

Improving Patient Safety through Medical Alert Management: An Automated Decision Tool to Reduce Alert Fatigue

Eva K. Lee, PhD^{*,1,2,3}, Amanda F. Mejia, BS^{1,2,3}, Tal Senior, RN⁴, James Jose, MD⁴.

¹Center for Operations Research in Medicine and HealthCare, ²NSF I/UCRC Center for Health Organization Transformation, ³Georgia Institute of Technology, Atlanta, GA, ⁴Children's HealthCare of Atlanta, Atlanta, GA

Abstract

Drug safety alerts, a feature of electronic medical records (EMRs), are increasingly recognized as valuable tools for reducing adverse drug events and improving patient safety. However, there has also been increased understanding that alert fatigue, a state in which users become overwhelmed and unresponsive to alerts in general, is a threat to patient safety. In this paper, we seek to mitigate alert fatigue by filtering superfluous alerts. We design a method of predicting alert overrides based on past alert override rate, range in override rate, and sample size. Using a dataset from a large pediatric network, we retroactively test and validate our method. For the test implementation, alerts are filtered with 91-96% accuracy, depending on the parameter values selected. By filtering these alerts, we reduce alert fatigue and allow users to refocus resources to potentially vital alerts, reducing the occurrence of adverse drug events.

*Corresponding author.

1. Introduction

Electronic medical records (EMRs) are being used by a growing number of healthcare providers to improve quality of care and patient safety. EMRs provide clinician support by offering customized, patient-specific decision support. When a provider orders a medication for a patient, a drug safety alert can notify the provider of any potential problem, including a patient allergy, incorrect dosage, or interaction with another drug. Alerts can be interruptive (i.e. a pop-up) or non-interruptive (i.e. highlighted text in a section of the order screen). When an alert appears, users have the option to cancel the medication order or override (dismiss) the alert.

Drug safety alerts can help avoid adverse drug events (ADEs), playing an important role in improving patient safety.^{1,2,3,4} However, alerts are only effective when users deem them to be relevant and therefore worth time and attention. If alerts are often inappropriate, they become ineffective and bothersome, leading to alert fatigue, a state in which the user becomes less responsive to alerts in general.

In this case, vital alerts can be missed, increasing the risk of ADEs and negatively impacting patient safety. In recent research, 49 to 96% of alerts are overridden by EMR users.^{5,6,7} This suggests that potentially, many existing alerts could be safely filtered (hidden from users).

Specific alert reduction strategies have been proposed by recent research. Shah et al. (2006) achieve significantly higher user compliance with interruptive alerts in ambulatory care by making non-interruptive all non-critical or low-severity alerts.⁸ Seidling et al. (2009) achieve positive results by incorporating dosage checking into alerts for dose-dependent drug-drug interaction alerts.⁹ Hsieh et al. (2004) and Swiderski et al. (2007) both studied doctors' decisions regarding drug allergy alerts and found that users usually overrode alerts based on non-exact matches to the patient's allergy list.^{10,11} They propose limiting such alerts.

However, some efforts to reduce alert fatigue have been ineffective. Lo et al. (2009) did not find improved physician action when they made all medication alerts in ambulatory care non-interruptive.¹² Van der Sijs et al. (2008) surveyed doctors about overridden alerts, but found that physicians did not agree on which alerts could be safely filtered.¹³ The failure of some interventions demonstrates the importance of validating and refining methods to achieve a beneficial result.

This research builds upon the work of Hsieh et al. (2004), Swiderski et al. (2009), and Van der Sijs et al. (2008).^{10,11,13} Similar to Van der Sijs et al. (2008), we identify highly overridden alerts based on past data.¹³ Building upon the prior work, we use a large dataset and incorporate cognitive and statistical methods to better predict future overrides. Finally, we present an automated decision tool to identify potentially filterable alerts in any healthcare system.

2. Methodology

Two assumptions are adopted throughout this paper. First, filtering alerts that are currently overridden by EMR users will improve patient safety, as users can

concentrate on relevant alerts to prevent harmful ADEs. Second, the alert quantity and frequency currently typical in most hospitals and other healthcare systems is extreme and should be carefully reduced to a manageable and effective level.

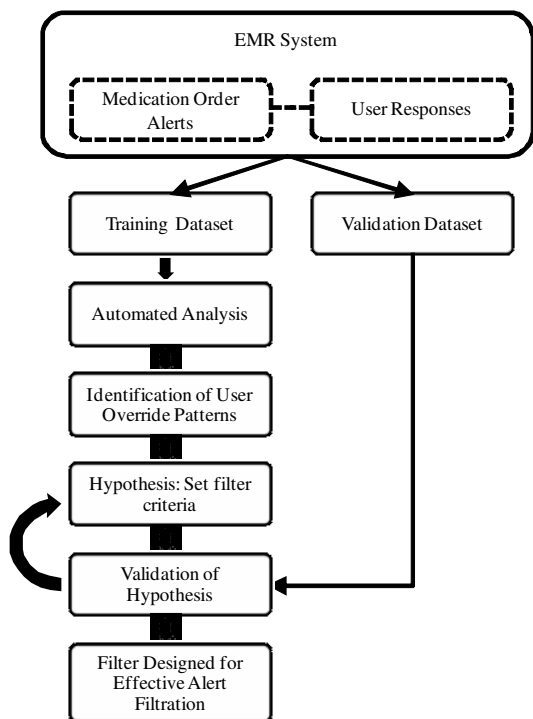


Figure 1. Methodology for filter design.

The objective of this research is to develop a method to identify consistently overridden alerts, which can then be automatically filtered to reduce alert fatigue and improve patient care and safety.

The steps described below are (1) divisional categorization of alerts, (2) hypothesis building by setting parameter values for three alert filtration criteria, and (3) hypothesis validation using independent alert data. The complete methodology is presented in Figure 1.

2.1 Definition of Alerts

The following alert characteristics are utilized in this analysis (Figure 2):

- **Alert Type:** dose alert, drug allergy alert, etc.
 - **Severity** of alert (e.g. moderate, severe)
 - **Drug or Drug Interaction** that triggered alert
 - **Filtered:** 1 if alert was filtered; 0 otherwise
 - **Override:** 1 if alert overridden by user; 0 otherwise
- An alert is uniquely defined by alert type, severity and drug or drug interaction.

2.2 Establishing Initial Filter Hypothesis

We use three criteria to analyze each alert. Alerts that meet all these criteria are expected to be highly overridden by users:

- **Sample Size (scalar):** number of alert instances in the data set
- **Override Rate (percentage):** the percentage of alert instances that were overridden by the users
- **Override Rate Range (percentage points):** the difference between the alert’s highest monthly override rate and lowest monthly override rate

We create alert filters by initializing with one or more specific values for each criteria.

2.3 Validation of Filter Hypothesis

Using the validation dataset, we retroactively apply the filters created in the previous step. The objective is to filter alerts that are often overridden by users. So in the validation step we test the success of this objective. We calculate the following validation measures:

- **Override Rate** of the alerts in the validation dataset that fell under each filter.
- **Accuracy of the filters** in filtering alerts that were overridden by users in the validation dataset.
- **Errors committed by the filters** in filtering non-overridden alerts in the validation set.

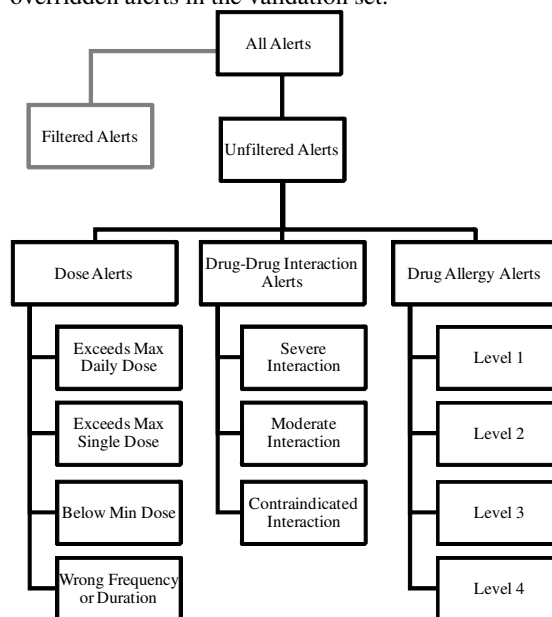


Figure 2. Alert categorization.

Alert Type	Severity	Drug Name/ Interaction	Alert Instances Training Dataset Jan-Jun 2009	Alert Instances Validation Dataset Jul-Sep 2009
Drug-Drug Interaction Alert	Moderate Interaction	137 interactions	1,244	457
	Severe Interaction	66 interactions	3,905	2,039
	Contraindicated Interaction	20 interactions	1,050	441
Dose Alert	Exceeds Max Daily Dose	210 drugs	4,480	2,309
	Exceeds Max Single Dose	112 drugs	672	293
	Below Min Daily/Single Dose	100 drugs	485	159
	Exceeds or Below Frequency/Duration	61 drugs	194	74
	Error Checking Dose	253 drugs	2,245	636
Drug Allergy Alert	Level 1	48 drugs	229	175
	Level 2	29 drugs	355	230
	Level 3	115 drugs	1,630	897
	Level 4	57 drugs	1,009	504
1,208 alerts			17,498 alert instances	8,214 alert instances

Table 1. (Unfiltered) alert identification and alert instances for training and validation datasets.

3. Results

In this section we describe a test implementation of this methodology at Children’s Healthcare of Atlanta (CHOA), a network of three pediatric hospitals in Atlanta, Georgia, which serves over half a million patients annually.

3.1 Medical Alert Data

Nine months of medication order alert data, January through September 2009, was provided by CHOA from the medication ordering system. The data source was two CHOA hospitals, CHOA at Egleston and CHOA at Scottish Rite. Data spanned all departments, including ED, NICU, PICU, cardiac, and general medicine. This dataset was separated into a training dataset of the first six months of data and a validation dataset of the final three months of data.

Based on the alert categorization described in Section 2.1, 1,208 different alerts were identified. The breakdown of these alerts is displayed in Table 1. The rankings of severity of ADE’s is a definition downloaded into the EMR system from the First DataBank. A high-level comparison of the training and validation datasets is given in Table 2.

3.2 Establishing Filters from the Training Dataset

In this implementation, based on the human cognitive decision for overriding preference, the initial values for the three alert filtering criteria are set as follows:

- Override rate: 90% or 95%

- Override rate variation: 10 or 15 percentage points
- Sample size: 50 alerts

Table 3 shows a two-by-two matrix that represents the four different combinations of the selected criteria values. Filter 1 is the most conservative filter and would filter alerts that had at least 50 or more occurrences, had 95-100% override rate, and 0-10 percentage point override rate variation. By sequentially adding filters 2, 3 and 4, the number of alerts filtered would increase.

	Training Dataset	Validatio n Dataset
Months of data	6	3
Number of Alerts	17,498	8,214
Alerts by Alert Type (Dose/Interaction/Allergy)	46/35/18	42/36/22
Override % of Alerts	63.7%	70.4%
Override % by Alert Type (Dose/Interaction/Allergy)	51/71/83	61/76/80

Table 2. Basic alert statistics of training and validation datasets.

Table 4 shows the percentage of alerts from the training set that fall under each filter. We note that no drug allergy alerts were filtered, and Filter 3 does not capture any alerts. Thus we exclude analysis of drug allergy alerts and Filter 3 from the validation described in Section 3.3.

Override % Variation	Override Rate		Sample Size
	≥ 95%	90-95%	
≤ 10 % pts	Filter 1	Filter 2	50+
≤ 15 % pts	Filter 3	Filter 4	

Table 3. Design of filtering criteria from training set.

	Drug Allergy	Dose	Drug-Drug	Total
Filter 1	0.0%	4.3%	1.9%	6.2%
Filter 2	0.0%	0.0%	11.8%	11.8%
Filter 3	0.0%	0.0%	0.0%	0.0%
Filter 4	0.0%	3.0%	1.1%	4.1%
Total	0.0%	7.4%	14.8%	22.2%

Table 4. Percentage of training set alerts that fall under each filter.

	Drug Allergy	Dose	Drug-Drug	Total
Filter 1		96.7%	93.7%	95.9%
Filter 2			91.2%	91.2%
Filter 3				
Filter 4		96.1%	94.2%	95.6%

Table 5. Override accuracy of unfiltered alerts in validation dataset that fell under each filter.

	Overridden	Not Overridden	Total
Filtered	1,568 25.7%	106 1.7%	1,674 27.4%
Filter 1	488 8.0%	21 0.3%	509 8.3%
Filter 2	709 11.6%	68 1.1%	777 12.7%
Filter 4	371 6.1%	17 0.3%	388 6.4%
Not Filtered	2,571 42.2%	1,849 30.3%	4,420 72.5%
Total	4,139 67.9%	1,955 32.0%	6,094 100%

Table 6. Confusion matrix comparing filters with user actions in the validation dataset.

3.3 Validation of Filters

To assess the performance of our filters in predicting the human decision to override, we apply them to a

separate, independent validation dataset, as described in Section 3.2. The override accuracy scores realized by applying our filters to the validation dataset are displayed in Table 5. For example, when applied to the validation dataset, Filter 1 was in agreement with the human decision whether or not to override dose alerts 96.7% of the time; was in agreement with the human decision pertaining to drug-drug alerts 93.7% of the time; and had overall agreement 95.9%.

Table 6 contrasts the performance of filters versus the user action (with actual numbers of alerts included). It offers a glimpse of the importance of an adaptive filtering schema, where filters can be established and continued to be refined through the time period for future prediction.

4. Discussion

In this research we have designed an automated filtering system that captures the decision patterns of EMR users to identify alerts that are highly overridden by users. This filtering system has shown a high degree of accuracy through a large-scale implementation and validation. However, we caution that a small percentage of alerts filtered by our decision tool are alerts that are not overridden by clinical users. Thus, we must analyze in the clinical setting the risk and benefits towards such usage. We emphasize that patient safety is affected by many factors, and while it is improved by the appearance of important alerts, it is hindered when the high quantity of alerts leads to alert fatigue and dismissal of important alerts.

In this case, 1,568 superfluous alert instances were avoided over three months at the expense of 106 “real” alert instances being filtered as well (Table 6). Though those 106 instances may have been important, so were the benefits of filtering 1,568 unnecessary instances: additional time, concentration and confidence in the relevance of EMR alerts. Therefore, part of the on-going research is to find a compromise between the quantity and quality of alerts and evaluate the risk and benefits of such an automatic filtering system.

The model we propose for identifying such alerts can be used by healthcare systems as a starting point to filter superfluous alerts. Before this analysis is converted into full implementation, each identified alert category should be carefully analyzed for its potential impact on patient safety. Doctors should be consulted on the risk of filtering each alert category and on their ability to continue treating patients

safely. Such analysis should focus on a key question: would filtering each alert category result in net benefit or net harm to patients? There is often no easy or clear answer to this question, as it will come down to achieving a balance between the potential benefit in avoiding an ADE of such an alert and its opportunity cost.

The overarching goal of alerts and alert filters is improved patient safety, hence critical analysis of tradeoffs must be performed to ensure that any alert filter would result in a net increase to patient safety.

Finally, each healthcare system should customize the criteria that are used to design filters: override rate, override rate variation and sample size. The validation step outlined in Section 3.3 should be used to assess and adjust the settings to achieve higher filter accuracy or broader alert filtration according to the goals and priorities of each healthcare system.

5. Conclusion

As implementation and utilization of EMR systems continue to grow, it becomes increasingly important to optimize the features and benefits offered through EMRs. There are several reasons to focus on optimizing the appearance of EMR drug safety alerts. First, drug safety alerts can improve patient safety by avoiding ADEs. More sophisticated alert systems that deliver more targeted, appropriate alerts would further improve patient safety by directing users' attention to necessary alerts and avoiding the potential conception that alerts tend to be of low importance or relevance. Furthermore, processing inappropriate or low-quality alerts wastes valuable time and resources.¹⁴ Finally, EMR adoption is a priority for many healthcare systems in the near future and is an important step in the continued modernization of healthcare systems. Developing more targeted and informative EMR systems will not only improve their effectiveness, but will result in more attractive products, aiding in the widespread adoption of EMR systems.

EMR drug safety alerts have great potential to reduce adverse drug effects and improve patient safety when properly implemented and utilized. This research aids in that process by providing a way to customize and improve an EMR alert system to meet the needs of individual healthcare systems and their patients.

Acknowledgement

We acknowledge funding from the National Science Foundation.

References

1. Schedlbauer A, Prasad V, Mulvaney C, et al. What evidence supports the use of computerized alerts and prompts to improve clinicians' prescribing behavior? *JAMIA*. 2009; 16:531-8.
2. Galanter WL, Didomenico RJ, Polikaitis A. A trial of automated decision support alerts for contraindicated medications using CPOE. *JAMIA*. 2005; 12:269-74.
3. Kazemi A, Ellenius J, Pourasghar F, et al. The effect of CPOE and decision support system on medication errors in the neonatal ward: experiences from an Iranian teaching hospital. *J. of Medical Systems*. 2009.
4. Kucher N, Puck M, Blaser J, et al. Physician compliance with advanced electronic alerts for preventing venous thromboembolism among hospitalized medical patients. *J. of Thrombosis and Haemostasis*. 2009; 7:1291-6.
5. Taylor LK, Kawasumi Y, Bartlett G, Tamblyn R. Inappropriate prescribing practices: the challenge and opportunity for patient safety. *Hlthcare Qtrly*. 2005; 8:81-5.
6. Van der Sijs H, Mulder A, Van Gelder T, et al. Drug safety alert generation and overriding in a large Dutch university medical centre. *Pharmacoepidemiology and Drug Safety*. 2009; 18:941-7.
7. Van der Sijs H, Aarts J, Vulto A, Berg M. Overriding of drug safety alerts in CPOE. *JAMIA*. 2006; 13:138-47.
8. Shah NR, Seger AC, Seger DL et al. Improving acceptance of computerized prescribing alerts in ambulatory care. *JAMIA*. 2006; 13:5-11.
9. Seidling HM, Storch CH, Bertsche T, et al. Successful strategy to improve the specificity of electronic statin-drug interaction alerts. *European J. of Clinical Pharmacology*. 2009; 65:1149-57.
10. Hsieh TC, Kuperman GJ, Jaggi T, et al. Characteristics and consequences of drug allergy alert overrides in a CPOE system. *JAMIA*. 2004; 11:482-491.
11. Swiderski SM, Pedersen CA, et al. A study of the frequency and rationale for overriding allergy warnings in a CPOE system. *J. of Patient Safety*. 2007; 3:91-6.
12. Lo HG, Matheny ME, Seger DL, et al. Impact of non-interruptive medication laboratory monitoring alerts in ambulatory care. *JAMIA*. 2009;16:66-71.
13. Van der Sijs H, Aarts J, van Gelder T, et al. Turning off frequently overridden drug alerts: limited opportunities for doing it safely. *JAMIA*. 2008; 15:439-448.
14. Judge J, Field TS, DeFlorio M, et al. Prescribers' responses to alerts during medication ordering in the long term care setting. *JAMIA*. 2006; 13:385-90.