Original Article

Communicating Calibrated Uncertainty at the Point of Dispensing

Lina Hassan[1]*, Omar Khalaf[2], Reem Jaber[1], Rania Mostafa[3]

[1]Department of Uncertainty Communication and Dispensing, Faculty of Pharmacy, University of Jordan, Amman, Jordan. [2]Department of Calibrated AI in Pharmacy Practice, Faculty of Pharmacy, Jordan University of Science and Technology, Irbid, Jordan. [3]Department of Point-of-Dispensing Decision Support, Faculty of Pharmacy, Al-Balqa Applied University, Salt, Jordan.

Abstract

Digital and artificial-intelligence-supported medication systems increasingly generate risk estimates, alerts, classifications, and recommendations that may appear more certain than their evidential basis permits. At the point of dispensing, pharmacists must interpret these outputs within brief encounters while considering incomplete data, medication-related consequences, patient preferences, and the limits of model applicability. Existing approaches commonly separate technical calibration, uncertainty display, clinical consequence, pharmacist judgement, and patient communication, leaving no coherent mechanism for translating uncertainty into a bounded medication decision. This article proposes a Point-of-Dispensing Uncertainty-Communication Framework that connects five functions: characterizing the source of uncertainty; appraising calibration and applicability; assessing the clinical consequence of error; constructing a decision-relevant communication package; and adapting that package through pharmacist mediation. The framework distinguishes confidence from calibration, uncertainty from recommendation, and model output from professional authority. It proposes five possible action boundaries: routine dispensing with counselling, enhanced uncertainty counselling, additional verification, temporary deferral, and escalation. The contribution is conceptual and non-empirical. Its proposed relationships require validation across technical performance, patient comprehension, appropriate human–AI reliance, medication-decision quality, workflow burden, safety, equity, and governance. Communication should preserve the source, direction, and practical meaning of uncertainty rather than merely simplify presentation. The framework does not establish universal thresholds, clinical effectiveness, regulatory acceptability, or deployment readiness. Its original contribution is to organize calibrated uncertainty as a pharmacist-mediated, consequence-sensitive communication and action problem rather than as a model metric or display feature alone.

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

Introduction

Medication-related artificial intelligence and clinical decision support may produce probabilities, risk categories, alerts, or recommended actions, yet these outputs rarely communicate all relevant uncertainty in a form that supports responsible human judgement. Clinical machine-learning outputs require usable uncertainty information because an apparently precise prediction may conceal limited evidence, model instability, missing inputs, or poor applicability to the current patient [1]. At dispensing, this problem is intensified by limited encounter time, competing workflow demands, incomplete access to clinical context, and the need to communicate information without delaying necessary treatment or creating unwarranted reassurance.

Machine-learning systems can prioritize prescriptions for pharmacist review, but medication-alert research remains heterogeneous and requires stronger external validation and implementation evidence [2, 3]. A prescription ranked as low risk may still involve a severe consequence if the system is wrong, whereas a high uncertainty estimate may reflect harmless variability rather than an immediate safety threat. Consequently, neither an alert score nor a confidence value should independently determine whether a medicine is dispensed, discussed, verified, deferred, or escalated.

Uncertainty communication is also relational. Shared decision-making models show that communicating options

Address for correspondence: Lina Hassan, Department of Uncertainty Communication and Dispensing, Faculty of Pharmacy, University of Jordan, Amman, Jordan.
lina.hassan@ju.edu.jo
Received: 05 February 2026; **Accepted:** 26 April 2026

How to cite this article: Hassan L, Khalaf O, Jaber R, Mostafa R. Communicating Calibrated Uncertainty at the Point of Dispensing. Arch Pharm Pract. 2026;17(2):30-8. https://doi.org/10.51847/VZ1sqYHTmD

requires recognition of choice, discussion of available evidence, and integration of patient preferences [4]. At the point of dispensing, however, uncertainty may concern several objects simultaneously: the reliability of a digital output, the completeness of patient data, the strength of medication evidence, the transferability of population findings, or the patient's unresolved preferences. These sources must not be compressed into a single unexplained qualifier such as "low confidence."

This article therefore proposes an original framework for moving from uncertain information to a bounded dispensing action. It treats communication as a structured professional function connecting technical evidence, clinical consequence, patient context, and pharmacist accountability. The framework is not a clinical guideline or validated intervention. It instead defines constructs, proposed relationships, evaluation requirements, and decision boundaries that can be tested in future pharmacy settings.

Calibration, Confidence, and Clinical Consequence

Confidence and calibration describe different properties. Confidence is the apparent strength assigned to an individual output, whereas calibration concerns agreement between predicted probabilities and observed event frequencies within a defined task, population, setting, and period [5]. A system may discriminate effectively between higher- and lower-risk prescriptions while assigning probabilities that are systematically too high or too low. Conversely, acceptable aggregate calibration does not prove that an estimate is applicable to a particular medicine, patient subgroup, clinical environment, or updated workflow.

Uncertainty quantification is particularly important when an incorrect output could produce a consequential medical decision [6]. Evidence from clinical machine learning also indicates that uncertainty information may help identify cases that warrant selective expert review [7]. Such findings support the principle of uncertainty-triggered review, but they do not establish a universal pharmacy threshold. The relationship between uncertainty and action depends on what is uncertain, whether additional information could resolve it, and what may happen if the pharmacist proceeds incorrectly.

Uncertainty can also be operationalized as a signal for targeted human review rather than displayed as a passive technical value [8]. At dispensing, this requires a distinction between the probability of error and the consequence of error. The article therefore proposes a consequence profile comprising severity, reversibility, detectability, urgency, duration of exposure, availability of safer alternatives, and the patient's vulnerability. These dimensions are not a validated score. They are a structured prompt for determining whether an uncertainty is immaterial, communicable, verifiable, or beyond the safe decision boundary.

The proposed relationship is consequence-sensitive rather than confidence-deterministic. High confidence should not suppress verification when calibration evidence is weak or the potential harm is severe. Low confidence should not automatically prevent dispensing when the uncertainty concerns a minor, reversible, or monitorable outcome. Calibration informs how much reliance a probability may warrant; clinical consequence informs how cautious the resulting action should be. The patient-facing communication must then explain not only that uncertainty exists but what it means for the immediate medication decision.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Communicating Calibrated Uncertainty at the Point of Dispensing.

Table 1. Definitions, components, relationships, and intended uses for the article Communicating Calibrated Uncertainty at the Point of Dispensing.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Confidence	A strong-looking output may be mistaken for correctness	Output probability, score, category, or verbal qualifier	Confidence is treated as an instance-level signal, not evidence of reliability	Prevents apparent precision from becoming automatic authority	Association between confidence and correctness in the intended setting	Overconfidence, false reassurance, or inappropriate automation	Confidence alone does not authorize a medication decision [1].
Calibration	Predicted probabilities may not correspond to observed frequencies	Validation population, task, time, site, and outcome definition	Calibration evidence qualifies how probabilities may be interpreted	Supports evidence-bounded probability communication	Local, temporal, and subgroup calibration assessment	Misleading probabilities under drift or population mismatch	Calibration does not establish usefulness, safety, or individual correctness [5].
Uncertainty source	Different limitations may require different responses	Variability, missing data, limited evidence, model scope, or patient preference	**Proposed:** classify uncertainty before selecting language or action	Avoids compressing heterogeneous limitations into one score	Content validation and inter-rater reliability of the proposed taxonomy	Wrong remedy applied to the wrong uncertainty	The taxonomy is proposed and requires empirical validation.

Selective review	Uniform automation may conceal difficult or atypical cases	Uncertainty estimate, case characteristics, data completeness	Uncertainty may trigger pharmacist review or information seeking	Directs professional attention to unresolved cases	Prospective evaluation of accepted, reviewed, deferred, and escalated cases	Excessive deferral or failure to identify consequential cases	Evidence from other clinical tasks does not validate dispensing thresholds [7].
Clinical consequence	Probability alone does not represent the seriousness of being wrong	Severity, reversibility, detectability, urgency, and alternatives	**Proposed:** communication and action become more cautious as consequence increases	Links uncertainty to medication-safety relevance	Prospective testing across medication and patient contexts	Minor uncertainties over-escalated or severe uncertainties normalized	The consequence profile is not a validated risk score [6].
Human review	A displayed uncertainty may not change behaviour appropriately	Pharmacist expertise, access to evidence, workflow, and escalation routes	Human review interrogates uncertainty rather than merely acknowledging it	Preserves professional challenge and accountability	Observation of review quality, overrides, and downstream actions	Automation bias, underreliance, or symbolic oversight	Human presence is not sufficient without authority and usable information [8].
Communication package	Technical outputs may not be interpretable to patients	Source, magnitude, practical implication, and next action	**Proposed:** translate uncertainty into a concise decision-relevant package	Connects evidence limitations to patient action	Comprehension, recall, reliance, and decision-quality studies	Distortion through oversimplification or excessive information	Adaptation must preserve the underlying meaning of uncertainty.
Action boundary	Systems may imply action without defining when not to proceed	Calibration, applicability, consequence, missing information, patient context	**Proposed:** select routine counselling, enhanced counselling, verification, deferral, or escalation	Makes non-action and escalation legitimate outcomes	Prospective safety, workflow, and equity evaluation	Unsafe dispensing, unnecessary delay, or unavailable escalation	No category constitutes a validated clinical rule.

Proposed Point-of-Dispensing Communication Framework

Clinical decision support operates within sociotechnical systems shaped by data quality, usability, workflow, alert burden, and professional roles [9]. The proposed framework therefore begins before information reaches the patient. It asks whether the uncertainty itself has been adequately characterized, whether the confidence estimate is calibrated and applicable, and whether the output addresses the clinical question that the pharmacist must resolve. Technical performance cannot by itself establish clinical usefulness, generalizability, or safe integration [10].

The framework comprises five linked functions. First, uncertainty characterization identifies whether the limitation arises from inherent variability, insufficient knowledge, data quality, applicability or distribution shift, conflicting clinical evidence, or unresolved patient preference. Second, calibration and applicability appraisal examines the population, medication, task, site, version, and period for which the estimate is supported. Third, clinical-consequence appraisal considers the seriousness, reversibility, detectability, and urgency of an incorrect interpretation. Fourth, communication-package construction converts the uncertainty into an essential message, its reason, its practical implication, and the required next action. Fifth, pharmacist mediation adapts language and format, checks understanding, verifies information, and determines whether the encounter can proceed.

Explanation must be matched to the needs and responsibilities of the intended stakeholder; technical interpretability is not equivalent to meaningful patient communication [11]. The pharmacist therefore does not merely repeat a model score. The pharmacist must distinguish uncertainty from recommendation, clarify whether more information could resolve the uncertainty, and communicate what remains professionally or clinically undecided.

The final action boundary contains five proposed outcomes: routine dispensing with ordinary counselling; enhanced uncertainty counselling; additional verification; temporary deferral; and escalation. Movement between these outcomes is conditional, not linear. New patient information may require reassessment of calibration, applicability, or consequence. Early clinical evaluation must examine human interaction, workflow, errors, and operational context rather than reporting model accuracy alone [12].

Figure 1 presents the original conceptual synthesis for point-of-dispensing uncertainty-communication framework, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

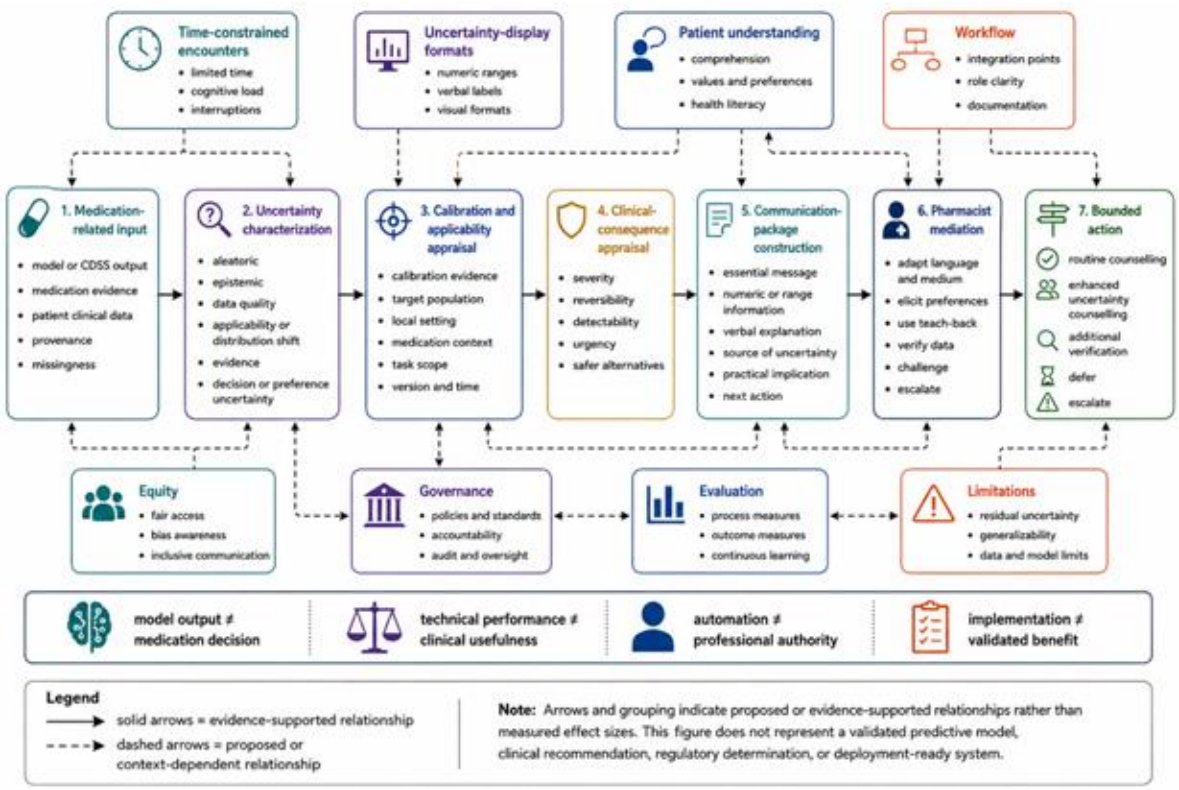

Figure 1. Point-of-Dispensing Uncertainty-Communication Framework

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Communicating Calibrated Uncertainty at the Point of Dispensing.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Communicating Calibrated Uncertainty at the Point of Dispensing.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Input and provenance	Establish what information generated the output	Patient data, medication data, evidence source, model version, and missingness enter the framework	Determines whether later calibration and explanation are meaningful	Completeness, traceability, timeliness, and version identification	Unknown source, stale data, conflicting records, or hidden missingness	Organization ensures provenance; pharmacist verifies decision-relevant gaps	No active use when material inputs cannot be traced.
Uncertainty characterization	Identify what is uncertain and whether it is reducible	Technical and clinical limitations are classified before communication	Directs calibration appraisal, verification, and display choice	Reproducible classification and recognition of mixed uncertainty	A single score obscures several limitations	System designers expose uncertainty source; pharmacist checks clinical relevance	Proposed taxonomy requires content and construct validation.
Calibration and applicability	Determine whether confidence is interpretable in the current context	Validation evidence is compared with the current patient, task, medication, and site	Modifies the weight given to the output and the need for review	Local, temporal, and subgroup calibration [5]	Miscalibration or distribution shift remains invisible	Developers validate; organizations monitor; pharmacists challenge obvious mismatch	Technical validation does not establish clinical usefulness [10].
Consequence appraisal	Determine the significance of an incorrect interpretation	Uncertainty is combined with severity, reversibility, urgency, and alternatives	Influences communication depth and action boundary	Consistency, clinical face validity, and prospective safety assessment	Severe consequences are normalized because confidence appears high	Pharmacist retains authority to verify, defer, or escalate	The proposed consequence profile is not a clinical score.

Communication construction	Translate uncertainty into decision-relevant content	Source, meaning, practical implication, and next action are selected	Connects technical appraisal to pharmacist–patient interaction	Objective comprehension and preservation of meaning	Simplification changes apparent certainty or hides limitations	Pharmacist adapts presentation; organization supplies accessible formats	No readiness without testing comprehension and distortion risk [11].
Pharmacist mediation	Interpret, personalize, verify, and challenge	Patient questions and contextual information return to earlier stages	Creates feedback loops rather than one-way delivery	Appropriate reliance, adaptation fidelity, verification, and escalation behaviour	Automation bias, symbolic review, or unsupported reassurance	Pharmacist exercises professional judgement; organization provides time and authority	Human involvement alone does not prove meaningful oversight.
Action boundary	Select whether and how the encounter proceeds	Appraised uncertainty and patient understanding inform action	Produces routine counselling, enhanced counselling, verification, deferral, or escalation	Decision quality, delay, safety, workflow burden, and equity	Unsafe dispensing, excessive deferral, or unavailable escalation	Local governance defines operational routes and accountability	Prospective clinical and workflow evaluation is required [12].
Monitoring and governance	Detect drift, incidents, and unequal effects	Outcomes and feedback return to model, workflow, and policy owners	Supports reassessment of every preceding layer	Incident review, subgroup monitoring, version control, and corrective action	Static deployment despite changing data or practice	Multidisciplinary governance owns surveillance and response	Implementation is not equivalent to validated benefit [9].

Displaying and Interpreting Different Uncertainty Types

Uncertainty should not be displayed as though every limitation had the same origin or implication. Communication format must be selected in relation to the audience, decision, reference class, and possibility of misunderstanding [13]. A probability may be appropriate when calibration is established and the patient can interpret the denominator. A range may better represent imprecision. A verbal qualifier may be necessary during a brief encounter, but it can conceal wide variation in meaning unless accompanied by an explanation or numerical anchor.

Numeric and verbal expressions of uncertainty are not interchangeable; they can produce different interpretations of magnitude, confidence, and credibility [14]. The proposed framework therefore separates the source of uncertainty from its display format. Aleatoric uncertainty concerns variability that may persist even with additional information. Epistemic uncertainty concerns limited knowledge that might be reduced by new evidence. Data-quality uncertainty arises from missing, delayed, conflicting, or unreliable information. Applicability uncertainty concerns whether a model or evidence base transfers to the present patient, medicine, setting, or time. Evidence uncertainty concerns limitations or disagreement in the clinical literature. Decision or preference uncertainty concerns unresolved trade-offs rather than a defect in a prediction.

The communication objective differs across these classes. Irreducible variability may require presentation of a range and monitoring plan. Missing data may require verification rather than further explanation. Applicability uncertainty may require an explicit statement that the evidence may not transfer. Conflicting evidence may require discussion of alternatives and trade-offs. Preference uncertainty requires elicitation of what matters to the patient rather than an attempt to calculate away the decision.

Communicating uncertainty does not necessarily destroy trust, although effects vary with the form and context of disclosure [15]. Trust preservation should not, however, become the primary design goal. A communication strategy that maintains confidence by suppressing limitations may increase inappropriate reliance. The more defensible objective is calibrated reliance: the patient and pharmacist should neither overvalue nor dismiss the information relative to its evidential support.

Visual aids may improve risk literacy when denominators, reference classes, comparisons, and visual scales are designed carefully [16]. At dispensing, visual or written formats may be useful for complex probabilities, but accessibility, time, language, and emotional burden must be considered. No format should be presumed universally superior. The same underlying uncertainty may require plain language, absolute frequencies, ranges, visual comparisons, staged explanation, or written reinforcement. These adaptations must preserve provenance, direction, magnitude, and practical implication.

The proposed taxonomy and display mappings are conceptual rather than validated. Their evaluation must determine whether patients understand what is uncertain, why it matters, what action is appropriate, and when additional professional advice is required. A concise display that obscures the decision boundary is not necessarily safer than a more detailed explanation, while excessive disclosure may overwhelm the encounter. The appropriate communication is therefore the least burdensome format that preserves decision-relevant meaning and supports safe action.

Pharmacist Mediation and Patient-Specific Adaptation

Uncertainty communication at dispensing cannot be reduced to selecting a numerical or verbal format. Patients differ in medication literacy, numeracy, language, prior knowledge, communication preferences, emotional state, sensory access, and willingness to participate. Medication literacy and desire for involvement are distinct; limited literacy should not be interpreted as limited interest in medication decisions [17]. Adaptation should therefore begin with an invitation to participate rather than an assumption about how much information a patient wants or can understand.

Structured pharmacist–patient conversations can support shared medication decision-making while remaining responsive to the encounter [18]. Within the proposed framework, the pharmacist first states the immediate medication decision, then explains what is known, what remains uncertain, why that uncertainty exists, and whether additional information could change the decision. The patient should also understand the practical consequence: whether treatment can proceed, whether monitoring is required, or whether verification or escalation is necessary.

Patient-centred pharmacy counselling is an interactive process rather than one-way transmission of information [19]. The pharmacist may adapt vocabulary, pace, sequence, medium, and level of detail, but should not alter the source, direction, or apparent magnitude of uncertainty. For example, replacing a poorly supported estimate with a reassuring phrase may make the explanation simpler while making it less faithful. The framework therefore proposes an adaptation-integrity rule: presentation may change, but provenance, evidential limitations, practical implications, and unresolved decision boundaries must remain visible.

Medication decisions may also involve family members, caregivers, prescribers, and other professionals whose knowledge is distributed across settings [20]. Pharmacist mediation may reveal missing information, conflicting goals, or patient concerns that require the framework to return to uncertainty characterization or consequence appraisal. Mediation is therefore represented as a feedback function rather than the final delivery of a fixed message.

Understanding should be checked through dialogue, patient questions, teach-back, or an equivalent method suited to the encounter. A correct repetition of words does not necessarily demonstrate understanding of the medication decision. The relevant test is whether the patient can explain what remains uncertain, what action is being taken, and when further advice is required. Pharmacist mediation may support this process, but it cannot correct invalid evidence, inaccessible provenance, severe distribution shift, or a prescription that exceeds the safe dispensing boundary.

Evaluation of Understanding and Medication Decisions

Evaluation must distinguish subjective trust from appropriate reliance. Clinical users can be influenced by incorrect artificial-intelligence advice, indicating that confidence in a system may persist even when its recommendation should be challenged [21]. A successful communication intervention should therefore reduce both overreliance on poorly supported outputs and underreliance on well-supported information.

Human–AI collaboration is not automatically superior to unaided professional judgement. Assistance can change decisions, improve selected classifications, or introduce new errors depending on the task, interface, and user response [22, 23]. Pharmacy evaluation should consequently examine what pharmacists do when an output conflicts with clinical information, whether uncertainty prompts verification, and whether escalation remains feasible under time pressure.

Medication-decision evaluation should include objective knowledge, understanding of probabilities or ranges, decisional conflict, informed choice, participation, and concordance between expressed values and the selected option [24]. Satisfaction alone is insufficient because a reassuring but misleading explanation may be rated favourably. Recall should also be assessed when the patient must monitor symptoms, follow conditional instructions, or seek additional care.

The proposed validation model therefore separates technical calibration, comprehension, appropriate reliance, medication-decision quality, workflow, medication safety, equity, and governance. These domains should not be collapsed into a single readiness score. A system may be well calibrated but poorly understood, understandable but inequitable, or efficient while producing inappropriate reliance. Prospective evaluation should compare routine communication with the proposed approach and should examine accepted, verified, deferred, and escalated cases separately. No individual metric can establish that the complete framework is safe or clinically effective.

Workflow and Equity Implications

Real-world integration requires more than adding an uncertainty display to an existing interface. Implementation depends on workflow redesign, stakeholder participation, monitoring, operational ownership, and clear escalation routes [25]. At dispensing, an uncertainty communication process may add verification work, prolong encounters, interrupt queues, or shift responsibility without providing the time or authority required for meaningful review. Evaluation must therefore include delays, duplicated work, workarounds, interruptions, documentation burden, and the consequences of unavailable escalation.

Equity failures may originate before communication begins. An apparently neutral output can reproduce inequity when its target or proxy reflects unequal access, utilization, or documentation [26]. Communication cannot repair a biased target merely by explaining it clearly. The framework therefore requires subgroup assessment of calibration, exposure to alerts, verification, deferral, escalation, and downstream medication decisions.

Communication for patients with limited health literacy is shaped by time, recognition, professional skill, information tailoring, and emotional context [27]. Equity should not be framed as simplification for selected patients while others receive the complete explanation. The aim is equivalent access to decision-relevant meaning through accessible language, interpretation support, suitable visual or written formats, and accommodations for sensory or cognitive needs.

Figure 2 presents the original conceptual synthesis for pharmacist-mediated adaptation of uncertainty for patient decisions, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

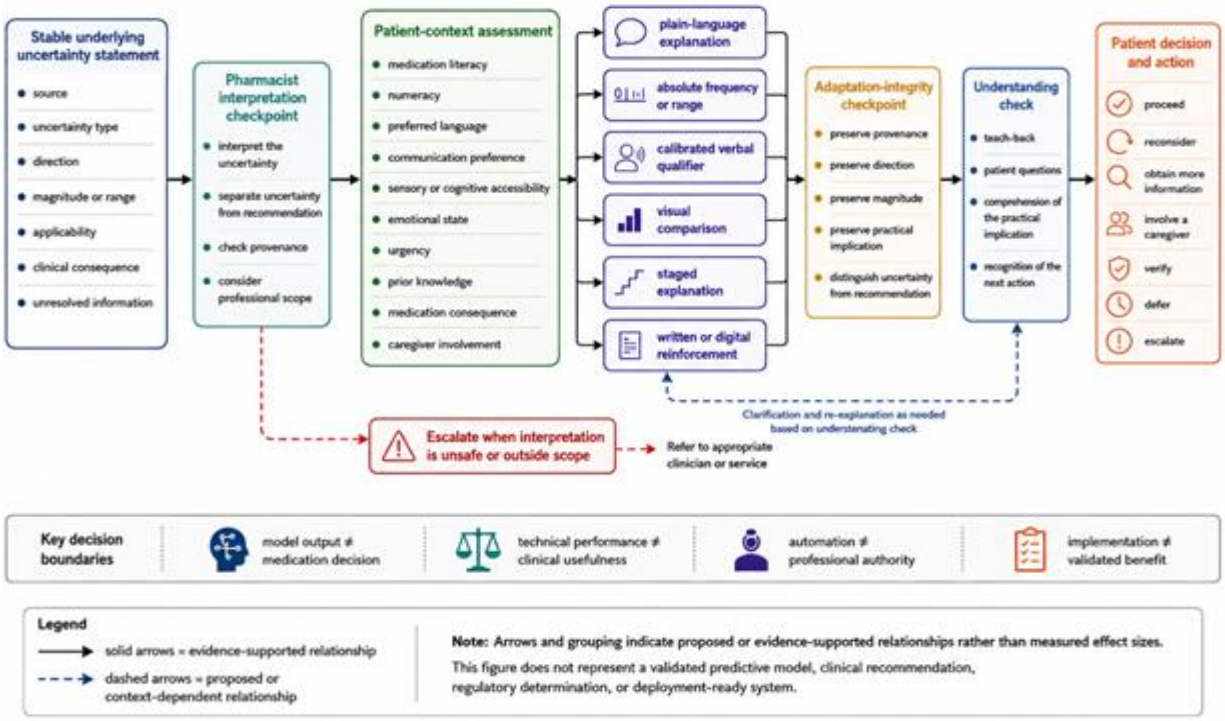

Figure 2. Pharmacist-Mediated Adaptation of Uncertainty for Patient Decisions

The adaptation pathway is proposed rather than prescriptive. Its feasibility may differ across community, hospital, remote, and automated dispensing settings. Equity evaluation should examine who receives uncertainty information, who understands it, whose encounters are delayed, whose prescriptions are escalated, and whether communication burdens are distributed differently across populations.

Limitations

The framework integrates evidence from pharmacy practice, risk communication, clinical decision support, and human–AI research, but several supporting studies arise from diagnostic or non-pharmacy settings. Their mechanisms require pharmacy-specific testing.

Post hoc explanations may appear plausible without faithfully representing how a model produced an output and cannot substitute for validation [28]. In some high-stakes tasks, intrinsically interpretable approaches may be preferable to explanations applied after an opaque prediction [29]. Neither interpretability nor pharmacist involvement guarantees safe use.

Machine-learning systems may also create unintended errors, dependence, deskilling, and new sociotechnical failure pathways [30]. The proposed taxonomy, consequence profile, adaptation-integrity rule, and action boundary have not been prospectively validated. The framework does not establish communication effectiveness, reduced medication error, improved adherence, regulatory acceptability, or deployment readiness.

Conclusion

Communicating calibrated uncertainty at dispensing is not merely a matter of displaying a confidence score. It requires

an explicit relationship among the source of uncertainty, evidence of calibration and applicability, the clinical consequence of error, the patient's informational and decisional needs, and the pharmacist's authority to verify, defer, or escalate.

The proposed framework organizes this work through uncertainty characterization, calibration and applicability appraisal, consequence assessment, communication-package construction, pharmacist mediation, and a bounded action decision. It distinguishes model output from medication decision, technical performance from clinical usefulness, explanation from validation, and patient trust from appropriate reliance.

Its professional contribution is to define uncertainty communication as active interpretive work rather than passive disclosure. Its organizational contribution is to make provenance, workflow capacity, escalation, monitoring, and accountability necessary conditions. Its safety contribution is to reject confidence-only action. Its equity contribution is to require evaluation of access, comprehension, burden, verification, delay, and escalation across patient groups.

These relationships remain proposed. The framework requires prospective, setting-specific validation across technical, communicative, clinical, human-factors, workflow, safety, equity, and governance outcomes. It should not be interpreted as a clinical protocol, regulatory standard, universal communication method, or deployment-ready system. Its value lies in providing a testable structure for deciding what uncertainty means, how it should be communicated, and when dispensing should proceed or pause.

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

References

1. 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
2. Corny J, Rajkumar A, Martin O, Dode X, Lajonchère JP, Billuart O, et al. A machine learning-based clinical decision support system to identify prescriptions with a high risk of medication error. J Am Med Inform Assoc. 2020;27(11):1688-94. doi:10.1093/jamia/ocaa154
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. Elwyn G, Durand MA, Song J, Aarts J, Barr PJ, Berger Z, et al. A three-talk model for shared decision making: Multistage consultation process. BMJ. 2017;359. doi:10.1136/bmj.j4891
5. 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
6. Begoli E, Bhattacharya T, Kusnezov D. The need for uncertainty quantification in machine-assisted medical decision making. Nat Mach Intell. 2019;1(1):20-3. doi:10.1038/s42256-018-0004-1
7. Leibig C, Allken V, Ayhan MS, Berens P, Wahl S. Leveraging uncertainty information from deep neural networks for disease detection. Sci Rep. 2017;7(1):17816. doi:10.1038/s41598-017-17876-z
8. Kang DY, DeYoung PN, Tantiongloc J, Coleman TP, Owens RL. Statistical uncertainty quantification to augment clinical decision support: A first implementation in sleep medicine. NPJ Digit Med. 2021;4(1):142. doi:10.1038/s41746-021-00515-3
9. 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
10. 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
11. 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
12. 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
13. Spiegelhalter D. Risk and uncertainty communication. Annu Rev Stat Appl. 2017;4:31-60. doi:10.1146/annurev-statistics-010814-020148
14. van der Bles AM, van der Linden S, Freeman ALJ, Mitchell J, Galvao AB, Zaval L, et al. Communicating uncertainty about facts, numbers and science. R Soc Open Sci. 2019;6(5):181870. doi:10.1098/rsos.181870
15. van der Bles AM, van der Linden S, Freeman ALJ, Spiegelhalter DJ. The effects of communicating uncertainty on public trust in facts and numbers. Proc Natl Acad Sci U S A. 2020;117(14):7672-83. doi:10.1073/pnas.1913678117
16. Garcia-Retamero R, Cokely ET. Designing visual aids that promote risk literacy: A systematic review of health research and evidence-based design heuristics. Hum Factors. 2017;59(4):582-627. doi:10.1177/0018720817690634
17. Cordina M, Hämeen-Anttila K, Lauri J, Tabone S, Enlund H. Health and medication literacy and the desire to participate in pharmacotherapy decision making—comparison of two countries. Res Social Adm Pharm. 2018;14(9):817-23. doi:10.1016/j.sapharm.2018.06.009
18. Weir KR, Bonner C, McCaffery K, Naganathan V, Carter SM, Rigby D, et al. Pharmacists and patients sharing decisions about medicines: Development and feasibility of a conversation guide. Res Social Adm Pharm. 2019;15(6):682-90. doi:10.1016/j.sapharm.2018.08.009
19. Fosgerau CF, Kaae S. Furthering patient-centered counseling: Exploring new aspects around pharmacists' experiences in pharmacy encounters through video-stimulated recall interviewing. Res Social Adm Pharm. 2021;17(4):723-32. doi:10.1016/j.sapharm.2020.06.018
20. Manias E, Hughes C, Woodward-Kron R, Ozavci G, Jorm C, Bucknall T. Decision-making about changing medications across transitions of care: Opportunities for enhanced patient and family engagement. Res Social Adm Pharm. 2024;20(5):520-30. doi:10.1016/j.sapharm.2024.02.002
21. 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
22. 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
23. Kiani A, Uyumazturk B, Rajpurkar P, Wang A, Gao R, Jones E, et al. Impact of a deep learning assistant on the histopathologic classification of liver cancer. NPJ Digit Med. 2020;3(1):23. doi:10.1038/s41746-020-0232-8
24. Stacey D, Légaré F, Lewis K, Barry MJ, Bennett CL, Eden KB, et al. Decision aids for people facing health treatment or screening decisions. Cochrane Database Syst Rev. 2017;4(4). doi:10.1002/14651858.CD001431.pub5

25. Sendak MP, Ratliff W, Sarro D, Alderton E, Futoma J, Gao M, et al. Real-world integration of a sepsis deep learning technology into routine clinical care: Implementation study. JMIR Med Inform. 2020;8(7). doi:10.2196/15182
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. Roodbeen R, Vreke A, Boland G, Rademakers J, van den Muijsenbergh M, Noordman J, et al. Communication and shared decision-making with patients with limited health literacy; helpful strategies, barriers and suggestions for improvement reported by hospital-based palliative care providers. PLoS One. 2020;15(6). doi:10.1371/journal.pone.0234926
28. 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
29. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206-15. doi:10.1038/s42256-019-0048-x
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