

ORIGINAL RESEARCH

Open access

# Interpretable Machine Learning Model for Predicting Laboratory Alert Fatigue Using Alert Frequency, Clinical Severity, Provider Specialty, Repeated Abnormal Results, Response Time, and Override Behavior

Chen Hao<sup>1\*</sup>, Liu Fang<sup>1</sup>, Zhao Lin<sup>2</sup>

## Abstract

Laboratory alert fatigue erodes the effectiveness of clinical decision support by making repeated abnormal result notifications less likely to prompt timely clinical attention. It is often recognised only after providers begin delaying, overriding, or ignoring alerts. Current alert reduction strategies are commonly based on broad thresholds or blanket suppression rules. These approaches do not explain which providers, specialties, alert types, or repeated result patterns are most susceptible to fatigue. This article proposes an interpretable machine learning model that could predict whether a provider will exhibit fatigued behaviour toward a specific laboratory alert. The model is intended to support transparent, provider-aware alert redesign rather than opaque automation. The proposed framework would use historical alert logs with features capturing alert frequency, clinical severity, provider specialty, repeated abnormal results, response time history, and override behaviour. A regularised logistic regression or gradient-boosted tree model with SHAP explanations would provide both prediction and interpretability. Conceptually, the model would generate a fatigue risk score for each provider–alert pair. It would also attribute the score to specific drivers, such as repeated low-severity results, accumulated alert burden, or recent override patterns. An interpretable fatigue prediction model could enable personalised alert suppression, escalation, or redesign before a clinically important laboratory result is missed. Such a system would support safer, more adaptive clinical decision support.

**Keywords** Clinical decision support, Alert fatigue, Explainable artificial intelligence, SHAP, Laboratory alerts, Override behaviour

\*Correspondence:

Chen Hao

chen.hao@outlook.com

<sup>1</sup> Department of Intelligent Healthcare Informatics, Faculty of Medicine, Zhejiang University, Hangzhou, China

<sup>2</sup> Department of Clinical Data Engineering, Faculty of Engineering, Nanjing University, Nanjing, China

## Introduction

Laboratory alert fatigue refers to the progressive desensitisation that occurs when clinicians receive frequent abnormal result notifications, many of which may be low severity, repetitive, or poorly contextualised. In hospital environments, this fatigue can weaken the safety purpose of laboratory alerting by slowing acknowledgement, increasing overrides, and reducing attention to genuinely critical results. Evidence from clinical decision support systems shows that workload, alert complexity, and repeated exposure can all shape clinician response behaviour [1]. Medication and intensive care alert

studies further demonstrate that override patterns and non-responsiveness can become safety-relevant signals when alerts are frequent and poorly targeted [2, 3].

Current mitigation strategies often depend on static rule tuning, generic threshold adjustment, or broad de-escalation of alert classes. Although these approaches can reduce volume, they rarely distinguish between a provider who is appropriately dismissing a clinically irrelevant alert and one who is becoming desensitised to repeated abnormal results. Reviews of clinical decision support quality and medication alert fatigue suggest that alert design, clinical role, and workflow fit are central to whether an alert is useful or burdensome [4, 5]. This indicates that laboratory alert optimisation should move beyond one-size-fits-all suppression toward context-aware and provider-sensitive modelling.

Electronic health record and laboratory information system logs create an opportunity to model fatigue as a behavioural risk pattern rather than a vague user complaint. Timestamped alerts, acknowledgement intervals, override clicks, repeat abnormal findings, and provider specialty can be converted into predictors of delayed or dismissive response. Prior work on alert overrides, medication-related clinical decision support, and machine learning-based alert filtering shows that audit data can support predictive modelling of alert response behaviour [6-8]. Interpretable methods are especially important because clinicians and governance committees need to understand why a model identifies a specific alert as fatigue-prone [9, 10].

This article proposes an explainable artificial intelligence framework for predicting laboratory alert fatigue at the provider-alert level. The model would use alert frequency, clinical severity, provider specialty, repeated abnormal result patterns, response time history, and override behaviour to estimate the likelihood of a fatigued response. Its purpose is not to replace clinical judgement, but to make alert burden visible, measurable, and modifiable through transparent explanations. By combining predictive modelling with SHAP-based interpretation, the framework could support targeted suppression, escalation, digesting, and alert redesign.

## Background

### Alert fatigue in clinical decision support

Alert fatigue in clinical decision support is commonly explained through repeated exposure to low-yield interruptions that weaken the perceived importance of future alerts. The “cry wolf” effect is particularly relevant when frequent alerts are clinically minor, duplicative, or poorly matched to provider context. Empirical studies of alert overrides and alert burden show that high-volume alerting can create a trade-off between safety coverage and clinician responsiveness [1, 6]. This dynamic is directly relevant to laboratory alerting because repeated abnormal values may be technically correct yet operationally unhelpful when they do not require immediate action.

### Laboratory alert characteristics and clinician response

Laboratory alerts typically arise from critical values, abnormal thresholds, delta checks, or repeated clinically significant deviations from expected ranges. Their usefulness depends not only on the numeric abnormality but also on whether the result is new, worsening, severe, or already known to the care team. Studies of medication-related and renal clinical decision support alerts show that context, appropriateness, and repeated presentation strongly influence whether providers accept or override alerts [11, 12]. For laboratory alerts, persistent mild abnormalities may become a source of desensitisation when the alert fires repeatedly without meaningful clinical change.

### Provider specialty and clinical context

Provider specialty can modify alert response because different specialties interpret the same abnormal result through different clinical priorities and baseline expectations. A mild creatinine increase may be more salient to a nephrologist than to a surgeon managing an unrelated postoperative issue, while electrolyte abnormalities may carry different urgency

in intensive care than on a general ward. Prior alert studies show that clinical role and care setting influence alert burden, appropriateness, and response patterns [2, 5, 13]. Therefore, a laboratory fatigue model should represent specialty not as a demographic label, but as a contextual modifier of alert relevance.

## Measuring alert fatigue from audit logs

Alert fatigue cannot usually be observed directly, so it must be inferred from behavioural proxies such as override rate, acknowledgement delay, repeated dismissal, and lack of downstream action. Audit logs provide the temporal structure needed to distinguish isolated non-response from cumulative fatigue over a shift or recent care episode. Prior predictive and descriptive studies have used alert logs to examine overrides, response patterns, and opportunities for filtering or redesign [8, 14, 15]. In a laboratory setting, the same logic could support fatigue labels based on delayed acknowledgement, repeated override, or low engagement with recurring abnormal result alerts.

## Explainable machine learning for alert optimisation


Explainable machine learning can make alert optimisation more trustworthy by showing which features drive fatigue risk predictions. SHAP, LIME, and partial dependence analyses can reveal whether predicted fatigue is mainly associated with alert frequency, low severity, repeat abnormality, specialty context, or recent override behaviour. Clinical AI scholarship emphasises that explanations must be usable in clinical workflows rather than merely technically accurate [9, 10, 16]. For laboratory alerting, this means explanations should help governance teams decide whether to suppress, escalate, merge, reword, or retime alerts.

# Model Development Overview

## High-level predictive framework

The proposed framework would aggregate alert activity at the provider–alert type–shift level while preserving enough event-level detail to score the next laboratory alert. The target would be a conceptual probability of fatigued response, represented by behaviours such as override without meaningful review, delayed acknowledgement, or repeated dismissal of similar abnormal results. This approach follows the broader direction of using clinical decision support audit data to model responsiveness and alert burden [7, 8]. The model would be designed as an early warning layer in the alert pipeline rather than as a replacement for laboratory critical value policy.

**Figure 1** illustrates the hierarchical explainable machine learning framework linking alert data, feature engineering, interpretable



**Figure 1.** Hierarchical Explainable AI Framework for Predicting Laboratory Alert Fatigue at the Provider–Alert Level

Core input features

The core feature set would include alert frequency per shift, cumulative alert count over a recent rolling window, clinical severity of the abnormal result, provider specialty, repeated-abnormality flags, recent median response time, and recent override proportion. These features reflect the major mechanisms described in alert fatigue literature: workload accumulation, low perceived usefulness, contextual mismatch, and repeated exposure to similar warnings. Prior studies of overrides and alert appropriateness support the value of combining alert content with provider response history [3, 11, 17]. In the laboratory context, these predictors would allow the model to distinguish a new critical result from the fifteenth low-severity repeat abnormality.

Design principles

The model should be intrinsically interpretable where possible and post-hoc explainable where necessary. A regularised logistic regression could provide transparent directional effects, while a gradient-boosted tree model with SHAP could capture nonlinear interactions between alert frequency, severity, and specialty. Prior work on explainable clinical machine learning cautions that interpretability must be aligned with clinical use rather than treated as a decorative add-on [10, 18, 19]. The system should also be computationally light enough for real-time preprocessing and flexible enough to update as provider behaviour and alert policies change.

Data Sources and Feature Engineering

Alert log extraction and response labelling

The primary data source would be timestamped alert events from the laboratory information system and electronic health record, linked to provider acknowledgement, order entry, override, and follow-up action where available. Response labels would be derived conceptually from time-to-acknowledge, time-to-action, repeated dismissal, or override behaviour rather than from subjective fatigue reports alone. Studies of clinical decision support overrides show that these behavioural traces can reveal both appropriate dismissal and potential fatigue-related non-responsiveness [2, 3, 14]. For laboratory

alerts, labelling should remain cautious because a delayed response may reflect competing clinical emergencies rather than disregard for the alert.

## Frequency, severity, and repeated abnormal features

Frequency features would use rolling counts of alerts by provider, alert type, shift, and recent time window, with higher weight assigned to recent exposure. Severity features would map laboratory result deviation into clinically meaningful categories, such as critical versus non-critical, worsening versus stable, and new versus persistent abnormality. Repeated abnormality features would flag whether the same or similar result has appeared recently without a clinically meaningful change. This design is consistent with evidence that alert burden, appropriateness, and repeated exposure influence override and response behaviour [1, 6, 12].

## Provider specialty and response behaviour history

Provider identity would be linked to specialty or service role in a privacy-conscious way, allowing the model to represent differences in alert relevance across clinical contexts. Individual response history would include recent acknowledgement speed, override tendency, and response to similar alert types, while avoiding punitive interpretation of provider behaviour. Role-tailored alerting has been proposed as one way to reduce fatigue because the same alert may be useful for one clinical role and distracting for another [5, 20]. Specialty features should therefore be used to personalise alert relevance, not to stereotype clinicians or obscure workflow constraints.

# Interpretable Machine Learning Architecture

## Model choice and interpretability strategy

A regularised logistic regression model would offer clear coefficients and straightforward governance review, making it attractive when transparency is prioritised over complex interaction modelling. A gradient-boosted tree model with SHAP explanations would be useful when alert fatigue depends on nonlinear combinations, such as high frequency becoming problematic only when severity is low and the abnormality is repeated. Prior explainable AI research in clinical settings supports additive explanations when they help clinicians understand why a prediction was generated [9, 16, 19]. The preferred architecture would therefore balance predictive flexibility with an explanation format that can be reviewed by laboratory, informatics, and clinical leadership.

## Input feature vector and preprocessing

Each alert-provider instance would be represented as a feature vector containing normalised frequency metrics, severity indicators, specialty encoding, repeat-abnormality status, recent response time summaries, and override history. Specialty could be represented through one-hot encoding or grouped service categories, while missing response history for new providers could be handled through service-level baselines. Prior studies of medication alert filtering and alert pattern analysis show the importance of structuring alert metadata and response history into usable modelling inputs [8, 21, 22]. Preprocessing should preserve clinical interpretability so that derived features remain understandable to governance committees.

## Output: fatigue risk score and explanation

The model output would be a calibrated fatigue risk score indicating the likelihood that a provider will respond to the current laboratory alert in a delayed, dismissive, or override-prone manner. A local SHAP explanation could show whether the score is driven mainly by accumulated alert count, low clinical severity, repeated abnormal result status, specialty context, or recent slow response to similar alerts. This mirrors the broader movement toward explainable clinical AI, where predictions must be accompanied by reasons that support safe action [10, 18, 23]. In practice, the explanation might justify digesting a non-critical repeat alert while escalating a high-severity alert that is at risk of being missed.

# Capturing Alert Frequency, Severity, and Repeated Abnormal Patterns

## Alert frequency metrics

Alert frequency should be represented through short-term and medium-term exposure measures, including alerts per shift, alerts by alert type, and rolling counts over recent days. A decay-weighted exposure feature would be useful because the most recent interruptions may have stronger influence on response behaviour than older alerts. Prior work links workload, repeated alerts, and national alert burden patterns to fatigue-related response changes [1, 6]. In laboratory alerting, this allows the model to distinguish ordinary background notification load from concentrated alert accumulation during a busy clinical period.

## Clinical severity encoding

Clinical severity encoding should map the degree of laboratory abnormality into clinically interpretable levels rather than treating all abnormal results as equivalent. A critical potassium value, a worsening creatinine trend, and a stable mild abnormality should contribute differently to fatigue risk because their urgency and perceived usefulness differ. Studies of alert appropriateness and renal medication-related decision support show that clinical context affects whether an alert is viewed as actionable or excessive [11, 12]. A severity feature could therefore help the model avoid suppressing high-risk alerts merely because provider workload is high.

## Repeated abnormal results and desensitisation

Repeated abnormal result features would identify alerts triggered by the same test, similar value range, or persistent abnormality over consecutive clinical periods. Such alerts may be clinically valid but become fatiguing when they repeat without a meaningful change in patient state or required action. Evidence on repeated alerts and override patterns suggests that repeated exposure can increase desensitisation and reduce responsiveness [1, 23]. In laboratory alerting, this feature would capture the nuisance effect of recurring mild abnormalities while preserving sensitivity to sudden deterioration.

## Interaction of frequency and severity

The model should represent the interaction between alert frequency and severity because fatigue risk is unlikely to rise uniformly across all alert types. High frequency of low-severity or repeated alerts would be expected to increase fatigue risk, whereas high-severity alerts may retain salience even during periods of heavy alert burden. Prior studies describe the safety trade-off between alert volume and responsiveness, as well as the need to filter alerts without losing clinically important signals [6, 21]. An interpretable interaction explanation would help governance teams understand when reducing frequency is safe and when escalation is more appropriate.

**Table 1** analytically links model input features to underlying fatigue mechanisms and their clinical interpretation to support transparent governance decisions.

**Table 1.** Analytical Mapping of Model Features to Mechanisms of Alert Fatigue and Clinical Interpretation

| Feature Domain  | Operational Definition                 | Underlying Fatigue Mechanism            | Expected Direction of Effect on Fatigue Risk | Clinical Interpretation Implication                                   |
|-----------------|----------------------------------------|-----------------------------------------|----------------------------------------------|-----------------------------------------------------------------------|
| Alert Frequency | Alerts per shift; rolling alert counts | Cognitive overload; interruption burden | Positive (higher frequency → higher fatigue) | Indicates accumulation pressure; supports alert batching or digesting |

|                                  |                                                      |                                              |                                                       |                                                                     |
|----------------------------------|------------------------------------------------------|----------------------------------------------|-------------------------------------------------------|---------------------------------------------------------------------|
| Clinical Severity                | Categorised abnormality (critical, worsening, mild)  | Perceived clinical value modulation          | Negative for high severity; positive for low severity | Prevents suppression of high-risk alerts; prioritises escalation    |
| Repeated Abnormal Results        | Recurrence of similar abnormal values without change | Desensitisation from redundancy              | Positive (more repetition → higher fatigue)           | Supports merging or suppressing non-actionable repeats              |
| Provider Specialty               | Clinical role or service category                    | Contextual relevance mismatch                | Context-dependent (varies by specialty-alert pairing) | Enables role-tailored alert routing and relevance adjustment        |
| Response Time History            | Median or recent acknowledgement delay               | Behavioural inertia; attentional fatigue     | Positive (longer delays → higher fatigue)             | Signals declining responsiveness; may justify escalation            |
| Override Behaviour               | Proportion of alerts overridden                      | Learned dismissal pattern                    | Positive (higher override rate → higher fatigue)      | Distinguishes systematic dismissal from appropriate filtering       |
| Frequency × Severity Interaction | Joint effect of high frequency and low severity      | Amplified fatigue under low-value repetition | Strong positive interaction effect                    | Identifies safe targets for suppression without compromising safety |

## Explainability Methods for Fatigue Insights

### Global explanations: drivers of alert fatigue across the institution

Global explanations would summarise which features most consistently contribute to predicted fatigue across the institution. SHAP summary plots, partial dependence curves, and grouped feature importance views could show whether alert frequency, specialty, repeated abnormality, or low severity is the dominant institutional driver. Clinical explainability scholarship emphasises that explanations should support organisational learning as well as individual prediction [10, 16, 19]. For laboratory governance, these global views could guide threshold review, alert consolidation, and redesign of high-burden alert categories.

### Local explanations: why will this provider likely override this alert?

Local explanations would show why a specific provider-alert pair receives a high fatigue risk score. A SHAP waterfall plot could attribute risk to accumulated alerts during the current shift, a repeated mildly abnormal result, recent slow acknowledgement, or provider-specialty context. Prior work on explainable clinical prediction and medication alert filtering supports the value of instance-level explanations when model outputs are used in care workflows [8, 9, 21]. This local rationale would make the system more transparent than a silent suppression rule.

### Counterfactual explanations for alert redesign

Counterfactual explanations could help alert designers test how changes in severity threshold, timing, routing, or wording might alter predicted fatigue risk. Instead of reporting fixed performance claims, the model would support conceptual simulations such as whether combining repeat non-critical alerts into a digest could reduce fatigue pressure. Studies on alert optimisation and clinical decision support redesign indicate that alert burden should be evaluated in relation to appropriateness, workflow fit, and clinician response [4, 15, 24]. Counterfactual reasoning would therefore make the model useful for governance decisions before live implementation.

### Transparent audit trail for alert governance

A transparent audit trail should record each fatigue prediction, the explanation features, and the alert action taken, such as suppression, escalation, digesting, or unchanged delivery. This record would allow clinical informatics teams to review whether model-informed decisions remain clinically appropriate over time. Prior research on alert overrides, explainability, and governance shows that predictions must be monitored because dismissal may be appropriate in some cases and unsafe in others [11, 17, 18]. For laboratory alerting, the audit trail would help ensure that explainable automation remains accountable and clinically reviewable.

## Clinical Integration and Alert System Redesign

### Real-time fatigue-aware alert suppression or escalation

A fatigue-aware alerting system would run before the alert is delivered and decide whether the notification should remain interruptive, be summarised, be delayed, or be escalated. For non-critical repeated alerts predicted to be ignored, the system could route information into a digest, while high-severity alerts at risk of being missed could be escalated to a covering provider or secondary channel. Prior studies of alert safety, override behaviour, and role-tailored design suggest that reducing burden must be balanced against the risk of missed clinically important alerts [2, 5, 13]. The goal is not silence, but smarter routing based on severity, context, and predicted responsiveness.

**Table 2** presents a conceptual decision matrix translating predicted fatigue risk and clinical context into actionable alert system responses.

**Table 2.** Conceptual Decision Matrix for Fatigue-Aware Alert Actions Based on Predicted Risk and Clinical Context

| Fatigue Risk Level | Clinical Severity | Alert Pattern             | Recommended System Action                    | Rationale                                         | Governance Consideration                               |
|--------------------|-------------------|---------------------------|----------------------------------------------|---------------------------------------------------|--------------------------------------------------------|
| High               | Low               | Repeated abnormal         | Digest or suppress                           | Minimises low-value interruption burden           | Requires audit to ensure no missed trend deterioration |
| High               | High              | New or worsening abnormal | Escalate (secondary channel or provider)     | Prevents critical alert from being ignored        | Must ensure escalation pathways are reliable           |
| Moderate           | Low               | Repeated abnormal         | Consolidate into periodic summary            | Reduces redundancy while preserving awareness     | Evaluate frequency of summaries for usability          |
| Moderate           | High              | Stable abnormal           | Maintain interruptive alert with explanation | Preserves visibility with contextual clarity      | Monitor for alert overload persistence                 |
| Low                | Low               | Isolated abnormal         | Maintain or downgrade alert priority         | Low fatigue risk; minimal intervention needed     | Avoid unnecessary system complexity                    |
| Low                | High              | Critical new abnormal     | Immediate interruptive alert                 | High clinical priority outweighs fatigue concerns | Align with critical value policy                       |

|                                     |     |                                       |                                     |                                     |                                   |
|-------------------------------------|-----|---------------------------------------|-------------------------------------|-------------------------------------|-----------------------------------|
| Variable<br>(context-<br>dependent) | Any | Specialty-<br>mismatched<br>relevance | Redirect to<br>appropriate provider | Improves contextual<br>fit of alert | Requires accurate role<br>mapping |
|-------------------------------------|-----|---------------------------------------|-------------------------------------|-------------------------------------|-----------------------------------|

## Personalised alert thresholds and feedback

Personalised alert thresholds could allow providers or services to receive fewer low-value interruptions during high-burden periods while preserving critical safety notifications. A feedback interface could show anonymised fatigue profiles, recurring alert types, and explanation patterns so clinicians understand why certain alerts are being modified. Evidence on alert fatigue solutions, knowledgebase transitions, and role-sensitive alert design supports the need for adaptive systems that account for clinical context and user behaviour [5, 20, 24]. Such feedback should be framed as system improvement rather than individual performance surveillance.

## Evaluation Strategy

### Predictive performance metrics

The model should be evaluated conceptually using discrimination, calibration, and minority-class performance measures suited to fatigued response prediction. These metrics would assess whether the model can separate likely delayed or override-prone responses from likely attentive responses without relying on unsupported performance claims. Prior alert prediction and filtering studies demonstrate that machine learning can be assessed against alert response outcomes derived from clinical decision support logs [7, 8, 21]. For laboratory alerting, evaluation should also examine whether predicted fatigue risk remains clinically interpretable across alert types and specialties.

### Explanation quality and clinical trust

Explanation quality should be assessed by whether clinicians and governance teams find the explanations plausible, actionable, and aligned with real workflow constraints. A provider-facing review could compare SHAP explanations with clinician judgement about why alerts are ignored, delayed, or overridden. Research on clinician expectations for explainable machine learning shows that explanations must be contextual, usable, and connected to practical decisions [10, 16]. This evaluation should therefore focus on whether explanations improve trust, not merely whether they appear technically sophisticated.

### Prospective impact on alert fatigue measures

Prospective evaluation should examine whether model-informed alert adjustments reduce fatigue indicators while preserving timely response to clinically important laboratory results. Outcomes could include aggregate override patterns, acknowledgement behaviour, critical result escalation review, and provider-reported alert burden, described without unsupported numerical claims. Prior work on alert appropriateness, alert pattern changes, and alert governance supports before-after evaluation as a way to assess whether redesign improves clinical decision support quality [15, 22, 24]. Any deployment should include safety monitoring to detect unintended suppression of important abnormal results.

## Limitations

### Data and labelling challenges

Fatigue labels based on override or delayed acknowledgement are imperfect because these behaviours may reflect appropriate clinical judgement, competing emergencies, or workflow barriers rather than desensitisation. A provider may override an alert because the abnormal result is already known, already treated, or clinically irrelevant in context. Prior studies of alert overrides emphasise that not all overrides are inappropriate and that harm assessment requires careful

clinical interpretation [11, 17]. Therefore, the proposed model should treat fatigue as a probabilistic risk signal rather than a definitive diagnosis of provider inattention.

## Specialty and context granularity

Specialty categories may be too broad to capture the full context of laboratory alert response. Providers may rotate across services, cover unfamiliar patients, or work in environments where alert relevance changes rapidly during a shift. Research on role tailoring and clinical alert patterns suggests that alert response depends on workflow, care setting, and knowledgebase design, not merely provider specialty labels [5, 14, 24]. Future versions of the model would need richer context features while avoiding excessive complexity that undermines interpretability.

## Conclusion

An interpretable machine learning model for laboratory alert fatigue could make alert burden measurable at the provider–alert level. By combining alert frequency, clinical severity, provider specialty, repeated abnormal result patterns, response time history, and override behaviour, the model could identify alerts that are likely to receive delayed or dismissive responses.

The central strength of this framework is that it links prediction with explanation. A fatigue risk score alone would be insufficient, but a SHAP-based rationale could show whether risk is driven by cumulative alert burden, repeated low-severity abnormalities, specialty context, or recent response behaviour.

Important challenges remain in defining valid fatigue labels, maintaining data quality, and ensuring that model-informed suppression does not hide clinically important laboratory results. Prospective validation would be essential before such a system could be trusted in live alert routing.

Collaborative pilot studies involving hospital laboratories, clinical informatics teams, frontline providers, and patient safety committees are needed. A shared repository of alert-fatigue prediction concepts, governance practices, and explanation templates could support safer and more transparent clinical decision support redesign.

## Acknowledgements

None

## Conflict of interest

None

## Financial support

None

## Ethics statement

None

## Rights and permissions

**Open Access** The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

## References

- Ancker JS, Edwards A, Nosal S, Hauser D, Mauer E, Kaushal R, et al. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Med Inform Decis Mak.* 2017;17(1):36.
- Wong A, Amato MG, Seger DL, Slight SP, Beeler PE, Dykes PC, et al. Evaluation of medication-related clinical decision support alert overrides in the intensive care unit. *J Crit Care.* 2017;39:156-61.
- Nanji KC, Seger DL, Slight SP, Amato MG, Beeler PE, Her QL, et al. Medication-related clinical decision support alert overrides in inpatients. *J Am Med Inform Assoc.* 2018;25(5):476-81.
- Carli D, Fahrni G, Bonnabry P, Lovis C. Quality of decision support in computerized provider order entry: systematic literature review. *JMIR Med Inform.* 2018;6(1):e3.
- Hussain MI, Reynolds TL, Zheng K. Medication safety alert fatigue may be reduced via interaction design and clinical role tailoring: a systematic review. *J Am Med Inform Assoc.* 2019;26(10):1141-9.
- Co Z, Holmgren AJ, Classen DC, Newmark L, Seger DL, Danforth M, et al. The tradeoffs between safety and alert fatigue: data from a national evaluation of hospital medication-related clinical decision support. *J Am Med Inform Assoc.* 2020;27(8):1252-8.
- Poly TN, Islam MM, Muhtar MS, Yang HC, Nguyen PA, Li YC. Machine learning approach to reduce alert fatigue using a disease medication-related clinical decision support system: model development and validation. *JMIR Med Inform.* 2020;8(11):e19489.
- Baron JM, Huang R, McEvoy D, Dighe AS. Use of machine learning to predict clinical decision support compliance, reduce alert burden, and evaluate duplicate laboratory test ordering alerts. *JAMIA Open.* 2021;4(1):ooab006.
- Lundberg SM, Nair B, Vavilala MS, Horibe M, Eisses MJ, Adams T, et al. Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nat Biomed Eng.* 2018;2(10):749-60.
- Amann J, Blasimme A, Vayena E, Frey D, Madai VI, Precise4Q Consortium, et al. Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC Med Inform Decis Mak.* 2020;20(1):310.
- Poly TN, Islam MM, Yang HC, Li YC. Appropriateness of overridden alerts in computerized physician order entry: systematic review. *JMIR Med Inform.* 2020;8(7):e15653.
- Shah SN, Amato MG, Garlo KG, Seger DL, Bates DW. Renal medication-related clinical decision support (CDS) alerts and overrides in the inpatient setting following implementation of a commercial electronic health record: implications for designing more effective alerts. *J Am Med Inform Assoc.* 2021;28(6):1081-7.
- Wong A, Amato MG, Seger DL, Rehr C, Wright A, Slight SP, et al. Prospective evaluation of medication-related clinical decision support overrides in the intensive care unit. *BMJ Qual Saf.* 2018;27(9):718-24.
- Yoo J, Lee J, Rhee PL, Chang DK, Kang M, Choi JS, et al. Alert override patterns with a medication clinical decision support system in an academic emergency department: retrospective descriptive study. *JMIR Med Inform.* 2020;8(11):e23351.
- Park H, Chae MK, Jeong W, Yu J, Jung W, Chang H, et al. Appropriateness of alerts and physicians' responses with a medication-related clinical decision support system: retrospective observational study. *JMIR Med Inform.* 2022;10(10):e40511.

Tonekaboni S, Joshi S, McCradden MD, Goldenberg A. What clinicians want: contextualizing explainable machine learning for clinical end use. In: *Mach Learn Healthc Conf*; 2019. p. 359-80.

Wong A, Rehr C, Seger DL, Amato MG, Beeler PE, Slight SP, et al. Evaluation of harm associated with high dose-range clinical decision support overrides in the intensive care unit. *Drug Saf.* 2019;42(4):573-9.

Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. *Lancet Digit Health.* 2021;3(11):e745-50.

Antoniadi AM, Du Y, Guendouz Y, Wei L, Mazo C, Becker BA, et al. Current challenges and future opportunities for XAI in machine learning-based clinical decision support systems: a systematic review. *Appl Sci (Basel).* 2021;11(11):5088.

Naeem A, Alwadie AF, Alshehri AM, Alharbi LM, Nawaz MU, AlHadidi RA, et al. The overriding of computerized physician order entry (CPOE) drug safety alerts fired by the clinical decision support (CDS) tool: evaluation of appropriate responses and alert fatigue solutions. *Cureus.* 2022;14(11):e31538.

Liu S, Kawamoto K, Del Fiore G, Weir C, Malone DC, Reese TJ, et al. The potential for leveraging machine learning to filter medication alerts. *J Am Med Inform Assoc.* 2022;29(5):891-9.

Chien SC, Chen YL, Chien CH, Chin YP, Yoon CH, Chen CY, et al. Alerts in clinical decision support systems (CDSS): a bibliometric review and content analysis. *Healthcare (Basel).* 2022;10(4):601.

Colicchio TK, Cimino JJ. Beyond the override: using evidence of previous drug tolerance to suppress drug allergy alerts: a retrospective study of opioid alerts. *J Biomed Inform.* 2023;147:104508.

Jung W, Yu J, Park H, Chae MK, Lee SS, Choi JS, et al. Effect of knowledgebase transition of a clinical decision support system on medication order and alert patterns in an emergency department. *Sci Rep.* 2023;13(1):21206.