Annals of Pharmacy Practice and Pharmacotherapy
ISSN: 3062-4436
2024, Volume 4, Issue 2, Page No: 47-56

Available online at: www.galaxypub.co/page/journals

Galaxy Publication

Predictive Medication-Safety Systems Can Change the Errors They Are Designed to Detect Through Automation Bias, Alert Migration, Data Drift, and Invisible Residual Risk

Naledi Mokoena[1]*, Peter Novak[2], Sofia Mendes[3]

[1]Department of Pharmacy Informatics and Medication Safety, Faculty of Pharmacy, University of Cape Town, Cape Town, South Africa.
[2]Department of Clinical Pharmacy and Health Informatics, Faculty of Pharmacy, Charles University, Prague, Czech Republic.
[3]Department of Pharmaceutical Care and Digital Health, Faculty of Pharmacy, University of Porto, Porto, Portugal.

***E-mail** ✉ naledi.mokoena@uct.ac.za

Received: 22 May 2024; **Revised:** 27 August 2024; **Accepted:** 01 September 2024

ABSTRACT

Predictive medication-safety systems are commonly evaluated as tools that identify patients, orders, or situations at elevated risk. Once embedded in clinical work, however, prediction becomes part of the system generating the outcomes on which later safety judgments depend. Recommendations can redistribute clinician attention, alerts can alter prescribing behavior, suppression and prioritization can change which hazards remain visible, and successful interventions can modify the labels and data streams used for subsequent monitoring or retraining. These effects complicate the interpretation of apparent improvement or deterioration after deployment. A model may retain acceptable technical performance while the surrounding clinical system develops new workload, reliance, or residual-risk problems; conversely, apparent statistical drift may partly reflect successful intervention. This Current Opinion examines predictive medication safety as a sociotechnical intervention rather than an isolated classification task. It distinguishes automation bias from appropriate reliance, alert fatigue from alert migration, model drift from intervention-induced feedback, and visible alert reduction from reduction in total medication risk. The resulting position is conditional: predictive systems can improve defined medication-safety outcomes, but their safety cannot be inferred from discrimination, calibration, alert counts, or override rates alone. Evaluation must include how the deployed system changes behavior, information production, error visibility, and the residual distribution of medication risk.

Keywords: Predictive medication safety, Clinical decision support, Automation bias, Alert fatigue, Model drift, Alert migration

How to Cite This Article: Mokoena N, Novak P, Mendes S. Predictive Medication-Safety Systems Can Change the Errors They Are Designed to Detect Through Automation Bias, Alert Migration, Data Drift, and Invisible Residual Risk. Ann Pharm Pract Pharmacother. 2024;4(2):47-56. https://doi.org/10.51847/nJu0NK3jgz

Introduction

Clinical decision support can improve the availability and timing of information at the point of care, but its effects depend on implementation, usability, workflow integration, and how clinicians respond to recommendations [1]. Medication safety therefore cannot be treated as a purely technical property of a rule, score, or prediction. Qualitative evidence from medication-related decision support shows that prescribers incorporate alerts into existing professional routines, practical workarounds, responsibility structures, and local safety practices [2]. Drug-safety reasoning similarly requires separation between evidence, interpretation, and the action ultimately taken [3].

These distinctions become more important when decision support is predictive. Responsible clinical machine-learning guidance has emphasized the full deployment lifecycle rather than model development alone [4], while analyses of artificial intelligence and clinical safety have identified bias, distributional instability, and implementation context as potential sources of harm [5]. More broadly, the translation of high algorithmic performance into clinical benefit remains dependent on generalizability, workflow fit, prospective evaluation, and real-world use [6]. For predictive medication safety, the central question is therefore not simply whether a model detects risk accurately. It is whether deployment changes clinician behavior, alert exposure, documentation, subsequent data, and the remaining distribution of medication errors in ways that alter what "performance" means.

Prediction as a sociotechnical intervention

A prediction becomes clinically consequential when it changes what somebody does. Once surfaced to a prescriber, pharmacist, nurse, or safety team, a risk estimate can alter attention, investigation, treatment choice, documentation, escalation, or monitoring. The distinction matters because the prediction is then no longer observing an untouched system. It participates in the processes that generate future medication orders and safety outcomes.

This interpretation follows naturally from the sociotechnical character of decision support [2]. The same model can produce different consequences when embedded in different interfaces, escalation pathways, staffing structures, or organizational norms. A warning that triggers pharmacist review has a different operational meaning from the same score presented passively in an electronic record. A model that changes prescribing before an adverse event occurs may also change the later availability of positive outcome labels, while a system that is routinely ignored can leave the underlying data-generating process largely unchanged.

Model performance and clinical-system performance must therefore remain analytically separate. Discrimination, calibration, sensitivity, specificity, or predictive value describe properties of a prediction task. They do not directly measure whether clinicians receive too many alerts, become over-reliant on recommendations, develop workarounds, encounter clinically irrelevant warnings, or shift errors into channels the system does not monitor. Predictive medication safety should consequently be evaluated as an intervention in a dynamic clinical environment, with the model treated as one component of a wider safety process rather than its sole determinant.

Automation bias and redistributed clinical attention

Automation bias describes inappropriate reliance on automated assistance when users accept, fail to interrogate, or insufficiently cross-check system recommendations [7]. In medication safety, the problem is not limited to an incorrect recommendation being followed. An apparently credible prediction can also redirect scarce attention toward the risks selected by the system and away from medication problems outside its target definition.

Experimental evidence demonstrates why the direction of effect cannot be assumed. In a randomized crossover study of artificial-intelligence-assisted diagnostic interpretation, incorrect AI advice could mislead clinicians even though AI assistance was also capable of improving decisions under other conditions [8]. The relevant lesson for medication safety is mechanistic rather than quantitative: the same assistive system can support correction and generate error depending on recommendation quality and how clinicians use it.

Recent generative-AI studies reinforce this conditionality. In a randomized clinical-management experiment, GPT-4 assistance improved physicians' task scores relative to conventional resources but increased task time, while the combined physician-plus-LLM condition was not significantly better than the LLM-alone comparison reported by the investigators [9]. In medication-safety scenarios spanning multiple specialties, performance also differed across pharmacist-only, LLM-only, and human–LLM configurations [10]. These results do not establish routine clinical benefit, but they show that system configuration alters the performance of the combined human–technology unit.

The clinically relevant safety variable is therefore not "trust in AI" in the abstract. It is calibration of reliance: when clinicians should accept, interrogate, ignore, or independently verify a recommendation. Poorly calibrated reliance can produce commission errors when an incorrect recommendation is followed and omission errors when attention is diverted from unmodeled risks. Under-reliance can also discard useful warnings. Predictive medication-safety evaluation must therefore examine how the system redistributes attention, not merely whether its output is statistically accurate.

Alert ecology: Burden, overrides, and workflow adaptation

Medication alerts operate within an ecology of repeated interruptions, competing priorities, local customization, and learned user responses. Override frequency, alert burden, and safety coverage are related but non-equivalent properties [11–13]. A high override rate may reflect poorly targeted alerts, but it may also include clinically appropriate decisions in which the prescriber possesses contextual information unavailable to the rule or model.

This distinction is supported by studies examining high-priority drug–drug interaction alerts in relation to override appropriateness and adverse drug events [14]. The safety meaning of an override therefore depends on the alert, clinical context, user reasoning, and downstream consequence. Treating every override as failure can reward systems that simply suppress difficult or context-dependent cases rather than improve decision quality.

Alert behavior is also modifiable. Tailored severe drug–drug interaction alerting in intensive care has demonstrated that implementation choices can change acceptance and override patterns [15]. Systematic examination of behavioral-design principles similarly indicates that presentation, salience, timing, and role-specific design can influence responses to alerts [16]. Usability assessment adds another distinct layer: formal tools for evaluating medication-alert usability do not measure the same construct as predictive accuracy or clinical relevance [17].

These findings argue against a single alert-performance endpoint. A system may reduce the number of interruptive warnings but make the remaining alerts more difficult to interpret. It may increase acceptance by improving specificity, by making warnings more salient, or by encouraging uncritical compliance. It may lower override rates without improving patient-relevant outcomes. Alert burden, acceptance, appropriateness, usability, and downstream safety therefore require separate observation before any change can be interpreted as improvement.

Alert migration and changed medication-safety data

Artificial intelligence is increasingly used to prioritize, filter, or suppress medication-related alerts. A recent scoping review found substantial interest in applying AI to optimize medication-alert generation and reduce low-value interruptions, although study designs and outcome definitions remain heterogeneous [18]. Applied work has likewise shown that machine-learning-based decision support can alter alert exposure and target specific wrong-drug or look-alike/sound-alike medication errors [19].

Optimization changes more than alert count. When a knowledgebase is replaced or recalibrated, the distribution of displayed warnings and medication orders can also change. An emergency-department study of a clinical decision-support knowledgebase transition documented changes in medication-order and alert patterns after implementation [20]. Such observations should not automatically be interpreted as clinician adaptation, because rule content itself has changed. They nevertheless demonstrate that the observable safety environment depends on the configuration used to detect and display risk.

Selective suppression can be beneficial. Optimization of interruptive alerts for duplicate antithrombotic prescribing reduced unnecessary warning exposure while preserving more targeted safety logic [21]. Similarly, a phased clinical decision-support intervention for anticoagulant duplication was associated with fewer and shorter duplication episodes while its workflow generated a considerably larger pool of candidate alerts for subsequent filtering and action [22]. This is an important disconfirming case against the proposition that predictive or algorithmic decision support necessarily increases medication risk. A system can improve a defined safety outcome.

The unresolved issue is what happens to risks not represented by the optimized surface. Alert suppression, ranking, and filtering determine which potential problems become salient to clinicians and which remain silent. They can therefore change the distribution of observable error even when total underlying medication risk is unknown. Here, alert migration denotes a change in where risk becomes visible and actionable after system redesign; it is not a synonym for alert fatigue. Fatigue concerns responsiveness under repeated alert exposure. Migration concerns redistribution of safety signals between highly visible, weakly visible, and unmonitored channels.

Deployment also changes the data used for later redesign. Free-text override comments can be analyzed to characterize clinician reasons for rejecting decision support and to identify opportunities for optimization [23]. Clinical-pharmaceutical review further shows that technically generated alerts may differ substantially in patient relevance [24]. A broader synthesis of medication-alert handling indicates that responses depend on interacting properties of clinicians, alerts, and implementation context rather than on a single determinant [25]. These data streams are valuable, but they are also selective. Recorded override comments represent documented reasoning,

not all reasoning; alert logs contain only events that the system chose to generate; and clinically irrelevant warnings that have been suppressed disappear from routine alert statistics.

The resulting measurement problem is central to predictive medication safety. A decrease in displayed alerts can reflect better targeting, successful prevention, narrower rule coverage, or migration of hazards outside the monitored alert surface. A lower override rate can reflect improved relevance or increased compliance. More complete documentation can reveal previously hidden problems and make apparent safety performance worse without any increase in true harm. The safety system therefore changes the evidence subsequently used to judge the safety system.

The relationship between system mechanism, clinician response, recorded data, and residual-risk visibility can be separated as follows.

Table 1. How predictive medication-safety interventions can change clinician behavior, recorded safety data, and the visibility of residual risk

System mechanism or change	Immediate clinician-facing effect	Resulting behavioral response that may occur	Safety data subsequently generated	Potential change in error visibility	Residual-risk issue requiring separate assessment
Risk prediction surfaced during prescribing	Directs attention toward selected medication risks	Acceptance, verification, rejection, or reliance on the recommendation	Alert response, order change, documentation, escalation	Modeled hazards become more visible than unmodeled hazards	Important medication problems outside the prediction target may receive less attention
Increased alert sensitivity or broader coverage	More interruptions and warnings	Greater review, selective override, habitual dismissal, or workaround development	Higher alert volume and response counts	More candidate hazards are recorded, but signal-to-noise may deteriorate	Clinically important alerts may compete with large numbers of low-value warnings
Alert prioritization or machine-learning filtering	Fewer or more selectively displayed alerts	Concentration of attention on the retained alert subset	Lower displayed-alert count and changed response distribution	Suppressed categories become less observable in routine alert metrics	Reduced alert burden cannot establish reduced total medication risk
Knowledgebase, threshold, or rule-set transition	Changes which orders trigger intervention	Altered prescribing, substitution, avoidance, or acceptance patterns	New alert and medication-order distributions	Apparent before/after differences partly reflect changed detection logic	Historical comparisons can confound real safety change with altered surveillance
Clinician override with documented rationale	Preserves clinician discretion while generating feedback	Override, modification, or escalation based on patient context	Structured override code or free-text comment	Reasons for rejecting alerts become partially visible	Undocumented reasoning and nonalerted hazards remain outside the feedback record
Pharmacist or expert filtering before escalation	Removes self-limiting or contextually irrelevant candidates	Prescribers receive a smaller set of higher-priority interventions	Candidate-alert, filtered-alert, and action logs	Intermediate filtering process becomes measurable if retained	Safety evaluation may be misleading if only final displayed alerts are counted

Repeated local optimization	Changes interaction patterns across successive system versions	Users adapt to new alert frequency, content, and perceived reliability	Version-specific response and outcome data	Error categories may move between visible, suppressed, and newly detected states	Comparisons across versions require preservation of the changing detection and workflow context

Note: The table separates evidence-supported classes of system change from the article's proposed residual-risk interpretation. It does not imply that every mechanism produces harm or that risk necessarily migrates in a particular direction. Patient-relevant outcomes, clinically important nonalerted errors, workflow effects, and the changing detection process require independent evaluation.

Drift, feedback, and performance instability

Predictive medication-safety systems are exposed to the same transportability problem as other clinical prediction models: performance established in one population, workflow, or period may not persist after deployment elsewhere. External validation of a widely implemented sepsis model demonstrated that operational performance can differ substantially from expectations derived from development settings [26]. Medication-related prediction is feasible at scale, including identification of patients at differing risk of adverse drug events, but predictive classification remains distinct from demonstrated prevention of those events [27].

Temporal change adds a second problem. Clinical practices, patient mix, coding, documentation, formularies, laboratory measurement, and organizational processes evolve after implementation. Empirical studies across clinical prediction tasks have documented temporal and system-shock-related drift [28–30]. Input monitoring may detect changes before conventional outcome-based performance measures become informative, although statistical drift is not automatically clinically important [31].

A more difficult case occurs when the intervention itself changes subsequent data. Feedback-aware monitoring research shows that model-triggered interventions can alter features and labels used for later evaluation or retraining [32]. If a prediction successfully prevents an adverse event, the future dataset may contain fewer positive labels among precisely the patients for whom the model predicted high risk. Naive surveillance can then interpret intervention success as declining calibration or model deterioration. Conversely, stable aggregate performance can conceal emerging workflow or reliance problems. Drift therefore requires causal and clinical interpretation before retraining is treated as the default response.

Figure 1 summarizes how predictive medication-safety systems can alter clinical decisions, subsequent data, and the visibility of residual risk.

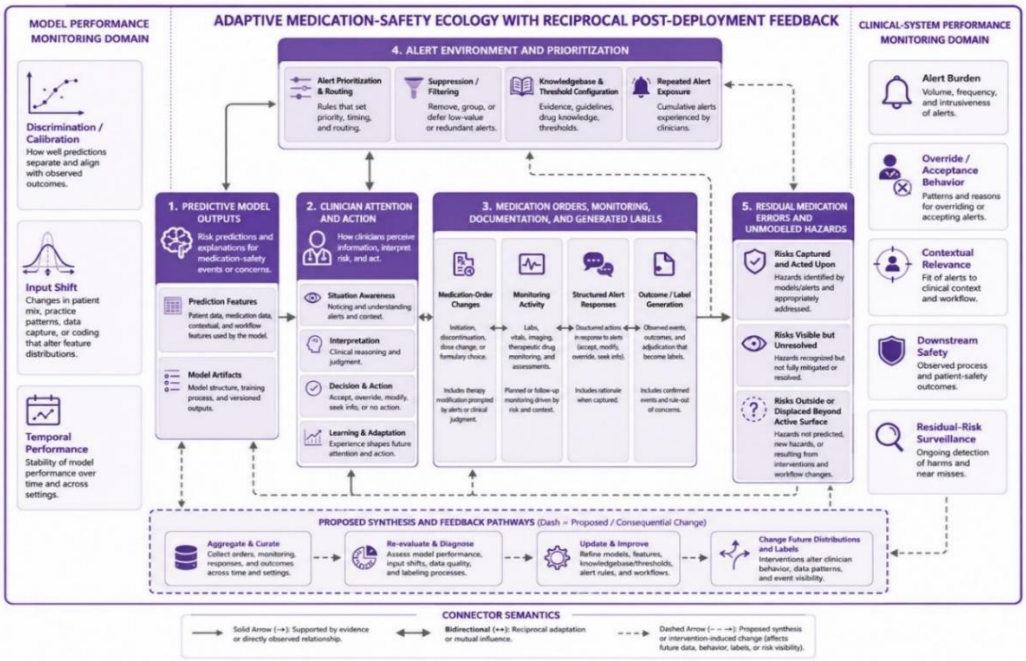

Figure 1. Feedback-coupled safety ecology after predictive medication-safety deployment

Invisible residual risk

Optimization can reduce specific medication errors while leaving uncertainty about the risk that remains outside the optimized alert surface. This residual component is "invisible" only in a relative sense: it is insufficiently represented by the routine measurements generated by the deployed system, not inherently unknowable. Risks may remain unmodeled, be filtered from interruptive presentation, occur downstream of an accepted recommendation, or emerge through inappropriate reliance on apparently credible predictions.

The concept therefore concerns risk visibility, not an assumed increase in total harm. Medication-alert optimization can improve defined outcomes, and successful interventions should be recognized as such. The analytical problem arises when falling alert counts or improving response rates are treated as evidence that the total residual medication-risk distribution has also improved. Alert filtering, automation bias, workflow adaptation, and drift can each change what is observed without producing the same change in what actually occurs.

This distinction yields a falsifiable implication. A strong risk-migration interpretation would be weakened where longitudinal evaluation demonstrates sustained improvement in patient-relevant medication outcomes together with reduced alert burden, stable or improved performance for clinically important nonalerted errors, preserved alert appropriateness, and no meaningful deterioration in residual-risk surveillance.

Figure 2 separates changes in visible alerting from changes in the broader distribution of medication risk.

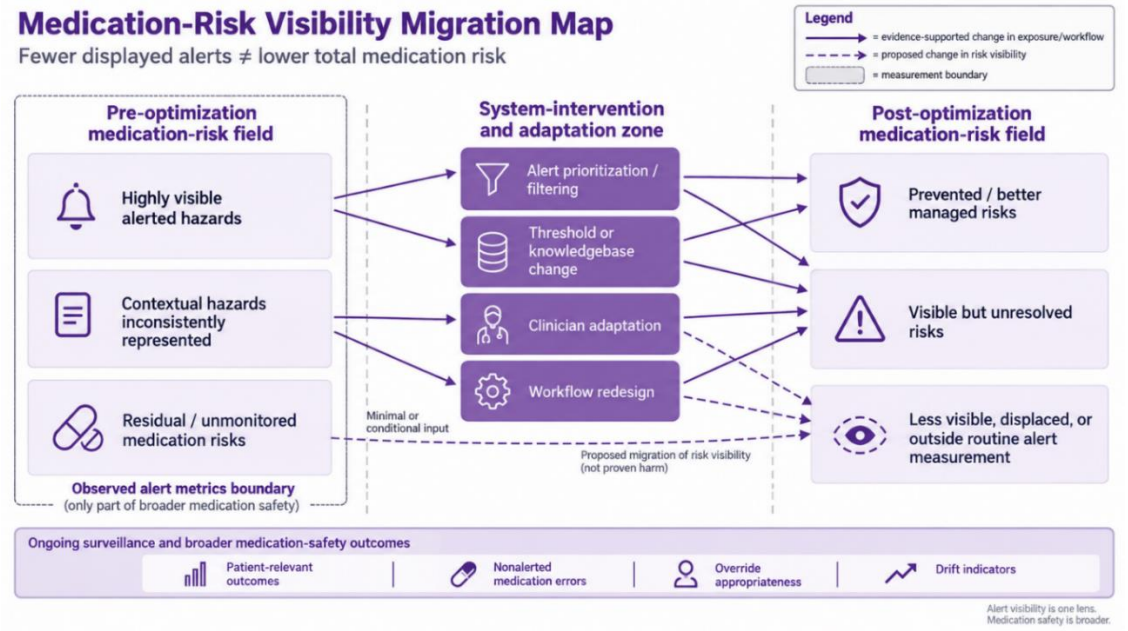

Figure 2. Risk visibility after alert optimization and workflow adaptation

Monitoring model performance and clinical-system performance separately

Post-deployment surveillance should distinguish technical degradation from deterioration in the surrounding safety system. Alert usability and patient relevance already demonstrate that clinically important properties can vary independently of predictive discrimination [17]. External validation likewise shows that technical performance itself can change across deployment contexts [26], while clinically generated alerts may be formally correct yet poorly matched to patient circumstances [24].

Monitoring should therefore combine model metrics with information about alert burden, override behavior, contextual relevance, downstream medication outcomes, and data drift [32–34]. No single monitoring stream resolves all failure modes. Input shift can precede measurable outcome deterioration; falling override rates can reflect better targeting or stronger automation bias; worsening calibration can reflect genuine drift or intervention-induced changes in labels.

The operational separation is summarized below.

Table 2. Distinct post-deployment signals for predictive medication-safety model performance and clinical-system performance

Monitoring domain	What is being observed	What deterioration may indicate	Important competing explanation	Why it cannot be interpreted alone
Discrimination and calibration	Statistical relationship between predictions and observed outcomes	Model degradation or transport failure	Successful intervention may alter observed labels	Does not capture burden, workflow, or unmodeled harms
Input-data shift	Change in feature distributions	Population, documentation, formulary, or workflow change	Statistically detectable change may be clinically irrelevant	Drift detection does not establish safety consequence
Alert burden	Number and frequency of clinician-facing warnings	Excess interruption or poor targeting	Broader legitimate safety coverage	Fewer alerts can reflect suppression rather than prevention
Override and acceptance behavior	Clinician response to generated alerts	Fatigue, distrust, inappropriate reliance, or poor relevance	Appropriate clinical disagreement	Response rate alone cannot determine correctness
Patient relevance	Fit between generated alert and current clinical context	Weak contextualization or overinclusive logic	Complex cases may legitimately require human adjudication	Technical validity and clinical usefulness differ
Patient-relevant medication outcomes	Actual medication-related harm; intercepted errors assessed separately	Failure of the combined clinical system	Outcome ascertainment can change after deployment	Rare events and intervention feedback complicate interpretation
Residual-risk surveillance	Important errors not captured by routine alert metrics	Blind spots or migration of visibility	Genuine reduction in the underlying hazard	Requires active surveillance beyond alert logs

Governance responses for adaptive medication-safety systems

Governance should begin with diagnosis of the observed failure mode rather than automatic retraining. Dataset shift, altered labels, alert fatigue, clinically irrelevant warnings, and inappropriate reliance require different responses. A detected statistical change may justify retraining, but it may instead require threshold review, workflow redesign, improved contextualization, or closer surveillance before the model is changed.

Drug-safety risk management already treats implementation and effectiveness evaluation as iterative activities [35]. Predictive medication safety extends this principle to systems whose operation can change the evidence later used for evaluation. Human escalation remains essential when automated monitoring cannot distinguish harmful degradation from benign shift or intervention-induced feedback.

Table 3. Governance responses matched to post-deployment failure mechanisms in predictive medication safety

Observed post-deployment problem	Diagnostic question before intervention	Candidate governance response	Human role that should remain explicit	Evidence needed before closure
Deteriorating discrimination or calibration	Has the patient, workflow, or label distribution changed?	Revalidation, recalibration, retraining, or model restriction	Clinical review of relevance and consequences	Temporal/external validation plus outcome review
Rising alert burden	Is burden driven by legitimate coverage or low-value triggering?	Threshold review, rule refinement, prioritization, or selective suppression	Pharmacist/clinician adjudication of alert value	Burden plus appropriateness and downstream safety
High override frequency	Are overrides appropriate, habitual, or caused by poor contextual fit?	Workflow redesign, contextualization, interface change, or targeted education	Review of representative override decisions	Override appropriateness rather than rate alone
Falling override frequency with unchanged outcomes	Has relevance improved or has reliance become less critical?	Focused automation-bias review and silent observation	Independent clinical verification in sampled cases	Reliance behavior and patient-relevant outcomes

Input-data drift without performance loss	Is the change clinically meaningful or transient?	Silent surveillance, targeted validation, or no immediate model change	Domain interpretation of changed variables	Persistence and clinical consequence of shift
Apparent performance decline after successful intervention	Could altered treatment or monitoring have changed outcome labels?	Feedback-aware analysis before retraining	Clinical interpretation of intervention pathways	Counterfactual or feedback-sensitive evaluation
Emergence of clinically important nonalerted errors	Has risk moved outside current model or alert scope?	Scope review, new detection strategy, workflow redesign, or human escalation	Safety-team investigation	Independent residual-risk surveillance

The governance objective is therefore adaptive control of the combined medication-safety system, not perpetual optimization of the predictor in isolation. Retraining is justified when model degradation is the diagnosed problem; it is insufficient when the dominant failure lies in workflow, alert ecology, clinical relevance, or residual-risk visibility.

Conclusion

Predictive medication-safety systems can improve defined safety tasks, but deployment also changes the environment in which their effectiveness is judged. Recommendations redistribute attention, alert configuration alters exposure to warnings, workflow adaptation changes recorded behavior, and successful intervention may modify the data and labels used for later surveillance. These mechanisms make fewer alerts, fewer overrides, or stable predictive metrics incomplete proxies for safer medication use.

The strongest defensible conclusion is conditional. Predictive systems change the errors they are designed to detect when their outputs materially alter clinical behavior, measurement, and subsequent system adaptation. Whether those changes reduce, preserve, or redistribute harm is an empirical question. Medication-safety evaluation should therefore follow both the predictor and the clinical system it changes, with explicit surveillance for residual risks that conventional model-performance metrics cannot reveal.

Acknowledgments: None

Conflict of Interest: None

Financial Support: None

Ethics Statement: None

References

1. Sutton RT, Pincock D, Baumgart DC, Sadowski DC, Fedorak RN, Kroeker KI. An overview of clinical decision support systems: Benefits, risks, and strategies for success. NPJ Digit Med. 2020;3:17. doi:10.1038/s41746-020-0221-y
2. Jeffries M, Salema NE, Laing L, Shamsuddin A, Sheikh A, Avery T, et al. Using sociotechnical theory to understand medication safety work in primary care and prescribers' use of clinical decision support: A qualitative study. BMJ Open. 2023;13(4):e068798. doi:10.1136/bmjopen-2022-068798
3. Hammad TA, Davies S. Navigating diverging perspectives: Reasoning, evidence, and decision-making in drug safety. Drug Saf. 2025;48(6):587-93. doi:10.1007/s40264-025-01537-6
4. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: A roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40. doi:10.1038/s41591-019-0548-6
5. Challen R, Denny J, Pitt M, Gompels L, Edwards T, Tsaneva-Atanasova K. Artificial intelligence, bias and clinical safety. BMJ Qual Saf. 2019;28(3):231-7. doi:10.1136/bmjqs-2018-008370
6. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17(1):195. doi:10.1186/s12916-019-1426-2

7. Khera R, Simon MA, Ross JS. Automation bias and assistive AI: Risk of harm from AI-driven clinical decision support. JAMA. 2023;330(23):2255-7. doi:10.1001/jama.2023.22557
8. Wang DY, Ding J, Sun AL, Liu SG, Jiang D, Li N, et al. Artificial intelligence suppression as a strategy to mitigate artificial intelligence automation bias. J Am Med Inform Assoc. 2023;30(10):1684-92. doi:10.1093/jamia/ocad118
9. Goh E, Gallo RJ, Strong E, Weng Y, Kerman H, Freed JA, et al. GPT-4 assistance for improvement of physician performance on patient care tasks: A randomized controlled trial. Nat Med. 2025;31(4):1233-8. doi:10.1038/s41591-024-03456-y
10. Ong JCL, Jin L, Elangovan K, Lim GYS, Lim DYZ, Sng GGR, et al. Large language model as clinical decision support system augments medication safety in 16 clinical specialties. Cell Rep Med. 2025;6(10):102323. doi:10.1016/j.xcrm.2025.102323
11. 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. doi:10.1093/jamia/ocx115
12. 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. doi:10.1093/jamia/ocz095
13. 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. doi:10.1093/jamia/ocaa098
14. Edrees H, Amato MG, Wong A, Seger DL, Bates DW. High-priority drug-drug interaction clinical decision support overrides in a newly implemented commercial computerized provider order-entry system: Override appropriateness and adverse drug events. J Am Med Inform Assoc. 2020;27(6):893-900. doi:10.1093/jamia/ocaa034
15. Karajizadeh M, Zand F, Vazin A, Saeidnia HR, Lund BD, Tummuru SP, et al. Design, development, implementation, and evaluation of a severe drug-drug interaction alert system in the ICU: An analysis of acceptance and override rates. Int J Med Inform. 2023;177:105135. doi:10.1016/j.ijmedinf.2023.105135
16. Hashemi S, Bai L, Gao S, Burstein F, Renzenbrink K. Sharpening clinical decision support alert and reminder designs with MINDSPACE: A systematic review. Int J Med Inform. 2024;181:105276. doi:10.1016/j.ijmedinf.2023.105276
17. Marcilly R, Zheng WY, Quindroit P, Pelayo S, Berdot S, Charpiat B, et al. Comparison of the validity, perceived usefulness, and usability of I-MeDeSA and TEMAS, two tools to evaluate alert system usability. Int J Med Inform. 2023;175:105091. doi:10.1016/j.ijmedinf.2023.105091
18. Graafsma J, Murphy RM, van de Garde EMW, Karapinar-Çarkit F, Derijks HJ, Hoge RHL, et al. The use of artificial intelligence to optimize medication alerts generated by clinical decision support systems: A scoping review. J Am Med Inform Assoc. 2024;31(6):1411-22. doi:10.1093/jamia/ocae076
19. Chen CY, Chen YL, Scholl J, Yang HC, Li YCJ. Ability of machine-learning based clinical decision support system to reduce alert fatigue, wrong-drug errors, and alert users about look alike, sound alike medication. Comput Methods Programs Biomed. 2024;243:107869. doi:10.1016/j.cmpb.2023.107869
20. 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. doi:10.1038/s41598-023-40188-4
21. Sundermann M, Clendon O, McNeill R, Doogue M, Chin PKL. Optimising interruptive clinical decision support alerts for antithrombotic duplicate prescribing in hospital. Int J Med Inform. 2024;186:105418. doi:10.1016/j.ijmedinf.2024.105418
22. Dahmke H, Cabrera-Diaz F, Heizmann M, Stoop S, Schuetz P, Fiumefreddo R, et al. Development and validation of a clinical decision support system to prevent anticoagulant duplications. Int J Med Inform. 2024;187:105446. doi:10.1016/j.ijmedinf.2024.105446
23. Liu S, McCoy AB, Wright AP, Nelson SD, Huang SS, Ahmad HB, et al. Why do users override alerts? Utilizing large language model to summarize comments and optimize clinical decision support. J Am Med Inform Assoc. 2024;31(6):1388-96. doi:10.1093/jamia/ocae041

24. Bauer J, Busse M, Koch S, Schmid M, Sommer J, Fromm MF, et al. Clinical-pharmaceutical assessment of medication CDSS alerts: Content appropriateness and patient relevance in clinical practice. Front Pharmacol. 2025;16:1510425. doi:10.3389/fphar.2025.1510425
25. Raven K, Zhang J, Vagliano I, Jaspers MWM, Peute LW, Medlock S. Determinants influencing medication-related alert handling in the electronic health record: A systematic review and meta-analysis. Int J Med Inform. 2025;203:106011. doi:10.1016/j.ijmedinf.2025.106011
26. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med. 2021;181(8):1065-70. doi:10.1001/jamainternmed.2021.2626
27. Langenberger B. Machine learning as a tool to identify inpatients who are not at risk of adverse drug events in a large dataset of a tertiary care hospital in the USA. Br J Clin Pharmacol. 2023;89(12):3523-38. doi:10.1111/bcp.15846
28. Lacson R, Eskian M, Licaros A, Kapoor N, Khorasani R. Machine learning model drift: Predicting diagnostic imaging follow-up as a case example. J Am Coll Radiol. 2022;19(10):1162-9. doi:10.1016/j.jacr.2022.05.030
29. Duckworth C, Chmiel FP, Burns DK, Zlatev ZD, White NM, Daniels TWV, et al. Using explainable machine learning to characterise data drift and detect emergent health risks for emergency department admissions during COVID-19. Sci Rep. 2021;11(1):23017. doi:10.1038/s41598-021-02481-y
30. Kagerbauer SM, Ulm B, Podtschaske AH, Andonov DI, Blobner M, Jungwirth B, et al. Susceptibility of AutoML mortality prediction algorithms to model drift caused by the COVID pandemic. BMC Med Inform Decis Mak. 2024;24(1):34. doi:10.1186/s12911-024-02428-z
31. Kore A, Abbasi Bavil E, Subasri V, Abdalla M, Fine B, Dolatabadi E, et al. Empirical data drift detection experiments on real-world medical imaging data. Nat Commun. 2024;15(1):1887. doi:10.1038/s41467-024-46142-w
32. Kim GYE, Corbin CK, Grolleau F, Baiocchi M, Chen JH. Monitoring strategies for continuous evaluation of deployed clinical prediction models. J Biomed Inform. 2025;168:104854. doi:10.1016/j.jbi.2025.104854
33. Silva GFDS, Barcellos Filho FN, Wichmann RM, da Silva Junior FC, Chiavegatto Filho ADP. Strategies for detecting and mitigating dataset shift in machine learning for health predictions: A systematic review. J Biomed Inform. 2025;170:104902. doi:10.1016/j.jbi.2025.104902
34. Andersen ES, Birk-Korch JB, Hansen RS, Fly LH, Röttger R, Arcani DMC, et al. Monitoring performance of clinical artificial intelligence in health care: A scoping review. JBI Evid Synth. 2024;22(12):2423-46. doi:10.11124/JBIES-24-00042
35. Mouchantaf R, Auth D, Moride Y, Raine J, Han SY, Smith MY. Risk management for the 21st century: Current status and future needs. Drug Saf. 2021;44(4):409-19. doi:10.1007/s40264-020-01033-z