Original Article

Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects

Saif Al-Hinai[1]*, Amal Al-Balushi[1], Mohammed Al-Maskari[2], Sultan Al-Mahruqi[3]

[1]Department of Reversible AI Medication Plans, College of Pharmacy, Sultan Qaboos University, Muscat, Oman. [2]Department of Auditable Clinical AI Objects, College of Pharmacy, University of Nizwa, Nizwa, Oman. [3]Department of Medication Plan Governance, College of Pharmacy, Dhofar University, Salalah, Oman.

Abstract

Artificial intelligence can generate medication plans that appear clinically coherent while remaining difficult to inspect, update, attribute, or reconstruct. When such plans persist only as prose, their assumptions, supporting evidence, model provenance, professional review, and subsequent modifications may become disconnected from the decisions they influence. This article proposes that an AI-generated medication plan should be represented as a reversible and auditable clinical-information object rather than treated as ephemeral text or an automatically executable medication order. The proposed Auditable Medication-Plan Object Model combines persistent identity, patient- and time-specific context, structured medication actions, rationale, evidence dependencies, uncertainty, model provenance, professional authorization, immutable versions, governance states, and append-only audit events. Material modification creates a successor version without erasing its predecessor. Reversal is separated into informational withdrawal or rollback and, where medication-related action has already occurred, clinical corrective or compensating action. Approval is represented as a version-specific professional decision rather than evidence of correctness or safety. Evaluation should test identity integrity, lineage completeness, state-transition validity, evidence traceability, recoverability, role comprehension, workflow burden, equity implications, and the separation of technical restoration from clinical recovery. The model is an original conceptual synthesis intended to organize future specification, human-factors assessment, workflow simulation, and prospective evaluation. It does not establish clinical effectiveness, medication-safety improvement, regulatory acceptability, legal sufficiency, universal applicability, or deployment readiness.

Keywords: Digital pharmacy, Artificial intelligence, Pharmacy practice, Medication safety, Human–AI collaboration, Clinical decision support

Introduction

Artificial intelligence is increasingly positioned as an augmentative component of clinical work rather than an independent replacement for professional judgment. Its potential value therefore depends not only on model capability but also on how outputs are integrated into accountable clinical relationships, information systems, and decisions [1].

Medication planning makes this integration unusually consequential. A generated plan may combine proposed initiation, discontinuation, dose modification, monitoring, interaction management, adherence support, and follow-up. Even when no recommendation is executed automatically, the plan may influence pharmacists, prescribers, patients, and subsequent digital systems.

Generative systems can produce medically plausible language while retaining clinically important limitations, uncertainty, and error risks [2]. Prose can obscure which patient data were available, which assumptions were made, which evidence was consulted, and which model version generated the response. Copying the text into a note does not necessarily preserve those dependencies. Later editing may remove the distinction between generated and professionally authored content, while a new output may silently replace an earlier plan without documenting why the change occurred. Fluency therefore cannot function as a substitute for provenance, authorization, or recoverability.

Address for correspondence: Saif Al-Hinai, Department of Reversible AI Medication Plans, College of Pharmacy, Sultan Qaboos University, Muscat, Oman.
saif.alhinai@squ.edu.om
Received: 09 March 2026; **Accepted:** 01 June 2026

How to cite this article: Al-Hinai S, Al-Balushi A, Al-Maskari M, Al-Mahruqi S. Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects. Arch Pharm Pract. 2026;17(2):75-83. https://doi.org/10.51847/OfbpAj2FTy

Medication-related decision support already operates within data-quality, alert-burden, workflow, and clinical-value constraints [3]. These constraints suggest that the safety problem is not limited to whether an individual recommendation is accurate. It also concerns whether the recommendation can be reviewed in context, challenged, modified without erasing history, linked to subsequent action, and reconstructed after a discrepancy or adverse event. A medication plan that persists only as unstructured text may be readable yet poorly governable.

Responsible clinical artificial intelligence requires safeguards across development, implementation, monitoring, and changing use contexts [4]. This article extends that lifecycle reasoning from the model to each patient-specific plan produced by or with the model. It proposes an Auditable Medication-Plan Object Model in which generated content becomes a persistent but non-executable clinical-information object. The object retains its identity across immutable versions, records evidence and assumptions, distinguishes review from approval, and preserves every material transition. "Reversible" refers to controlled withdrawal, replacement, or restoration of information states and not to the literal undoing of medication exposure. "Auditable" refers to reconstructable content, lineage, actors, timing, reasons, and dependencies and not merely to the availability of an explanation. The contribution is conceptual and requires technical, sociotechnical, clinical, and governance validation.

Requirements for a Governed Clinical Object

A governed clinical object requires more than a text field with a timestamp. Healthcare artificial intelligence has been associated with separable requirements involving reliability, transparency, traceability, responsibility, user control, review, monitoring, and audit trails [5]. Applied to medication planning, these requirements imply that the object must expose what was generated, the context in which it was generated, the intended use, the person or system responsible for each subsequent decision, and the evidence needed to assess whether its status remains valid. The article therefore treats governance as a property of the object's entire lifecycle rather than a final approval label.

Information needs also differ by role. A pharmacist may need access to medication history, interactions, laboratory trends, contraindications, uncertainty, and unresolved questions. A prescriber may need the proposed therapeutic change, rationale, alternatives, and authorization consequences. A patient may require an understandable explanation of purpose, choices, risks, and follow-up. Developers and governance teams may require model version, input provenance, error reports, and monitoring metadata. Multidisciplinary work on explainability supports the principle that information should be matched to the user and decision context [6]. The proposed object therefore permits role-sensitive views while retaining a common underlying record.

Explainability and auditability must remain distinct. A model-generated rationale may help a reviewer understand the apparent basis of a recommendation, but current explanation approaches cannot by themselves establish evidence validity, model reliability, causal correctness, or clinical safety [7]. An auditable medication-plan object must consequently preserve source data, evidence dependencies, assumptions, model provenance, human decisions, and version history even when an explanation is available. Explanations become inspectable components of the object rather than substitutes for governance.

Structured documentation frameworks demonstrate how provenance, composition, intended uses, assumptions, and limitations can be made visible [8]. This article proposes transferring that documentation logic to patient-specific medication plans while recognizing that a clinical plan carries more immediate decision consequences than a dataset description. Required properties include persistent identity, immutable versions, explicit lineage, attributable events, evidence-linked plan elements, defined governance states, and recoverability at a specified time. These properties are proposed requirements for evaluation; satisfying them would not establish clinical correctness, legal adequacy, regulatory acceptance, or safety.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects.

Table 1. Definitions, components, relationships, and intended uses for the article Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Governed clinical object	Generated prose may lack persistent control and accountability.	Plan content, context, provenance, review, state, and event records.	**Proposed:** bind content and governance metadata under one persistent identity.	Inspectable lifecycle rather than isolated output.	Formal specification and workflow evaluation informed by responsible-AI requirements [5].	A complete object may still contain an unsafe plan.	Object integrity does not establish clinical validity.

Role-sensitive representation	Different users require different information.	User role, decision purpose, authorization, and communication needs.	**Proposed:** generate role-specific views from a common record.	Reduced information mismatch during review and communication.	Pharmacist, prescriber, patient, and governance usability testing informed by multidisciplinary explainability work [6].	Important details may be hidden by oversimplified views.	Role adaptation must not alter the underlying record.
Evidence dependency	Recommendations may become detached from their basis.	Sources, patient data, assumptions, and plan elements.	**Proposed:** link each material plan element to supporting or constraining evidence.	Targeted reassessment when evidence changes.	Traceability and reviewer-reconstruction testing.	Links may be present but scientifically weak or outdated.	Traceability does not prove evidence quality.
Assumption and uncertainty record	Fluent language may conceal incomplete information.	Missing data, inferred conditions, model uncertainty, and unresolved questions.	**Proposed:** store assumptions separately from observed facts and clinical decisions.	More explicit review and escalation.	Human-factors assessment of recognition and interpretation.	Users may ignore, misunderstand, or ritualistically accept uncertainty fields.	Documentation does not eliminate uncertainty.
Explanation component	Explanation may be mistaken for validation.	Model rationale, reviewer rationale, and evidence links.	Treat explanation as one inspectable field within a broader audit structure.	Supports scrutiny without assigning explanation excessive authority.	Comparison of explanation use, error detection, and decision quality, recognizing known limitations [7].	Persuasive explanations may amplify automation bias.	Explanation is not proof of correctness or safety.
Provenance record	Model, data, or software changes may be invisible.	Generator identity, model version, software version, source data, and generation time.	**Proposed:** preserve provenance for every immutable version.	Reconstruction of the conditions under which a plan was produced.	Completeness, integrity, and cross-system preservation tests informed by structured documentation principles [8].	Provenance can be complete while inputs remain biased or erroneous.	Provenance records origin, not validity.
Approval state	Generated, reviewed, and authorized content may be conflated.	Reviewer role, authorization, intended use, conditions, and time.	**Proposed:** use explicit state transitions rather than an undifferentiated "accepted" status.	Clearer authority and use boundaries.	Simulation and prospective evaluation of state interpretation.	Approval may become superficial or ambiguous.	Approval is not evidence of adequate review.
Version and lineage	Editing may erase the history of a plan.	Prior version, modification reason, actor, and successor relationship.	**Proposed:** material changes create immutable successor versions.	Non-destructive comparison and reconstruction.	Version-integrity, lineage, and recoverability testing.	Excessive versioning may burden users or obscure the current plan.	Technical recoverability does not restore completed clinical actions.

Proposed Medication-Plan Object Model

The proposed model begins with an interoperable object boundary. Digital medicine depends on the ability to preserve meaning when information moves across systems and professional settings [9]. The medication-plan object is therefore conceptualized as a persistent information layer rather than a model-specific response format. It does not replace a prescription, medication order, administration record, clinical note, or pharmacist intervention record. Instead, it links generated planning content to the data, evidence, professional decisions, and subsequent clinical records required to interpret its use.

Structured standards such as HL7 FHIR demonstrate that granular longitudinal clinical and medication information can be represented in exchangeable forms [10]. The proposed object could therefore be implemented through profiles or linked resources, but this article does not define a complete FHIR implementation guide. Its conceptual layers are: object identity and lineage; patient–encounter–time context; medication actions and intended purpose; rationale and evidence dependencies; assumptions, exclusions, uncertainty, and unresolved questions; generator and model provenance; professional review; approval state and permitted use; and version, audit, monitoring, and validation metadata.

Interoperability does not guarantee semantic uniformity. FHIR-based initiatives differ in scope, resource selection, extensions, terminology, validation, and implementation context [11]. The model consequently requires explicit definitions for each plan element and transition. A "medication action," for example, should identify whether it proposes initiation, discontinuation, substitution, dose adjustment, monitoring, reconciliation, adherence support, or referral. Each action should remain linked to the patient context and evidence on which it depends. Local implementations could represent these elements differently,

but their meaning and lineage would need to survive exchange.

The object remains subordinate to professional authority. Responsible AI-enabled decision support requires controlled development, testing, supervision, use, and accountability [12]. A generated object therefore begins as a draft. It may be rejected, revised, reviewed, or approved for a specified purpose, but it cannot become clinically authoritative merely because it is complete or technically valid. Material alteration creates a successor version; it does not silently overwrite the approved record. This architecture is proposed to improve inspectability and recoverability, not to establish safer decisions.

Figure 1 presents the original conceptual synthesis for auditable medication-plan object model, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

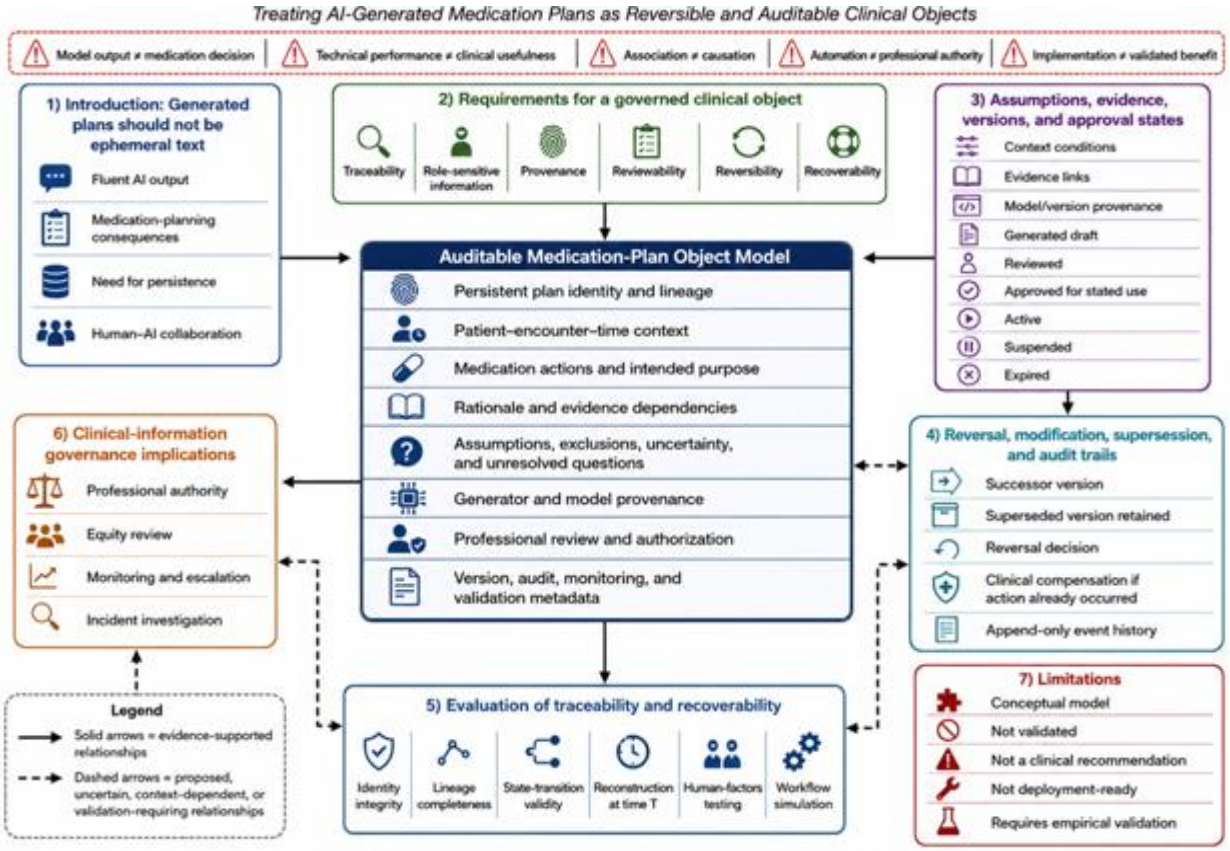

Figure 1. Auditable Medication-Plan Object Model

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Identity and lineage	Maintain one plan identity across immutable versions.	Initial object identifier → successor-version links.	Connects provenance, state, evidence, and audit events.	Unique identity, valid parent links, no destructive overwrite.	Duplicate, orphaned, or circular lineage.	Informatics team defines integrity controls; clinical governance approves their use.	Passing integrity tests does not establish clinical correctness.

Patient–encounter–time context	Bind the plan to the circumstances in which it was generated.	Clinical data snapshot → plan elements and validity conditions.	Conditions evidence interpretation and approval.	Completeness, temporal consistency, and contextual plausibility.	Stale, missing, or mismatched patient data.	Clinical reviewer confirms relevance; system owner preserves provenance.	Structured representation is feasible [10], but local clinical adequacy requires validation.
Medication-action layer	Represent proposed medication changes and intended purpose.	Context and evidence → discrete proposed actions.	Links to rationale, uncertainty, approval, and subsequent records.	Action specificity and unambiguous interpretation.	Ambiguous dose, route, duration, purpose, or conditional instruction.	Pharmacist verifies medication meaning and identifies unresolved ambiguity.	The object is not an executable order.
Evidence and rationale	Make the basis of each material action inspectable.	Evidence source and patient datum → dependent plan element.	Triggers reassessment when supporting information changes.	Reviewer can reconstruct why an action was proposed.	Citation without relevance, outdated evidence, or persuasive unsupported rationale.	Clinical reviewer evaluates evidence fit; governance defines minimum documentation.	Traceability does not prove scientific validity.
Provenance	Preserve generator, model, software, source data, and time.	Generation environment → immutable version metadata.	Supports monitoring, comparison, and incident investigation.	Complete and internally consistent provenance.	Unknown model version or lost generation context.	Model owner maintains provenance; organization verifies preservation.	Provenance does not establish model performance.
Interoperability layer	Preserve structure and meaning across systems.	Object fields → mapped exchange resources and terminologies.	Connects the object to medication, observation, encounter, and practitioner records.	Syntactic conformance and semantic equivalence.	Valid exchange with altered clinical meaning.	Informatics and terminology teams maintain mappings.	Interoperability is necessary but insufficient for safe use [9].
Review and authorization	Separate generation, review, revision, and approval.	Draft → review decision → state transition.	Controls permitted use and version activation.	Authorized actor, reason, conditions, and resulting state are recorded.	Rubber-stamp review or unclear authority.	Pharmacy and medical leadership define role boundaries.	Responsible oversight is required [12], but human presence alone does not ensure safety.
Monitoring and change control	Detect changes affecting validity and create successors.	New evidence, patient change, or model change → impact assessment.	Interacts with assumptions, versions, suspension, and supersession.	Affected dependencies are identifiable and material changes create new versions.	Silent alteration, unassessed drift, or excessive version burden.	Governance body defines triggers; clinical teams assess consequences.	Implementation heterogeneity requires local validation [11].

Assumptions, Evidence, Versions, and Approval States

The proposed object is conditional on assigned responsibility. Local deployment frameworks for clinical prediction models emphasize multidisciplinary oversight across development, validation, implementation, and monitoring [13]. A medication-plan object similarly requires named authority for generation controls, pharmacy review, prescriber or authorized-clinician decisions, technical maintenance, incident investigation, and governance escalation. These functions may be distributed differently across organizations, but the allocation should be explicit. The object should not imply that accountability belongs solely to the model developer or to the pharmacist who encounters the output.

Plan validity is also time-dependent. Clinical models can experience calibration drift as populations, practice, and data relationships change [14]. For a generated medication plan, model provenance must therefore include the model and software version, generation time, relevant configuration, and source-data context. Plan-object versions and model versions are distinct: the same model can generate different plan versions as patient information changes, while the same plan content may require reassessment after the underlying model or evidence changes. Version metadata makes that relationship inspectable but does not quantify whether a new model is better.

Assumptions should be recorded as conditions rather than hidden premises. Clinical translation of artificial intelligence is constrained by dataset shift, generalizability, confidence, workflow integration, and sociocultural change [15]. A medication-plan object should therefore distinguish observed patient facts from inferred conditions, missing information, applicability assumptions, exclusions, unresolved questions, and uncertainty. Each material plan element should identify the evidence and patient-data dependencies whose change could invalidate it. This dependency structure is proposed; empirical testing must determine whether reviewers can interpret it without excessive cognitive or documentation burden.

Approval states provide governance meaning to versions. Early clinical evaluation of AI decision support should examine actual use, errors, human factors, workflow, and safety rather than relying only on offline model performance [16]. The proposed states are: generated draft, review pending, revision required, clinically reviewed, approved for stated use, active, suspended, superseded, reversed, expired, and archived. A version cannot move directly from generated draft to active. Approval applies only to the specified version, context, intended use, conditions, and validity period. A successor version does not automatically inherit approval. These states are original propositions for formal, simulation, human-factors, and prospective testing. They do not establish clinical readiness, and an approved state cannot be interpreted as proof that the plan is correct, complete, acceptable to the patient, or safe beyond its recorded conditions.

Reversal, Modification, Supersession, and Audit Trails

An auditable medication-plan object requires non-destructive change control. Electronic health record audit-log research shows that clinical-system activity can be reconstructed only when events, measures, users, and timestamps are defined consistently [17]. The proposed model therefore prohibits overwriting an approved or active version. A material modification creates a successor version linked to its predecessor and records the actor, role, reason, evidence, affected fields, previous state, and resulting state.

Timestamped activity records can support analysis of workflow and accountability [18]. Within the proposed model, every generation, review, revision, approval, activation, suspension, reversal, supersession, expiry, and archival event enters an append-only audit trail. The event record complements rather than replaces the immutable version snapshot. A log showing that a change occurred is insufficient unless the content before and after the change can also be reconstructed.

Audit trails remain incomplete representations of clinical work. Event logs may omit verbal communication, handwritten information, reasoning outside the system, patient preferences, and work performed through unrecorded channels [19]. The object should therefore include structured reason fields and links to associated clinical documentation without claiming to capture the complete decision process. Apparent inactivity in a digital log cannot be interpreted automatically as absence of professional work.

Recoverability also depends on record quality. Completeness, conformance, and plausibility are distinct dimensions of electronic health record data quality [20]. A plan may be technically retrievable yet clinically uninterpretable because its context, evidence links, or state history are missing. Evaluation should therefore test whether an authorized reviewer can reconstruct which version was valid at a specified time and why.

Reversal must be distinguished from deletion. Before execution, reversal may withdraw a plan or restore a previously valid version. After a medication has been supplied, administered, discontinued, or otherwise affected care, a digital rollback cannot undo the clinical consequence. The object must instead record the reversal decision, assess whether an affected action occurred, document any corrective or compensating action, create a new plan version where required, and preserve the complete history. Supersession similarly makes a successor current without erasing the prior version. These rules are proposed governance mechanisms and have not been shown to reduce medication-related harm.

Evaluation of Traceability and Recoverability

Evaluation should begin with object-level tests rather than clinical-effect claims. Reports of artificial-intelligence interventions should describe inputs, outputs, errors, human interaction, and intervention handling sufficiently for independent appraisal [21]. For the proposed object, corresponding tests include identifier uniqueness, version immutability, lineage completeness, permitted state transitions, actor attribution, timestamp integrity, evidence-link reconstruction, and preservation across system exchange.

Prospective protocols should prespecify how unavailable inputs, contradictory evidence, system failure, user disagreement, and unsafe output will be handled [22]. Scenario testing should therefore include concurrent modification, missing laboratory data, changed renal function, revised evidence, new contraindications, model replacement, unauthorized approval, failed exchange, and reversal after partial execution. Recoverability should be measured by whether reviewers can reconstruct the plan active at a selected time, not merely whether historical records exist.

Where plan content depends on predictive or classification models, object integrity and model validity must be assessed separately. Transparent reporting of development, intended use, performance, validation, and limitations remains necessary [23]. An object may preserve an inaccurate prediction perfectly. Conversely, a well-performing model may generate a poorly governed plan if its output lacks context, authorization, or recoverable lineage.

Patient-benefit claims require transparency, replication, ethical scrutiny, effectiveness evidence, and implementation evaluation beyond technical performance [24]. Evaluation should consequently progress from formal specification and synthetic cases to retrospective reconstruction, pharmacist and prescriber human-factors testing, workflow simulation, prospective limited use, multisite evaluation, equity assessment, longitudinal monitoring, and only then comparative clinical-outcome research. Evaluation should jointly report the intervention, human interaction, model limitations, and intended-use boundaries [21-23].

Successful traceability would show that information dependencies can be followed. Successful recoverability would show that the valid object state can be reconstructed. Neither result would demonstrate medication safety, clinical effectiveness, workflow benefit, patient acceptability, or regulatory readiness.

Clinical-Information Governance Implications

Governance must examine how generated plans affect different patient groups, professional decisions, and resource allocation rather than treating fairness as a single technical property [25]. **Figure 2** presents the original conceptual synthesis for versioning, reversal, supersession, approval states, and audit trails, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

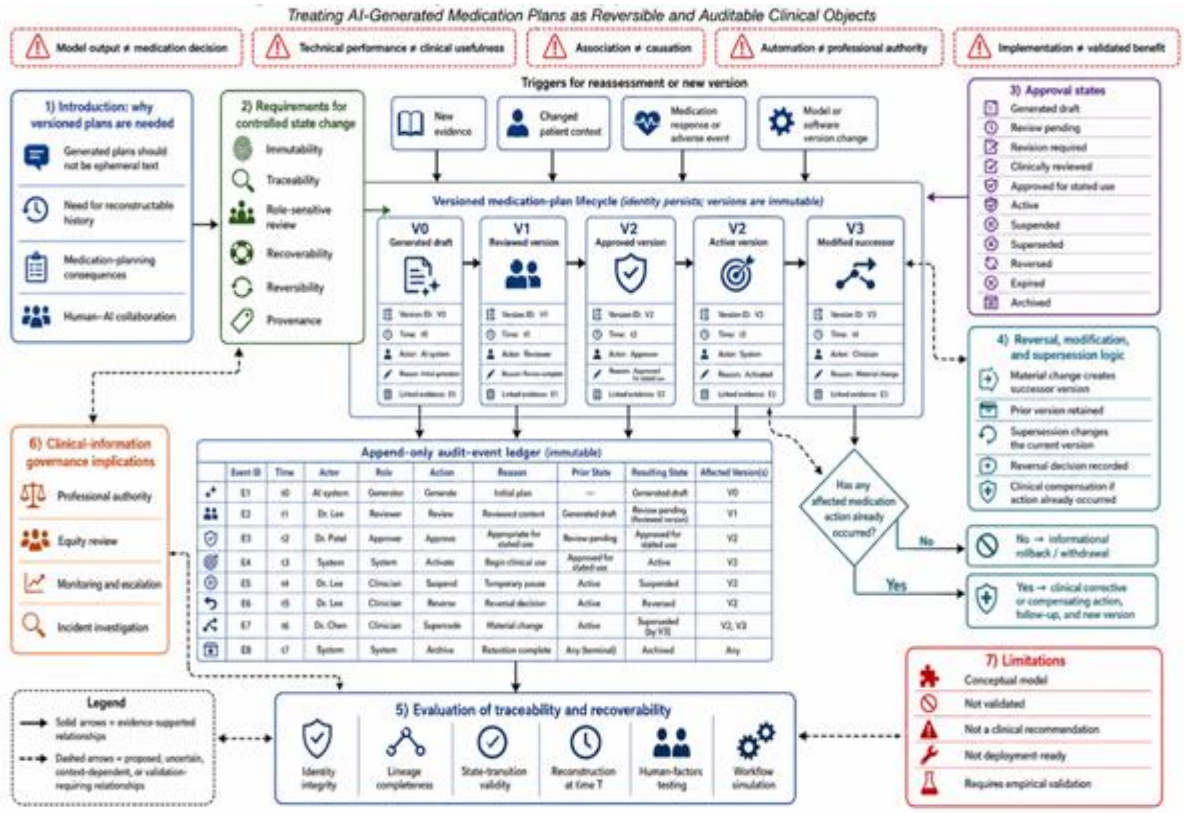

Figure 2. Versioning, Reversal, Supersession, Approval States, and Audit Trails

Apparently neutral targets or proxies can reproduce structural inequity when they do not represent the clinical need they are assumed to measure [26]. The object should therefore retain model targets, proxy definitions, relevant subgroup information, intended use, and downstream decision consequences where lawful and appropriate. These records may support investigation but cannot make the plan equitable by themselves.

Technical fairness adjustments also cannot settle every normative or structural problem in healthcare [27]. Equity governance must address model design, proxy selection, access, professional interpretation, patient communication, and downstream allocation rather than relying on a single fairness metric [25-27]. Human decision makers retain epistemic and normative responsibilities even when algorithmic support is used [28]. Approval rights, escalation routes, override authority, and responsibility for corrective action must therefore be explicit.

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for governance responsibilities, implementation conditions, and monitoring requirements for the article Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects.

Table 3. Governance responsibilities, implementation conditions, and monitoring requirements for the article Treating AI-Generated Medication Plans as Reversible and Auditable Clinical Objects.

Governance domain	Responsible actor	Required authority	Documentation requirement	Monitoring indicator	Escalation trigger	Corrective action	Accountability boundary

Generation controls	Model and system owner	Control model, configuration, and permitted use.	Model version, software version, source context, and generation event.	Missing provenance or use outside scope.	Unknown generator or unauthorized use.	Suspend generation and investigate affected objects.	Technical ownership does not confer clinical authority.
Medication review	Pharmacist or authorized medication specialist	Review medication meaning, evidence, feasibility, and uncertainty.	Findings, unresolved issues, proposed revisions, and rationale.	Repeated ambiguity, missing context, or unsupported actions.	Potential medication harm or inability to complete review.	Reject, revise, or escalate the plan.	Review does not transfer sole responsibility to the pharmacist.
Clinical approval	Authorized clinician under local policy	Approve a specified version for a stated use.	Actor, role, conditions, validity period, and approval reason.	Approval without recorded review or required evidence.	Conflict, contraindication, or authority uncertainty.	Withhold or revoke approval.	Approval is not proof of correctness or patient acceptance.
Activation and monitoring	Clinical service and governance body	Permit operational use and define monitoring.	Activation event, local conditions, deviations, overrides, and incidents.	Drift, missing fields, abnormal override patterns, or workflow burden.	Safety signal, contextual change, or failed control.	Suspend use, reassess dependencies, and create corrective versions.	Local implementation does not establish generalizability.
Reversal and supersession	Authorized clinical decision maker with informatics support	Withdraw, replace, or suspend a version.	Reason, affected version, execution status, successor, and clinical compensation.	Reversal without preserved history or unlinked corrective action.	Executed medication action or unresolved patient risk.	Initiate clinical follow-up and preserve a new version.	Digital reversal cannot erase completed exposure.
Equity oversight	Multidisciplinary governance group	Review proxies, subgroup effects, access, and allocation consequences.	Fairness assumptions, subgroup analyses, complaints, and corrective decisions.	Unequal errors, exclusions, or downstream effects.	Material disparity or structurally inappropriate proxy.	Redesign, restrict, suspend, or investigate use.	Metric compliance does not establish substantive equity.
Incident investigation	Independent safety and governance function	Access complete object history and require remediation.	Reconstructed versions, events, actors, evidence, and system conditions.	Missing lineage or inability to reconstruct the active version.	Harm, near miss, unexplained state change, or audit failure.	Containment, disclosure where required, remediation, and monitoring.	Auditability supports investigation but does not prevent harm.

The proposed allocation defines functional responsibilities rather than legal liability. Organizations must adapt authority, retention, access, and escalation rules to their professional, technical, and jurisdictional context. Human authorization should not become a ceremonial confirmation step, and organizations should not transfer systemic responsibility entirely to individual practitioners.

Limitations

The article synthesizes evidence from clinical artificial intelligence, interoperability, audit logs, medication decision support, reporting standards, ethics, and governance. Much clinical-AI evidence remains retrospective, task-specific, or weakly connected to real-world comparisons [29]. Evidence from prediction, imaging, and general decision support may not transfer directly to generated medication plans.

The proposed object may also create unintended consequences, including documentation burden, automation bias, deskilling, duplicated records, false reassurance, and diffusion of responsibility [30]. Greater visibility may produce ritualized field completion rather than meaningful scrutiny.

Bias and safety risk may arise through data, model objectives, workflow integration, professional interpretation, and changing clinical environments [31]. Auditability can reveal a harmful process without preventing it. The model does not determine clinical correctness, legal sufficiency, regulatory classification, retention requirements, professional liability, or patient acceptability. Its constructs, transition rules, and evaluation criteria require formal, human-factors, workflow, equity, and prospective clinical validation.

Conclusion

AI-generated medication plans should not be treated as disposable prose or as self-authorizing clinical instructions. This article proposes representing them as persistent, versioned, and auditable clinical-information objects linked to patient context, medication actions, evidence, assumptions, uncertainty, model provenance, professional review, approval states, and append-only events.

The model's central contribution is the integration of object structure with clinical-information governance. Material changes create successor versions rather than overwriting history. Supersession preserves prior decisions. Reversal distinguishes informational withdrawal from the corrective clinical work required when medication-related action has

already occurred. Approval is version-specific and conditional rather than evidence of correctness or safety.

The model may organize future specification, interoperability work, human-factors testing, workflow simulation, incident reconstruction, equity review, and prospective evaluation. It does not demonstrate that greater auditability improves medication safety, professional decision making, patient outcomes, or organizational accountability. Those effects remain empirical questions. Clinical authority must remain with appropriately authorized professionals and institutions, while technical controls must remain open to scrutiny, correction, and contextual limitation.

ACKNOWLEDGMENTS: None
CONFLICT OF INTEREST: None
FINANCIAL SUPPORT: None
ETHICS STATEMENT: None

REFERENCES

1. Topol EJ. High-performance medicine: The convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7
2. Lee P, Bubeck S, Petro J. Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. N Engl J Med. 2023;388(13):1233-9. doi:10.1056/NEJMsr2214184
3. 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
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. Garcia-Gomez JM, Blanes-Selva V, Alvarez Romero C, de Bartolome Cenzano JC, Pereira Mesquita F, Pazos A, et al. Mitigating patient harm risks: A proposal of requirements for AI in healthcare. Artif Intell Med. 2025;167:103168. doi:10.1016/j.artmed.2025.103168
6. Amann J, Blasimme A, Vayena E, Frey D, Madai VI; Precise4Q consortium. Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. BMC Med Inform Decis Mak. 2020;20(1):310. doi:10.1186/s12911-020-01332-6
7. 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). doi:10.1016/S2589-7500(21)00208-9
8. Gebru T, Morgenstern J, Vecchione B, Vaughan JW, Wallach H, Daumé H III, et al. Datasheets for datasets. Commun ACM. 2021;64(12):86-92. doi:10.1145/3458723
9. Lehne M, Sass J, Essenwanger A, Schepers J, Thun S. Why digital medicine depends on interoperability. NPJ Digit Med. 2019;2:79. doi:10.1038/s41746-019-0158-1
10. Bennett AM, Ulrich H, van Damme P, Wiedekopf J, Johnson AEW. MIMIC-IV on FHIR: Converting a decade of in-patient data into an exchangeable, interoperable format. J Am Med Inform Assoc. 2023;30(4):718-25. doi:10.1093/jamia/ocad002
11. Duda SN, Kennedy N, Conway D, Cheng AC, Nguyen V, Zayas-Cabán T, et al. HL7 FHIR-based tools and initiatives to support clinical research: A scoping review. J Am Med Inform Assoc. 2022;29(9):1642-53. doi:10.1093/jamia/ocac105
12. Labkoff S, Oladimeji B, Kannry J, Solomonides A, Leftwich R, Koski E, et al. Toward a responsible future: Recommendations for AI-enabled clinical decision support. J Am Med Inform Assoc. 2024;31(11):2730-9. doi:10.1093/jamia/ocae209
13. Bedoya AD, Economou-Zavlanos NJ, Goldstein BA, Young A, Jelovsek JE, O'Brien C, et al. A framework for the oversight and local deployment of safe and high-quality prediction models. J Am Med Inform Assoc. 2022;29(9):1631-6. doi:10.1093/jamia/ocac078
14. 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
15. 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
16. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
17. Rule A, Chiang MF, Hribar MR. Using electronic health record audit logs to study clinical activity: A systematic review of aims, measures, and methods. J Am Med Inform Assoc. 2020;27(3):480-90. doi:10.1093/jamia/ocz196
18. Adler-Milstein J, Adelman JS, Tai-Seale M, Patel VL, Dymek C. EHR audit logs: A new goldmine for health services research? J Biomed Inform. 2020;101:103343. doi:10.1016/j.jbi.2019.103343
19. Rule A, Melnick ER, Apathy NC. Using event logs to observe interactions with electronic health records: An updated scoping review shows increasing use of vendor-derived measures. J Am Med Inform Assoc. 2023;30(1):144-54. doi:10.1093/jamia/ocac177
20. Lewis AE, Weiskopf N, Abrams ZB, Foraker R, Lai AM, Payne PRO, et al. Electronic health record data quality assessment and tools: A systematic review. J Am Med Inform Assoc. 2023;30(10):1730-40. doi:10.1093/jamia/ocad120
21. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK, Chan AW, et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. Nat Med. 2020;26(9):1364-74. doi:10.1038/s41591-020-1034-x
22. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ, Ashrafian H, et al. Guidelines for clinical trial protocols for interventions involving artificial intelligence: The SPIRIT-AI extension. Nat Med. 2020;26(9):1351-63. doi:10.1038/s41591-020-1037-7
23. 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
24. Vollmer S, Mateen BA, Bohner G, Király FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. BMJ. 2020;368. doi:10.1136/bmj.l6927
25. Rajkomar A, Hardt M, Howell MD, Corrado G, Chin MH. Ensuring fairness in machine learning to advance health equity. Ann Intern Med. 2018;169(12):866-72. doi:10.7326/M18-1990
26. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-53. doi:10.1126/science.aax2342
27. McCradden MD, Joshi S, Mazwi M, Anderson JA. Ethical limitations of algorithmic fairness solutions in health care machine learning. Lancet Digit Health. 2020;2(5). doi:10.1016/S2589-7500(20)30065-0
28. Grote T, Berens P. On the ethics of algorithmic decision-making in healthcare. J Med Ethics. 2020;46(3):205-11. doi:10.1136/medethics-2019-105586
29. Nagendran M, Chen Y, Lovejoy CA, Gordon AC, Komorowski M, Harvey H, et al. Artificial intelligence versus clinicians: Systematic review of design, reporting standards, and claims of deep learning studies. BMJ. 2020;368. doi:10.1136/bmj.m689
30. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. JAMA. 2017;318(6):517-8. doi:10.1001/jama.2017.7797
31. 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