

A Potential Answer to the Alert Override Riddle: Using Patient Attributes to Predict False Positive Alerts

Timothy I. Kennell Jr., BS and James J. Cimino, MD
Informatics Institute, University of Alabama at Birmingham, Birmingham, AL

Abstract

Electronic health records (EHRs) use alerts to help prevent medical errors, yet clinicians override many of these alerts due to desensitization from constant exposure (alert fatigue). We hypothesize that a clinician might override an alert warning about the dangers of a treatment if the patient's health is so poor that the treatment is worth the risk or if a patient's health suggests the treatment is not needed. We used logistic regression with general estimating equations to determine if the Early Warning Score (EWS), a measurement used to predict critical care need, could be used to predict alert overrides. EWS was a significant predictor of overrides for three alerts. Although EWS could not predict overrides for all alert rules, these results suggest that EWS may be helpful for some alerts, but that additional EHR data will be needed for predicting override behavior to a useful degree.

Background and Significance

A hallmark of electronic health record (EHR) systems is the implementation of logic to issue automated alerts and reminders as a form clinical decision support (CDS). Many of these alerts provide clinicians two options: an override option and a non-override option (typically an order cancellation). While this type of CDS can provide useful information when relevant, the messages it provides (referred to in this paper as “alerts”) are often judged as false positives by the receiving clinicians, leading to high override percentages.¹ One study showed that primary care physicians can receive up to 56 alerts per day just from an email notification system² and several studies point to override percentages in the range of 49 – 96%.^{1,3,4} Excessive override rates are problematic because they disrupt workflow, reduce the clinician's quality of life, and can produce “alert fatigue” that may result in “pseudo-false positive” alerts that are inappropriately overridden, to the possible detriment of the patient.⁵

Methods of reducing the alert volume and the override rate typically focus on improving alert logic to increase its specificity (that is, replace false positives with true negatives).¹ Such approaches are complicated by the fact that clinicians do not always accurately report their reasons for overrides.⁶ Other approaches have included manual review of the literature to prioritize alerts,⁷ maintenance of a drug-drug interaction database containing alert priorities,⁸ and dashboard construction, allowing easier manual intervention.⁹ One study proposed an automated method for predicting overrides per clinician based on previous actions by the specific provider in question and suppressing such alerts going forward.¹⁰ Each of these methods is either time-consuming or not patient-specific.

While alerts typically use basic criteria to determine whether an alert should fire, studies indicate that physicians regularly consider more than what the EHR is using.¹¹ We presume that a significant number of alerts are overridden, not because of fatigue, disagreement with the alert logic, or general stubbornness, but because the clinician has some awareness of the patient's particular situation that is not covered by the logic. Indeed, although an override of an alert is not the same as a false positive alert, studies have shown many alert overrides are appropriate.¹² We suspect that information about the situation may be available in the EHR at the time of the alert and could be used to suppress alerts that are predicted to be overridden. To our knowledge, the use of just-in-time patient-specific data outside that used in the alert logic itself has not been studied. Despite the limited criteria used by alerts, the EHR is overflowing with additional information that might reflect a physician's decision when overriding an alert. One such example might be a patient's overall health status. For example, a clinician might ignore a reminder to administer an influenza vaccine if the patient is currently undergoing cardiopulmonary resuscitation. Such an overall health status of a patient is readily available in most EHRs called the Early Warning Score (EWS) (see **Table 1** for details on the score calculation.). The EWS is an aggregate score calculated from a patient's vital signs each time they are recorded and applies increasing values to a patient's worsening vital signs in order to identify patients at risk of deteriorating.^{13,14} The score is used for both prognosis of disease and indication for intensive care admission.¹³

Table 1. Early Warning Score Calculation. The column on the left indicates the parameter measured and the remaining column headers indicate the points added based on the range that the parameter falls into. Table format from Subbe, et. al.¹²

Physiological Parameter	Points						
	3	2	1	0	1	2	3
Systolic Blood Pressures	≤ 90	91–100	101–110	111–219			≥ 220
Temperature in °F	≤ 95		95.1–96.8	96.9–100.4	100.5 - 102.2	≥ 102.3	
Heart Rate	≤ 40		41–50	51–90	91–110	111–130	≥ 131
Respiratory Rate	≤ 8		9–11	12–20		21–24	≥ 25
Oxygen Saturation	≤ 91	92–93	94–95	≥ 95			
Oxygen Therapy	Responses not equal to Room air or CPAP adds 2 points						
Level of Consciousness	Not Awake, alert; Awake, sedated; Obeys commands; Responds to pain; Sleeping, difficult to arouse; Sleeping, easy to arouse; or Unable to assess adds 3 points						

Objective

We reason that extenuating information present in the record, such as the EWS, could be used to suppress irrelevant alerts, even when the alert logic does not consider it. Furthermore, we hypothesize that clinicians override alerts, in part, because of something they know about the patient that is not represented in the alert logic. This paper describes a study of the correlation of patient characteristics *unrelated to specific alert logic that is nevertheless present in the EHR*, with the intent of using such correlations to suppress alerts that are likely to be overridden anyway, thus reducing their adverse effects. We report the application of the above approach with the use of a general measure of patient condition severity.

Methods

Data Source

All data used in the study came from the University of Alabama at Birmingham’s (UAB) University Hospital. Specifically, the information came from both the Cerner Millennium EHR and PowerInsight, a research copy of the EHR’s database. All data were acquired with IRB approval for data reuse with waiver of consent.

Alerts

Initially, we acquired all alerts events ever triggered in the EHR from the time of its implementation in 2010 to 2017. Events that resulted from alerts having no override option were removed from the set as they did not pertain to the objective of the study. The remaining alerts are listed in **Table 2** with a brief description to supplement the name of the alert. Despite the apparent duplication of some alerts, these function in different, nuanced capacities in UAB’s EHR.

Table 2. Descriptions of alerts used in the analysis of alert overrides and the early warning score.

Alert Name	Alert Description
ALLERGYDRUG	Patient allergic to ordered medication
AMB_RULE10_FEC_SPEC_ALERT	Order requires a fecal specimen not yet collected
AMB_RULE11_TIMED_UR_ALERT	Order requires a urine sample not yet collected
AMB_RULE30_LUNG_CANCER	Prior, negative lung cancer CT screen within 24 hrs.
AMB_RULE8_FIT_TEST_ALERT	Fecal iron test ordered before specimen is collected
DRUGALLERGY	Patient allergic to ordered medication

DRUGDRUG	Ordered medication interacts with current medication
DRUGDUP	Ordered medication is a duplicate of current medication
DRUGFOOD	Ordered medication interacts with patient's diet
LH_ADVSR_VTE_2	Prompts for a venous thromboembolism screening
UAB_ADE16_NITRO_ACIDOSIS	Nitroprusside medication ordered on a patient with contraindicated labs
UAB_ADE1_HT_WT_ALLERGY	Requires documentation of a patient's allergies, height, and weight
UAB_ADE2_RCMDMKIDNEY	Radiology study with contrast ordered on a patient with contraindicated labs
UAB_ADE6_DIGOXINLAB	Digoxin-type medication ordered on a patient with contraindicated labs
UAB_ADE8_METFORM_ACIDOSIS	Metformin prescribed to a patient with contraindicated labs in the last 14 days
UAB_ADE9_ANTIHIISTELD_V2	Anticholinergic or antihistamine medication ordered on a patient over 65
UAB_RULE101_MEDS_10_DAYS	Medication order signed with start date more than 10 days in the past
UAB_RULE101_MED_10DAYS_MM	Medication added to scratch pad with start date more than 10 days in the past
UAB_RULE105_DUP_PSYCH_MED	Two or more psychotropic agents ordered on scratch pad
UAB_RULE105_PSYMEDS_DUP	Two or more psychotropic agents ordered and signed
UAB_RULE111_SULF_UREA_ALT	Sulfonylurea ordered on a patient >50 years old with creatinine clearance <60
UAB_RULE158_ED_FLOOR_ADM	Suggests admission to ICU for patients matching lab values or vital signs
UAB_RULE175_TRANSFUSION	Transfusion ordered on a patient with a hemoglobin ≥ 7 g/dL
UAB_RULE184_CHESTAPALERT	Chest x-ray ordered with existing one less than 3 days old
UAB_RULE186_PLATELET_TRAN	Platelet order placed without results within 30 days or without indication
UAB_RULE209_METHOTREX_CRC	Methotrexate ordered on a patient with creatinine clearance \leq to 50
UAB_RULE22_RAD_HVC_WT_LMT	Radiology study is ordered on a patient weighing > 225 kg
UAB_RULE22_RAD_NM_WT_LMT	Nuclear medicine study on a patient weighing over 158.76 kg
UAB_RULE22_RAD_PET_WT_LMT	PET scan ordered on a patient over a pre-defined weight limit
UAB_RULE22_RAD_WT_HIGHLAN	Radiology study ordered on a patient over equipment weight limit
UAB_RULE22_UABH_C_CT_WTLM	Radiology study ordered on a patient over equipment weight limit
UAB_RULE232_CDIFF_TOXIN	C diff toxin test ordered on a patient taking a laxative
UAB_RULE255_MYASTHENIAGRA	Contraindicated medication ordered on a patient with myasthenia gravis
UAB_RULE257_NAC_TYL	Tylenol ordered on a patient taking acetylcysteine
UAB_RULE268_HALOPERDL65YO	Haloperidol is prescribed to a patient > 65 years old
UAB_RULE27_HIT_AB	Heparin ordered on a patient with anti-heparin antibodies
UAB_RULE282_RAD_CTLUNG_SC	Requires a completed form before performing a lung cancer CT screen
UAB_RULE28_STERILE_WATER	Sterile water without an electrolyte additive ordered
UAB_RULE35_WAR_INF_SIGN	Warfarin ordered and signed without an INR within 7 days
UAB_RULE35_WAR_INR_SCRATC	Warfarin ordered on scratchpad without an INR within 7 days
UAB_RULE35_WAR_INR_SR_AMB	Warfarin ordered and signed without an INR within 7 days (ambulatory rule)
UAB_RULE37_PO_CONV_ALERT	Recommends IV to oral medication conversion if a patient's diet order suggests the action would be appropriate.
UAB_RULE61_MRI_SAFETY_SCR	Requires completed safety form before signing MRI order
UAB_RULE62_MEPPERIDINECRCL	Meperidine ordered on a patient with a previous creatinine clearance or glomerular filtration rate < 50 mL/min.
UAB_RULE63_US_GUIDED_LAB	Ultrasound procedure ordered without platelet lab results
UAB_RULE63_US_GUIDED_LB_C	Ultrasound procedure ordered without platelet lab results
UAB_RULE90_DOPAMINERGIC	Contraindicated medication ordered for a patient taking dopamine
UAB_RULE90_DOPAMINERGIC_2	Contraindicated medication ordered for a patient taking dopamine
UAB_RULE90_DOPAMIN_MEDM	Contraindicated medication ordered for a patient taking dopamine
UAB_RULE90_DOPAMIN_MEDM2	Contraindicated medication ordered for a patient taking dopamine
UAB_RULE92_ACE_PHER_ALERT	ACE inhibitor is ordered for a patient scheduled for plasmapheresis
UAB_RULE92_ACE_PHER_MEDM	ACE inhibitor is ordered for a patient scheduled for plasmapheresis

Patient Data

In addition to acquiring all alerts for the seven-year period, all patient data recorded during that time period were retrieved. The information included the date of birth for age calculation along with basic demographic information for the analysis such as race and sex. Duplicate patients (based on possession of identical medical record numbers) in the data set were combined while retaining the most specific demographic information. For example, if two duplicate patient records had a race of “Unknown” and “African American”, the final aggregated patient record’s race was set to “African American”.

Early Warning Score

After combining the appropriate patient information with each alert event, the early warning score was added. Since the patient’s vitals, and by extension the EWS, are not guaranteed to be measured at the time an alert event triggers, a reasonable estimate was acquired using the following procedure. For each alert event, a time window consisting of 12 hours before and 2 hours after was created. Then, the EWS within that window and occurring closest to the time the alert triggered was taken to be the patient’s EWS at the time of the event.

Using 12 hours or less prior to the alert increased the likelihood that the EWS would be from the same hospital stay as the alert event and hence more likely to accurately reflect the patient’s current status. Likewise, searching two hours after the alert made it possible to detect EWSs occurring slightly after the alert that would more accurately reflect the patient’s status compared to an EWS taken ten hours prior to the alert. Any alert events without an EWS within the time window were excluded from analysis.

Statistical Analysis

In order to determine if a patient’s overall health status as indicated by the EWS was predictive of an alert’s override likelihood, a logistic regression model with general estimating equations (GEE) was employed. Logistic regression allows modeling of the binary outcome: override or non-override. While capable of modeling this type of outcome, logistic regression expects independence of each alert event. However, an alert can occur on the same patient, breaking the independence assumption of the model. The addition of GEE to the analysis corrected for this non-independence by estimating correlation between alert events.

Following de-identification of the data, the logistic regression model with GEE was performed using SAS 9.4. Each alert was tested in a separate GEE model in order to make interactions between a specific alert and the EWS levels more easily interpretable (i.e. allow negative versus positive correlations between EWS and an alert to be more recognizable). Patient demographics served as covariates in the model in order to account for any variation that might occur due to these categories. Additionally, the amount of time between the EWS measurement and the alert event might affect the accuracy of EWS and, as a result, the override likelihood. Therefore, both the time between the EWS and the alert along with a time / EWS interaction term were included. A conservative Type III analysis was used for each categorical variable. Finally, due to the multiple models being run simultaneously, a Bonferroni correction was applied to all p-values. As 46 models were run, an alpha value of 0.0011 (0.05 / 46) was considered to be significant for Type III analyses.

Results

The UAB EHR contains 52 alerts that offer physicians an override option (see **Table 2**). Over the seven-year period (starting at the launch of UAB’s EHR), these alert rules represent a total of 4,043,431 unique alert events with approximately 17.37% overridden on average. The events per alert rule range from 1 to 1,998,811 (see **Figure 1**). Even after excluding rules that occurred less than 1000 times, override percentages ranged from 0.02% or 100% (see **Figure 1**).

Additionally, the basic demographic information of the UAB patient population was retrieved for inclusion in the models. This information included patient sex, race, and age at the time of the alert. The breakdown of the patient population during the time that alert event data was pulled is shown in **Table 3**.

The appropriate patient demographic information was combined with each alert event. Then, each alert event was combined with a patient’s EWS closest to the event and occurring within 12 hrs. before the event or 2 hrs. after the

event (see Methods for details). Alert events without an EWS within this time window were excluded leaving a total of 46 alert rules spanning 443,019 events from the original 52 alert rules and 4,043,431 events.

These 46 alerts were then individually analyzed with logistic regression incorporating GEE. Due to limitations of the data for some alerts, 20 were unable to be analyzed for various reasons. The analysis removed eight alerts due to one of the variables having only one value (e.g. all alert events in the data set were overridden), and it removed twelve alerts for insufficient variation in one or more variables. The alerts events in the data with only one type of override status had either a low number of events in the combined data set, extreme override frequencies (high or low), or both.

Table 3. UAB Patient demographics

Demographic	Category	N (%)
Gender	Female	921,874 (54.84%)
	Male	754,637 (44.89%)
	Other	4,384 (0.26%)
Race	Other	1,517,369 (90.27%)
	White	109,209 (6.5%)
	African American	54,317 (3.23%)
Age	Mean (yr.)	54.7
	Standard Deviation (yr.)	17.1

The results of the remaining alerts are shown in **Table 4**. After using a Bonferroni correction ($\alpha = 0.05/46 = 0.0011$) to account for using multiple models (one per alert), EWS only predicts override frequency for three alerts: DRUGDRUG, DRUGDUP, and UAB_RULE184_CHESTAPALERT. For two of these alerts (DRUGDRUG and UAB_RULE184_CHESTAPALERT), the time term is also significant in addition to the separate EWS and time terms, suggesting that the time between the EWS measurement and alert event affects the utility of EWS in predicting an alert's override. Additionally, the interaction term between time and EWS reached significance for a few of the alerts, including the DRUGDRUG and the UAB_RULE184_CHESTAPALERT. This information suggests that both the time between the alert event and the EWS measurement and the EWS itself combine to produce different likelihoods of an alert overrides.

These results can further be visualized in **Figure 2**. This representation depicts the override percentage of an alert at each EWS level. Although there is not a drastic increase or decrease in override percentage with increasing EWS for the three significant alerts, the amount of data analyzed allows detection of smaller effects. Additionally, as in the UAB_RULE184_CHESTAPALERT alert, the baseline override percentage is already high, and does slightly increase with increasing EWS.

Discussion

EHRs use alerts to warn providers of potential errors or provide suggestions on patient care. The large number of alert events that occur have potentially decreased their utility in certain instances and introduced other issues including alert fatigue. Although most EHRs give the provider a default list of override reasons, these do not always accurately reflect the reasoning behind the override decision.⁶ Additionally, appropriate overrides are typically the result of the clinician considering an aspect of the clinical context not considered by the alert logic.^{11,15}

This study investigated the potential of using a patient's health status (EWS) to predict an alert's override. Presumably, clinicians might override minor workflow suggestions as a patient's health status deteriorates and adhere to severe alerts (e.g. drug-renal alerts) more frequently as a patient's EWS increases. Although overrides and false positive alerts are not identical, many alerts are appropriately overridden.¹² Therefore, targeting override prediction is a good first step to removing potentially unnecessary alerts. In order to be as comprehensive as possible, all overridable alerts at UAB were analyzed; however, a patient's health status might not affect a clinician's decision for all alerts. For example, while the UAB_ADE8_METFORM_ACIDOSIS alert warning of a contraindicated prescription of metformin might cause a clinician to consider the severity of the patient's condition, the

UAB_ADE1_HT_WT_ALLERGY alert requiring patient allergy, height, and weight information be added to the record might not have the same effect.

Although EWS did not achieve significance for explaining overrides for many of the alerts, it did for three: drug-drug interaction (DRUGDRUG), drug duplicate orders (DRUGDUP), and a chest X-ray advisory (UAB_RULE184_CHESTAPALERT). Regarding the two drug alerts, a clinician might consider the medication more important for the deteriorating patient than the potential of an interaction. Indeed, a previous study showed that the two most common override reasons for drug-drug interaction alerts were ‘clinically irrelevant alert’ and ‘benefit assessed to be greater than the risk’,¹⁵ and this result is reflected in the UAB EHR as the most common override reasons for both the DRUGDRUG and DRUGDUP alert are ‘interaction noted, will take precautions’ and ‘essential therapy, will take precautions’. Similarly, the UAB_RULE184_CHESTAPALERT alert warns clinicians that a chest x-ray has already been performed on the patient within 3 days and suggests not ordering another one. A deteriorating patient would likely increase the likelihood of overriding such an alert as shown in the model results (see **Figure 2**). It should be noted that the increase override percentage shown **Figure 2** for these alerts is small and might not offer a threshold for EWS at which the alert completely loses utility. However, this result is likely due to the already high override percentage.

Table 4. Logistic Regression with GEE results for the EWS predictor and the EWS * Time interaction term.

Alert Name	EWS p-value from Type 3 GEE	Time p-value from Type 3 GEE	EWS * Time p-value from Type 3 GEE	EWS Coefficient from GEE
ALLERGYDRUG	0.12	<.0001	<.0001	-0.01
DRUGALLERGY	0.01	0.03	0.03	0.01
DRUGDRUG	<.0001	<.0001	<.0001	0.04
DRUGDUP	<.0001	0.05	0.05	0.11
DRUGFOOD	0.04	0.67	0.67	0.17
UAB_ADE16_NITRO_ACIDOSIS	0.02	0.18	0.18	0.24
UAB_ADE2_RCMDMKIDNEY	0.2	0.98	0.98	-0.03
UAB_ADE6_DIGOXINLAB	0.02	0.99	0.99	0.08
UAB_ADE8_METFORM_ACIDOSIS	0.98	0.75	0.75	0
UAB_ADE9_ANTIHISTELD_V2	0.03	0.01	0.01	0.05
UAB_RULE101_MEDS_10_DAYS	0.52	0.82	0.82	-0.05
UAB_RULE111_SULF_UREA_ALT	0.17	0.51	0.51	0.24
UAB_RULE158_ED_FLOOR_ADM	0.07	0.67	0.67	-0.06
UAB_RULE175_TRANSFUSION	0.54	0.06	0.06	0.01
UAB_RULE184_CHESTAPALERT	<.0001	<.0001	<.0001	0.17
UAB_RULE186_PLATELET_TRAN	0.99	0.04	0.04	0
UAB_RULE22_RAD_NM_WT_LMT	0.45	0.54	0.54	-0.51
UAB_RULE268_HALOPERDL65YO	0.03	0.3	0.3	0.24
UAB_RULE35_WAR_INF_SIGN	0.36	0.01	0.01	0.04
UAB_RULE35_WAR_INR_SCRATC	0.41	0.63	0.63	-0.06
UAB_RULE37_PO_CONV_ALERT	0.01	0.24	0.24	-0.03
UAB_RULE61_MRI_SAFETY_SCR	0.37	0.02	0.02	0.02
UAB_RULE62_MEPPERIDINECRCL	0.16	0.19	0.19	-0.1
UAB_RULE90_DOPAMINERGIC	0.37	0.05	0.05	0.09
UAB_RULE90_DOPAMINERGIC_2	0.51	0.63	0.63	0.11
UAB_RULE92_ACE_PHER_ALERT	0.83	0.55	0.55	0.1

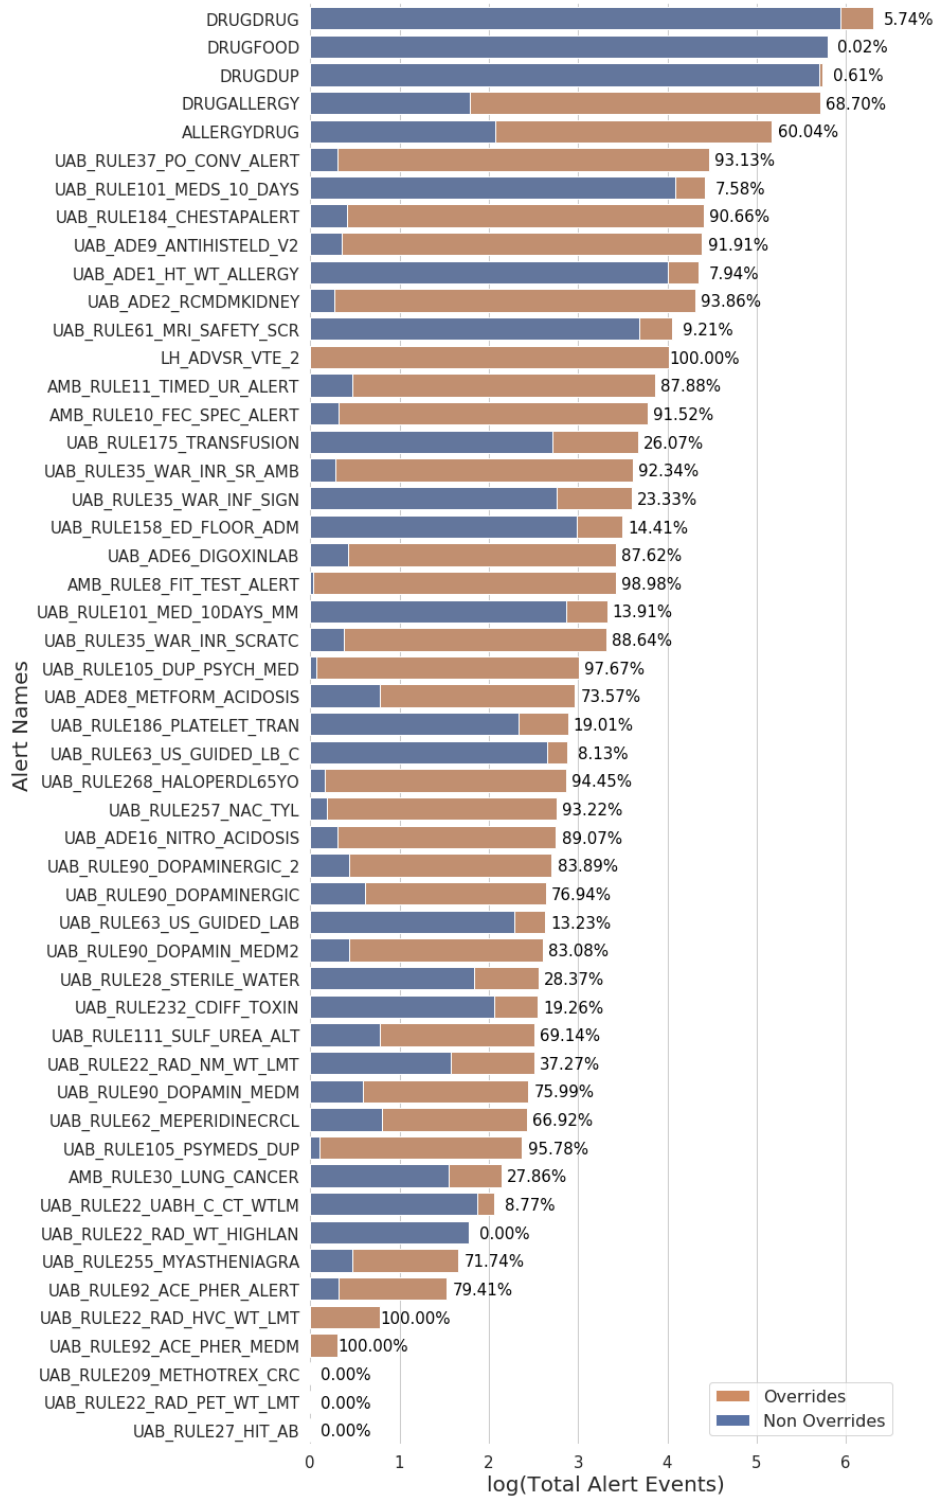


Figure 1. The distribution of alert events in the UAB EHR broken down by override status. Alert rules were limited to those that provided clinicians an override option. The length of the bars indicates the log transformed number of the total alert events triggered by the alert rule on the y-axis. The two colors of the bar reflect the ratio of overridden to non-overridden alert events (alerts events that clinicians chose the alternative option to overriding) with the percentage of each alert rule’s overrides displayed to the right of the bar. Many of these alerts have override rates above 70% showing that the UAB institution has similar override rates to the 49 – 96% reported by other institutions.

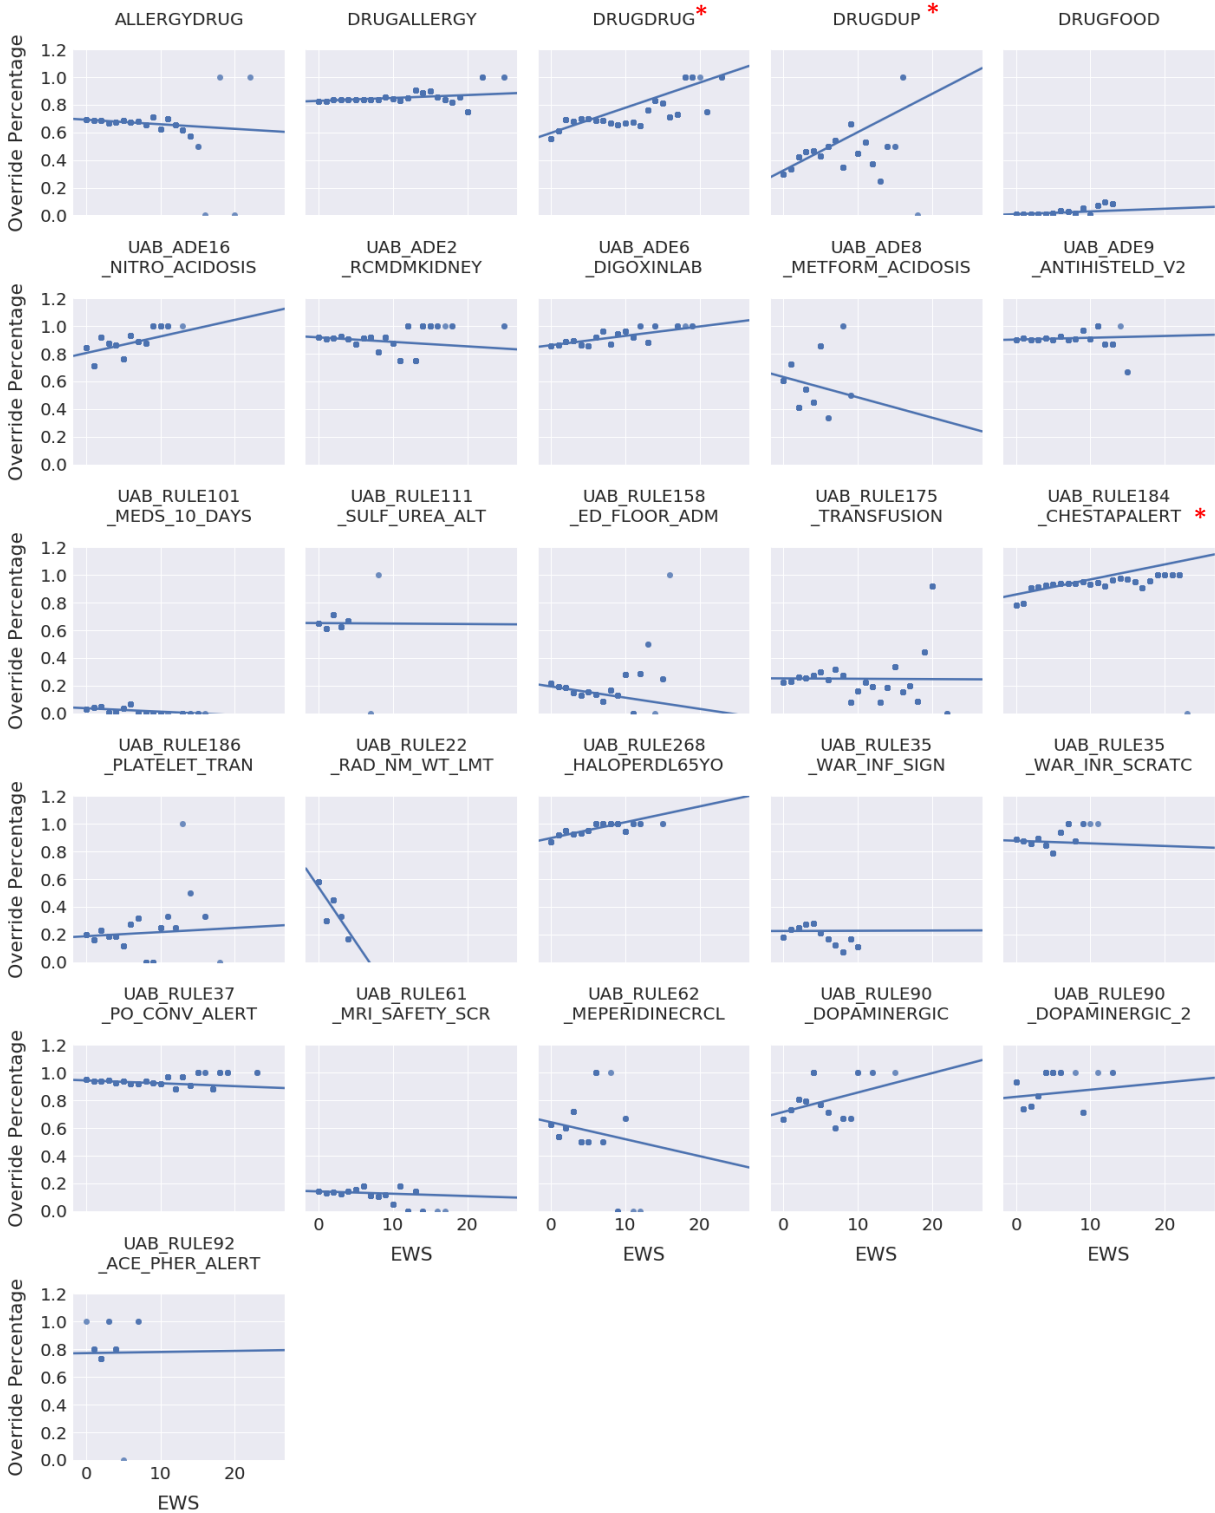


Figure 2. The effect of EWS on the override frequency of the 26 alerts analyzed by logistic regression with GEE. Override frequency per alert per EWS level was calculated and plotted with an estimated line using simple linear regression. Each data point for a specific EWS and override frequency was duplicated to reflect the number of alert events and weight the information. Thus, the trend lines indicate whether an alert is more or less likely to be overridden with deteriorating patient status. The asterisks indicate alerts in which EWS is a significant predictor of alert overrides. Other alerts have some increase or decrease in override frequency with an increasing EWS suggesting that patient status might play a small part in their override likelihood.

Additionally, the interaction term between time and EWS reached significance for a few of the alerts. This result might be an indication of the accuracy of EWS at the time of the alert. Although the method of selecting the EWS was intended to choose the measurement most reflective of the patient's status at the time of the alert event, this result might not have been achieved as closely as desired. Thus, the time between the alert event and the EWS measurement might cause EWS to differentially predict alert overrides.

However, EWS did not significantly correlate with the likelihood of an override for many of the alerts. Some of the reasons for this might be due to the fact that a patient's health status simply does not factor into a physician's decision to override the alert as discussed above. However, another reason might be that many of the alerts already have an override rate above 75% (see **Figure 1**) limiting the ability of the models to detect an effect. Additionally, it should be noted that medical alerts widely vary in the decision support they provide. Some simply suggest cost-saving measures as in the chest x-ray alert (UAB_RULE184_CHESTAPALERT). Others provide minor suggestions for patient care such as warning of false positives on a lab test (UAB_RULE232_CDIFF_TOXIN). Some provide critical care advice like the alert warning of a methotrexate prescription on a patient under a specified creatinine clearance (UAB_RULE209_METHOTREX_CRC). Given this variation, a single patient characteristic is unlikely to provide a strong predictor for the overrides of all alerts.

As some of the alerts did approach significance for EWS as a predictor, more covariates or additional predictors are needed to account for the variance in the models. For example, the UAB_ADE9_ANTIHIISTELD_V, which warns about prescribing an anticholinergic or antihistamine medication to a patient over 65, might benefit from accounting for previously prescribed and tolerated medications. For other alerts, such as the UAB_RULE184_CHESTAPALERT, previous diagnoses or procedures might be effective at determining whether a clinician would override the alert. Newly recorded procedures might prompt a physician to order a patient X-ray in order to track the individual's current status.

Alternatively, given there is a need to remove clinically unnecessary alerts, another more effective approach might be to implement a machine learning model to predict overrides. This methodology would incorporate multiple EHR variables along with complex interactions and might come closer to solving the immediate need of removing unnecessary alert events.

Conclusion

High rates of false positive alerts continue to plague clinicians despite decades of work to reduce them. We have explored a methodology for predicting such alerts based on patient data in the EHR available just prior to the time of the alert. The use of the EWS for this purpose may be one small step toward alert reduction. The EHR is replete with many other possible targets for the methodology we describe here, offering the potential for a giant leap forward in addressing this decades-old problem.

Acknowledgements

Drs. Kennell and Cimino are supported by research funds from the UASOM Informatics Institute. Dr. Cimino is also supported in part by research funds from the Center for Clinical and Translational Sciences (CCTS) under grant UL1TR001417 from the National Center for the Advancement of Translational Science (NCATS).

References

1. van der Sijs, H., Aarts, J., Vulto, A. & Berg, M. Overriding of drug safety alerts in computerized physician order entry. *J. Am. Med. Inform. Assoc. JAMIA* **13**, 138–147 (2006).
2. Murphy, D. R., Reis, B., Sittig, D. F. & Singh, H. Notifications received by primary care practitioners in electronic health records: a taxonomy and time analysis. *Am. J. Med.* **125**, 209.e1–7 (2012).
3. Castillo-Páramo, A. *et al.* Inappropriate prescribing according to the STOPP/START criteria in older people from a primary care setting. *Eur. J. Gen. Pract.* **20**, 281–289 (2014).
4. van der Sijs, H. *et al.* Drug safety alert generation and overriding in a large Dutch university medical centre. *Pharmacoepidemiol. Drug Saf.* **18**, 941–947 (2009).
5. McCoy, A. B. *et al.* A framework for evaluating the appropriateness of clinical decision support alerts and responses. *J. Am. Med. Inform. Assoc. JAMIA* **19**, 346–352 (2012).

6. Dekarske, B. M., Zimmerman, C. R., Chang, R., Grant, P. J. & Chaffee, B. W. Increased appropriateness of customized alert acknowledgement reasons for overridden medication alerts in a computerized provider order entry system. *Int. J. Med. Inf.* **84**, 1085–1093 (2015).
7. Riedmann, D. *et al.* Development of a context model to prioritize drug safety alerts in CPOE systems. *BMC Med. Inform. Decis. Mak.* **11**, 35 (2011).
8. Phansalkar, S. *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.* **20**, 489–493 (2013).
9. Simpao, A. F. *et al.* Optimization of drug–drug interaction alert rules in a pediatric hospital’s electronic health record system using a visual analytics dashboard. *J. Am. Med. Inform. Assoc.* **22**, 361–369 (2015).
10. Lee, E. K., Mejia, A. F., Senior, T. & Jose, J. Improving Patient Safety through Medical Alert Management: An Automated Decision Tool to Reduce Alert Fatigue. *AMIA. Annu. Symp. Proc.* **2010**, 417–421 (2010).
11. Taylor, L. K. & Tamblyn, R. Reasons for physician non-adherence to electronic drug alerts. *Stud. Health Technol. Inform.* **107**, 1101–1105 (2004).
12. Nanji, K. C. *et al.* Medication-related clinical decision support alert overrides in inpatients. *J. Am. Med. Inform. Assoc.* **25**, 476–481 (2018).
13. Akgun, F. S., Ertan, C. & Yucel, N. The prognostic efficiencies of modified early warning score and mainz emergency evaluation score for emergency department patients. *Niger. J. Clin. Pract.* **21**, 1590–1595 (2018).
14. Subbe, C. P., Kruger, M., Rutherford, P. & Gemmel, L. Validation of a modified Early Warning Score in medical admissions. *QJM Int. J. Med.* **94**, 521–526 (2001).
15. Ahn, E. K., Cho, S.-Y., Shin, D., Jang, C. & Park, R. W. Differences of Reasons for Alert Overrides on Contraindicated Co-prescriptions by Admitting Department. *Healthc. Inform. Res.* **20**, 280–287 (2014).