Original Article

Safety Boundaries for Autonomous Sterile Compounding Cells

Ethan Wright[1]*, Chloe Bennett[2], Jack Turner[1], Olivia Harris[3]

[1]Department of Sterile Compounding Safety and AI, Faculty of Pharmacy, University of Leeds, Leeds, United Kingdom. [2]Department of Autonomous Compounding Boundaries, Faculty of Pharmacy, University of Sheffield, Sheffield, United Kingdom. [3]Department of Safety-Critical Pharmacy Automation, Faculty of Pharmacy, University of York, York, United Kingdom.

Abstract

Autonomous sterile compounding combines digital decision systems with a high-consequence physical process in which errors may affect product identity, dose, sterility, occupational exposure, workflow integrity, and medication release. Existing technologies can automate selected preparation, measurement, documentation, and verification tasks, but no single sensor, accuracy measure, artificial-intelligence model, or professional check is sufficient to establish system safety. This article proposes a cyber-physical safety architecture for autonomous sterile compounding cells. The architecture treats autonomy as conditional authority bounded by an authorized operating envelope, material and order controls, environmental and process sensing, independent verification, interlocks, hazard-specific safe states, accountable human supervision, evidence preservation, and controlled incident recovery. It distinguishes technical performance from medication-use decisions, automated execution from professional release authority, and successful implementation from validated benefit. The proposed relationships are organized as testable safety claims rather than assumed properties of automation. Evaluation would require component verification, aseptic and environmental qualification, fault injection, human-factors testing, workflow assessment, traceable safety-case evidence, and prospective monitoring under defined operating conditions. Changes in products, software, equipment, environment, staffing, or workflow would require reassessment of the applicable safety boundary. The architecture is conceptual and does not establish regulatory conformity, clinical effectiveness, universal applicability, or deployment readiness. Its original contribution is an integrated account of how physical containment, computational control, professional authority, and organizational governance may be connected without treating any individual component as a substitute for system-level safety assurance.

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

Introduction

Sterile compounding automation should be understood as a system-level change rather than the replacement of an isolated manual step. International experience with robotic intravenous compounding indicates that operating procedures, responsibilities, maintenance, workflow integration, and local risk controls must be specified around the technology [1]. An autonomous sterile compounding cell therefore includes not only a robot but also the order interface, materials, sensors, software, containment environment, verification pathway, operators, pharmacists, maintenance functions, and product-release decision.

Implementation evidence from an oncology satellite pharmacy further shows that introducing a compounding robot changes staffing, production organization, exception handling, and local evaluation requirements [2]. Autonomy is consequently defined here as delegated capability to perform specified physical and informational actions without continuous direct manipulation. It is not unrestricted

Address for correspondence: Ethan Wright, Department of Sterile Compounding Safety and AI, Faculty of Pharmacy, University of Leeds, Leeds, United Kingdom.
ethan.wright@leeds.ac.uk
Received: 11 March 2026; **Accepted:** 04 June 2026

How to cite this article: Wright E, Bennett C, Turner J, Harris O. Safety Boundaries for Autonomous Sterile Compounding Cells. Arch Pharm Pract. 2026;17(2):84-92. https://doi.org/10.51847/WzmggKdwbq

authority to interpret every order, continue through every anomaly, or release a medication.

Technology-assisted sterile compounding can improve traceability and alter checking, error-detection, and task-time patterns, but these consequences remain dependent on the workflow and setting in which the system operates [3]. Technical performance must therefore be distinguished from clinical usefulness and medication safety. A cell may execute an intended movement accurately while receiving an incorrect input, operating in an unsuitable environment, failing to detect contamination, or producing evidence insufficient for professional release.

The broader medical-AI literature frames high performance as a complementary relationship between human and computational capabilities rather than autonomous sufficiency [4]. In sterile compounding, human–AI collaboration should similarly be designed around differences in capability. Machines may provide repeatable execution, continuous sensing, and traceable records, whereas pharmacists and trained operators retain contextual interpretation, exception management, product-disposition, and professional-accountability functions.

Responsible health-care machine learning also requires prospective harm analysis, monitoring, stakeholder responsibility, and controls extending beyond model discrimination or accuracy [5]. This article accordingly proposes a cyber-physical safety architecture in which autonomy remains conditional on evidence, operating context, and human authority. The architecture is an original conceptual synthesis. It does not establish that any current robotic compounding platform satisfies the proposed requirements.

Hazards and System Boundaries in Sterile Compounding

A safety boundary specifies the conditions, components, interfaces, and decisions within which an autonomous cell may operate. Its purpose is not merely to enclose equipment, but to identify where hazards enter, propagate, become detectable, and require control. Gravimetric workflow can strengthen measurement and checking for hazardous sterile preparations, yet such control addresses only selected preparation deviations and cannot establish global safety [6].

Nonhazardous sterile compounding also presents distinct workflow and error-capture conditions [7]. Hazard classification should therefore consider product characteristics, route, formulation, container, compounding complexity, environmental requirements, and consequences of failure. A single authorization category for all compounded sterile preparations would obscure clinically and operationally important differences.

Sterility cannot be inferred from robotic enclosure or repeatable motion. Microbiological performance requires qualification of aseptic processes and the environment in which compounding occurs [8]. The proposed boundary consequently includes the isolator or cleanroom state, equipment surfaces, material transfer, interventions, airflow-related conditions, cleaning status, and evidence supporting continued aseptic control.

Hazardous-drug contamination may extend across pharmacy and care environments rather than remaining within the nominal robotic enclosure [9]. The relevant system boundary must therefore include receipt, preparation, transfer, waste, decontamination, maintenance, and downstream handling. Boundary crossings should be treated as control points at which identity, containment, evidence, and responsibility may be lost.

Surface-contamination patterns support environmental surveillance as a safety function distinct from dose verification [10]. Environmental signals should not be interpreted as proof that every released preparation is contaminated or uncontaminated. Instead, they may indicate loss of process control, the need for investigation, or restriction of operation pending evaluation.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Safety Boundaries for Autonomous Sterile Compounding Cells.

Table 1. Definitions, components, relationships, and intended uses for the article Safety Boundaries for Autonomous Sterile Compounding Cells.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Authorized operating envelope (proposed)	Unbounded autonomous action	Product class, equipment, environment, validated configuration	Permits operation only under predefined conditions	Restricts unsupported use	Local qualification and change evidence	Conditions may drift unnoticed	Does not authorize use outside validated conditions

Order and material identity	Wrong input or component	Prescription, barcode, ingredient, container	Identity agreement precedes execution	Prevents propagation of input errors	Traceable order–material reconciliation	Correct execution of incorrect input	Technical matching does not replace pharmacist review
Dose and volume verification	Quantitative preparation error	Target amount, gravimetry, process record	Independent comparison with expected preparation [6]	Detects selected deviations	Measurement-system validation	Shared calibration or data failure	Accuracy alone does not establish safety
Product-specific workflow boundary	Inappropriate uniform control	Hazard class, formulation, workflow complexity	Controls vary by preparation context [7]	Improves contextual fit	Product- and setting-specific evaluation	Overgeneralization across preparations	No universal pathway is assumed
Aseptic and environmental state	Microbial or particulate control loss	Qualification, interventions, cleaning, environmental status	Environmental acceptability conditions continued operation [8]	Supports aseptic-process assurance	Microbiological and environmental qualification	Intermittent or undetected excursions	Enclosure does not prove sterility
Hazardous-drug containment	Occupational and environmental exposure	Drug properties, surfaces, transfer, waste	Boundary extends across the handling chain [9]	Supports containment and escalation	Exposure and contamination monitoring	Contamination outside the cell	Robot boundaries are not exposure boundaries
Environmental surveillance	Loss of process control	Surface and environmental signals	Trends trigger investigation or restriction [10]	Detects possible control deterioration	Valid sampling and response procedures	False reassurance or unnecessary shutdown	Signals require contextual interpretation
Human authority (proposed synthesis)	Ambiguous release and exception responsibility	Evidence package, anomaly, professional judgment	Human authority governs release, disposition, and escalation	Preserves accountable medication decisions	Human-factors and workflow validation	Nominal oversight without usable control	Automation does not acquire professional authority

Proposed Cyber-Physical Safety Architecture

Dose accuracy and precision are measurable attributes of robotic compounding, but they do not represent the complete safety state of a sterile product or production system [11]. The proposed architecture therefore places the autonomous cell inside an authorized operating envelope and surrounds execution with identity control, environmental sensing, process sensing, independent verification, interlocks, human authority, evidence preservation, and governance.

Figure 1 presents the original conceptual synthesis for cyber-physical safety architecture for autonomous sterile compounding, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

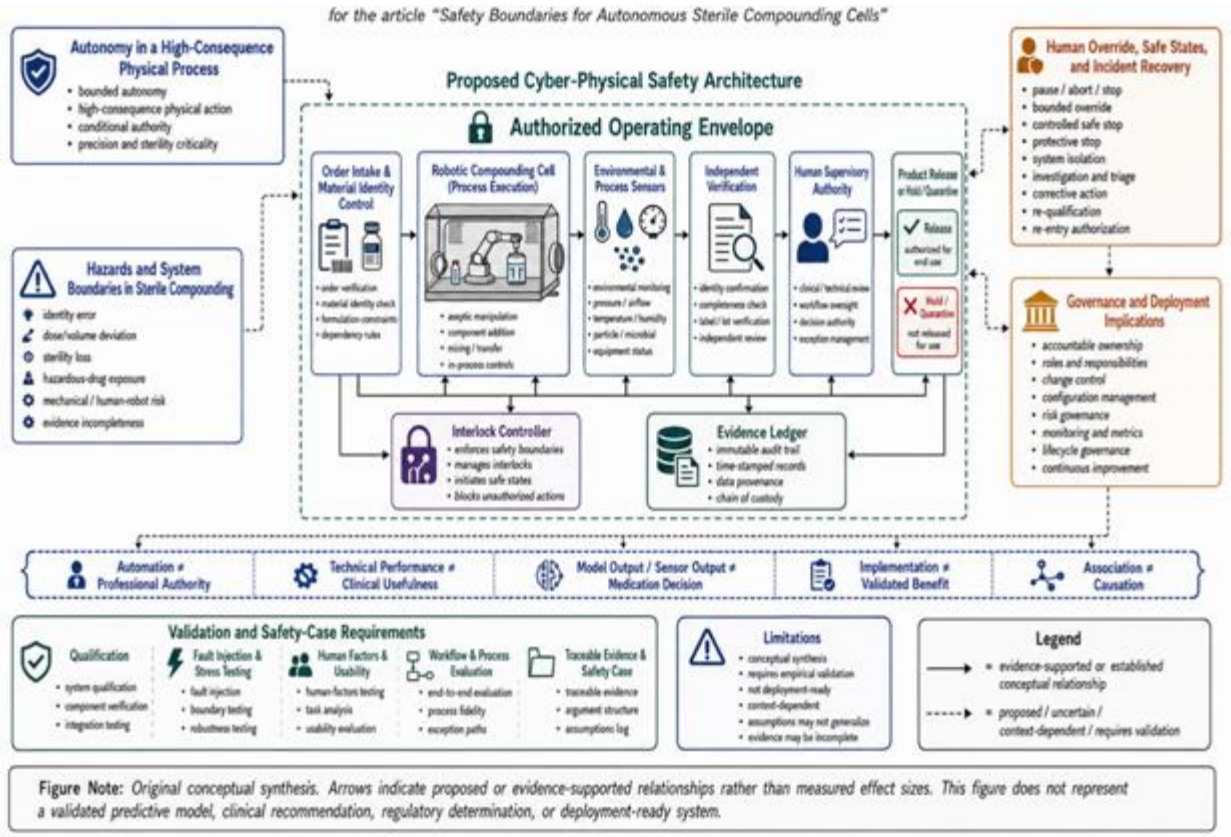

Figure 1. Cyber-Physical Safety Architecture for Autonomous Sterile Compounding

Technology-assisted workflow studies indicate that automation redistributes detectable errors, checking activities, traceability, and process time [12]. The architecture therefore rejects a single final accuracy gate. It proposes layered controls in which input validation, execution monitoring, environmental acceptability, quantitative verification, exception detection, and product-release review provide distinguishable evidence.

Remote sterile-product checking is feasible only when the pharmacist receives an adequate evidence package and responsibility remains explicit [13]. In the proposed architecture, the evidence ledger connects sensor observations, material identity, executed actions, verification results, interventions, configuration state, and unresolved anomalies. Remote access is an interface to that evidence, not a transfer of accountability to the software.

Robotic and gravimetric-assisted workflows have different processing and human-intervention profiles [14]. The architecture accordingly defines separate operating modes rather than assuming that all automation levels are interchangeable. A degraded mode may permit limited operation only when its hazards, evidence requirements, and human responsibilities have been prospectively specified; otherwise, the cell should enter hold or a hazard-specific safe state.

Sterile compounding is ultimately a sociotechnical work system involving people, tasks, technologies, environments, organizational structures, and transitions [15]. The proposed architecture therefore treats staffing, training, workload, communication, maintenance, and escalation capacity as safety-relevant conditions. These relationships require empirical validation and should not be interpreted as evidence that the integrated architecture is already effective.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Safety Boundaries for Autonomous Sterile Compounding Cells.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Safety Boundaries for Autonomous Sterile Compounding Cells.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary

Operating-envelope gate (proposed)	Authorize bounded operation	Configuration and context → permission	Constrains all layers	Correct acceptance and rejection of conditions	Operation after unassessed change	System owner approves scope	Not ready without defined configuration limits
Input and material control	Reconcile order and components	Order, barcode, material data → cell	Feeds execution and ledger	Traceable identity agreement	Incorrect but internally consistent data	Pharmacist defines discrepancy response	Matching does not equal clinical appropriateness
Physical execution	Perform preparation	Authorized command → robot action	Monitored by process sensors	Reproducibility and fault detectability	Misalignment, obstruction, wrong action	Technical team maintains equipment	Motion accuracy does not prove product safety
Quantitative verification	Check dose and volume	Expected and measured values → comparison	Independent of primary execution where feasible	Validated measurement and discrepancy handling [11]	Common-mode calibration error	Pharmacy validates measurement process	No release based on one metric
Environmental control	Maintain acceptable compounding conditions	Environmental state → interlock and record	Conditions execution and release	Qualified detection and response	Undetected excursion	Pharmacy and facilities share control	Qualification is context-specific
Evidence ledger (proposed)	Preserve traceable decision evidence	All layers → time-linked record	Supports review, recovery, and safety case	Completeness, integrity, retrievability	Missing or altered evidence	Information-governance owner	Incomplete evidence prevents autonomous release
Human supervisory layer	Review anomalies and disposition	Evidence and alerts → human decision	May pause, abort, quarantine, or escalate	Timely, usable, independent oversight	Automation bias or inaccessible evidence	Pharmacist retains release authority	Override cannot bypass product verification
Change and recovery gate (proposed)	Control return after fault or modification	Incident/change evidence → reauthorization	Connects governance to operation	Requalification and closure of unresolved hazards	Premature restart	Named authority approves re-entry	Implementation does not establish validated benefit

Interlocks, Environmental Sensing, and Process Verification

Interlocks are controls that prevent initiation, interrupt execution, or restrict continued operation when specified conditions are not satisfied. Occupational biological-monitoring evidence for antineoplastic exposure supports treating containment failure and operator exposure as safety concerns requiring monitoring and escalation, although it does not define a real-time sensor specification for a compounding cell [16]. The proposed interlock model therefore includes physical, environmental, identity, quantitative, software, and evidence-completeness conditions.

Incident recovery cannot be represented as an undifferentiated cleaning step. Comparative evidence indicates that removal effectiveness varies across antineoplastic agents and decontamination solutions [17]. Recovery logic should consequently identify the agent, affected surfaces, exposure pathway, selected method, verification requirement, product disposition, and conditions for re-entry.

Contamination control requires complementary microbiological qualification, environmental surveillance, and validated decontamination rather than reliance on enclosure alone [8, 9, 17].

Process verification should produce a structured safety argument, not merely a collection of successful test results. Assurance-case methods link claims to arguments, evidence, assumptions, and context [18]. Adapted to sterile compounding, the top-level claim might state that a defined product was compounded within an authorized state. Supporting evidence would then address identity, environment, execution, quantitative checks, equipment configuration, interventions, and unresolved anomalies. This structure is proposed and requires testing for usability, completeness, and resistance to documentation burden.

Human–robot collaboration introduces physical hazards that require context-sensitive sensing and separation or collision controls [19]. Industrial evidence can inform hazard identification, but direct transfer to sterile-compounding cells would be inappropriate without pharmacy-specific analysis. Interlocks must account for gloved access, material transfer, maintenance, cleaning, emergency intervention, and the possibility that stopping motion may preserve personnel safety while compromising aseptic conditions or product integrity.

A safe state is therefore hazard-specific. Depending on the initiating event, the appropriate response may be controlled pause, material isolation, product quarantine, restricted degraded mode, protective stop, environmental containment, or emergency shutdown. The safest state for personnel may differ from the state that best protects the product, environment, or evidence trail. The controller should prioritize predefined hazard logic while preserving information needed for human assessment.

Autonomous-system safety evidence can become outdated when software, configuration, equipment, environment, products, or operating context changes [20]. The proposed architecture links interlocks to a dynamic safety-case function: material change, recurring anomaly, sensor degradation, unexplained disagreement, or altered workflow should trigger reassessment of permissible operation. Continued operation after change would remain a governed decision rather than an automatic consequence of prior qualification.

Human Override, Safe States, and Incident Recovery

Human override is not inherently protective. Automation bias and verification complexity can cause users to accept automated outputs or perform only superficial checks, particularly when independent assessment is difficult [21]. Meaningful override therefore requires accessible evidence, clear authority, sufficient time, trained personnel, and an interface that supports independent judgment.

Explainability should be matched to the information needs of pharmacists, operators, engineers, and incident investigators; an explanation cannot replace independent verification or accountable product disposition [22]. The proposed architecture consequently separates system-generated explanations from the professional decision to release, quarantine, discard, or investigate a preparation.

Conditional deferral may be useful when system confidence, sensor agreement, or operating conditions fall outside predefined competence boundaries [23]. Incorrect automated recommendations can nevertheless influence expert performance [24]. Human intervention must therefore be triggered before evidence becomes irrecoverable and must not permit an unresolved discrepancy to be bypassed.

Robotic work is also a team-system activity involving communication, coordination, training, workload, and recovery capacity [25]. Meaningful oversight requires independent verification, calibrated deferral, and protection against automation-biased acceptance of incorrect machine output [21, 23, 24].

The proposed safe states are hazard-specific: controlled pause, product hold, quarantine, degraded operation, protective stop, environmental containment, or emergency shutdown. Override cannot connect directly to product release. Incident recovery requires evidence preservation, product disposition, investigation, corrective action, requalification, and authorized re-entry. These transitions remain proposed and require sterile-compounding-specific validation.

Validation and Safety-Case Requirements

Validation should test safety claims across components, interactions, users, environments, and failure conditions. Reporting guidance for artificial-intelligence interventions emphasizes description of system inputs, human interaction, error handling, and evaluation procedures [26]. Prospective protocols should likewise prespecify intended use, operating conditions, interaction pathways, anticipated failures, and analysis plans [27].

Early-stage evaluation should examine usability, workflow integration, human factors, safety events, and iterative refinement before effectiveness or deployment claims are made [28]. Where learned models are used, data provenance, model construction, validation, and reproducibility require explicit documentation [29]. Predictive components also require transparent reporting of intended use, calibration, performance, and limitations [30].

For an autonomous compounding cell, the proposed safety case should connect each claim to evidence from equipment qualification, media fills, environmental monitoring, identity and dose verification, sensor testing, fault injection, cybersecurity and data-integrity assessment, human-factors simulation, workflow evaluation, and longitudinal monitoring. Evidence should be configuration-specific. A prior validation result should not automatically cover modified software, ingredients, containers, sensors, environmental conditions, staffing arrangements, or operating modes.

Governance and Deployment Implications

Deployment complexity extends beyond technical installation. Health technologies may be rejected, abandoned, or fail to scale when technology, users, organizations, value propositions, and wider systems are misaligned [31]. Implementation determinants should therefore be assessed across the innovation, inner and outer settings, involved individuals, and implementation process [32].

Lifecycle governance must assign authority for design, procurement, configuration, validation, production authorization, release, override, monitoring, incident response, change control, re-entry, and retirement [33]. Ethical governance must also address transparency, responsibility, bias, privacy, trust, and distributive consequences [34].

Deployment readiness depends jointly on technology complexity, implementation context, and explicit lifecycle governance [31-33].

Figure 2 presents the original conceptual synthesis for interlocks, safe states, human override, verification, and incident recovery, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

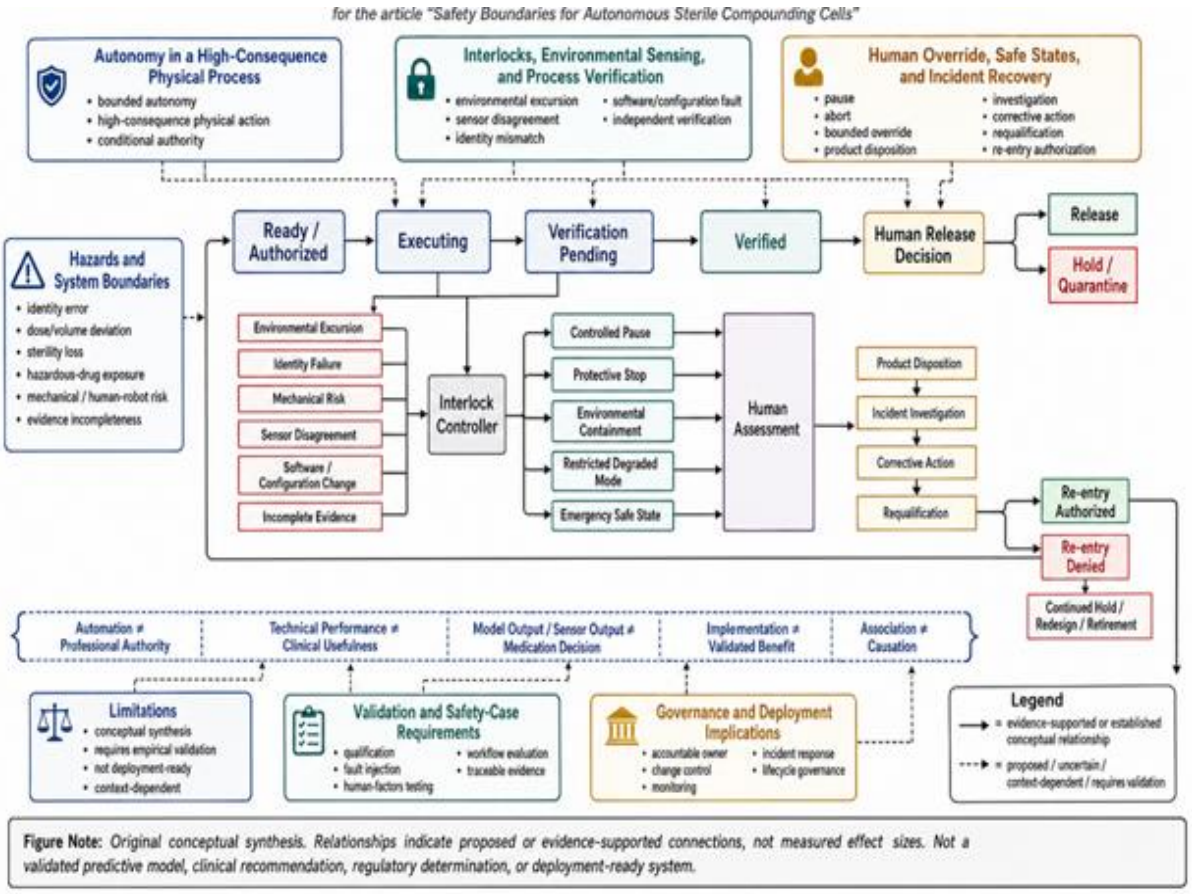

Figure 2. Interlocks, Safe States, Human Override, Verification, and Incident Recovery

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for governance responsibilities, implementation conditions, and monitoring requirements for the article Safety Boundaries for Autonomous Sterile Compounding Cells.

Table 3. Governance responsibilities, implementation conditions, and monitoring requirements for the article Safety Boundaries for Autonomous Sterile Compounding Cells.

Governance domain	Responsible actor	Required authority	Documentation requirement	Monitoring indicator	Escalation trigger	Corrective action	Accountability boundary
Design and procurement	Pharmacy, engineering, procurement	Approve intended use and requirements	Hazard analysis and supplier evidence	Unresolved requirement or interface risk	Safety-critical requirement not demonstrated	Redesign, additional evidence, or rejection	Supplier evidence does not replace local validation
Local configuration	System owner and pharmacy lead	Approve validated configuration	Version, interface, material, and environment record	Configuration deviation	Unauthorized or unassessed change	Restrict operation and reassess	Prior qualification covers only the defined configuration
Production authorization	Pharmacist and operational lead	Permit defined operating mode	Qualification and readiness record	Failed prerequisite or incomplete evidence	Operating-envelope violation	Hold production	Automation cannot authorize itself
Product release	Authorized pharmacist	Release, quarantine, or reject	Complete product evidence package	Unresolved discrepancy	Missing, conflicting, or unreliable evidence	Quarantine and investigate	Technical completion is not medication release
Human override	Trained authorized user	Pause, abort, or enter bounded override	Reason, action, evidence, and outcome	Override frequency and unresolved exceptions	Repeated or unsafe override pattern	Suspend mode, retrain, or redesign	Override cannot bypass verification

Monitoring and change control	Governance and quality team	Restrict operation and require revalidation	Trend, incident, maintenance, and change records	Drift, recurring alarms, contamination, or sensor disagreement	Safety case no longer supports operation	Update controls and requalify	Monitoring does not prove continuing benefit
Incident recovery	Incident lead and quality authority	Control disposition and re-entry	Preserved evidence, investigation, corrective action	Incomplete cause analysis or failed requalification	Residual or unexplained hazard	Continue hold, redesign, or retire	Re-entry requires independent authorization
Equity and access	Institutional governance	Review distributive consequences	Access, delay, error-distribution, and service-impact record	Unequal burden or service displacement	Persistent differential impact	Modify workflow or allocation	Technical efficiency does not establish equitable benefit

Limitations

Aggregate technical performance may conceal biased or unsafe failure modes [35]. Data-dependent components may also deteriorate under distribution shift, making operating envelopes and revalidation triggers necessary [36]. Optimization targets or proxies can embed structural inequity even when computational performance appears acceptable [37]. Weak design, leakage, limited external validation, or unrepresentative data can similarly make apparently strong evidence unsuitable for deployment [38].

The architecture remains conceptual. Its applicability may differ across products, robots, isolators, institutions, staffing models, jurisdictions, and risk classes. Evidence from medical AI, industrial robotics, and other clinical technologies informs relevant mechanisms but does not directly validate autonomous sterile compounding cells. The proposal does not replace applicable professional standards, local validation, or regulatory requirements.

Conclusion

Autonomous sterile compounding should be governed as a bounded cyber-physical and sociotechnical system rather than evaluated through a single accuracy measure, artificial-intelligence model, sensor, or professional checkpoint. The proposed architecture connects an authorized operating envelope with identity controls, environmental and process sensing, independent verification, interlocks, hazard-specific safe states, accountable human authority, evidence preservation, incident recovery, and lifecycle governance.

Its principal scientific contribution is a testable organization of relationships that are often assessed separately. Its professional contribution is the preservation of pharmacist authority over medication release and exception disposition. Its organizational contribution is the assignment of responsibility across configuration, operation, monitoring, change, and recovery. Its safety contribution is the explicit treatment of failure transitions and incomplete evidence. Its equity contribution is recognition that efficiency and automation may distribute access, delay, and error unevenly.

These elements remain proposed. Their value depends on empirical evaluation within defined products, environments, workflows, and institutions. The architecture does not establish clinical benefit, regulatory conformity, universal applicability, or deployment readiness. It provides a bounded research framework through which those claims may later be tested.

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

References

1. Yaniv AW, Orsborn A, Bonkowski JJ, Chew L, Krämer I, Calabrese S, et al. Robotic i.v. medication compounding: Recommendations from the international community of APOTECAchemo users. Am J Health Syst Pharm. 2017;74(1). doi:10.2146/ajhp151027
2. Bhakta SB, Colavecchia AC, Coffey W, Curlee DR, Garey KW. Implementation and evaluation of a sterile compounding robot in a satellite oncology pharmacy. Am J Health Syst Pharm. 2018;75(11 Suppl 2). doi:10.2146/ajhp170461
3. Fan M, Yang D, Ng B, Jackson J, Bouris K, Eng S, et al. Impact of technology-assisted versus manual sterile compounding on safety and efficiency in a Canadian community hospital. Am J Health Syst Pharm. 2022;79(19):1685-96. doi:10.1093/ajhp/zxac167
4. 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
5. 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
6. Roberts PA, Willoughby IR, Barnes N, Eckel SF, Paruscio A, Valgus JM, et al. Evaluation of a gravimetric-based technology-assisted workflow system on hazardous sterile product preparation. Am J Health Syst Pharm. 2018;75(17):1286-92. doi:10.2146/ajhp170564
7. Bucci TG, Hedrick TL, Roberts PA, Amerine LB. Evaluation of gravimetric-based technology-assisted workflow for nonhazardous sterile product preparation. Am J Health Syst Pharm. 2019;76(14):1071-7. doi:10.1093/ajhp/zxz097
8. Geersing TH, Franssen EJF, Pilesi F, Crul M. Microbiological performance of a robotic system for aseptic compounding of cytostatic drugs. Eur J Pharm Sci. 2019;130:181-5. doi:10.1016/j.ejps.2019.01.034
9. Sottani C, Grignani E, Cornacchia M, Negri S, della Cuna FSR, Cottica D, et al. Occupational exposure assessment to antineoplastic drugs in nine Italian hospital centers over a 5-year survey program. Int J Environ Res Public Health. 2022;19(14):8601. doi:10.3390/ijerph19148601
10. Salch SA, Zamboni WC, Zamboni BA, Eckel SF. Patterns and characteristics associated with surface contamination of hazardous drugs in hospital pharmacies. Am J Health Syst Pharm. 2019;76(9):591-8. doi:10.1093/ajhp/zxz033
11. Geersing TH, Klous MG, Franssen EJF, van den Heuvel JJG, Crul M. Robotic compounding versus manual compounding of chemotherapy: Comparing dosing accuracy and precision. Eur J Pharm Sci. 2020;155:105536. doi:10.1016/j.ejps.2020.105536

12. Eckel SF, Higgins JP, Hess E, Cerbone T, Civiello JB, Conley C, et al. Multicenter study to evaluate the benefits of technology-assisted workflow on i.v. room efficiency, costs, and safety. Am J Health Syst Pharm. 2019;76(12):895-901. doi:10.1093/ajhp/zxz067
13. Jean SJ, Francart SJ, Eckel SF, Schenkat D, Eberwein S, Lamm M, et al. Evaluation of telepharmacy and the use of a gravimetric technology-assisted workflow system for remote sterile product pharmacist checks. Am J Health Syst Pharm. 2020;77(7):560-7. doi:10.1093/ajhp/zxaa015
14. Pang B, Earl M, Knoer S, Yaniv AW, Boyd AM. Comparison of IV oncology infusions compounded via robotics and gravimetrics-assisted workflow processes. Am J Health Syst Pharm. 2021;78(2):122-34. doi:10.1093/ajhp/zxaa366
15. Carayon P, Wooldridge A, Hoonakker P, Schoofs Hundt A, Kelly MM. SEIPS 3.0: Human-centered design of the patient journey for patient safety. Appl Ergon. 2020;84:103033. doi:10.1016/j.apergo.2019.103033
16. Leso V, Sottani C, Santocono C, Russo F, Grignani E, Iavicoli I. Exposure to antineoplastic drugs in occupational settings: A systematic review of biological monitoring data. Int J Environ Res Public Health. 2022;19(6):3737. doi:10.3390/ijerph19063737
17. Simon N, Guichard N, Odou P, Decaudin B, Bonnabry P, Fleury-Souverain S. Efficiency of four solutions in removing 23 conventional antineoplastic drugs from contaminated surfaces. PLoS One. 2020;15(6). doi:10.1371/journal.pone.0235131
18. Wei R, Kelly TP, Dai X, Zhao S, Hawkins R. Model based system assurance using the Structured Assurance Case Metamodel. J Syst Softw. 2019;154:211-33. doi:10.1016/j.jss.2019.05.013
19. Patalas-Maliszewska J, Łosyk H, Dudek A. Improving safety in human–robot collaboration towards sustainable production in Industry 5.0. J Intell Manuf. 2026;37(8):3069-104. doi:10.1007/s10845-025-02676-4
20. Buysse L, Habli I, Vanoost D, Pissoort D. Safe autonomous systems in a changing world: Operationalising dynamic safety cases. Saf Sci. 2025;191:106965. doi:10.1016/j.ssci.2025.106965
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. Amann J, Blasimme A, Vayena E, Frey D, Madai VI. 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
23. Dvijotham KD, Winkens J, Barsbey M, Ghaisas S, Stanforth R, Pawlowski N, et al. Enhancing the reliability and accuracy of AI-enabled diagnosis via complementarity-driven deferral to clinicians. Nat Med. 2023;29(7):1814-20. doi:10.1038/s41591-023-02437-x
24. Dratsch T, Chen X, Rezazade Mehrizi M, Kloeckner R, Mähringer-Kunz A, Püsken M, et al. Automation bias in mammography: The impact of artificial intelligence BI-RADS suggestions on reader performance. Radiology. 2023;307(4). doi:10.1148/radiol.222176
25. Catchpole K, Cohen T, Alfred M, Lawton S, Kanji F, Shouhed D, et al. Human factors integration in robotic surgery. Hum Factors. 2024;66(3):683-700. doi:10.1177/00187208211068946
26. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK; CONSORT-AI and SPIRIT-AI Working Group. 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. 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
29. Norgeot B, Quer G, Beaulieu-Jones BK, Torkamani A, Dias R, Gianfrancesco M, et al. Minimum information about clinical artificial intelligence modeling: The MI-CLAIM checklist. Nat Med. 2020;26(9):1320-4. doi:10.1038/s41591-020-1041-y
30. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385. doi:10.1136/bmj-2023-078378
31. Greenhalgh T, Wherton J, Papoutsi C, Lynch J, Hughes G, A'Court C, et al. Beyond adoption: A new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. J Med Internet Res. 2017;19(11). doi:10.2196/jmir.8775
32. Damschroder LJ, Reardon CM, Widerquist MAO, Lowery J. The updated Consolidated Framework for Implementation Research based on user feedback. Implement Sci. 2022;17(1):75. doi:10.1186/s13012-022-01245-0
33. 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
34. 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
35. Challen R, Denny J, Pitt M, Gompels L, Edwards T, Tsaneva-Atanasova K. Artificial intelligence, bias and clinical safety. BMJ Qual Saf. 2019;28(3):231-7. doi:10.1136/bmjqs-2018-008370
36. Subbaswamy A, Saria S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics. 2020;21(2):345-52. doi:10.1093/biostatistics/kxz041
37. 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
38. Roberts M, Driggs D, Thorpe M, Gilbey J, Yeung M, Ursprung S, et al. Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. Nat Mach Intell. 2021;3(3):199-217. doi:10.1038/s42256-021-00307-0