Available online www.ijpras.com

International Journal of Pharmaceutical Research & Allied Sciences, 2026, 15(2):40-49
https://doi.org/10.51847/tqeRGEASpD

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Adaptive Pharmacovigilance under Model Drift, Product Evolution, Data-Source Change, and Emerging Patterns of Medicine Use across Global Health Systems

Sven Larsson[1]*, Erik Johansson[2], Anna Nilsson[3], Lars Andersson[4]

[1]Department of Adaptive Pharmacovigilance and Model Drift, Faculty of Pharmacy, Swedish University of Agricultural Sciences, Uppsala, Sweden.
[2]Department of Product Evolution and Data-Source Change, Faculty of Pharmacy, KTH Royal Institute of Technology, Stockholm, Sweden.
[3]Department of Emerging Medicine Use Patterns, Faculty of Pharmaceutical Sciences, Lund University, Lund, Sweden.
[4]Department of Global Health Systems and Safety, Faculty of Pharmacy, University of Gothenburg, Gothenburg, Sweden.

***Email:** sven.larsson@slu.se*

ABSTRACT

Pharmacovigilance systems operate within changing medicinal-product, population, data, analytical, and organizational environments. Static surveillance configurations may consequently become misaligned with the products being monitored, the populations exposed, the information captured, or the decisions supported. This article develops an original adaptive-governance framework for pharmacovigilance under model drift, product evolution, data-source change, and emerging patterns of medicine use. The framework conceptualizes surveillance drift as a multidomain change in the relationship among a medicinal product, its use context, observed safety data, analytical processes, and decision environment. It distinguishes product drift, use-pattern drift, data-source drift, model-performance drift, and governance drift without assuming that these categories are independently observable or causally separable. The proposed Adaptive Pharmacovigilance Governance Lifecycle connects versioned baseline specification, multidomain change sensing, change-impact hypothesis formation, consequence-proportionate evidence assignment, targeted revalidation, human-authorized action, rollback, and post-change monitoring. Emerging use patterns and off-label exposure are treated as changing exposure contexts requiring stratified interpretation rather than presumptive evidence of harm. Change-impact assessment is separated from revalidation, and both are distinguished from causal assessment, clinical recommendation, and regulatory acceptability. Human oversight is defined through attributable decision ownership, auditability, override documentation, and recoverable prior states rather than nominal reviewer involvement. The framework requires retrospective temporal evaluation, external and subgroup validation, data-lineage assessment, source-transition analysis, prospective silent-mode testing, human-factors research, and controlled rollback exercises. It is an original conceptual synthesis rather than a validated model, universal standard, regulatory pathway, or deployment-ready system.

Key words: *Pharmacovigilance, Drug safety, Safety signal, Causal inference, Adverse event, Real-world data*
Received: 11 February 2026; **Revised:** 23 April 2026; **Accepted:** 29 April 2026

INTRODUCTION

Pharmacovigilance encompasses the detection, assessment, understanding, and prevention of adverse effects and other medicine-related problems. Its contemporary infrastructure increasingly includes automated case

processing, statistical signal detection, real-world data analysis, active surveillance, and machine-learning-supported prioritization. These tools may improve information handling, but their outputs remain components of medicine-safety evaluation rather than autonomous pharmacovigilance conclusions [1].

The value of real-world data similarly depends on whether evidence generation is connected to interpretation, communication, and proportionate safety action. A technically sophisticated analysis that is detached from product identity, exposure context, causal assessment, or decision responsibility may produce information without completing the pharmacovigilance function [2]. Adaptation must therefore concern the entire evidence-to-action pathway rather than only analytical model maintenance.

Machine-learning applications have been investigated across adverse-event identification, case processing, literature surveillance, signal detection, and risk prediction. However, the supporting evidence is heterogeneous, task-specific, and uneven in its treatment of external validation, operational consequences, and human oversight [3]. The prevailing framing often evaluates individual algorithms under relatively stable assumptions while giving less attention to changes in products, coding systems, data-generating processes, medicine use, exposed populations, and organizational practice.

This Original Adaptive-Governance Framework Article proposes the Adaptive Pharmacovigilance Governance Lifecycle. Its objective is to explain how material change should be identified, interpreted, documented, and connected to proportionate revalidation and accountable action. The contribution is explicitly non-empirical. It organizes established concepts and selected evidence into a proposed architecture whose categories, relationships, decision rules, and implementation requirements remain subject to empirical evaluation.

Why pharmacovigilance systems require continuous adaptation

Pharmacovigilance systems are generally designed around known products, established data streams, specified analytical procedures, and anticipated patterns of reporting. Unexpected or structurally novel events may fall outside these learned or institutionalized patterns, particularly when automation has been optimized for recurrent processing rather than epistemic surprise [4]. Adaptation is therefore required not because every unusual observation represents harm, but because unusual observations may invalidate assumptions embedded in routine surveillance.

Temporal change can also alter the meaning of otherwise familiar safety data. Change-point analysis may reveal shifts in reporting patterns that are not apparent from static disproportionality measures alone [5]. Such changes may reflect an emerging safety concern, stimulated reporting, altered market exposure, media attention, coding revision, or operational disruption. Continuous adaptation must therefore preserve uncertainty about competing explanations rather than converting temporal novelty directly into causal interpretation.

Analytical capability is only one condition of adaptive pharmacovigilance. Network coordination, harmonized procedures, workforce development, institutional learning, and shared operational standards influence whether new information can be recognized and acted upon consistently [6]. A surveillance process may remain statistically functional while becoming organizationally ineffective because responsibilities, escalation practices, expertise, or communication pathways have changed.

Continuous adaptation is accordingly defined here as the governed reassessment of a pharmacovigilance system after a potentially material change in its product, population, use context, data, analytical method, or operational environment. It does not mean unrestricted continual learning, automatic model replacement, or repeated intervention after every fluctuation. The proposed objective is controlled responsiveness: preserving a stable, versioned baseline while creating explicit conditions under which change requires investigation, revalidation, restriction, or rollback.

Sources of surveillance drift

Dataset shift in health systems can arise from changing populations, clinical workflows, professional behaviour, technology adoption, and interactions between users and analytical systems [7]. Pharmacovigilance is exposed to comparable sociotechnical changes. A model may receive different data not only because patients have changed, but because reporting incentives, coding practices, data linkage, case-processing procedures, or reviewer responses have altered the observed evidence.

Distributional change should therefore be interpreted in relation to the causal processes generating the data rather than treated only as a statistical discrepancy [8]. A changed association between a medicine and an adverse event may result from altered exposure, outcome recording, channeling, confounding, stimulated reporting, product

substitution, or a genuine change in risk. Detection of drift identifies a difference requiring explanation; it does not determine which explanation is correct.

Model behaviour may also deteriorate without an explicit change in software or decision rules. Temporal validation has shown that predictive calibration can change as populations and clinical practices evolve [9]. In pharmacovigilance, calibration drift may affect risk-ranking or prediction tools, while other signal-detection systems may exhibit changes in sensitivity, alert burden, subgroup error, or prioritization stability. These properties must be assessed according to intended use rather than reduced to one universal performance metric.

This article proposes surveillance drift as a higher-order construct describing a material change in the relationship among the monitored medicine, exposed population, observed data, analytical process, and safety-decision context. Five forms are distinguished: product drift, use-pattern drift, data-source drift, model-performance drift, and governance drift. They may occur separately or jointly, and one form must not be inferred solely from another. Their analytical separation is a governance device that requires validation rather than an established causal taxonomy.

Proposed adaptive pharmacovigilance framework

The proposed framework begins from the premise that pharmacovigilance technologies require explicit intended-use definitions, evidentiary qualification, and readiness boundaries before operational reliance [10]. The same analytical output may be appropriate for prioritizing case review but inadequate for changing a product label, restricting medicine use, or inferring causality. Intended use therefore determines what change matters, what evidence is required, and who must authorize action.

Task-level automation provides a second foundation. Automated methods may support adverse-event coding, duplicate detection, case triage, information extraction, and other bounded processing functions [11]. Performance within one task does not establish validity across the wider surveillance pathway. The framework consequently treats every automated component as a versioned dependency whose inputs, outputs, error modes, and downstream consequences must remain identifiable.

The third foundation is accountable professional involvement. Automation changes the responsibilities and competencies of drug-safety professionals rather than eliminating the need for judgement [12]. Human oversight must include authority to question outputs, request additional evidence, document disagreement, restrict use, and initiate rollback. A reviewer retained only to endorse an automated recommendation does not constitute meaningful governance.

These foundations indicate that intended-use control, task-level evaluation, and attributable human responsibility must operate together; none is sufficient as an isolated safeguard [10-12]. The framework therefore organizes adaptation around the pharmacovigilance system rather than around a single model. Its unit of governance includes medicinal-product identity, exposure context, data provenance, coding versions, analytical methods, decision rules, professional roles, and the consequences attached to system outputs.

The proposed Adaptive Pharmacovigilance Governance Lifecycle contains seven stages. A versioned baseline first records the product, population, data sources, mappings, models, intended uses, performance references, escalation rules, and recoverable prior states. Multidomain sensing then identifies potentially material change. An explicit change-impact hypothesis specifies affected components, plausible mechanisms, alternative explanations, and uncertainty. Consequence and uncertainty determine a proportionate evidence tier, followed by targeted revalidation. An attributable human gate authorizes continued monitoring, recalibration, retraining, remapping, intended-use restriction, pause, or rollback. Post-change surveillance then establishes a new versioned baseline without implying that adaptation has produced validated safety improvement.

Model-performance and calibration monitoring

Calibration monitoring can identify when the relationship between predicted and observed outcomes has changed sufficiently to justify investigation of model updating [13]. Such an alert should initiate diagnosis rather than automatic retraining because deterioration may originate in the outcome definition, population, data source, workflow, or exposure context rather than the model alone.

Discrimination and calibration represent different properties. A model may preserve ranking while producing poorly aligned probabilities, or remain calibrated overall while performing inadequately within clinically or demographically relevant subgroups [14]. Monitoring requirements must therefore reflect the model's intended pharmacovigilance use and the consequences attached to its output.

Performance assessment should be embedded within wider quality controls covering data provenance, development methods, external validation, reporting, implementation, and lifecycle management [15]. Statistical stability cannot demonstrate that product identity, coding practices, reporting behaviour, or medicine use has remained stable.

The framework proposes monitoring discrimination, calibration, subgroup error, missing-input patterns, alert burden, source coverage, reviewer disagreement, and decision consequences. No indicator is treated as universally decisive. Trigger levels require empirical justification for the relevant product, population, data environment, analytical task, and safety decision.

Figure 1 presents the original conceptual synthesis for adaptive pharmacovigilance lifecycle under changing products, data, and use, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

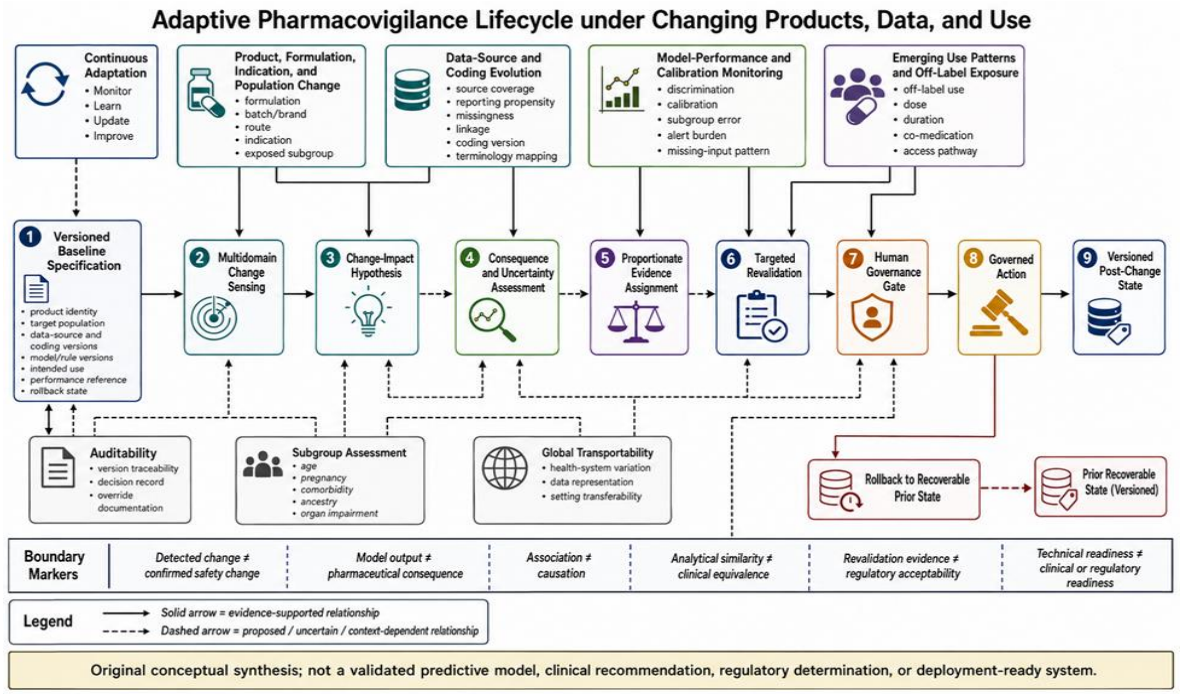

Figure 1. Adaptive Pharmacovigilance Lifecycle under Changing Products, Data, and Use.

The figure is an original conceptual synthesis developed for "Adaptive Pharmacovigilance under Model Drift, Product Evolution, Data-Source Change, and Emerging Patterns of Medicine Use across Global Health Systems." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

Product, formulation, indication, and population change

Product evolution may alter the object being monitored even when the active substance name appears unchanged. For biological medicinal products, meaningful attribution may depend on brand, manufacturer, presentation, and batch traceability [16]. A safety observation cannot be interpreted reliably when the relevant product identity is unavailable.

Empirical examination of biological-product reporting has shown that product and batch identification may be incomplete in clinical records and adverse-event reports [17]. Traceability failure does not demonstrate a product-specific safety problem, but it limits the capacity to investigate one.

Population and indication changes can similarly modify the surveillance context. Population-specific analyses may reveal reporting patterns obscured within aggregated data, although a subgroup signal remains distinct from confirmed subgroup harm [18]. Relevant changes include age, pregnancy, organ impairment, comorbidity, ancestry, concomitant therapy, and access pathway.

The proposed framework therefore treats formulation, route, indication, identifiable product, batch, and exposed population as separate versioned attributes. **Table 1** organizes the evidence, constructs, relationships,

uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in adaptive pharmacovigilance under model drift, product evolution, data-source change.

Table 1. Definitions, components, relationships, and intended uses in Adaptive Pharmacovigilance under Model Drift, Product Evolution, Data-Source Change

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Versioned surveillance baseline	Unclear reference state after change	Product, population, data, model, workflow, intended use	**Proposed:** records the state against which later change is interpreted	Traceable comparison and recoverability	Version completeness and reconstruction testing	Missing or inconsistent versions	A documented baseline does not establish validity
Model-performance monitoring	Undetected analytical degradation	Calibration, discrimination, subgroup error, alert burden	**Proposed:** performance change triggers investigation, not automatic updating	Earlier recognition of possible analytical misalignment	Temporal, external, and subgroup evaluation [13]	False alarms or missed deterioration	Performance change does not establish safety change
Product and batch identity	Inability to attribute observations accurately	Brand, manufacturer, formulation, presentation, batch	Traceability links observations to the relevant medicinal product [16]	More specific investigation of product evolution	Product-identification completeness and audit testing	Misclassification or untraceable exposure	Product identity does not establish causality
Population and indication context	Aggregate results conceal heterogeneous exposure	Age, indication, comorbidity, physiology, concomitant treatment	Population-specific surveillance may reveal distinct patterns [18]	Better contextualization of subgroup signals	Stratified validation and adequate representation	Sparse data, multiplicity, unstable estimates	Subgroup association is not confirmed subgroup harm
Data-source and coding state	Apparent change caused by altered observation	Source coverage, coding version, mapping, linkage, missingness	**Proposed:** source and terminology changes are assessed before clinical interpretation	Separation of measurement discontinuity from safety change	Bridging, overlap, provenance, and mapping studies	Semantic loss and artificial discontinuity	Standardization does not guarantee equivalence
Human governance and rollback	Unaccountable or irreversible system modification	Decision authority, audit record, override, prior states	**Proposed:** attributable review selects continuation, restriction, pause, or rollback	Controlled action under uncertainty	Human-factors testing and rollback exercises	Nominal oversight or unrecoverable states	Human involvement does not compensate for invalid evidence

Data-source and coding-system evolution

Real-world-data quality is conditional on fitness for the intended purpose rather than an intrinsic property that can be summarized by one universal score [19]. A source suitable for descriptive surveillance may be inadequate for causal assessment, subgroup analysis, calibration evaluation, or product-level attribution.

Standardized data structures and analytical frameworks can improve reproducibility across observational databases [20]. Nevertheless, structural standardization does not ensure equivalent clinical meaning, completeness, coding practice, capture timing, or transportability across institutions and health systems.

Coding-system evolution may introduce discontinuities that resemble epidemiological change. Transitions between classification systems can alter disease or outcome frequencies without corresponding biological change [21]. Similar problems may arise through terminology updates, mapping revisions, changed extraction rules, or modifications to case-processing practice.

Adaptive pharmacovigilance should therefore version data sources, mappings, coding systems, linkage procedures, and extraction logic. Where feasible, source transitions should include overlap periods, dual coding, bridging analyses, and explicit examination of missingness and reporting propensity. Residual discontinuity must remain analytically indeterminate until measurement and safety explanations have been assessed separately.

Emerging use patterns and off-label exposure

Off-label and unlicensed medicine use are related but distinct concepts and should not be used interchangeably [22]. Their interpretation depends on jurisdiction, authorization status, patient group, indication, dose, route, and formulation. Clear terminology is necessary before emerging use can be monitored coherently.

Spontaneous-report data may identify safety concerns within off-label exposure contexts, including populations insufficiently represented in preauthorization development [23]. However, such reports do not provide exposure denominators, incidence estimates, or independent causal confirmation.

Use-pattern drift may involve new indications, altered doses, longer treatment duration, polypharmacy, online access, shortages, substitution, or migration into different demographic and geographic populations. These changes may modify both true risk and the probability that an event is observed or reported.

The framework proposes that medicine-use context be versioned alongside the product and analytical system. **Table 2** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for adaptive pharmacovigilance under model drift, product evolution, data-source change.

Table 2. Testable propositions, evidence requirements, and validation criteria for Adaptive Pharmacovigilance under Model Drift, Product Evolution, Data-Source Change

Architecture layer or process stage	Core function	Information or material flow	Interaction with other components	Validation criterion	Potential failure mode	Human or experimental responsibility	Readiness boundary
Versioned baseline	Defines the reference surveillance state	Product, population, source, coding, model, workflow, intended use	Supplies comparison state for all later stages	Prior states can be reconstructed accurately	Incomplete lineage or undocumented changes	Data stewards and safety-system owners verify versions	Documentation alone does not establish performance
Multidomain sensing	Detects potentially material change	Product, use, data, model, and workflow indicators	Feeds the change-impact hypothesis	Detection is reproducible and distinguishable from ordinary variation	Alert fatigue or undetected change	Analysts define and prospectively test triggers	A trigger is not a safety signal or causal conclusion
Change-impact hypothesis	States how a change could affect surveillance meaning	Affected components, mechanisms, populations, alternatives	Determines evidence scope and priority	Competing explanations are explicit and testable	Premature attribution	Multidisciplinary reviewers document alternatives	Plausibility is not empirical confirmation
Consequence-proportionate evidence	Matches evaluation intensity to potential consequence	Exposure reach, uncertainty, reversibility, decision impact	Determines targeted revalidation design	Evidence burden is justified by intended use	Under-testing or disproportionate resource use	Governance body approves evidence level	No universal evidence tier is established
Targeted revalidation	Tests affected	Temporal, external, subgroup,	Informs human	Results are relevant to the changed	Validation on unchanged or	Independent analysts or investigators perform testing	Revalidation does not confer

	system properties	provenance, workflow evidence	authorization	component and intended setting	unrepresentative data		regulatory acceptance
Source-transition assessment	Separates measurement change from safety change	Parallel coding, overlap data, mappings, completeness	Interacts with model and signal interpretation	Discontinuity is explained or bounded	Artificial trends from coding change	Data and terminology specialists assess comparability	Apparent trend is not presumed clinical change
Population and use-pattern assessment	Tests whether validity differs across exposure contexts	Indication, dose, duration, subgroup, co-medication	Interacts with product and model monitoring	Aggregate and subgroup behaviour are compared [18]	Sparse data or hidden heterogeneity	Pharmacoepidemiologists define subgroup analyses	Stable overall results do not establish subgroup validity
Governed action and rollback	Selects continuation, modification, restriction, pause, or recovery	Revalidation evidence, uncertainty, audit record	Produces a new versioned state	Decisions are attributable, reversible where feasible, and monitored	Irreversible change or nominal review	Authorized safety leadership documents the decision	Operational action does not establish improved outcomes

Change-impact assessment and revalidation

Revalidation should begin with transparent documentation of the changed data, model, outcome, population, setting, and intended use [24]. Without this information, apparently favourable performance may reflect incomparable evaluation conditions rather than preserved validity.

Evidence should also be assessed for methodological quality, risk of bias, and applicability to the intended decision context [25]. A technically correct evaluation may remain unsuitable when its population, data source, outcome definition, or operational setting differs materially from the proposed use.

Transparent reporting and structured assessment of quality, bias, and applicability are complementary requirements for interpretable revalidation evidence [24, 25]. Neither requirement establishes that a signal is causal, that a model is clinically beneficial, or that a modified system is acceptable for regulatory reliance.

The framework proposes consequence-proportionate revalidation. Documentary verification may be sufficient for a low-consequence metadata correction, whereas model, coding, product, population, or intended-use changes may require temporal, external, subgroup, prospective, or human-factors evaluation. Evidence intensity should reflect plausible consequence, uncertainty, exposure reach, detectability, and reversibility rather than a universal threshold.

Human oversight, auditability, and rollback

Explainability requirements should depend on the pharmacovigilance task, intended user, uncertainty, and consequence of error [26]. The explanation needed for case prioritization may differ from that required for signal escalation, causal assessment, or a decision affecting medicine use.

Ethical governance also requires transparency, accountability, fairness, privacy protection, and identifiable responsibility [27]. These principles should be translated into operational controls, including role definitions, access restrictions, decision records, conflict documentation, and review of differential effects across populations.

The proposed framework defines human oversight as an attributable decision function. Reviewers must be able to interrogate inputs, reject outputs, request further evidence, record disagreement, restrict intended use, and authorize pause or rollback. Merely retaining a person at the end of an automated process does not ensure meaningful control.

Auditability requires reconstruction of the data, product, coding, model, rule, reviewer, and evidence versions supporting each action. Rollback requires a recoverable prior state and a defined post-rollback assessment. It should not be interpreted as proof that safety has been restored, because the underlying product, population, or data-generating environment may also have changed.

Limitations and implementation priorities

Responsible machine-learning implementation requires attention to development, validation, workflow integration, monitoring, and potential harm throughout the system lifecycle [28]. Applying these principles to pharmacovigilance requires task-specific adaptation because signal detection, causal assessment, case processing, and regulatory decision support have different evidentiary requirements.

Global applicability is constrained by health-data poverty and unequal representation of populations, products, coding systems, and health-service environments [29]. A framework calibrated in data-rich settings may perform differently where reporting infrastructures, access pathways, diagnostic practices, or product traceability are less complete.

Responsible implementation therefore requires both harm-oriented technical governance and explicit attention to populations and health systems that are poorly represented in available data [28, 29]. Local analytical success cannot be assumed to establish fairness, transportability, or global decision readiness.

The principal priorities are prospective multicentre evaluation, subgroup and transportability studies, product- and source-transition analyses, silent-mode monitoring, human-factors research, resource assessment, and controlled rollback exercises. Research should examine missed changes, false escalation, reviewer burden, response latency, equity effects, and whether adaptive governance improves decisions compared with well-designed static surveillance.

CONCLUSION

The Adaptive Pharmacovigilance Governance Lifecycle provides an original conceptual architecture for governing pharmacovigilance when medicinal products, populations, medicine-use patterns, data sources, coding systems, analytical models, and organizational practices change. Its central contribution is to connect versioned baselines, multidomain change sensing, explicit impact hypotheses, proportionate revalidation, accountable human authorization, rollback, and post-change monitoring while preserving distinctions among adverse events, safety signals, causal conclusions, model outputs, and regulatory decisions. The framework does not establish universal thresholds, validated performance, clinical benefit, regulatory acceptance, or deployment readiness. Its categories and propositions require empirical testing across products, populations, surveillance methods, and global health systems.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Bate A, Hobbiger SF. Artificial intelligence, real-world automation and the safety of medicines. Drug Saf. 2021;44(2):125-32. doi:10.1007/s40264-020-01001-7
2. Morales DR, Gardarsdottir H. From evidence to impact: Closing the loop in pharmacovigilance with real-world data. Pharmacoepidemiol Drug Saf. 2025;34(11). doi:10.1002/pds.70244
3. Kompa B, Hakim JB, Palepu A, Kompa KG, Smith M, Bain PA, et al. Artificial intelligence based on machine learning in pharmacovigilance: A scoping review. Drug Saf. 2022;45(5):477-91. doi:10.1007/s40264-022-01176-1
4. Kjoersvik O, Bate A. Black swan events and intelligent automation for routine safety surveillance. Drug Saf. 2022;45(5):419-27. doi:10.1007/s40264-022-01169-0
5. Trinh NTH, Solé E, Benkebil M. Benefits of combining change-point analysis with disproportionality analysis in pharmacovigilance signal detection. Pharmacoepidemiol Drug Saf. 2019;28(3):370-6. doi:10.1002/pds.4613
6. Radecka A, Loughlin L, Foy M, de Ferraz Guimaraes MV, Macolic Sarinic V, Di Giusti MD, et al. Enhancing pharmacovigilance capabilities in the EU regulatory network: The SCOPE Joint Action. Drug Saf. 2018;41(12):1285-302. doi:10.1007/s40264-018-0708-5

7. Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The clinician and dataset shift in artificial intelligence. N Engl J Med. 2021;385(3):283-6. doi:10.1056/NEJMc2104626
8. Subbaswamy A, Saria S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics. 2020;21(2):345-52. doi:10.1093/biostatistics/kxz041
9. Davis SE, Lasko TA, Chen G, Siew ED, Matheny ME. Calibration drift in regression and machine learning models for acute kidney injury. J Am Med Inform Assoc. 2017;24(6):1052-61. doi:10.1093/jamia/ocx030
10. Ball R, Dal Pan G. "Artificial intelligence" for pharmacovigilance: Ready for prime time? Drug Saf. 2022;45(5):429-38. doi:10.1007/s40264-022-01157-4
11. Schmider J, Kumar K, LaForest C, Swankoski B, Naim K, Caubel PM. Innovation in pharmacovigilance: Use of artificial intelligence in adverse event case processing. Clin Pharmacol Ther. 2019;105(4):954-61. doi:10.1002/cpt.1255
12. Danysz K, Cicirello S, Mingle E, Assuncao B, Tetarenko N, Mockute R, et al. Artificial intelligence and the future of the drug safety professional. Drug Saf. 2019;42(4):491-7. doi:10.1007/s40264-018-0746-z
13. Davis SE, Greevy RA Jr, Lasko TA, Walsh CG, Matheny ME. Detection of calibration drift in clinical prediction models to inform model updating. J Biomed Inform. 2020;112:103611. doi:10.1016/j.jbi.2020.103611
14. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW, Bossuyt P, et al. Calibration: The Achilles heel of predictive analytics. BMC Med. 2019;17(1):230. doi:10.1186/s12916-019-1466-7
15. de Hond AAH, Leeuwenberg AM, Hooft L, Kant IMJ, Nijman SWJ, van Os HJA, et al. Guidelines and quality criteria for artificial intelligence-based prediction models in healthcare: A scoping review. NPJ Digit Med. 2022;5(1):2. doi:10.1038/s41746-021-00549-7
16. Klein K, Stolk P. Challenges and opportunities for the traceability of (biological) medicinal products. Drug Saf. 2018;41(10):911-8. doi:10.1007/s40264-018-0678-7
17. Klein K, Hazell L, Stolk P, Shakir S. The UK BIO-TRAC study: A cross-sectional study of product and batch traceability for biologics in clinical practice and electronic adverse drug reaction reporting in the UK. Drug Saf. 2020;43(3):255-63. doi:10.1007/s40264-019-00891-6
18. Star K, Sandberg L, Bergvall T, Choonara I, Caduff-Janosa P, Edwards IR. Paediatric safety signals identified in VigiBase: Methods and results from Uppsala Monitoring Centre. Pharmacoepidemiol Drug Saf. 2019;28(5):680-9. doi:10.1002/pds.4734
19. Reynolds MW, Bourke A, Dreyer NA. Considerations when evaluating real-world data quality in the context of fitness for purpose. Pharmacoepidemiol Drug Saf. 2020;29(10):1316-8. doi:10.1002/pds.5010
20. Reps JM, Schuemie MJ, Suchard MA, Ryan PB, Rijnbeek PR. Design and implementation of a standardized framework to generate and evaluate patient-level prediction models using observational healthcare data. J Am Med Inform Assoc. 2018;25(8):969-75. doi:10.1093/jamia/ocy032
21. Khera R, Dorsey KB, Krumholz HM. Transition to the ICD-10 in the United States: An emerging data chasm. JAMA. 2018;320(2):133-4. doi:10.1001/jama.2018.6823
22. Aronson JK, Ferner RE. Unlicensed and off-label uses of medicines: Definitions and clarification of terminology. Br J Clin Pharmacol. 2017;83(12):2615-25. doi:10.1111/bcp.13394
23. Mascolo A, Scavone C, Bertini M, Brusco S, Punzo F, Pota E, et al. Safety of anticancer agents used in children: A focus on their off-label use through data from the spontaneous reporting system. Front Pharmacol. 2020;11:621. doi:10.3389/fphar.2020.00621
24. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385. doi:10.1136/bmj-2023-078378
25. Moons KGM, Damen JAA, Kaul T, Hooft L, Andaur Navarro C, Dhiman P, et al. PROBAST+AI: An updated quality, risk of bias, and applicability assessment tool for prediction models using regression or artificial intelligence methods. BMJ. 2025;388. doi:10.1136/bmj-2024-082505
26. Hauben M. Artificial intelligence in pharmacovigilance: Do we need explainability? Pharmacoepidemiol Drug Saf. 2022;31(12):1311-6. doi:10.1002/pds.5501
27. Jain A, Salas M, Aimer O, Adenwala Z. Safeguarding patients in the AI era: Ethics at the forefront of pharmacovigilance. Drug Saf. 2025;48(2):119-27. doi:10.1007/s40264-024-01483-9
28. 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

29. Ibrahim H, Liu X, Zariffa N, Morris AD, Denniston AK. Health data poverty: An assailable barrier to equitable digital health care. Lancet Digit Health. 2021;3(4). doi:10.1016/S2589-7500(20)30317-4