Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

LEARNING WITHOUT SHARING THE DATA: AN EVIDENCE MAP OF PRIVACY-PRESERVING COLLABORATION IN PHARMACEUTICAL RESEARCH

Carlos Almeida[1]*, Gabriela Mendez[2], Ricardo Fonseca[3], Ivan Pereira[4]

1. ***Department of Privacy-Preserving Collaboration and Evidence Map, Faculty of Pharmacy, Federal University of Rio Grande do Sul, Porto Alegre, Brazil.***
2. ***Department of Federated Learning and Data Privacy, Faculty of Pharmaceutical Sciences, Federal University of Pampa, São Gabriel, Brazil.***
3. ***Department of Collaborative Research without Data Sharing, Faculty of Pharmacy, National University of Uruguay, Montevideo, Uruguay.***
4. ***Department of Privacy-Preserving AI in Pharma, Faculty of Pharmacy, University of Buenos Aires, Buenos Aires, Argentina.***

ARTICLE INFO

Received:
17 June 2025
Received in revised form:
21 October 2025
Accepted:
25 October 2025
Available online:
28 October 2025

Keywords: Federated learning, Privacy-enhancing technologies, Secure multiparty computation, Synthetic data, Distributed analytics, Pharmaceutical data science

ABSTRACT

Pharmaceutical research increasingly depends on evidence distributed across companies, healthcare institutions, research networks, and jurisdictions, yet legal restrictions, commercial sensitivity, technical incompatibility, and patient-privacy obligations often prevent centralized data pooling. Privacy-preserving collaboration has therefore emerged as a possible means of learning across organizational boundaries while retaining data at their sources or limiting disclosure during analysis. However, the evidence base combines conceptually distinct technologies, heterogeneous validation designs, and inconsistent claims about privacy, analytical utility, and implementation readiness. This evidence-mapping review critically organizes recent peer-reviewed scholarship on federated learning, secure computation, synthetic data, and distributed analytics across pharmaceutical discovery, development, and safety. The review maps evidence according to the protected asset, exchanged artifact, application context, threat model, validation setting, implementation maturity, utility limitation, and translational boundary. The synthesis indicates that data locality can enable collaboration but does not itself constitute a complete privacy guarantee. Formal privacy mechanisms provide stronger protection for specified adversarial conditions, although their usefulness remains conditional on computational feasibility, privacy parameters, model performance, and organizational infrastructure. Evidence is most developed for retrospective modeling, molecular-property prediction, distributed clinical analysis, and cross-company quantitative structure–activity relationship learning. By contrast, prospective decision impact, independent security testing, semantic interoperability, governance evaluation, and routine pharmaceutical integration remain comparatively underdeveloped. The article proposes a bounded maturity interpretation that separates conceptual plausibility, technical feasibility, retrospective validation, operational piloting, and decision-integrated readiness. This synthesis is an organizing framework rather than a validated scoring instrument. Trustworthy pharmaceutical collaboration will require privacy claims to be linked explicitly to threat models, utility assessments, data suitability, reproducibility, governance, and the scientific consequences of model use.

To Cite This Article: Almeida C, Mendez G, Fonseca R, Pereira I. Learning without Sharing the Data: An Evidence Map of Privacy-Preserving Collaboration in Pharmaceutical Research. Pharmacophore. 2025;16(5):119-29. https://doi.org/10.51847/mS21cxN933

Introduction

Pharmaceutical research is increasingly computational, data-intensive, and distributed across organizations that hold complementary but non-interchangeable evidence. Molecular structures, assay measurements, omics profiles, imaging data, trial records, real-world observations, and safety reports may each contribute to a scientific or development decision, yet they are frequently separated by institutional ownership, contractual restrictions, jurisdictional rules, incompatible data models, and commercial sensitivity. Machine learning has expanded across target identification, molecular design, property prediction,

Corresponding Author: Carlos Almeida; Department of Privacy-Preserving Collaboration and Evidence Map, Faculty of Pharmacy, Federal University of Rio Grande do Sul, Porto Alegre, Brazil. E-mail: carlos.almeida@ufrgs.br.

clinical development, and post-authorization evidence generation, but the scientific value of these applications remains dependent on the suitability, representativeness, and interpretability of the underlying data rather than on algorithmic sophistication alone [1]. Centralizing all relevant records is therefore neither technically straightforward nor always legally, ethically, or commercially acceptable.

Federated learning has become a prominent response to this constraint because it permits participating organizations to train or evaluate models while retaining source records within local environments. Its importance lies not merely in avoiding a large central repository, but in enabling coordinated analysis across otherwise inaccessible datasets. Nevertheless, decentralization does not eliminate data heterogeneity, governance obligations, model leakage, infrastructure vulnerabilities, or uncertainty about whether collaborating institutions are solving sufficiently comparable tasks [2]. Privacy-preserving machine learning consequently includes a broader family of approaches, such as differential privacy, secure aggregation, homomorphic encryption, multiparty computation, trusted execution environments, synthetic-data generation, and distributed query systems. These mechanisms protect different objects against different adversaries and should not be described as equivalent forms of privacy protection [3].

The unresolved problem is therefore not whether collaboration without conventional data pooling is technically imaginable. It is how the evidence should be interpreted when reported studies use different meanings of privacy, different validation environments, and different measures of utility. In pharmaceutical research, the protected information may include proprietary chemical structures, confidential assay labels, target hypotheses, patient-level observations, genomic records, trial data, safety signals, model parameters, or organizational participation itself. A method that protects one of these assets may expose another through gradients, outputs, metadata, membership signals, or repeated queries. Drug-discovery scholarship has accordingly emphasized the need to evaluate decentralized learning together with model leakage, cryptographic protection, computational burden, trust assumptions, and the distinctive commercial sensitivity of molecular and assay data [4].

The available literature also spans markedly different levels of implementation. Some studies partition public datasets to simulate multiple partners, whereas others execute analyses across institutions or companies holding genuinely distributed information. Cross-pharmaceutical federated quantitative structure–activity relationship modeling demonstrates that large-scale collaboration can be operationalized without transferring proprietary source records, yet the resulting performance improvements remain dependent on endpoint characteristics, partner distributions, data quality, and the validation design [5]. Such evidence does not by itself establish experimental success, candidate quality, clinical utility, regulatory acceptability, or routine organizational readiness. The objective of this evidence-mapping review is therefore to classify privacy-preserving pharmaceutical collaboration by technology, protected object, exchanged artifact, application, threat model, validation quality, maturity, utility limitation, and translational boundary. Its central contribution is an evaluative map that separates demonstrated capability from benchmark performance, plausible pharmaceutical usefulness, prospective decision value, and operational deployment.

Evidence-Mapping Questions and Conceptual Boundaries

The primary conceptual boundary concerns what is meant by *learning without sharing the data*. The phrase is used here to denote collaborative computation in which participating organizations do not conventionally transfer their complete source datasets into a common analytical repository. It does not imply that nothing is exchanged, that the resulting model is non-disclosive, or that every participant receives the same protection. The evidence map therefore asks what remains local, what crosses organizational boundaries, which asset is intended to be protected, which adversary is considered, and what scientific decision the collaboration is intended to support. Within molecular discovery, this distinction is essential because predictive performance on molecular-property tasks must be separated from experimental confirmation, biological mechanism, synthesizability, developability, safety, and therapeutic value [6]. The map consequently treats pharmaceutical usefulness as a sequence of bounded evidentiary claims rather than as a direct consequence of collaborative model training.

A second boundary separates architectural data locality from privacy guarantees. In conventional federated learning, participants commonly exchange parameters, gradients, predictions, or other model-derived artifacts while raw records remain local. Multi-institutional demonstrations show that this architecture can facilitate collaboration without conventional patient-data transfer, but they also show that the resulting system still depends on server design, client behavior, update visibility, harmonized tasks, and local data quality [7]. Synthetic data occupy a different position because generated records may be transferred or released, sometimes more freely than their source data. Their acceptability depends on whether they preserve properties needed for the intended analysis while limiting disclosure about source individuals, organizations, or proprietary observations [8]. The evidence map therefore classifies federation and synthetic release separately rather than treating both as equivalent examples of decentralization.

A third boundary concerns the meaning of utility. Distributional similarity, model accuracy, calibration, analytical validity, clinical usefulness, and pharmaceutical decision value are related but distinct outcomes. A synthetic dataset may reproduce marginal distributions yet distort rare safety events or relationships among covariates. A federated model may improve average prediction while performing poorly for one organization, chemical series, patient group, or assay family. Evaluation must therefore be aligned with the downstream purpose and should examine fidelity, analytical utility, disclosure risk, stability, subgroup performance, and sensitivity to alternative validation designs. Methodological evaluations of synthetic patient data support this multidimensional approach and show why resemblance to source records cannot function as a sufficient indicator

of either privacy or usefulness [9]. The same principle applies to collaborative modeling: a result is scientifically meaningful only when the evaluated metric corresponds to the decision the model is expected to inform.

A fourth boundary separates formal protection from practical security and organizational trustworthiness. Secure multiparty computation and homomorphic encryption can provide mathematically specified confidentiality for particular computations, but the guarantee remains conditional on the protocol, key distribution, implementation, adversary model, collusion assumptions, communication architecture, and information revealed through outputs. Multiparty homomorphic federated analytics illustrate how distributed biomedical analyses can be executed under stronger cryptographic protection while also exposing practical dependencies on computational overhead and system design [10]. Accordingly, the present review maps privacy claims according to their stated assumptions instead of assigning a generic "privacy-preserving" label. It also distinguishes technical confidentiality from governance, interoperability, accountability, legal responsibility, and scientific validity. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop evidence-mapping questions and conceptual boundaries within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Evidence-mapping questions and conceptual boundaries

Evidence domain	Eligible evidence or method	Reported capability or claim	Strength of support	Reproducibility or bias concern	Unresolved gap
Collaborative architecture	Federated learning, split computation, distributed query execution, secure aggregation	Enables coordinated analysis while retaining source records within participating environments	Strong support for technical feasibility; variable support for pharmaceutical usefulness	Public-data partitioning may be presented as equivalent to real organizational federation; client distributions and local workflows may be incompletely described	Independent validation across genuinely distinct companies, laboratories, data models, and decision environments
Formal privacy protection	Secure multiparty computation, homomorphic encryption, differential privacy, trusted execution, cryptographic aggregation	Reduces specified disclosure risks under defined threat and trust assumptions	Strong theoretical support for bounded guarantees; implementation evidence remains context-dependent	Privacy parameters, collusion assumptions, implementation details, and residual output leakage may be underreported	Standardized threat-model reporting and independent adversarial testing under realistic infrastructure
Synthetic-data release	Generative models, disclosure controls, distributional and task-based validation	Produces transferable records intended to preserve selected analytical properties while limiting disclosure	Moderate and heterogeneous support across datasets and downstream tasks	Fidelity metrics may obscure rare-event distortion, memorization, subgroup error, or task-specific utility loss	Fit-for-purpose release criteria combining disclosure attacks, analytical validity, and intended-use restrictions
Pharmaceutical discovery	Molecular-property prediction, quantitative structure–activity relationships, phenotypic modeling, computational toxicology	Expands collaborative training data and may improve selected predictive tasks	Strongest for retrospective modeling and benchmarked prediction; weaker for experimental consequences	Simulated partners, inconsistent endpoint definitions, chemical-series leakage, and limited prospective testing	Evidence that collaborative models alter experiment selection, candidate prioritization, or reproducible laboratory outcomes
Development and safety	Distributed clinical analysis, trial-support analytics, pharmacoepidemiology, pharmacovigilance	Supports multi-source evidence generation without conventional record pooling	Moderate support for distributed analysis; uneven support for decision-integrated machine learning	Variable coding, local data quality, missingness, confounding, and inconsistent outcome definitions may limit comparability	Prospective validation of scientific, operational, and safety-decision consequences
Utility and uncertainty	Discrimination, calibration, task utility, subgroup analysis,	Demonstrates whether protected collaboration	Frequently reported for benchmark	Aggregate metrics can hide partner-specific	Joint reporting of privacy parameters,

	robustness, external validation	retains adequate analytical performance	accuracy; less consistently reported for calibration and decision utility	degradation, distribution shift, or privacy-induced uncertainty	calibration, transportability, subgroup performance, and decision thresholds
Governance and interoperability	FAIR data practices, common data models, provenance, access control, accountability, audit and incident procedures	Enables technically valid and organizationally accountable collaboration	Strong conceptual support; limited comparative evaluation of governance arrangements	Governance may be treated as administrative context rather than part of system validity	Evaluated governance models covering contribution, withdrawal, liability, monitoring, versioning, and benefit allocation

Search, Classification, and Charting Methodology

The review was designed as an evidence map rather than an effectiveness meta-analysis. Search concepts combined privacy-preserving computation with pharmaceutical discovery, molecular modeling, biomedical data science, clinical development, distributed evidence generation, and drug safety. Eligible evidence was restricted to peer-reviewed journal articles with sufficient methodological or conceptual detail to support a preassigned manuscript claim. Articles were excluded when they mentioned privacy-preserving collaboration only superficially, lacked a direct connection to a planned evidence domain, or reported a narrow benchmark that could not be interpreted in relation to validity, privacy assumptions, utility, or pharmaceutical relevance. Reporting was structured around the transparency principles used for scoping reviews, including explicit questions, eligibility logic, information sources, charting fields, and synthesis boundaries [11]. No pooled effectiveness estimate was attempted because the selected literature differed substantially in technology, data type, application, comparison, privacy mechanism, and validation design.

Classification dimensions were defined before narrative synthesis to reduce post hoc category drift. Each article was charted by application context, pharmaceutical stage, technology family, protected asset, data location, exchanged artifact, collaboration topology, threat model, validation setting, privacy claim, utility measure, reproducibility concern, maturity, and translational limitation. The technology categories comprised federated learning, secure computation, synthetic-data generation, distributed analytics, and layered combinations of privacy-enhancing technologies. Maturity was interpreted qualitatively through conceptual or protocol rationale, simulated proof-of-concept, retrospective multi-site or multi-dataset validation, cross-organizational operational piloting, and routine decision-integrated use. These categories were informed by updated methodological guidance recommending that scoping-review concepts, boundaries, and charting decisions be established in advance and refined transparently when the evidence requires clarification [12]. They are proposed organizing categories, not validated scales, and they do not predict implementation success.

Evidence extraction preserved differences among what a study attempted, what it demonstrated, and what could reasonably be concluded. Reported findings were charted together with application context, data distribution, technical architecture, validation design, quality concern, and contribution to the relevant review question. “Not clearly reported” was retained where an article did not provide enough information, rather than filling missing categories by inference. Critical appraisal focused on risks especially relevant to privacy-preserving collaboration: simulated rather than authentic organizational separation, unclear adversary assumptions, inadequate attack testing, unreported privacy parameters, utility assessment limited to aggregate accuracy, missing calibration, weak external validation, semantic incompatibility, and absence of prospective decision-impact evaluation. The synthesis then organized convergent findings, divergent findings, contextual explanations, uncertainties, and evidence gaps without converting the map into an unsupported comparative ranking. This approach follows recommendations that scoping-review results should preserve heterogeneity and present evidence distributions and uncertainties rather than imply effectiveness where comparative evidence is unavailable [13].

Federated Learning, Secure Computation, Synthetic Data, and Distributed Analytics

Federated learning coordinates model development across data holders by moving model-derived information rather than conventionally pooling complete source datasets. In pharmaceutical research, this architecture is attractive when compound structures, assay outcomes, clinical observations, or safety records are commercially or legally difficult to centralize. A federated quantitative structure–activity relationship prototype demonstrated that molecular-property models can be trained across separated data partitions, establishing technical feasibility for collaborative prediction [14]. Nevertheless, simulated client partitions do not reproduce every feature of company-to-company collaboration. Endpoint definitions, chemical-series composition, assay protocols, missingness, infrastructure, and incentives may differ substantially among actual partners. Federated feasibility should therefore be interpreted as an early evidence layer rather than as proof of cross-company usefulness.

Secure computation addresses a different problem by protecting inputs or intermediate values during jointly specified calculations. Secure multiparty computation can allow organizations to evaluate collaborative drug-discovery procedures

without revealing their original inputs to one another, although the feasible computation is constrained by the selected protocol and its communication and runtime requirements [15]. This approach may provide stronger confidentiality for defined operations than ordinary parameter exchange, but it does not automatically establish correctness of the source data, relevance of the analytical task, or usefulness of the output. Secure computation should consequently be charted according to the protected object, participating parties, collusion assumptions, permitted outputs, computational burden, and whether the analysis was independently reproduced under realistic infrastructure.

Data heterogeneity is a central limitation across decentralized learning. Pharmaceutical partners rarely hold random samples from one common distribution; instead, they may investigate different chemical series, targets, populations, assays, indications, or development stages. Federated learning on non-identically distributed drug-discovery data shows that aggregation behavior and local performance can change when partner datasets differ systematically [16]. A globally improved model may therefore coexist with deterioration for one contributor or one scientific domain. Appropriate evaluation requires partner-level performance, calibration, transportability, uncertainty, and sensitivity to the aggregation rule. Data locality cannot compensate for incompatible labels, inconsistent measurements, or an analysis that combines scientifically non-comparable tasks.

Synthetic data and layered privacy-enhancing architectures extend the design space beyond federated parameter exchange. Comparative benchmarking of synthetic health-record generators indicates that privacy, fidelity, and analytical utility vary by model and evaluation context, with no single metric establishing safe or useful release [17]. End-to-end privacy-preserving learning may combine federation with encryption, secure aggregation, and differential privacy, thereby reducing several exposure routes while increasing computational complexity and potentially altering model utility [18]. These technology families should not be ranked through a single privacy label. Their suitability depends on whether the intended activity is model training, aggregate analysis, external data release, exploratory collaboration, or regulated evidence generation, as well as on the adversary, output, and scientific consequences of error.

Pharmaceutical Use Cases across Discovery, Development, and Safety

Within discovery, molecular-property prediction is the most developed application cluster. Federated graph neural networks have been evaluated for learning molecular properties from heterogeneous client datasets, supporting the possibility that organizations can contribute complementary chemical evidence without disclosing their complete molecular collections [19]. However, the dominant evidence remains retrospective and benchmark-centered. Predictive improvement does not establish that generated rankings identify experimentally active compounds, preserve applicability across chemical neighborhoods, or improve synthesis and testing decisions. Discovery studies should therefore report chemical split design, partner heterogeneity, applicability domain, calibration, local benefit, and the relation between the modeled endpoint and the next experimental decision.

Industrial collaboration provides stronger implementation evidence than simulated federation because genuine partners contribute proprietary tasks, infrastructure, and governance constraints. Federated learning with knowledge distillation has been applied to collaborative drug-discovery modeling, expanding the range of architectures available when conventional shared-parameter training is undesirable or impractical [20]. Its value remains conditional on task similarity, teacher quality, endpoint definition, baseline selection, and the information transferred through predictions or distilled representations. Such work supports operational feasibility, but prospective evidence is still needed to determine whether collaboration changes compound selection, experiment prioritization, project attrition, or reproducibility across companies.

In development, privacy-preserving collaboration may connect trial, biomarker, sensor, imaging, and observational evidence held by different stakeholders. Proposed federated approaches for Parkinson's disease drug development illustrate how distributed datasets could support disease characterization, endpoint development, patient stratification, and longitudinal analysis without requiring a single shared repository [21]. The evidence nevertheless remains uneven across these uses. A technically successful model does not establish that a biomarker is clinically meaningful, that a treatment effect is causal, or that an analytical output improves trial design. Development applications require prospective protocols, predefined decision roles, external validation, uncertainty assessment, and documentation of how outputs would influence scientific or operational choices.

Drug safety offers a comparatively mature setting for distributed analysis because multi-database networks already rely on common protocols and local execution. Machine learning may extend these networks, but its interpretability depends on harmonized definitions, appropriate confounding control, transportability, and transparent local data processing [22]. Federated analyses across multiple drug-safety sources further demonstrate that participating databases can contribute results without transferring complete patient-level records [23]. These implementations support practical collaboration, yet they do not eliminate coding differences, missingness, outcome misclassification, or variation in healthcare practice. Safety claims must therefore distinguish signal detection, association, causal assessment, and decision-ready evidence.

Privacy Guarantees, Utility Loss, Governance, and Attack Surfaces

The first interpretive requirement is to avoid equating local data retention with complete privacy. Federated systems remain exposed to model poisoning, malicious updates, inference attacks, compromised aggregation, membership disclosure, and infrastructure failure [24]. Risks differ according to whether the server is trusted, clients can collude, updates are visible, participation is observable, or repeated queries are permitted. A defensible privacy claim must identify the protected asset and adversary, rather than merely state that raw data were not centralized. Security evaluation should also consider insider threats,

implementation errors, model outputs, logs, and metadata that may reveal organizational behavior or sensitive scientific activity.

Gradient leakage illustrates why exchanged model information can remain sensitive. Under some conditions, gradients or updates can be exploited to infer features of training records or reconstruct aspects of local inputs [25]. The practical severity of this risk depends on model architecture, batch size, update frequency, optimization, attacker access, and available auxiliary information. Consequently, an attack demonstrated in one configuration cannot be generalized automatically to every pharmaceutical federation, but neither can the risk be dismissed because source files remain local. Collaboration agreements should define update visibility, aggregation protection, retention, monitoring, red-team testing, and response procedures for suspected leakage.

Differential privacy provides a formal framework for limiting the influence of individual records or contributions, but the resulting protection is parameter-dependent and can reduce analytical utility. Evaluations in medical deep learning demonstrate that privacy-preserving noise may affect model performance differently across tasks and privacy settings [26]. Reporting only final accuracy is therefore inadequate. Studies should specify the privacy accounting method, protected unit, privacy parameters, training duration, clipping procedure, calibration, subgroup effects, and uncertainty. In pharmaceutical contexts, the protected unit may be a patient, compound, assay, experimental batch, institution, or project; these choices are not interchangeable and determine what the guarantee actually means.

Technical safeguards must be embedded within organizational governance. Healthcare-focused analyses of federated privacy emphasize that secure aggregation, differential privacy, encryption, access control, monitoring, and policy should be treated as complementary layers rather than substitutes for one another [27]. Governance must specify participation, authorization, permitted uses, model ownership, contribution, benefit allocation, withdrawal, incident handling, auditing, versioning, and responsibility for harmful or misleading outputs. These requirements also interact with utility: a highly protected system that cannot answer the scientific question is not useful, whereas a high-performing system without defensible controls is not trustworthy. Privacy and utility should therefore be evaluated jointly within the intended decision context.

Evidence Map of Implementation Maturity and Validation Quality

Implementation maturity should be distinguished from algorithmic sophistication. Reviews of federated-learning implementation identify persistent barriers involving data harmonization, infrastructure, workflow integration, security, law, incentives, maintenance, and organizational coordination [28]. The proposed evidence map therefore separates conceptual rationale, simulated proof-of-concept, retrospective multi-site validation, authentic cross-organizational piloting, and routine decision-integrated use. These categories describe the reported setting rather than assigning a validated readiness score. Progression between them is not automatic: a technically strong method may remain unsuitable because the collaborating data are incompatible, the threat model is incomplete, or the analytical result lacks a defined pharmaceutical use.

The mapped literature is heterogeneous in data type, model architecture, collaboration pattern, privacy mechanism, and validation practice. A scoping review of federated and distributed learning for structured medical data found substantial variation in application context and technical design, limiting direct comparison across studies [29]. Similar heterogeneity is expected in pharmaceutical research, where molecular, clinical, omics, and safety datasets differ fundamentally. Validation quality must therefore be interpreted relative to the claim. A public benchmark may support algorithmic comparison, whereas cross-company usefulness requires genuine organizational separation, protected execution, partner-level evaluation, reproducibility, and evidence that the collaborative result is relevant to an actual decision.

Clinical and technical reviews further indicate that retrospective performance should not be treated as a substitute for external validity, fairness, security, or prospective usefulness [30]. For pharmaceutical applications, validation should examine whether an approach performs across organizations, populations, chemical spaces, instruments, laboratories, and time. It should also test whether privacy protections alter calibration, rare-event detection, subgroup reliability, or local benefit. The highest evidentiary layer would require prospective evaluation of how the collaborative analysis affects experiment selection, trial conduct, safety assessment, or another predefined decision. Even then, success in one use case would not establish universal readiness.

The available evidence supports a broad technical opportunity space, but surveys continue to identify unresolved needs in reproducibility, security testing, standardization, and cross-site validation [31]. The article therefore proposes a privacy–utility collaboration airlock in which a claim must pass through linked checks for question definition, data suitability, protected execution, utility, generalizability, operational integration, and accountable governance. The airlock is an original conceptual synthesis and has not been empirically validated as a predictive maturity instrument. **Figure 1** presents the privacy–utility collaboration airlock, showing how the article's key components and boundaries are connected within evidence map of implementation maturity and validation quality.

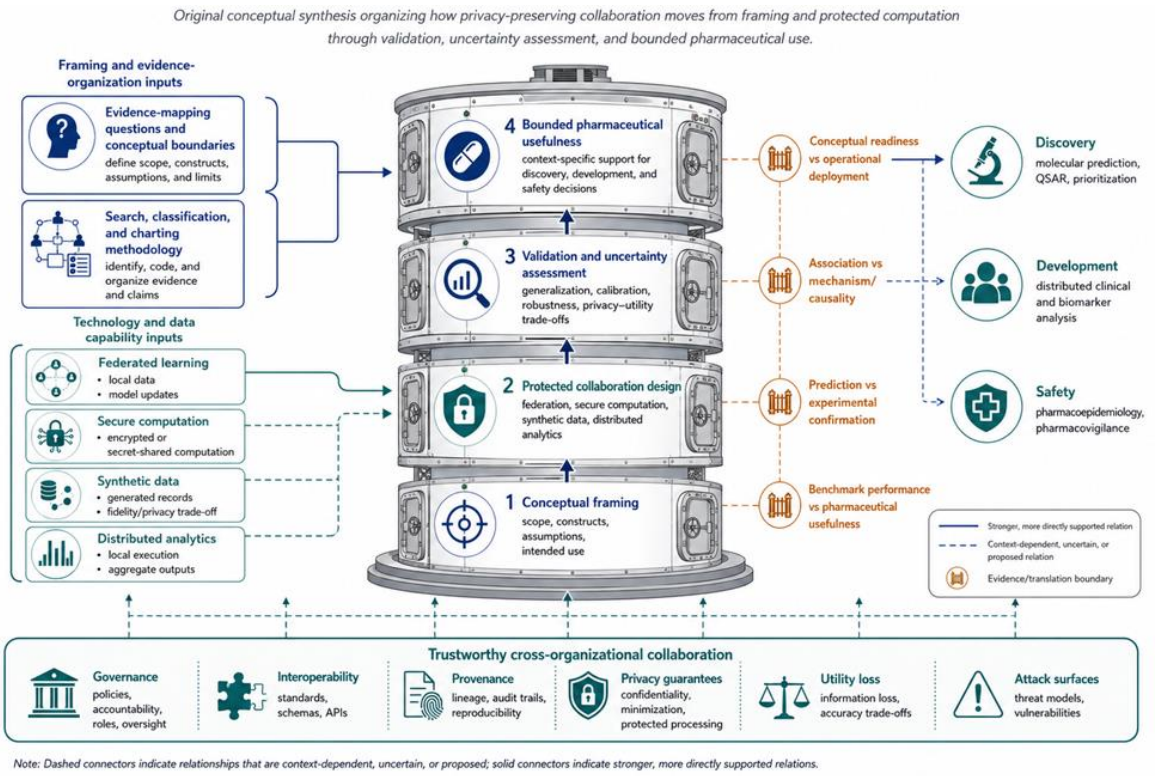

Figure 1. Privacy–Utility Collaboration Airlock

The figure is an original conceptual synthesis that organizes the article's central contribution across evidence-mapping questions and conceptual boundaries, search, classification, and charting methodology, charting methodology, federated learning, secure computation, synthetic data, and distributed analytics, federated learning, secure computation. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Priorities for Trustworthy Cross-Organizational Collaboration

The first priority is to treat governance as part of analytical validity rather than as an external administrative layer. A scoping review of federated-learning governance in healthcare identifies recurring requirements involving roles, accountability, control, participation, access, oversight, and lifecycle management [32]. Pharmaceutical consortia additionally require arrangements for commercially sensitive endpoints, contribution asymmetry, intellectual property, model ownership, withdrawal, and benefit allocation. Governance should be specified before training begins and tested through realistic scenarios, including partner departure, security incidents, disputed outputs, model updates, and changes in permitted use. Comparative evaluation of governance models remains an important research need.

The second priority is to improve data suitability, metadata, and provenance. FAIR principles in biopharmaceutical research emphasize that data should be findable, accessible under appropriate authorization, interoperable, and reusable [33]. Privacy-preserving computation does not correct mislabeled assays, incompatible units, undocumented transformations, population imbalance, or missing provenance. Before collaboration, partners should establish common definitions, quality rules, transformation records, endpoint semantics, and procedures for handling unresolved incompatibility. Where full harmonization is impossible, the model should preserve site-specific meaning and report uncertainty rather than forcing apparently uniform variables into an invalid pooled representation.

The third priority is semantic and technical interoperability. Digital medicine depends on shared representations and communication standards because data held in different systems cannot support valid joint interpretation merely by remaining local [34]. Pharmaceutical collaboration requires interoperable molecular identifiers, assay vocabularies, clinical concepts, temporal definitions, outcome models, and provenance structures. Common data models should be evaluated for information loss and local mapping error, particularly when rare events or specialized laboratory measurements are important. Interoperability should therefore be validated scientifically, not inferred from successful data transmission or software connectivity.

The fourth priority is alignment between privacy-enhancing technology, legal responsibility, and scientific purpose. Reviews of federated machine learning, privacy-enhancing technologies, and data-protection law show that technical decentralization does not remove obligations associated with sensitive model artifacts, processing purposes, participant rights, accountability, and cross-organizational control [35]. A trustworthy collaboration should document its protected objects, lawful and permitted uses, adversaries, safeguards, residual risks, validation plan, model lineage, monitoring, and withdrawal criteria. The proposed

priorities are not a universal compliance framework; they organize the evidence needed for context-specific assessment. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop priorities for trustworthy cross-organizational collaboration within the article's central argument.

Table 2. Evaluation, implementation, and research priorities arising from Learning without Sharing the Data: An Evidence Map of Privacy-Preserving Collaboration in Pharmaceutical Research

Approach or application	Data and task context	Evaluation practice	Evidence maturity	Translational limitation	Review interpretation
Federated molecular-property learning	Distributed chemical structures, assay endpoints, and molecular-property tasks	Chemical-series-aware splits, partner-level performance, calibration, applicability-domain analysis, and external validation	Predominantly proof-of-concept to retrospective multi-partner validation, with selected operational pilots	Benchmark improvement does not establish experimental activity, synthesizability, developability, or decision impact	Promising for collaborative prediction when endpoint semantics and chemical distributions are compatible
Cross-company quantitative structure–activity relationship modeling	Proprietary assay collections held by pharmaceutical partners	Authentic organizational execution, local and global baselines, endpoint-specific analysis, and reproducibility checks	Operational cross-organizational piloting	Benefits may be uneven across endpoints and partners; prospective project impact remains uncertain	Strong implementation evidence, but not proof of routine discovery readiness
Secure multiparty and homomorphic computation	Joint analysis of confidential molecular, clinical, genomic, or safety data	Formal protocol verification, adversary specification, runtime, communication, collusion testing, and output-leakage analysis	Technical proof-of-concept to retrospective distributed validation	Computational burden and restricted analytical flexibility may limit practical use	Stronger bounded confidentiality than data locality alone when assumptions are explicit
Differentially private collaborative learning	Patient-, compound-, assay-, or institution-level contribution protection	Privacy accounting, protected-unit definition, calibration, subgroup utility, and sensitivity to privacy parameters	Retrospective technical validation	Utility loss and uncertainty are task-dependent; privacy parameters may be difficult to interpret operationally	Useful when the protected unit and acceptable utility loss are defined prospectively
Synthetic-data generation and release	Transferable patient, trial, omics, or safety-like records	Distributional fidelity, downstream analytical validity, rare-event preservation, and disclosure attacks	Predominantly retrospective methodological and benchmarking evidence	Similarity does not establish privacy, and privacy does not establish fitness for every downstream use	Release should be restricted to validated tasks and supported by attack testing
Distributed pharmacoepidemiology and pharmacovigilance	Multiple healthcare databases using common protocols and local execution	Semantic validation, local quality control, confounding assessment, transportability, and reproducibility	Retrospective multi-source implementation	Coding variation, missingness, healthcare-system differences, and outcome misclassification remain	Among the more operationally grounded collaboration models, but causal and decision claims remain bounded
Governance, FAIR data, and interoperability	Cross-company, institutional, and jurisdictional collaboration	Role definition, provenance auditing, mapping validation, incident simulation, withdrawal testing, and lifecycle monitoring	Conceptual guidance with emerging implementation evidence	Governance outcomes and interoperability failures are rarely evaluated comparatively	Necessary components of scientific validity and sustainable collaboration
Prospective decision-integrated validation	Experiment selection, trial	Prospective protocol,	Limited	Cost, organizational	Highest-priority gap for separating

	operations, safety assessment, or other predefined pharmaceutical decisions	predefined success and failure criteria, independent replication, uncertainty reporting, and documented decision consequences	complexity, and ethical or legal constraints may restrict evaluation	plausible utility from demonstrated pharmaceutical value

Figure 2 presents the evidence, validation, and translation boundary observatory for learning without sharing the data an evidence, showing how the article's key components and boundaries are connected within priorities for trustworthy cross-organizational collaboration.

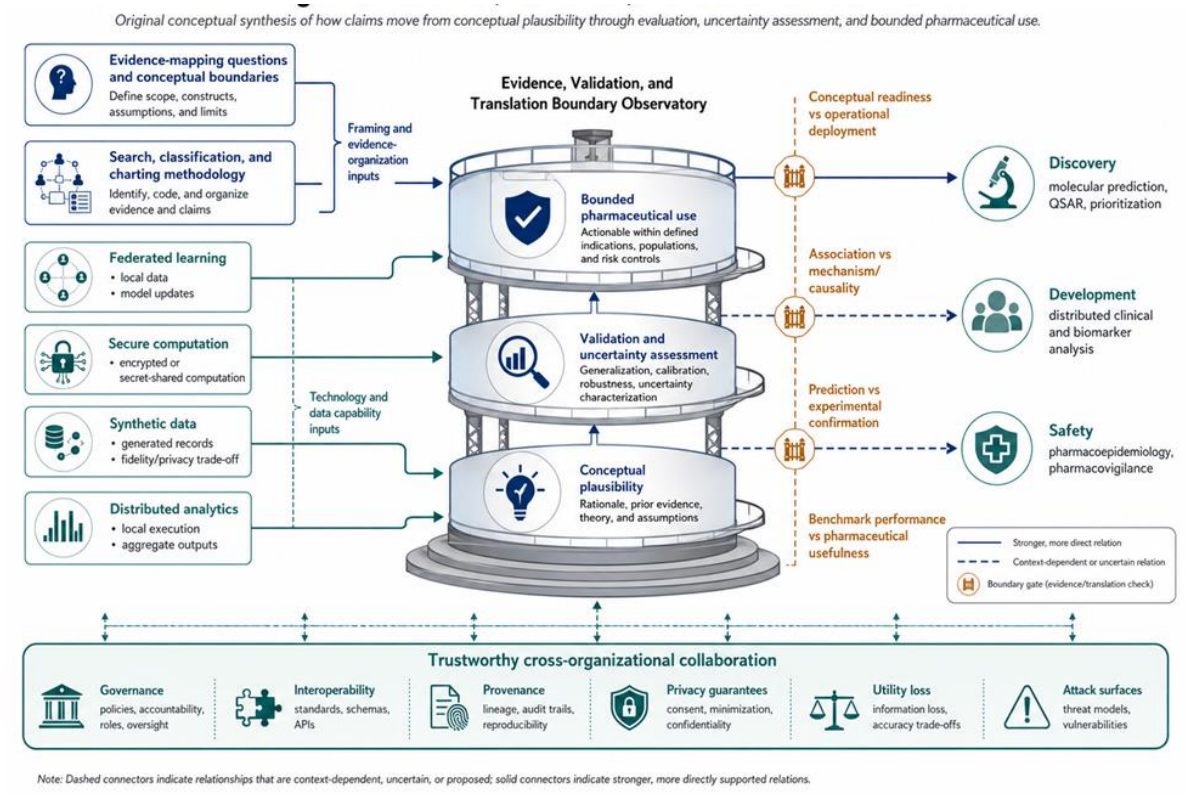

Figure 2. Evidence, Validation, and Translation Boundaries

The figure is an original conceptual synthesis that shows how claims move from conceptual plausibility through evaluation, uncertainty assessment, and bounded pharmaceutical use. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Conclusion

Privacy-preserving collaboration can expand the evidence available to pharmaceutical research without requiring conventional pooling of every source dataset, but decentralization is not synonymous with privacy, scientific validity, or readiness. This evidence map distinguishes federated learning, secure computation, synthetic data, and distributed analytics according to what remains local, what is exchanged, what is protected, how utility is evaluated, and which risks remain. The strongest evidence supports technical feasibility, retrospective distributed modeling, and selected operational collaborations. More limited evidence addresses prospective pharmaceutical decisions, independent attack testing, calibrated uncertainty, semantic compatibility, governance performance, lifecycle monitoring, and routine integration. The proposed maturity map, privacy–utility collaboration airlock, and translation-boundary observatory are organizing syntheses rather than validated instruments. Their purpose is to prevent benchmark performance or data locality from being interpreted as proof of experimental, clinical, regulatory, or operational value. Trustworthy collaboration will depend on aligning the scientific question, data suitability, threat model, privacy mechanism, utility requirement, interoperability, governance, and decision consequences before claims of pharmaceutical readiness are made.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5
2. Rieke N, Hancox J, Li W, Milletarì F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119. doi:10.1038/s41746-020-00323-1
3. Kaissis GA, Makowski MR, Rückert D, Braren RF. Secure, privacy-preserving and federated machine learning in medical imaging. Nat Mach Intell. 2020;2(6):305-11. doi:10.1038/s42256-020-0186-1
4. Smajić A, Grandits M, Ecker GF. Privacy-preserving techniques for decentralized and secure machine learning in drug discovery. Drug Discov Today. 2023;28(12):103820. doi:10.1016/j.drudis.2023.103820
5. Heyndrickx W, Mervin L, Morawietz T, Sturm N, Friedrich L, Zalewski A, et al. MELLODDY: Cross-pharma federated learning at unprecedented scale unlocks benefits in QSAR without compromising proprietary information. J Chem Inf Model. 2024;64(7):2331-44. doi:10.1021/acs.jcim.3c00799
6. Hanser T. Federated learning for molecular discovery. Curr Opin Struct Biol. 2023;79:102545. doi:10.1016/j.sbi.2023.102545
7. Sheller MJ, Edwards B, Reina GA, Martin J, Pati S, Kotrotsou A, et al. Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. Sci Rep. 2020;10(1):12598. doi:10.1038/s41598-020-69250-1
8. Chen RJ, Lu MY, Chen TY, Williamson DFK, Mahmood F. Synthetic data in machine learning for medicine and healthcare. Nat Biomed Eng. 2021;5(6):493-7. doi:10.1038/s41551-021-00751-8
9. Goncalves A, Ray P, Soper B, Stevens J, Coyle L, Sales AP. Generation and evaluation of synthetic patient data. BMC Med Res Methodol. 2020;20(1):108. doi:10.1186/s12874-020-00977-1
10. Froelicher D, Troncoso-Pastoriza JR, Raisaro JL, Cuendet MA, Sousa JS, Cho H, et al. Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption. Nat Commun. 2021;12(1):5910. doi:10.1038/s41467-021-25972-y
11. Tricco AC, Lillie E, Zarin W, O'Brien KK, Colquhoun H, Levac D, et al. PRISMA extension for scoping reviews (PRISMA-ScR): Checklist and explanation. Ann Intern Med. 2018;169(7):467-73. doi:10.7326/M18-0850
12. Peters MDJ, Marnie C, Tricco AC, Pollock D, Munn Z, Alexander L, et al. Updated methodological guidance for the conduct of scoping reviews. JBI Evid Synth. 2020;18(10):2119-26. doi:10.11124/JBIES-20-00167
13. Pollock D, Peters MDJ, Khalil H, McInerney P, Alexander L, Tricco AC, et al. Recommendations for the extraction, analysis, and presentation of results in scoping reviews. JBI Evid Synth. 2023;21(3):520-32. doi:10.11124/JBIES-22-00123
14. Chen S, Xue D, Chuai G, Yang Q, Liu Q. FL-QSAR: A federated learning-based QSAR prototype for collaborative drug discovery. Bioinformatics. 2020;36(22-23):5492-8. doi:10.1093/bioinformatics/btaa1006
15. Ma R, Li Y, Li C, Wan F, Hu H, Xu W, et al. Secure multiparty computation for privacy-preserving drug discovery. Bioinformatics. 2020;36(9):2872-80. doi:10.1093/bioinformatics/btaa038
16. Huang D, Ye X, Zhang Y, Sakurai T. Collaborative analysis for drug discovery by federated learning on non-IID data. Methods. 2023;219:1-7. doi:10.1016/j.ymeth.2023.09.001
17. Yan C, Yan Y, Wan Z, Zhang Z, Omberg L, Guinney J, et al. A multifaceted benchmarking of synthetic electronic health record generation models. Nat Commun. 2022;13(1):7609. doi:10.1038/s41467-022-35295-1
18. Kaissis G, Ziller A, Passerat-Palmbach J, Ryffel T, Usynin D, Trask A, et al. End-to-end privacy preserving deep learning on multi-institutional medical imaging. Nat Mach Intell. 2021;3(6):473-84. doi:10.1038/s42256-021-00337-8
19. Zhu W, Luo J, White AD. Federated learning of molecular properties with graph neural networks in a heterogeneous setting. Patterns (N Y). 2022;3(6):100521. doi:10.1016/j.patter.2022.100521
20. Hanser T, Ahlberg E, Amberg A, Anger LT, Barber C, Brennan RJ, et al. Data-driven federated learning in drug discovery with knowledge distillation. Nat Mach Intell. 2025;7(3):423-36. doi:10.1038/s42256-025-00991-2
21. Khanna A, Adams J, Antoniades C, Bloem BR, Carroll C, Cedarbaum J, et al. Accelerating Parkinson's disease drug development with federated learning approaches. NPJ Parkinsons Dis. 2024;10(1):225. doi:10.1038/s41531-024-00837-5
22. Wong J, Prieto-Alhambra D, Rijnbeek PR, Desai RJ, Reps JM, Toh S. Applying machine learning in distributed data networks for pharmacoepidemiologic and pharmacovigilance studies: Opportunities, challenges, and considerations. Drug Saf. 2022;45(5):493-510. doi:10.1007/s40264-022-01158-3

23. Gedeborg R, Igl W, Svennblad B, Wilén P, Delcoigne B, Michaëlsson K, et al. Federated analyses of multiple data sources in drug safety studies. Pharmacoepidemiol Drug Saf. 2023;32(3):279-86. doi:10.1002/pds.5587
24. Mothukuri V, Parizi RM, Pouriyeh S, Huang Y, Dehghantanha A, Srivastava G. A survey on security and privacy of federated learning. Future Gener Comput Syst. 2021;115:619-40. doi:10.1016/j.future.2020.10.007
25. Yang H, Ge M, Xue D, Xiang K, Li H, Lu R. Gradient leakage attacks in federated learning: Research frontiers, taxonomy, and future directions. IEEE Netw. 2024;38(2):247-54. doi:10.1109/MNET.001.2300140
26. Ziller A, Usynin D, Braren R, Makowski M, Rueckert D, Kaissis G. Medical imaging deep learning with differential privacy. Sci Rep. 2021;11(1):13524. doi:10.1038/s41598-021-93030-0
27. Pati S, Kumar S, Varma A, Edwards B, Lu C, Qu L, et al. Privacy preservation for federated learning in health care. Patterns (N Y). 2024;5(7):100974. doi:10.1016/j.patter.2024.100974
28. Li M, Xu P, Hu J, Tang Z, Yang G. From challenges and pitfalls to recommendations and opportunities: Implementing federated learning in healthcare. Med Image Anal. 2025;101:103497. doi:10.1016/j.media.2025.103497
29. Li S, Liu P, Nascimento GG, Wang X, Leite FRM, Chakraborty B, et al. Federated and distributed learning applications for electronic health records and structured medical data: A scoping review. J Am Med Inform Assoc. 2023;30(12):2041-9. doi:10.1093/jamia/ocad170
30. Teo ZL, Jin L, Liu N, Li S, Miao D, Zhang X, et al. Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture. Cell Rep Med. 2024;5(2):101419. doi:10.1016/j.xcrm.2024.101419
31. Madathil NT, Dankar FK, Gergely M, Belkacem AN, Alrabaee S. Revolutionizing healthcare data analytics with federated learning: A comprehensive survey of applications, systems, and future directions. Comput Struct Biotechnol J. 2025;28:217-38. doi:10.1016/j.csbj.2025.06.009
32. Eden R, Chukwudi I, Bain C, Barbieri S, Callaway L, de Jersey S, et al. A scoping review of the governance of federated learning in healthcare. NPJ Digit Med. 2025;8(1):427. doi:10.1038/s41746-025-01836-3
33. Wise J, Grebe de Barron A, Splendiani A, Balali-Mood B, Vasant D, Little E, et al. Implementation and relevance of FAIR data principles in biopharmaceutical R&D. Drug Discov Today. 2019;24(4):933-8. doi:10.1016/j.drudis.2019.01.008
34. Lehne M, Sass J, Essenwanger A, Schepers J, Thun S. Why digital medicine depends on interoperability. NPJ Digit Med. 2019;2:79. doi:10.1038/s41746-019-0158-1
35. Brauneck A, Schmalhorst L, Kazemi Majdabadi MM, Bakhtiari M, Völker U, Baumbach J, et al. Federated machine learning, privacy-enhancing technologies, and data protection laws in medical research: Scoping review. J Med Internet Res. 2023;25. doi:10.2196/41588