Original Article

Translating Algorithmic Uncertainty into Actionable Medication Advice

Omar Zayed[1]*, Leila Ibrahim[1], Ahmed Mansour[2]

[1]Department of Clinical Pharmacy and Uncertainty Communication, Faculty of Pharmacy, Cairo University, Cairo, Egypt. [2]Department of AI Actionability and Medication Advice, Faculty of Pharmacy, Alexandria University, Alexandria, Egypt.

Abstract

Artificial intelligence can attach probabilities, confidence scores, intervals, disagreement measures, or out-of-distribution signals to medication-related outputs, yet these numerical representations do not independently specify what a pharmacist, prescriber, or patient should do. The problem is not only whether uncertainty is estimated, but whether its source, decision relevance, consequences, and ownership are translated into a proportionate response. This article develops a proposed Uncertainty-Translation Framework for actionable medication advice. The framework separates algorithmic output from medication decision-making through five connected functions: uncertainty characterization; calibration, coverage, data-completeness, and scope checks; interpretation of the affected medication decision; selection of a bounded action class; and recipient-specific communication supported by an auditable governance record. Proposed action classes include proceeding with qualified advice, verifying data, checking an independent source, delaying action, communicating uncertainty, or escalating to an accountable professional. The framework treats these classes as conditional design hypotheses rather than clinical rules. Validation would require technical assessment of calibration and coverage; human-factors testing of comprehension, workload, and inappropriate reliance; evaluation of medication decisions and downstream consequences; subgroup and equity analyses; and prospective governance of versioning, overrides, incidents, and withdrawal. The framework cannot correct biased targets, invalid data, unsuitable model purposes, or deficient professional judgment. Its original contribution is to define uncertainty translation as an intermediate sociotechnical function between prediction and advice while preserving the boundaries between technical performance, clinical usefulness, professional authority, and validated benefit.

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

Introduction

Medication-related algorithms can produce risk estimates, classifications, rankings, alerts, or treatment suggestions accompanied by measures labelled as confidence or uncertainty. Such outputs may appear decision-ready, but clinical impact depends on more than predictive performance and requires alignment with the intended task, workflow, users, and evaluation context [1]. A probability of harm, a wide interval, or disagreement among models does not state whether a dose should be withheld, a laboratory value rechecked, an interaction investigated, a prescriber contacted, or a patient advised to wait.

This gap is not resolved simply by adding an explanation. Post hoc explanations can describe influential features or provide a plausible account of model behaviour, but they do not establish that the recommendation is reliable, that uncertainty has been validly estimated, or that the explanation supports a safe medication decision [2]. Advice requires a normative and operational step: uncertainty must be related

Address for correspondence: Omar Zayed, Department of Clinical Pharmacy and Uncertainty Communication, Faculty of Pharmacy, Cairo University, Cairo, Egypt.
omar.zayed@cu.edu.eg
Received: 10 May 2025; **Accepted:** 30 August 2025

How to cite this article: Zayed O, Ibrahim L, Mansour A. Translating Algorithmic Uncertainty into Actionable Medication Advice. Arch Pharm Pract. 2025;16(3):53-70. https://doi.org/10.51847/A5cYIK22eM

to the particular medication choice, the consequence of error, the availability of corrective information, the urgency and reversibility of action, and the professional who retains authority.

The article therefore distinguishes three objects. The algorithmic output is the model-produced result. The uncertainty object describes what is not known about that result, why it is not known, and under which conditions the uncertainty estimate is meaningful. The advice object is a proposed, bounded statement about what action is permissible next and what conditions must accompany that action. These objects should not be collapsed because a well-performing model can enter an unsuitable workflow, a calibrated probability can be irrelevant to the immediate decision, and a clearly worded message can still direct attention towards the wrong action. Clinical usefulness therefore depends on sociotechnical integration and lifecycle safeguards rather than model output alone [3, 4].

The proposed contribution is an uncertainty-translation function positioned between algorithmic output and medication advice. It does not prescribe universal thresholds or claim that uncertainty communication improves outcomes. Instead, it specifies the information that must be connected before an output may be converted into a check, delay, communication, escalation, or qualified recommendation. This framing treats uncertainty as a decision-governance problem as well as a statistical one.

Forms and Sources of Algorithmic Uncertainty

A useful taxonomy begins with aleatoric and epistemic uncertainty. Aleatoric uncertainty concerns irreducible variability or ambiguity in observations and outcomes, whereas epistemic uncertainty concerns incomplete knowledge about the model, parameters, representation, or relevant data. Additional observations may reduce some epistemic uncertainty but cannot necessarily remove aleatoric variability; in practice, the two may be difficult to separate perfectly [5]. For medication advice, the distinction matters because obtaining more information is rational only when that information could meaningfully reduce the uncertainty relevant to the decision.

Uncertainty quantification must also be separated from uncertainty usefulness. A mathematically defined interval, probability distribution, or confidence measure can be technically coherent yet fail to indicate which medication decision is affected or which response is warranted. Clinical use therefore requires uncertainty information to be interpreted against the decision context rather than treated as meaningful merely because it is available [6]. The same numerical value could justify routine confirmation in one setting and immediate escalation in another because urgency, potential harm, reversibility, and available alternatives differ.

The estimate itself is method-dependent. Bayesian approximations, ensembles, evidential approaches, conformal methods, and other techniques rely on different assumptions and may represent different aspects of uncertainty. Their outputs are not automatically equivalent, and apparent confidence can be affected by data quality, model specification, calibration, and evaluation design [7]. The framework therefore avoids ranking methods abstractly; it asks whether a chosen estimator is appropriate for the target task and whether its output has demonstrated interpretable behaviour under relevant conditions.

A further class is distributional and contextual uncertainty. Population, setting, time, measurement practice, prevalence, and workflow can differ from development conditions, making generalization uncertain even when internal performance appears strong [8]. In medication use, contextual uncertainty may arise from missing renal-function trends, incomplete medication histories, altered formulary practices, unusual combinations of diseases, or patients outside the represented population. Detection of shift does not itself explain the correct response, but it should limit the authority granted to the output.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Translating Algorithmic Uncertainty into Actionable Medication Advice.

Table 1. Definitions, components, relationships, and intended uses for the article Translating Algorithmic Uncertainty into Actionable Medication Advice.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Aleatoric uncertainty	Variable or ambiguous outcomes may be mistaken for model ignorance.	Outcome variability, measurement noise, label ambiguity, and stochastic processes.	Evidence-derived distinction: variability that additional model knowledge may not remove [5].	Prevents reflexive demands for more data when variability is not readily reducible.	Task-specific characterization of variability and measurement processes.	False reassurance that variability has been fully quantified.	The category does not prove that a particular estimator isolates aleatoric uncertainty.

Epistemic uncertainty	Missing knowledge may be hidden behind a precise output.	Sparse representation, model specification, parameter uncertainty, and unobserved states.	Evidence-derived distinction: uncertainty may decline when relevant knowledge or data increase [5].	Supports consideration of verification, data acquisition, or deferral.	External and stress testing showing responsiveness to knowledge gaps.	Treating every wide estimate as remediable.	Additional data are useful only when they address the relevant knowledge gap.
Numerical uncertainty representation	A score or interval may be technically valid but operationally silent.	Probability, interval, disagreement, uncertainty score, or confidence label.	Proposed relationship: representation becomes useful only after linkage to a medication decision and response [6].	Separates quantification from action.	Calibration, comprehension, and decision-consequence studies.	False precision or action inferred from an unlabelled number.	Numerical display alone is not medication advice.
Estimator and model dependence	Different methods may produce non-equivalent uncertainty signals.	Modelling family, assumptions, training procedure, calibration method, and test design.	Evidence-derived relationship: uncertainty behaviour depends on method and evaluation conditions [7].	Requires method-specific validation rather than generic claims.	Comparative evaluation under relevant perturbations and failure cases.	Unsupported ranking or interchange of uncertainty methods.	No method is presumed universally superior.
Distributional and contextual uncertainty	Development conditions may not match the patient, site, time, or workflow.	Population, prevalence, measurement practice, setting, temporal change, and workflow.	Evidence-derived relationship: generalization can fail under population or context shift [8].	Limits output authority when intended-use assumptions are violated.	Local external validation, shift testing, and subgroup analysis.	A shift flag may be absent, insensitive, or clinically uninterpretable.	Shift detection does not itself specify the correct medication action.
Decision-relevance translation—**proposed**	Not every uncertainty source can change the immediate medication decision.	Affected decision, consequence of error, urgency, reversibility, and available alternatives.	A proposed gate determines whether uncertainty is material and which response class may be considered.	Prevents unnecessary interruption while preserving escalation for consequential uncertainty.	Prospective human–AI and medication-safety evaluation.	Hidden value judgements or unvalidated thresholds.	The gate is a conceptual construct, not a clinical rule.
Recipient and governance layer—**proposed**	One message cannot serve all users or preserve accountability by itself.	Recipient role, information need, model version, provenance, and override authority.	A proposed one-record/multiple-view architecture preserves common evidence while tailoring presentation.	Supports role-appropriate communication and auditability.	Comprehension, workflow, accountability, and implementation studies.	Tailoring may omit material information or blur responsibility.	Presentation must not transfer professional authority to the system.

Proposed Uncertainty-Translation Framework

The framework begins with an output-registration layer that records the model result together with the uncertainty measure, model version, input recency, missingness, intended population, and use context. For probabilistic outputs, calibration is essential because discrimination alone does not show whether predicted probabilities correspond to observed risk [9]. Calibration nevertheless remains only an input condition: it does not determine whether the prediction is relevant to the current medication choice or whether acting on it produces benefit.

The second layer is decision-relevance translation. It asks: What is uncertain? Which medication decision could change? What is the consequence of a wrong action or an inappropriate delay? Can the uncertainty be reduced by checking data or obtaining an independent source? Who has authority to resolve it? This layer converts uncertainty from a numerical property into a reason for a bounded next step. Communication of uncertainty is particularly important when the appropriate response may be corroboration or a second opinion rather than acceptance of the model output [10].

Figure 1 presents the original conceptual synthesis for uncertainty-translation framework for actionable medication advice, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

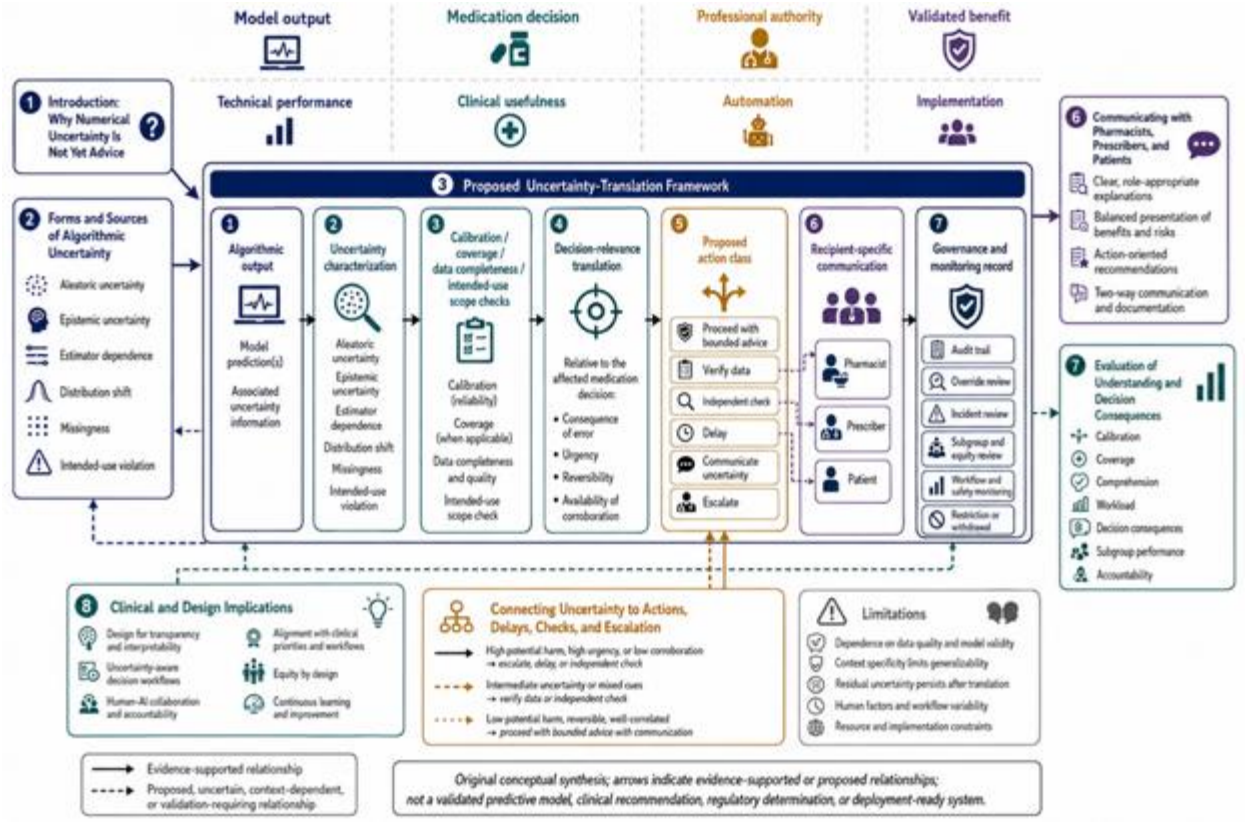

Figure 1. Uncertainty-Translation Framework for Actionable Medication Advice

The third layer is an uncertainty-translation record. Structured model information can make intended use, inputs, performance, limitations, and governance-relevant facts visible to clinical users [11]. The proposed record extends that principle by linking the model fact to the affected medication decision, the selected action class, the rationale for that class, unresolved assumptions, the accountable professional, and the outcome data needed for monitoring. It is intended to make an action traceable, not merely explainable.

The fourth layer provides recipient-specific views of the same record. Pharmacists may need data provenance, interaction mechanisms, alternatives, and override history; prescribers may need the affected treatment choice and required clinical confirmation; patients may need a plain-language statement of what is uncertain, why it matters, who will review it, and what happens next. Explanation needs differ across stakeholders, making a single undifferentiated display unlikely to be sufficient [12]. Tailoring, however, is proposed as an interface principle rather than assumed to improve comprehension.

A governance envelope surrounds all layers. It assigns ownership for model versioning, translation rules, local configuration, overrides, incident review, monitoring, and retirement. No layer authorizes autonomous prescribing or dispensing. The framework's central boundary is that algorithmic uncertainty may inform professional action but does not replace professional authority.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Translating Algorithmic Uncertainty into Actionable Medication Advice.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Translating Algorithmic Uncertainty into Actionable Medication Advice.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Output registration	Preserve the model result and conditions of production.	Inputs, output, uncertainty measure, version, recency,	Supplies the characterization and decision-relevance layers.	Calibration and coverage are assessed for the	A precise-looking result is recorded without valid	The model owner verifies version, inputs, and	No advice translation when provenance or

		missingness, and intended use.		intended context [9].	calibration or scope.	documented limitations.	intended-use scope is unknown.
Uncertainty characterization	Identify the source and meaning of uncertainty.	Aleatoric, epistemic, estimator-dependent, and contextual signals.	Constrains which corrective actions are plausible.	The estimator responds appropriately to known perturbations and knowledge gaps [7].	Different uncertainty types are collapsed into one label.	The technical team documents estimator assumptions; the clinical team reviews relevance.	Characterization is insufficient when the signal has no demonstrated task meaning.
Decision-relevance gate—**proposed**	Determine whether uncertainty could alter the medication decision.	Affected choice, consequence, urgency, reversibility, and alternatives.	Connects characterization to action mapping.	Users can identify the affected decision and reason for uncertainty.	Hidden value judgements or thresholds direct action without validation.	An accountable pharmacist or prescriber confirms materiality.	No automated gate is treated as clinically authoritative.
Action mapping—**proposed**	Select a bounded response class.	Proceed with qualification, verify data, independent check, delay, communicate, or escalate.	Feeds communication and governance records.	The selected action is proportionate and prompts corroboration when needed [10].	Over-escalation creates burden; under-escalation permits unsafe reliance.	Clinical governance defines local rules and override authority.	Action classes remain hypotheses until task-specific evaluation is completed.
Provenance record—**proposed extension**	Make the rationale and conditions auditable.	Model facts, decision rationale, unresolved assumptions, action, override, and outcome.	Supports recipient views, monitoring, and incident review.	Intended use and limitations are visible and retrievable [11].	Documentation is incomplete, stale, or disconnected from the decision.	The organization maintains version control and audit access.	Disclosure does not establish comprehension or effectiveness.
Recipient-specific communication—**proposed**	Present material information at an appropriate level.	A shared record rendered for a pharmacist, prescriber, or patient.	Receives action and rationale and returns questions or corrections.	Essential uncertainty and next steps are understood without obscuring authority [12].	Simplification omits material risk or tailoring reinforces assumptions about users.	Designers and clinical teams conduct user-centred testing.	No communication format is presumed universally suitable.
Governance and monitoring—**proposed**	Control change, accountability, incidents, and withdrawal.	Overrides, errors, subgroup signals, workflow effects, and outcome observations.	Encloses all layers and feeds revision or retirement.	Rules, ownership, review triggers, and stopping conditions are explicit.	Responsibility is diffused or failures persist without corrective action.	Named organizational owners review performance and consequences.	Implementation remains conditional on local validation and continuing oversight.

Connecting Uncertainty to Actions, Delays, Checks, and Escalation

The framework proposes six action classes: proceed with bounded advice, verify data, check an independent source, delay pending information, communicate uncertainty, and escalate to an accountable professional. These are not arranged as a universal ladder. Their appropriateness depends on whether the uncertainty is reducible, whether the decision is time-sensitive, whether action is reversible, and whether error could cause meaningful medication harm.

Proceed with bounded advice is reserved for cases in which the output remains within intended use, the uncertainty is characterized, and the remaining uncertainty is unlikely to change the immediate decision. The advice should state its conditions rather than convert residual uncertainty into false certainty. Defer or abstain is different: it applies when the system lacks sufficient support to issue an interpretable recommendation. Selective-prediction research demonstrates that an algorithm can be designed to withhold output in cases it cannot support adequately, providing a basis for a proposed deferral branch [13]. The threshold and consequences of abstention must be established for each medication task rather than transferred from another clinical domain.

Verify data is appropriate when missing, stale, contradictory, or implausible inputs could be corrected, such as an outdated renal-function measurement or incomplete medication list. Check an independent source is appropriate when uncertainty concerns evidence, a drug interaction, dose constraints, or a model inference that can be corroborated without delaying urgent care. Delay pending information is a distinct action because time can be clinically meaningful: delay may reduce risk when a nonurgent decision depends on confirmatory information, but it may increase risk when treatment is time-critical.

The framework also constrains interruption. Workload, work complexity, and repeated alerts can reduce attention to decision-support messages [14]. Consequently, not every uncertainty signal should generate an interruptive alert. A proposed interruption rule should consider immediacy,

potential harm, reversibility, and whether the recipient can act on the message at that point in the workflow. Low-consequence uncertainty may be documented for later review, whereas unresolved high-consequence uncertainty may justify direct escalation.

Medication-alert systems illustrate both the opportunity and the risk. Artificial intelligence may help prioritize or suppress medication alerts, but heterogeneity in the evidence and limited evaluation of clinical relevance, workflow effects, and safety prevent a general assumption of benefit [15]. The proposed translation framework therefore separates prioritization from action: an alert may prompt data verification, targeted evidence checking, communication, delay, or escalation rather than a binary accept-or-dismiss response.

Finally, action uptake is not equivalent to improved care. A pharmacist opening an alert, a prescriber changing an order, or a team escalating a case are process events. Evidence concerning drug–drug interaction alerts indicates that changes in clinician behaviour do not necessarily establish patient-important benefit [16]. Evaluation must therefore trace the pathway from uncertainty display to understanding, action, workflow consequence, medication-safety outcome, and patient outcome. Until that chain is demonstrated, the proposed action classes remain conditional design hypotheses rather than validated recommendations.

Communicating with Pharmacists, Prescribers, and Patients

Uncertainty communication should support a specific medication decision rather than display a detached probability or generic warning. User-centred design evidence indicates that AI information is more appropriately embedded within clinical context, interpretive needs, and professional conversation [17]. The proposed framework therefore communicates four elements: what is uncertain, why the uncertainty matters, which action is proposed, and who remains responsible for resolving or accepting it.

For patients, communication should avoid presenting the algorithm as an independent authority. Resistance to medical AI may arise when patients believe that a system cannot account for their individual circumstances [18]. Patient-facing advice should therefore explain how personal factors were considered, identify information that may be incomplete, and describe the next human action. This does not require disclosing every technical detail; it requires preserving the distinction between simplification and concealment.

Patients may also need information about professional oversight, data use, accountability, and the role assigned to AI [19]. The framework consequently proposes a concise patient view containing the affected medication issue, the uncertainty in plain language, any immediate precaution, and the pharmacist or prescriber responsible for review. Communication should not imply that uncertainty necessarily indicates danger or that professional review guarantees a favourable outcome.

For pharmacists and prescribers, the professional view should include the source of uncertainty, relevant inputs, intended-use limitations, plausible consequences, recommended check, and override route. Clinicians have identified concerns involving trust, responsibility, professional roles, and effects on patient relationships when AI is introduced into care [20]. Communication must therefore preserve decision authority explicitly: the system may recommend verification, delay, or escalation, but the accountable professional determines the medication action.

Evaluation of Understanding and Decision Consequences

Evaluation should begin with human–AI interaction rather than end with model accuracy. Studies of clinical collaboration show that combined performance depends on how users integrate algorithmic information, while incorrect advice can influence decisions and behavioural observation may reveal responses not captured by self-report [21–23]. Testing should therefore expose pharmacists and prescribers to correct, incorrect, overconfident, uncertain, and abstaining outputs. Relevant outcomes include recognition of uncertainty, identification of the affected decision, appropriate checking, resistance to unsafe advice, and correct escalation.

Understanding should be distinguished from satisfaction, trust, and action. A user may report that a display is clear while misunderstanding its limitations or following an inappropriate recommendation. Evaluation must also separate decision effects, care-delivery effects, and patient outcomes because evidence across these levels is heterogeneous [24]. Proposed criteria include calibration and coverage, appropriateness of abstention, comprehension, action concordance, workload, interruption burden, medication-error consequences, subgroup performance, and accountability. Thresholds should be prespecified for each task and setting rather than treated as universal properties of the framework.

Clinical and Design Implications

Evaluation should progress from formative usability and simulation to prospective workflow testing and, where justified, comparative clinical assessment. Early-stage and trial-focused AI guidance emphasizes explicit reporting of intended use, algorithm version, inputs, human interaction, workflow integration, errors, and analysis procedures [25–27]. For uncertainty translation, these requirements should extend to action rules, overrides, delays, escalations, communication failures, and changes made after incidents.

Governance must assign named responsibility for translation thresholds, local configuration, monitoring, subgroup review, version changes, and withdrawal. A healthcare AI governance model similarly emphasizes transparency, safety, accountability, and continuing organizational oversight [28]. Translation rules should therefore be treated as controlled clinical-system components rather than informal interface wording.

Figure 2 presents the original conceptual synthesis for from algorithmic uncertainty to checks, delay, communication, and escalation, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

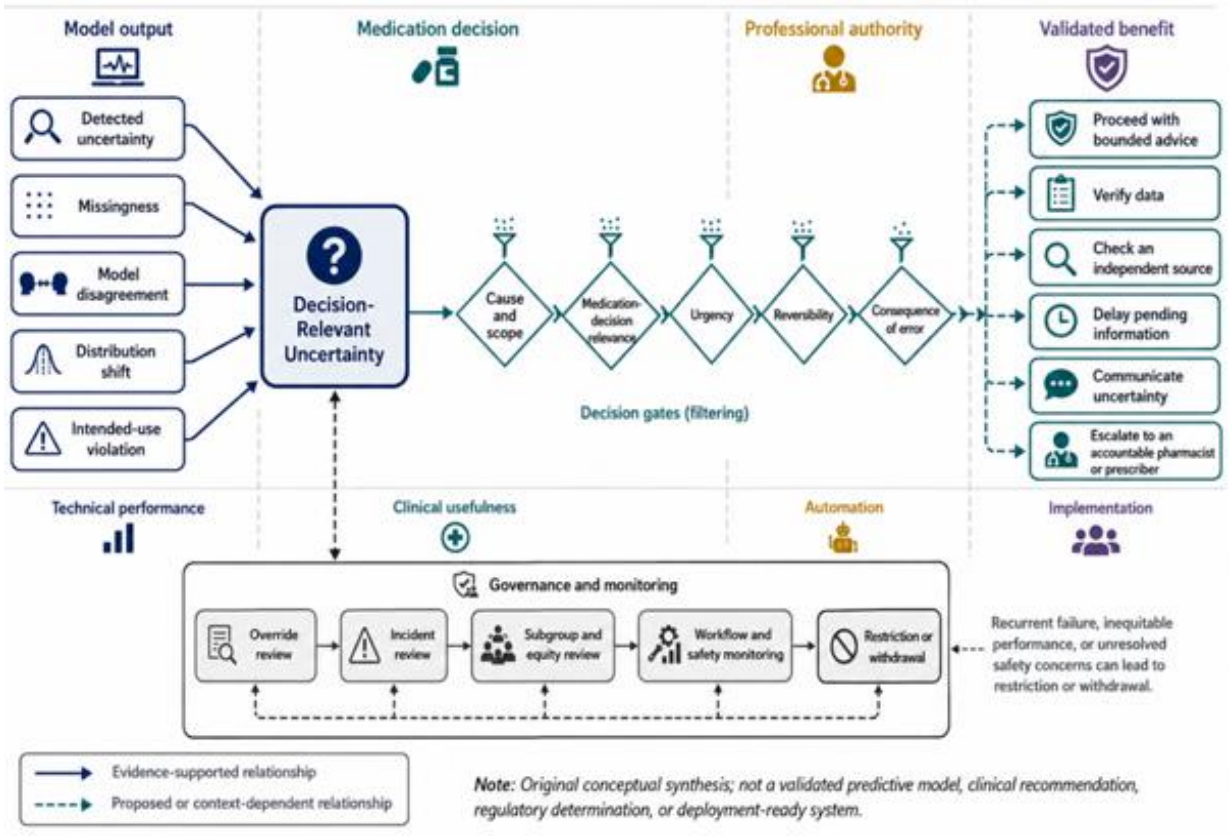

Figure 2. From Algorithmic Uncertainty to Checks, Delay, Communication, and Escalation

Limitations

The framework cannot correct an invalid target, biased proxy, or structurally inequitable objective. Empirical evidence shows that an algorithm may reproduce racial inequity when the variable used as a proxy does not represent the intended clinical need [29]. Uncertainty presentation cannot repair that defect.

Aggregate performance may also conceal systematic underperformance in underserved groups [30]. Validation must therefore examine subgroup calibration, abstention, action recommendations, communication effects, and downstream consequences.

Finally, much of the supporting evidence comes from clinical domains outside pharmacy. Methodological weaknesses, limited external validation, and dataset problems can undermine apparent readiness [31]. Medication-specific prospective testing is therefore required before any proposed translation rule is used to direct care.

Conclusion

Numerical uncertainty is not actionable medication advice until it is interpreted against the medication decision, the source and reducibility of uncertainty, the urgency and reversibility of action, the consequences of error, and the authority of the responsible professional. The proposed Uncertainty-Translation Framework inserts this interpretive and governance function between model output and advice.

Its principal contribution is a bounded architecture connecting uncertainty characterization, calibration and scope checks, decision relevance, proportionate action classes, recipient-specific communication, and auditable oversight. It also defines checks, delay, communication, escalation, restriction, and withdrawal as possible responses rather than automatic consequences of a score.

The framework remains a conceptual proposal. Its value must be tested through technical validation, human-factors evaluation, medication-workflow studies, patient communication research, equity analysis, prospective safety assessment, and continuing governance. It should not be interpreted as transferring clinical authority to an algorithm or as establishing that uncertainty displays improve decisions or outcomes. Its intended role is narrower: to make explicit the reasoning, evidence, responsibility, and validation required before algorithmic uncertainty can be converted into medication advice.

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

REFERENCES

1. 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
2. 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
3. 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(1):17. doi:10.1038/s41746-020-0221-y
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. Hüllermeier E, Waegeman W. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. Mach Learn. 2021;110(3):457-506. doi:10.1007/s10994-021-05946-3
6. Begoli E, Bhattacharya T, Kusnezov D. The need for uncertainty quantification in machine-assisted medical decision making. Nat Mach Intell. 2019;1(1):20-23. doi:10.1038/s42256-018-0004-1
7. Abdar M, Pourpanah F, Hussain S, Rezazadegan D, Liu L, Ghavamzadeh M, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. Inf Fusion. 2021;76:243-97. doi:10.1016/j.inffus.2021.05.008
8. Goetz L, Seedat N, Vandersluis R, van der Schaar M. Generalization—a key challenge for responsible AI in patient-facing clinical applications. NPJ Digit Med. 2024;7(1):126. doi:10.1038/s41746-024-01127-3
9. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW; Topic Group 'Evaluating diagnostic tests and prediction models' of the STRATOS initiative. Calibration: The Achilles heel of predictive analytics. BMC Med. 2019;17(1):230. doi:10.1186/s12916-019-1466-7
10. Kompa B, Snoek J, Beam AL. Second opinion needed: Communicating uncertainty in medical machine learning. NPJ Digit Med. 2021;4(1):4. doi:10.1038/s41746-020-00367-3
11. Sendak MP, Gao M, Brajer N, Balu S. Presenting machine learning model information to clinical end users with model facts labels. NPJ Digit Med. 2020;3(1):41. doi:10.1038/s41746-020-0253-3
12. 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
13. Shashikumar SP, Wardi G, Malhotra A, Nemati S. Artificial intelligence sepsis prediction algorithm learns to say "I don't know". NPJ Digit Med. 2021;4(1):134. doi:10.1038/s41746-021-00504-6
14. Ancker JS, Edwards A, Nosal S, Hauser D, Mauer E, Kaushal R; HITEC Investigators. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. BMC Med Inform Decis Mak. 2017;17(1):36. doi:10.1186/s12911-017-0430-8
15. 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
16. Holbrook AM, Matos Silva J, Faruque JAY, Deng J, Schneider T, Jaffer A. Effect of electronic drug-drug interaction alerts on patient and clinician outcomes: A systematic review. J Am Med Inform Assoc. 2025;32(10):1617-28. doi:10.1093/jamia/ocaf139
17. Staes CJ, Beck AC, Chalkidis G, Scheese CH, Taft T, Guo JW, et al. Design of an interface to communicate artificial intelligence-based prognosis for patients with advanced solid tumors: A user-centered approach. J Am Med Inform Assoc. 2024;31(1):174-87. doi:10.1093/jamia/ocad201
18. Longoni C, Bonezzi A, Morewedge CK. Resistance to medical artificial intelligence. J Consum Res. 2019;46(4):629-50. doi:10.1093/jcr/ucz013
19. Ongena YP, Haan M, Yakar D, Kwee TC. Patients' views on the implementation of artificial intelligence in radiology: Development and validation of a standardized questionnaire. Eur Radiol. 2020;30(2):1033-40. doi:10.1007/s00330-019-06486-0
20. Blease C, Kaptchuk TJ, Bernstein MH, Mandl KD, Halamka JD, DesRoches CM. Artificial intelligence and the future of primary care: Exploratory qualitative study of UK general practitioners' views. J Med Internet Res. 2019;21(3). doi:10.2196/12802
21. Tschandl P, Rinner C, Apalla Z, Argenziano G, Codella N, Halpern A, et al. Human-computer collaboration for skin cancer recognition. Nat Med. 2020;26(8):1229-34. doi:10.1038/s41591-020-0942-0
22. Gaube S, Suresh H, Raue M, Merritt A, Berkowitz SJ, Lermer E, et al. Do as AI say: Susceptibility in deployment of clinical decision-aids. NPJ Digit Med. 2021;4(1):31. doi:10.1038/s41746-021-00385-9
23. Nagendran M, Festor P, Komorowski M, Gordon AC, Faisal AA. Eye tracking insights into physician behaviour with safe and unsafe explainable AI recommendations. NPJ Digit Med. 2024;7(1):202. doi:10.1038/s41746-024-01200-x
24. Susanto AP, Lyell D, Widyantoro B, Berkovsky S, Magrabi F. Effects of machine learning-based clinical decision support systems on decision-making, care delivery, and patient outcomes: A scoping review. J Am Med Inform Assoc. 2023;30(12):2050-63. doi:10.1093/jamia/ocad180
25. 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
26. 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
27. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ; SPIRIT-AI and CONSORT-AI Working Group. 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
28. Reddy S, Allan S, Coghlan S, Cooper P. A governance model for the application of AI in health care. J Am Med Inform Assoc. 2020;27(3):491-7. doi:10.1093/jamia/ocz192
29. 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
30. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. 2021;27(12):2176-82. doi:10.1038/s41591-021-01595-0
31. Varoquaux G, Cheplygina V. Machine learning for medical imaging: Methodological failures and recommendations for the future. NPJ Digit Med. 2022;5(1):48. doi:10.1038/s41746-022-00592-y