Original Article

The Digital Pharmacist Proxy and the Clinical Acts That Must Remain Human

Elena Petrova[1]*, Ivan Georgiev[1], Nikolay Stoyanov[2], Petar Kolev[3]

[1]Department of Digital Proxies and Pharmacy Practice, Faculty of Pharmacy, Sofia University, Sofia, Bulgaria. [2]Department of Human-Only Clinical Acts, Faculty of Pharmacy, University of Plovdiv, Plovdiv, Bulgaria. [3]Department of AI Boundaries in Pharmacy, Faculty of Pharmacy, University of Ruse, Ruse, Bulgaria.

Abstract

Artificial intelligence is moving from supporting pharmacy work toward producing outputs that can appear to represent the pharmacist. This shift creates a boundary problem: technical capability to retrieve, classify, draft, or recommend does not by itself establish authority to perform a professional clinical act. This article proposes a Digital Pharmacist Proxy–Professional Boundary Model for distinguishing acts that may be delegated to a digital proxy, acts that require conditional human–AI co-production, and acts that must remain human-led. The model begins with the act rather than the technology and evaluates verifiability, reversibility, contextual sufficiency, relational dependence, ethical or normative load, authority and consequence, and the operational management of uncertainty and accountability. These dimensions are assembled within a proposed permission envelope specifying scope, users, population, inputs, outputs, autonomy, pharmacist oversight, fallback, monitoring, and expiry. Informational and administrative acts may be delegable when they are bounded, auditable, reversible, and low in normative consequence. Medication-related interpretation may require co-production when uncertainty, contextual incompleteness, or meaningful clinical consequence demands pharmacist review and authorization. Relational engagement, value elicitation, ethical exception handling, professional commitment, and final authority-bearing decisions remain human-led under the proposed model. Validation would require task-level technical assessment, human-factors testing, workflow evaluation, equity analysis, prospective safety study, and governance testing. The model is an original non-empirical synthesis. It does not establish professional equivalence, clinical benefit, regulatory acceptance, or deployment readiness, and its boundaries require empirical and normative validation across pharmacy settings.

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

Introduction

Clinical artificial intelligence is commonly framed as augmenting rather than replacing professional judgment, because computational scale and pattern recognition do not reproduce the full clinical, relational, and accountable work of a health professional [1]. In pharmacy, however, the language and interface of contemporary systems increasingly make them appear to speak, explain, prioritize, or recommend as though they were performing the pharmacist's role. The relevant question is therefore no longer only whether a tool assists a pharmacist, but when its output represents an act that users may attribute to a pharmacist.

Available clinical-pharmacy applications remain heterogeneous and are concentrated in bounded functions such as prediction, information handling, screening, and decision support rather than validated representation of the integrated professional role [2]. This matters because a technically competent output can cross a professional boundary when it is presented as individualized medication

Address for correspondence: Elena Petrova, Department of Digital Proxies and Pharmacy Practice, Faculty of Pharmacy, Sofia University, Sofia, Bulgaria.
elena.petrova@uni-sofia.bg
Received: 01 March 2026; **Accepted:** 23 May 2026

How to cite this article: Petrova E, Georgiev I, Stoyanov N, Kolev P. The Digital Pharmacist Proxy and the Clinical Acts That Must Remain Human. Arch Pharm Pract. 2026;17(2):67-74. https://doi.org/10.51847/pN0iDcYpWH

advice, when it alters a medication decision, or when patients and clinicians reasonably interpret it as carrying professional authority.

Pharmacy-facing AI now spans multiple practice functions, including generative systems proposed for drug-information, therapy-support, and hospital-workflow tasks [3, 4]. Expansion of function does not establish equivalence of role. A system may generate a plausible counseling script without building trust, identify a medication issue without negotiating priorities, or propose an action without possessing authority to accept responsibility for its consequences. Alternative explanations for apparent professional performance include imitation of professional language, access to structured knowledge, and successful completion of narrow benchmarks.

This article therefore treats a digital pharmacist proxy as a system that does more than assist: it produces or communicates an output that is presented, interpreted, or operationalized as representing pharmacist work. The literature describes many pharmacy applications [2]. It also reveals a representation gap between discrete tools and the pharmacist's integrated contribution [3]. The original contribution proposed here is an act-level boundary model. It does not classify an entire technology as acceptable or unacceptable; instead, it asks which act is being represented, what makes that act professional, what evidence supports delegation, who retains authority, and how permission can be restricted or withdrawn.

Defining Representation of Pharmacist Work

A defensible boundary requires a definition of the work being represented. Pharmacist work includes observable tasks, but it also includes the reasoning that connects patient information, medication knowledge, uncertainty, and possible action. Efforts to make medication-review reasoning explicit show that professional performance cannot be reduced to whether a final recommendation was produced [5]. Representation therefore concerns both output and the cognitive or interpretive process claimed by that output.

Professional representation also has a relational dimension. Patients develop trust in community pharmacists through interpersonal and professional conduct, not solely through receipt of correct information [6]. Trust should not be romanticized as proof that every interaction must be face-to-face, but it does show that accuracy is only one determinant of whether an act is recognizably professional.

Medication counseling further illustrates the distinction. Patient participation is shaped by patient, pharmacist, and environmental factors, making counseling a reciprocal and context-sensitive process rather than one-way information transmission [7]. A digital system may prepare education, translate instructions, or identify questions; it does not thereby become the author of a shared understanding or the accountable negotiator of a medication plan.

Finally, pharmacist work is organizationally situated. Integration into teams involves mediation, coordination, and negotiated professional relationships [8]. A proxy can consequently alter more than a task: it can redistribute who interprets information, who communicates uncertainty, who authorizes action, and who is expected to respond when an output fails. In this article, representation of pharmacist work is therefore defined as the production or communication of an act attributed to pharmacist reasoning, relationship, ethical judgment, or professional authority. This definition is proposed as an analytic boundary and is not a universal taxonomy of pharmacy practice.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article The Digital Pharmacist Proxy and the Clinical Acts That Must Remain Human.

Table 1. Definitions, components, relationships, and intended uses for the article The Digital Pharmacist Proxy and the Clinical Acts That Must Remain Human.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Digital assistant	Confusion between support and representation	User, task, output, and workflow position	Supports a human act without claiming authorship or authority; proposed distinction	Preserves a narrow support role	Task accuracy, usability, and workflow fit	Users may overattribute competence	Assistance does not equal professional representation [1].
Digital pharmacist proxy	Outputs appear to stand for pharmacist work	Presentation, personalization, actionability, and attribution	Proposed proxy status arises when an output is treated as pharmacist-authored or pharmacist-authorized	Makes hidden delegation visible	User-interpretation and workflow studies	Proxy status may arise unintentionally	The proxy is an analytic category, not a licensed professional.

Clinical reasoning	Final outputs can conceal how cases were interpreted	Patient data, medication knowledge, alternatives, and uncertainty	Representation must account for reasoning, not only task completion	Prevents output-only equivalence claims	Process tracing and case-based validation	Plausible answers may mask faulty reasoning	One reasoning model is not an exhaustive taxonomy [5].
Relational work	Correct information may not create trust or understanding	Continuity, communication, responsiveness, and patient perception	Human conduct shapes trust; the proposed model treats relational dependence as a boundary test	Identifies acts requiring human presence or leadership	Patient-reported and interaction evidence	Digital fluency may be mistaken for trustworthiness	Accuracy alone does not establish relational competence [6].
Participatory counseling	Information transfer can be mistaken for shared understanding	Patient goals, questions, literacy, and setting	Counseling emerges through participation and contextual adaptation	Separates educational preparation from relational authorship	Observation, participation, comprehension, and concordance evidence	Standardized output may suppress patient priorities	Digital preparation may support, but not automatically replace, counseling [7].
Organizational mediation	Task automation can redistribute professional work invisibly	Team roles, escalation routes, local norms, and accountability	Proxy use changes coordination and negotiated role boundaries	Makes workflow consequences assessable	Sociotechnical and implementation studies	Responsibility may become fragmented	A task cannot be evaluated outside its organizational setting [8].
Permission envelope	Broad labels do not define safe use	Task, users, population, inputs, outputs, autonomy, fallback, monitoring, and expiry	Proposed envelope limits what the proxy may do and under which conditions	Converts abstract boundaries into governable permissions	Prospective task- and setting-specific validation	Scope may expand without evidence	Permission is act-specific, conditional, monitored, and revocable.
Human-retained act	Some acts carry irreducible relational, ethical, or authority content	Values, conflict, consequence, and professional responsibility	Proposed rule retains human leadership when an act constitutes relationship, judgment, or accountable commitment	Protects professional and patient-facing boundaries	Normative analysis plus empirical testing	Overrestriction may block useful support; underrestriction may displace responsibility	Human retention concerns authority over the act, not exclusion of digital support.

Proposed Proxy–Professional Boundary Model

Human–AI combinations are not inherently better than either humans or AI alone; benefit depends on task structure, interaction design, and comparative capabilities [9]. The proposed model therefore rejects both blanket automation and the weak reassurance that any "human in the loop" is sufficient. Its unit of analysis is a candidate pharmacy act: for example, retrieving interaction evidence, drafting instructions, prioritizing a medication problem, explaining uncertainty, negotiating a preference-sensitive choice, or authorizing a change.

The act is assessed through seven proposed boundary tests: verifiability, reversibility, contextual sufficiency, relational dependence, ethical or normative load, authority and consequence, and uncertainty and accountability. These tests assign the act to one of three zones. A delegable proxy act is bounded, auditable, reversible, contextually adequate, low in normative load, and linked to clear accountability. A conditional co-produced act may use AI-generated analysis, but a pharmacist must interpret, check, contextualize, and authorize the consequential output. Experimental evidence that collaboration can work under specified arrangements supports the plausibility of an intermediate co-production zone, not universal superiority or pharmacy validation [10]. A human-retained act is one whose performance constitutes a therapeutic relationship, value judgment, ethical exception, or professional commitment rather than merely supplying information.

Figure 1 presents the original conceptual synthesis for digital pharmacist proxy–professional boundary model, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

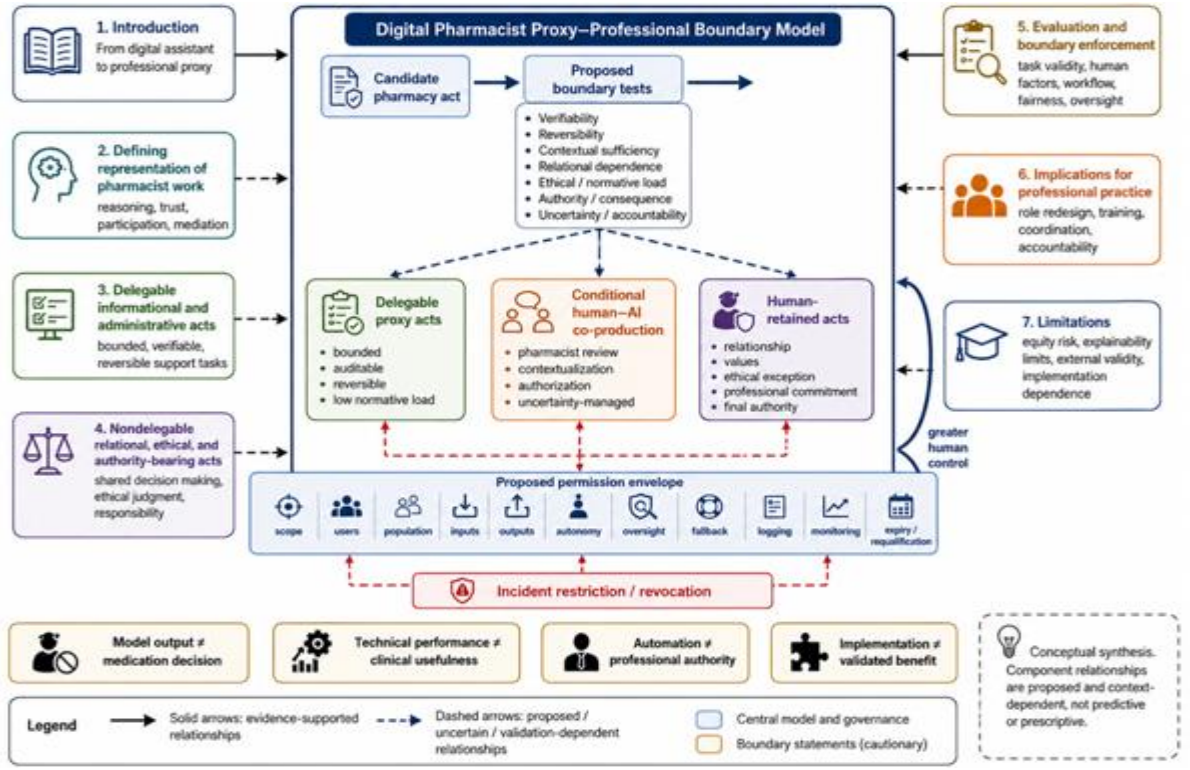

Figure 1. Digital Pharmacist Proxy–Professional Boundary Model

Permissions within the model are asymmetric. Escalation toward greater human control is always allowed, whereas movement toward greater delegation requires new evidence and explicit approval. This asymmetry responds to the possibility that machine-learning use can create new errors, dependencies, and sociotechnical consequences [11]. A high-performing system may therefore remain restricted if its failures are difficult to detect, its outputs are hard to reverse, or its use obscures responsibility.

The three zones operate inside a proposed permission envelope. The envelope specifies the task, authorized users, population and setting, accepted inputs, permitted outputs, autonomy, required pharmacist review, fallback, monitoring, and expiry. Professional AI literacy and oversight remain necessary because users must understand both capability and limitation [12]. The envelope is not a declaration that the system is safe; it is a governance object defining what must be evaluated, what is prohibited, and when use must stop.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article The Digital Pharmacist Proxy and the Clinical Acts That Must Remain Human.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article The Digital Pharmacist Proxy and the Clinical Acts That Must Remain Human.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Candidate-act specification	Define the exact act before evaluating the system	Intended act, user, setting, recipient, and consequence	Initiates all boundary tests	Act is observable, bounded, and distinguishable from the whole pharmacist role	Technology-level approval substitutes for act-level analysis	Service owner and pharmacist define scope	No assessment until the act and consequence are explicit.
Capability comparison	Determine whether AI, human performance, or combination is appropriate	Comparative task evidence and error profiles	Informs zone assignment	Evidence matches the intended act and context	Assumed synergy or benchmark overgeneralization	Multidisciplinary evaluation team	Human–AI combination is not presumed superior [9].

Boundary testing	Apply the seven proposed tests	Task properties, context, relationship, ethics, consequence, and uncertainty	Produces provisional zone assignment	Each test has task-specific evidence and a documented rationale	Important dimensions are omitted or collapsed into one score	Pharmacist-led governance body	No composite score may offset a critical boundary failure.
Zone assignment	Classify an act as delegable, co-produced, or human-retained	Boundary-test results	Determines oversight intensity	Classification is reproducible and understandable	Similar-looking acts receive inconsistent permissions	Accountable clinical owner	Assignment is provisional and setting-specific.
Conditional co-production	Combine machine analysis with pharmacist interpretation and authorization	AI output, uncertainty, patient context, and pharmacist judgment	Sits between delegation and human retention	Pharmacist can meaningfully verify and alter the output	Human review becomes ceremonial	Pharmacist and workflow designer	Collaboration requires demonstrated task-specific complementarity [10].
Permission envelope	Limit use to approved conditions	Scope, users, inputs, outputs, autonomy, fallback, and monitoring	Constrains every zone	Controls are technically and operationally enforceable	Scope creep, unlogged use, or missing fallback	Organization assigns owner, escalation route, and expiry	Permission is conditional, monitored, and revocable [12].
Incident response	Escalate, restrict, or revoke permission	Errors, near misses, complaints, and drift signals	Can move any act toward greater human control	Signals trigger timely investigation and containment	Harmful patterns persist because failures are normalized	Safety and governance leads	Unresolved safety or accountability failure blocks continued use [11].
Requalification	Reassess after change or expiry	Model, data, workflow, population, or policy change	Returns the act to specification and testing	Material changes are detected and reassessed	Prior evidence is reused after conditions change	Designated clinical and technical owners	No automatic expansion or renewal of permission.

Delegable Informational and Administrative Acts

Delegation should attach to an act, not to the label "pharmacist." Informational and administrative acts may enter the delegable zone when they are narrowly specified, auditable, reversible, and separable from final clinical authority. Candidate examples include retrieving formulary information, organizing medication histories, identifying missing fields, preparing comparison summaries, translating standardized instructions, drafting documentation, routing requests, and generating questions for pharmacist review. These are proposed categories, not an assertion that every instance is safe to automate.

Evidence that generative AI can perform bounded clinical-pharmacy question tasks supports evaluation of informational delegation, but it does not establish equivalence to integrated pharmacist practice or authorize direct patient-specific action [13]. The same nominal task may also move between zones. A general summary of administration instructions may be delegable; adapting those instructions to cognitive impairment, conflicting goals, uncertain adherence, or a high-consequence regimen may require co-production or human retention.

Three conditions are central. First, verifiability requires a competent reviewer or reliable check to determine whether the output is correct. Second, reversibility asks whether an error can be corrected before meaningful harm or commitment. Third, contextual sufficiency asks whether the available data represent the patient, medication trajectory, setting, and decision purpose. These conditions are proposed as necessary but not sufficient. A task can be verifiable yet impose excessive checking burden, or reversible in theory yet operationally difficult to intercept.

Medication verification illustrates the conditional zone. AI helpfulness and uncertainty presentation shape pharmacists' verification behavior [14, 15]. The implication is not that uncertainty displays guarantee safe reliance. Rather, interface design, uncertainty communication, time pressure, and the pharmacist's ability to contest the output become part of the act's evidence requirements. A system that supplies a recommendation without enabling meaningful verification should not be treated as safely co-produced merely because a pharmacist remains nominally responsible.

The proposed delegability test therefore combines low normative load, limited consequence, reliable verification, reversibility, contextual adequacy, transparent uncertainty, and an identified accountable owner. Task allocation remains conditional because human–AI advantage depends on the particular interaction [9]. The permission envelope must also state what the proxy may not do, when a pharmacist must take over, and how the system responds to missing or conflicting information [12]. Even within an approved act, unexpected uncertainty should trigger escalation rather than confident completion [14].

Delegation is consequently best understood as restricted permission to perform a defined contribution. It does not

transfer professional identity, ethical standing, or final authority. The pharmacist remains responsible for deciding whether the act belongs in the delegable zone, whether its preconditions are present, and whether accumulated exceptions require restriction or withdrawal.

Nondelegable Relational, Ethical, and Authority-Bearing Acts

Some pharmacy acts should remain human-led not because a digital system can never contribute information, but because the act itself constitutes a relationship, ethical judgment, or exercise of professional authority. An AI system may organize options for discussion; it cannot independently establish that a patient has been heard, that conflicting priorities have been reconciled, or that a preference-sensitive medication plan reflects meaningful participation. Evidence on AI-enabled decision aids indicates that users continue to value clinician involvement and shared decision processes [16]. Value elicitation, negotiation of treatment burdens, and resolution of disagreement therefore belong in the proposed human-retained zone.

Ethical acts also exceed prediction or information synthesis. Healthcare AI raises concerns at individual, interpersonal, institutional, and societal levels, including autonomy, fairness, responsibility, and distribution of benefit and burden [17]. A digital proxy may identify that an ethical conflict exists, but selecting among competing obligations requires an accountable human decision-maker who can explain and defend the judgment.

Professional authority is distinct from technical expertise. Clinical action involves responding to a person's concern, particularizing general knowledge to an individual situation, and accepting responsibility for the resulting commitment [18]. Under the proposed model, final authorization of consequential medication changes, acceptance of unresolved uncertainty, departure from standard pathways, and communication of ethically significant recommendations remain pharmacist-led.

"Human-retained" does not mean digitally unsupported. AI may retrieve evidence, structure alternatives, document discussions, or reveal inconsistencies. The boundary concerns authorship and authority over the completed act. A pharmacist must remain able to reject the output, reformulate the problem, explain uncertainty, and assume responsibility without being constrained by a system-generated recommendation.

Evaluation and Boundary Enforcement

Evaluation must examine more than model discrimination, linguistic quality, or agreement with expert answers. Early clinical evaluation of AI decision support requires attention to safety, human factors, workflow integration, and the conditions under which users interact with the system [19]. The proposed boundary model should therefore be tested at the act level: whether the intended contribution is accurately specified, whether the proxy remains within its permission envelope, and whether pharmacists can identify and correct unsafe outputs.

Trustworthiness is multidimensional. Fairness, universality, traceability, usability, robustness, and explainability should be assessed separately rather than collapsed into a single readiness judgment [20]. A technically accurate proxy may still be unsuitable if it performs unequally across populations, obscures its information sources, creates excessive checking burden, or cannot be restricted after a workflow change.

Boundary enforcement also requires effective rather than nominal human oversight. Automation bias and verification complexity can lead users to accept incorrect advice or make review so burdensome that the human role becomes ceremonial [21]. Evaluation should therefore test disagreement behavior, escalation, time pressure, missing-data handling, uncertainty communication, and the pharmacist's practical ability to override the system.

The proposed governance sequence comprises act specification, evidence review, permission approval, monitored use, incident-triggered restriction, and expiry or requalification. Permission should narrow when uncertainty increases, contextual assumptions fail, or accountability becomes unclear. Expansion toward greater delegation requires new evidence; prior approval should not automatically transfer across populations, interfaces, models, or pharmacy settings.

Implications for Professional Practice

The proposed model reframes digital transformation as professional redesign rather than simple substitution. Pharmacy's response to changing technologies and labor conditions requires deliberate decisions about which activities should be automated, which should be co-produced, and which express the profession's distinctive contribution [22]. Efficiency gains should therefore be evaluated alongside effects on judgment, patient access, continuity, responsibility, and professional development.

Introducing AI can create new forms of work, including checking outputs, translating recommendations, resolving discrepancies, documenting overrides, and coordinating escalation. Qualitative evidence from clinical AI implementation shows that integration depends on such adaptation and coordination rather than software installation alone [23]. Pharmacy organizations should assign this work explicitly and avoid treating oversight as an invisible addition to existing workloads.

Education should combine technical literacy with strengthened relational, ethical, and authority-bearing capabilities. Preparing pharmacists for digital practice requires understanding system limitations while retaining the

competencies through which pharmacists interpret patient circumstances and assume professional responsibility [24]. Training should therefore include uncertainty communication, automation-bias recognition, contestability, incident reporting, and decisions about when not to use a proxy.

Limitations

The proposed model is a conceptual synthesis and has not been empirically validated. Algorithms may encode inequity through inappropriate proxy targets, while apparently helpful explanations may create false assurance; neither predictive performance nor explainability alone establishes a legitimate basis for professional delegation [25, 26]. The model also draws some mechanisms from healthcare fields outside pharmacy, and their transfer requires testing.

External validity, workflow fit, prospective evaluation, and implementation evidence remain necessary before decision use [27]. Boundary assignments may vary by setting, patient population, consequence, and available oversight. The model is not a clinical recommendation, regulatory standard, universal taxonomy, or deployment authorization.

Conclusion

The Digital Pharmacist Proxy–Professional Boundary Model proposes that decisions about pharmacy AI should begin with the clinical act rather than with the technology. It distinguishes delegable informational and administrative contributions, conditional human–AI co-production, and human-retained relational, ethical, and authority-bearing acts. The central boundary is not whether a system can generate a plausible answer. It is whether the act can be verified, reversed, contextually supported, separated from relationship and values, and governed through identifiable human accountability. Where these conditions are absent, apparent technical capability should not be treated as professional authority.

The model also makes delegation conditional and revocable. Permission is limited by a defined envelope, monitored in practice, and narrowed when uncertainty, inequity, workflow disruption, or accountability failures emerge. Its scientific value lies in converting a broad replacement debate into testable act-level questions. Its professional value lies in preserving the pharmacist's responsibility for interpretation, relationship, ethical judgment, and consequential commitment while allowing bounded digital support. Empirical, normative, organizational, and equity-focused validation is required before the proposed boundaries can guide implementation.

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. Ranchon F, Chanoine S, Lambert-Lacroix S, Bosson JL, Moreau-Gaudry A, Bedouch P. Development of artificial intelligence powered apps and tools for clinical pharmacy services: A systematic review. Int J Med Inform. 2023;172:104983. doi:10.1016/j.ijmedinf.2022.104983
3. Hatzimanolis J, Riley B, El-Den S, Aslani P, Zhou J, Chaar BB. Applications of artificial intelligence in current pharmacy practice: A scoping review. Res Social Adm Pharm. 2025;21(3):134-41. doi:10.1016/j.sapharm.2024.12.007
4. Li X, Guo H, Li D, Zheng Y. Engine of innovation in hospital pharmacy: Applications and reflections of ChatGPT. J Med Internet Res. 2024;26. doi:10.2196/51635
5. Guignard B, Crevier F, Charlin B, Audétat MC. A graphical model to make explicit pharmacist clinical reasoning during medication review. Res Social Adm Pharm. 2024;20(12 Pt A):1142-50. doi:10.1016/j.sapharm.2024.09.005
6. Gregory PAM, Austin Z. How do patients develop trust in community pharmacists? Res Social Adm Pharm. 2021;17(5):911-20. doi:10.1016/j.sapharm.2020.07.023
7. Qudah B, Thakur T, Chewning B. Factors influencing patient participation in medication counseling at the community pharmacy: A systematic review. Res Social Adm Pharm. 2021;17(11):1863-76. doi:10.1016/j.sapharm.2021.03.005
8. Nabhani-Gebara S, Fletcher S, Shamim A, May L, Butt N, Chagger S, et al. General practice pharmacists in England: Integration, mediation and professional dynamics. Res Social Adm Pharm. 2020;16(1):17-24. doi:10.1016/j.sapharm.2019.01.014
9. Vaccaro M, Almaatouq A, Malone T. When combinations of humans and AI are useful: A systematic review and meta-analysis. Nat Hum Behav. 2024;8(12):2293-303. doi:10.1038/s41562-024-02024-1
10. Reverberi C, Rigon T, Solari A, Hassan C, Cherubini P, Antonelli G, et al. Experimental evidence of effective human-AI collaboration in medical decision-making. Sci Rep. 2022;12(1):14952. doi:10.1038/s41598-022-18751-2
11. 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
12. Meskó B, Görög M. A short guide for medical professionals in the era of artificial intelligence. NPJ Digit Med. 2020;3:126. doi:10.1038/s41746-020-00333-z
13. Huang X, Estau D, Liu X, Yu Y, Qin J, Li Z. Evaluating the performance of ChatGPT in clinical pharmacy: A comparative study of ChatGPT and clinical pharmacists. Br J Clin Pharmacol. 2024;90(1):232-8. doi:10.1111/bcp.15896
14. Tsai CC, Kim JY, Chen Q, Rowell B, Yang XJ, Kontar R, et al. Effect of artificial intelligence helpfulness and uncertainty on cognitive interactions with pharmacists: Randomized controlled trial. J Med Internet Res. 2025;27. doi:10.2196/59946
15. Lester C, Rowell B, Zheng Y, Co Z, Marshall V, Kim JY, et al. Effect of uncertainty-aware AI models on pharmacists' reaction time and decision-making in a web-based mock medication verification task: Randomized controlled trial. JMIR Med Inform. 2025;13. doi:10.2196/64902
16. Hassan N, Slight R, Bimpong K, Bates DW, Weiand D, Vellinga A, et al. Systematic review to understand users perspectives on AI-enabled decision aids to inform shared decision making. NPJ Digit Med. 2024;7:332. doi:10.1038/s41746-024-01326-y
17. Morley J, Machado CCV, Burr C, Cowls J, Joshi I, Taddeo M, et al. The ethics of AI in health care: A mapping review. Soc Sci Med. 2020;260:113172. doi:10.1016/j.socscimed.2020.113172
18. Heyen NB, Salloch S. The ethics of machine learning-based clinical decision support: An analysis through the lens of professionalisation theory. BMC Med Ethics. 2021;22:112. doi:10.1186/s12910-021-00679-3
19. 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
20. Lekadir K, Frangi AF, Porras AR, Glocker B, Cintas C, Langlotz CP, et al. FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388. doi:10.1136/bmj-2024-081554
21. Lyell D, Coiera E. Automation bias and verification complexity: A systematic review. J Am Med Inform Assoc. 2017;24(2):423-31. doi:10.1093/jamia/ocw105
22. Traulsen JM, Druedahl LC. Shifting perspectives—Planning for the future of the pharmacy profession taking current labor market trends into consideration. Res Social Adm Pharm. 2018;14(12):1189-94. doi:10.1016/j.sapharm.2018.02.006
23. Farič N, Hinder S, Williams R, Ramaesh R, Bernabeu MO, van Beek E, et al. Early experiences of integrating an artificial intelligence-based diagnostic decision support system into radiology settings: A qualitative study. J Am Med Inform Assoc. 2024;31(1):24-34. doi:10.1093/jamia/ocad191
24. Gregório J, Cavaco A. The pharmacist's guide to the future: Are we there yet? Res Social Adm Pharm. 2021;17(4):795-8. doi:10.1016/j.sapharm.2020.05.029
25. 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
26. 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
27. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17:195. doi:10.1186/s12916-019-1426-2