alerts. *AMIA Annu Symp Proc* 2011;2011:339–48.
- 12 Juurlink DN, Mamdani M, Kopp A, *et al.* Drug-drug interactions among elderly patients hospitalized for drug toxicity. *JAMA* 2003;289:1652–8.
- 13 Uijtendaal EV, Zwart-van Rijkom JEF, van Solinge WW, *et al.* Frequency of laboratory measurement and hyperkalaemia in hospitalised patients using serum potassium concentration increasing drugs. *Eur J Clin Pharmacol* 2011;67:933–40.
- 14 RANDOM.ORG—Integer Generator. <http://www.random.org/integers> (accessed 29 Feb 2012).
- 15 Zeger SL, Liang KY, Albert PS. Models for longitudinal data: a generalized estimating equation approach. *Biometrics* 1988;44:1049–60.
- 16 R Development Core Team. *R: A language and environment for statistical computing*. Vienna:R Foundation for Statistical Computing, 2010.

- 17 Judge J, Field TS, DeFlorio M, *et al*. Prescribers' responses to alerts during medication ordering in the long term care setting. *J Am Med Inform Assoc* 2006;13:385–90.
- 18 Hsieh TC, Kuperman GJ, Jaggi T, *et al*. Characteristics and consequences of drug allergy alert overrides in a computerized physician order entry system. *J Am Med Inform Assoc* 2004;11:482–91.
- 19 Phansalkar S, Edworthy J, Hellier E, *et al*. A review of human factors principles for the design and implementation of medication safety alerts in clinical information systems. *J Am Med Inform Assoc* 2010;17:493–501.
- 20 Zachariah M, Phansalkar S, Seidling HM, *et al*. Development and preliminary evidence for the validity of an instrument assessing implementation of human-factors principles in medication-related decision-support systems—I-MeDeSA. *J Am Med Inform Assoc* 2011;18(Suppl 1): i62–72.
- 21 Paterno MD, Maviglia SM, Gorman PN, *et al*. Tiering drug-drug interaction alerts by severity increases compliance rates. *J Am Med Inform Assoc* 2009;16:40–6.