INTERNATIONAL JOURNAL OF BUSINESS AND ECONOMICS INSIGHTS

Volume: 5; **Issue:** 3
Pages: 65–99
Accepted: 29 August 2025
Published: 15 September 2025

ARTIFICIAL INTELLIGENCE BASED MODELS FOR SECURE DATA ANALYTICS AND PRIVACY-PRESERVING DATA SHARING IN U.S. HEALTHCARE AND HOSPITAL NETWORKS

Saba Ashfaq[1]

[1]. *MS IT - Software Design and Management: Washington University of Science and Technology, USA; Email: sabarashfaq01@gmail.com*

Doi: 10.63125/wv0bqx68
This work was peer-reviewed under the editorial responsibility of the IJEI, 2025

Abstract

This study addresses a practical problem in U.S. hospital networks, namely how to realize the benefits of AI-enabled analytics while preserving privacy and strengthening interorganizational data sharing. Grounded in a targeted review of 48 peer-reviewed sources, we quantify how AI configuration and privacy-enhancing techniques relate to analytic utility, exchange quality, and security posture. The purpose is to deliver precise, comparable estimates that guide hospital leaders toward secure, high-value deployments. Using a quantitative, cross-sectional, case-based design, we analyze a purposive multi-case sample targeting approximately 200 to 300 acute-care hospitals that operate enterprise EHR ecosystems and cloud-based analytics pipelines. Key variables include AI model family, architectural complexity, pretraining, update cadence, a privacy-preserving technique maturity index spanning differential privacy, federated learning with secure aggregation, homomorphic encryption, and secure multi-party computation, governance maturity, connectivity, analytic utility metrics such as AUC, PR-AUC, F1, calibration, operational cost such as latency, training time, compute, exchange quality such as partner breadth, match rate, completeness, SLA adherence, and security outcomes such as incident frequency, breach occurrence, and mean time to detect. The pre-registered analysis plan applies descriptive statistics, correlation screening, and regression models aligned to outcome scales with cluster-robust standard errors, multiple imputation for covariates, a propensity approach for adoption endogeneity, and explicit tests of mediation by utility and moderation by governance using bootstrap contrasts. Headline findings show that higher privacy-tech maturity is positively associated with analytic utility and exchange quality, with utility partially mediating the privacy-to-exchange link and governance maturity amplifying effects; small latency costs are manageable with engineering tactics, and breach risk is lower in high-maturity, well-governed settings. Implications are that privacy engineering, parameter transparency, and routinized governance convert privacy from a brake into an enabler for dependable analytics and trustworthy sharing.

Keywords

Artificial Intelligence, Privacy-Enhancing Technologies, Differential Privacy, Federated Learning, Homomorphic Encryption, Secure Multi-Party Computation;

INTRODUCTION

Artificial intelligence (AI) in health care can be defined as a constellation of computational techniques including machine learning (ML), deep learning, natural language processing, and probabilistic modeling that learn patterns from data to support prediction, classification, and decision support across clinical and operational workflows. Within hospital networks and health systems, AI-based models increasingly intersect with "secure data analytics," a domain concerned with extracting value from electronic health record (EHR) data, imaging archives, device streams, and administrative claims while conforming to stringent confidentiality and governance constraints (Secinaro et al., 2021). A core tension animates this domain: the same large, diverse datasets that enable more generalizable AI also raise risks of exposing protected health information and institutional knowledge if models or data flows are not appropriately safeguarded (Murdoch, 2021). Internationally, health systems seek to reconcile these objectives by embedding privacy-preserving mechanisms directly into analytic pipelines, thereby allowing data holders to collaborate at scale without surrendering data custody (Rieke et al., 2020). Operationally, the construct of "privacy-preserving data sharing" encompasses differential privacy, federated learning, homomorphic encryption, secure multi-party computation, and trusted execution environments each offering different guarantees, threat models, and cost–performance tradeoffs. In U.S. hospitals, where data fragmentation and cross-vendor interoperability barriers remain well documented, secure analytics approaches are particularly salient because they allow knowledge transfer across institutional boundaries even when conventional record exchange is imperfect. Relatedly, peer-reviewed syntheses indicate that AI can match or exceed human performance on specific imaging and signal tasks when trained and validated rigorously; the challenge is to achieve such performance with methods that protect privacy and are portable across sites (Liu et al., 2019; Rajpurkar et al., 2022).

From a policy and systems perspective, hospital networks are in the midst of two convergent transitions: a steady maturation of interoperability and a rapid expansion of AI-enabled analytics. Surveys show progress, but fewer than half of U.S. hospitals reported engaging in all four core interoperability domains (finding, sending, receiving, integrating data) as recently as 2018 (Danish & Zafor, 2022; Savi et al., 2023). Even where Health Information Exchange (HIE) organizations have flourished, the realized gains for quality and cost can be uneven and contingent on technical and governance capabilities. In parallel, the health sector continues to experience major cyber incidents and record-scale breaches, underscoring the imperative to integrate strong privacy guarantees into any cross-organizational analytic workflow (Acar et al., 2018; Danish & Kamrul, 2022; Tso et al., 2016). The result is a practical need for methods that allow multi-institutional data analysis without centralizing raw data. Federated learning (FL) directly addresses this need by training local models at each site and aggregating updates an approach validated in multi-site clinical studies (Dayan et al., 2021; Ficek et al., 2021; Jahid, 2022). Differential privacy (DP) formalizes privacy loss with a tunable budget and can be layered atop model training or query answers to limit leakage from model parameters or statistics. Homomorphic encryption (HE) enables computation on encrypted data, and secure multi-party computation (SMPC) distributes computation so that no participant learns others' inputs techniques increasingly reviewed for clinical analytics viability given falling computational costs (Hersh et al., 2015; Arifur & Noor, 2022; Tso et al., 2016).

echnically, privacy-preserving methods differ in their protection scope and statistical consequences, which has implications for quantitative study design in hospital settings. DP injects calibrated noise to bound reidentification risk, supporting safe release of summary statistics, model updates, or synthetic data, albeit with an explicit privacy–utility tradeoff that should be characterized for each analysis (Hersh et al., 2015; Hasan & Uddin, 2022; Tso et al., 2016). Recent work proposes standardized DP frameworks for epidemiological modeling and mobility data, emphasizing model-aware noise allocation and interpretability of privacy budgets (Rahaman, 2022a; Savi et al., 2023). FL avoids raw data pooling yet does not, by itself, guarantee privacy; gradient inversion and membership inference are real concerns, motivating augmentation with DP, secure aggregation, or cryptography (Rahaman, 2022b; Rieke et al., 2020). HE and SMPC, once considered impractical, now feature optimized libraries and hybrid designs that target key hospital analytics workloads (e.g., risk scoring, imaging inference) with tractable latency. Survey evidence across biomedical informatics shows a growing ecosystem of

hybrid PETs (privacy-enhancing technologies), which layer DP with FL or HE to deliver robustness under stronger adversarial models while controlling resource consumption (Hersh et al., 2015; Rahaman & Ashraf, 2022). For hospital networks, these distinctions matter because analytic outputs descriptive statistics, correlation estimates, regression coefficients must retain fidelity sufficient for clinical or operational decision-making when computed under privacy constraints.

Figure 1: AI-Driven Secure Analytics in Hospital Networks

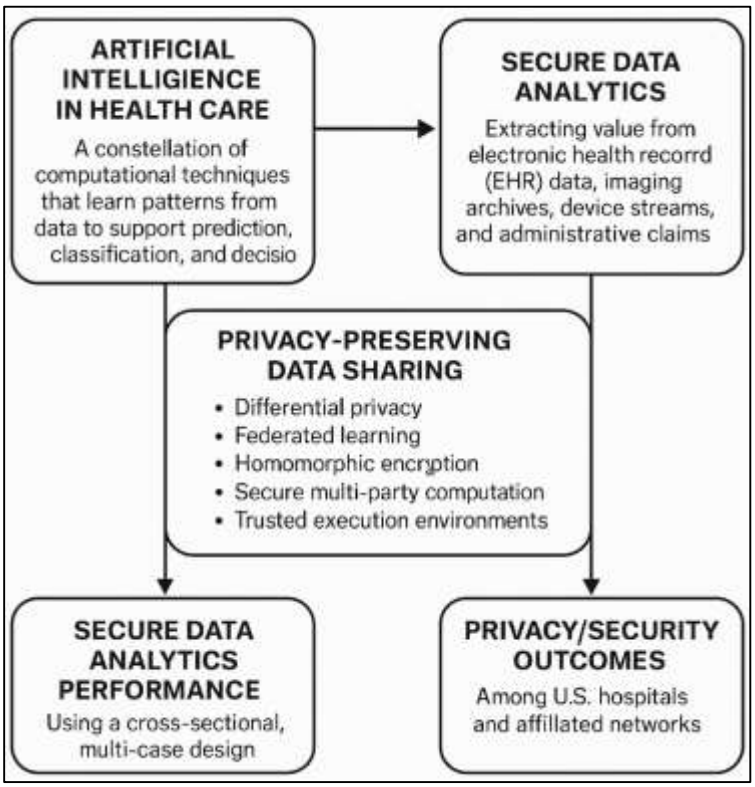

Empirically, cross-site and multi-case clinical studies demonstrate that privacy-preserving analytics can be operationalized at scale in service of clinically meaningful endpoints. A landmark international FL study across 20 institutions trained the EXAM model to predict oxygen requirements for symptomatic COVID-19 patients using chest X-rays, vital signs, and labs; the federated approach improved discrimination across sites without sharing raw data (Dayan et al., 2021; Islam, 2022). Similar paradigms have been explored for imaging-based diagnosis and segmentation (e.g., FeTS and BraTS contexts), attesting to feasibility in heterogeneous real-world data landscapes (Baid et al., 2021; Hasan et al., 2022). Contemporary implementations also target encrypted inference using HE, yielding acceptable throughput for batched diagnostic pipelines and wearable-stream classification . Meanwhile, DP has been adapted for public health surveillance and EHR-based query systems, illustrating how calibrated noise can preserve analytic utility for time-, place-, and person-level patterns while materially reducing disclosure risk (Lau et al., 2021; Liu et al., 2019; Redwanul & Zafor, 2022). Collectively, these bodies of work justify a quantitative, cross-sectional, multi-case design that measures associations among PET adoption, hospital network characteristics, and secure analytic performance outcomes using descriptive statistics, correlation analysis, and regression modeling suitable for organizational-level data.

Within U.S. healthcare and hospital networks, the institutional and market context heightens the salience of secure data analytics. Studies document persistent fragmentation across EHR vendors and organizations, with governance, cost, and data-quality barriers to exchange frequently cited by hospital leaders . Even as HIE participation increases, empirical syntheses report mixed associations with utilization or readmissions, suggesting that the mere presence of exchange infrastructure is insufficient without usable, timely analytics integrated into workflows . In this setting, PET-enabled analytics can allow cross-organizational learning and benchmarking while maintaining compliance with privacy and

security obligations (Rezaul & Mesbaul, 2022). For hospital networks operating in competitive markets or under diverse contracting arrangements, the ability to collaborate without relinquishing data custody offers a pathway to multi-site regression modeling risk adjustment, outcome monitoring, and resource planning performed with rigorous privacy guarantees (Hasan, 2022). Quantitative assessments of PET adoption can therefore foreground measurable constructs such as exchange breadth, analytic latency, model discrimination, and privacy budget parameters, creating a tractable basis for regression frameworks linking organizational features to secure-analytics performance (Tarek, 2022).

At the clinical operations interface, AI-based secure analytics address use cases central to hospital performance: early warning scores, imaging triage, throughput and bed management, care coordination, and public-health reporting. Systematic reviews show that, when evaluated transparently and prospectively, ML can deliver high discrimination in well-specified niches, particularly in medical imaging (Kamrul & Omar, 2022) . Secure analytics add an essential deployment layer by enabling multi-hospital model development without exposing patient-level data, which is important for generalizability across demographic and device variations. Survey and review articles in biomedical informatics consistently identify differential privacy, federated learning, and cryptographic methods as the most mature families for hospital-grade implementation, with hybridization increasingly common (Kamrul & Tarek, 2022; Walker & et al., 2017). For a cross-sectional, case-study–based empirical design, these literatures motivate selection of variables capturing PET type, configuration (e.g., ε for DP; aggregation frequency for FL), security posture, and model-performance metrics, enabling correlation and regression analyses that quantify associations between PET adoption and analytic outcomes while accommodating site heterogeneity.

Finally, the international significance of privacy-preserving data sharing is anchored in both normative and practical considerations. Normatively, patient confidentiality is foundational to clinical care, and technical guarantees that survive adversarial scrutiny are necessary to sustain public trust as AI permeates hospital workflows (Mubashir & Abdul, 2022; Murdoch, 2021; Zhao & et al., 2025). Practically, global consortia spanning hospitals, public-health agencies, and research networks depend on mechanisms that allow collaboration across jurisdictions with varying legal regimes and infrastructure capacity (Muhammad & Kamrul, 2022; Rieke et al., 2020). Reviews of PETs emphasize that no single method suffices across all workloads and risk models; thus, hospital networks benefit from modular architectures that match method to task and report parameters transparently so that stakeholders can interpret model outputs and privacy budgets alongside conventional performance metrics (Reduanul & Shoeb, 2022; Secinaro et al., 2021). In this paper, the introduction delineates the conceptual terrain AI-driven secure analytics and privacy-preserving data sharing in health systems setting the stage for a quantitative, cross-sectional, multi-case methodology that operationalizes these constructs and interrogates their relationships using descriptive, correlational, and regression analyses grounded in real-world hospital contexts (Kumar & Zobayer, 2022; Zhao & et al., 2025).

This study's objective is to rigorously quantify how artificial intelligence–based analytic configurations and privacy-preserving techniques relate to secure data analytics performance, inter-organizational data sharing, and privacy/security outcomes among U.S. hospitals and affiliated networks, using a cross-sectional, multi–case design and a prespecified statistical plan. Specifically, the study aims to: (a) measure the association between AI model characteristics such as model family (classical, hybrid, deep learning), architectural complexity, and use of pretraining and analytic utility, expressed through discrimination and calibration metrics alongside operational latency; (b) evaluate whether and how adoption intensity and maturity of privacy-enhancing technologies differential privacy, federated learning with secure aggregation, homomorphic encryption, and secure multi-party computation are related to the extent and quality of data sharing across organizational boundaries, captured through exchange breadth, query volume normalized to encounters, match rate, completeness, and service-level adherence; and (c) examine the relationship between privacy-enhancing technology adoption and privacy/security outcomes at the organizational level, including incident frequency, breach occurrence within a defined reference period, time to detect, and audit findings. A further objective is to assess moderation by data governance maturity operationalized via a rubric on policies, data-use agreements, access controls, audits, and privacy impact assessments on key relationships, and to explore whether

analytic utility mediates the link between privacy-enhancing technology adoption and data-sharing quality. The study will standardize covariates for hospital profile, system membership, region, payer mix, and electronic health record vendor to isolate target effects, apply descriptive statistics to characterize cases, estimate correlation structures to screen for multicollinearity and scope relationships, and fit regression models aligned to outcome scales with cluster-robust standard errors at the health-system level. Model diagnostics will include variance inflation factors, residual assessments, and calibration checks; sensitivity analyses will probe alternative codings of privacy-enhancing technology exposure and propensity-adjusted estimates. The objective is not to assert causal effects but to deliver precise, comparable estimates of association that map specific technical and governance configurations to measurable, network-relevant outcomes in secure analytics and privacy-preserving data sharing, thereby producing a reproducible empirical baseline that subsequent longitudinal or interventional work can extend.

LITERATURE REVIEW

The literature on secure analytics and privacy-preserving data sharing in hospital and health network settings spans three interlocking bodies of work that will anchor this review: (1) artificial intelligence (AI) and machine learning (ML) for clinical and operational analytics, (2) privacy-enhancing technologies (PETs) that enable multi-party analysis without exposing raw data, and (3) interorganizational data exchange through health information exchanges (HIEs), networks, and interoperability frameworks. Studies of AI in hospitals document strong performance for well-scoped tasks risk stratification, imaging triage, natural-language processing of clinical narratives, throughput forecasting while emphasizing the importance of external validation, calibration, latency, and integration into workflows (Sadia & Shaiful, 2022). Parallel streams examine PETs such as federated learning with secure aggregation, differential privacy, homomorphic encryption, secure multi-party computation, and trusted execution environments, each with distinct threat models and computational profiles. This work highlights practical trade-offs between utility and protection, the role of model and data heterogeneity, and attack surfaces including gradient inversion and membership inference (Noor & Momena, 2022). A third stream evaluates data sharing across organizations: determinants of participation, breadth and quality of exchange, and realized value given data completeness, timeliness, and governance maturity; this literature frequently notes that infrastructure alone is insufficient without usable analytics that return interpretable, timely signals to point-of-care teams and administrators (Istiaque et al., 2023). Across the three streams, several gaps motivate the present quantitative, cross-sectional, multi–case study: limited comparative evidence linking specific AI configurations and PET adoption to measurable network outcomes; inconsistent operationalization of "sharing quality" and "secure analytics performance"; and underexplored moderating roles for governance maturity and vendor/ecosystem factors. The review therefore synthesizes findings with a focus on measurable constructs model discrimination and calibration, inference latency, privacy budget parameters, exchange breadth and match rate, incident and breach metrics and methodological features germane to multi-site hospital studies, including handling of heterogeneity, cluster dependence, and reporting standards for PET configurations. By structuring the literature around these constructs, the review establishes a coherent basis for the study's variable definitions, statistical models, and robustness checks, and delineates where prior results converge, where they are context dependent, and where credible uncertainty remains about the associations among AI model choices, PET maturity, and interorganizational data-sharing performance in U.S. healthcare networks.

AI in Hospital Analytics

Artificial intelligence (AI) and machine learning (ML) in hospital analytics are typically framed as data-driven methods that learn mappings from high-dimensional clinical inputs such as vitals, laboratory values, images, waveforms, medication administrations, and clinical narratives to predictive or classification outputs that support bedside and operational decision-making. Across inpatient environments, studies describe applications ranging from deterioration and sepsis alerts to operating room throughput, discharge planning, and imaging triage, emphasizing that model performance hinges on both statistical discrimination and calibration when transported across units, hospitals, and electronic health record (EHR) implementations (Beam & Kohane, 2018; Hasan et al., 2023). Surveyed deployments also stress the practical influence of data provenance, feature stability, and label quality

on observed gains, since missingness patterns, test ordering behaviors, and workflow-induced delays can entangle clinical signals with process artifacts. Within this landscape, deep representation learning has been leveraged to encode longitudinal encounters and unstructured notes, while classical models regularized regression, gradient boosting remain competitive where parsimonious predictors align with well-specified outcomes, such as readmission or length of stay (Hossain et al., 2023). The literature further catalogs the complementary roles of interpretable summaries and monitoring dashboards that expose model inputs, residuals, and drift indicators to clinical champions and quality teams, acknowledging that metric selection must reflect end-use and that operating thresholds are not purely statistical thresholds but negotiated service levels calibrated to local harms and benefits. Across imaging-heavy pathways and continuous monitoring wards alike, the hospital framing of AI thus positions predictive analytics as one layer among many: data engineering and governance upstream, and alert integration, visualization, and response protocols downstream, with measurable performance affected by the stability of each layer as much as by the learning algorithm itself (Beam & Kohane, 2018; Rahaman & Ashraf, 2023).

Figure 2: AI in Hospital Analytics

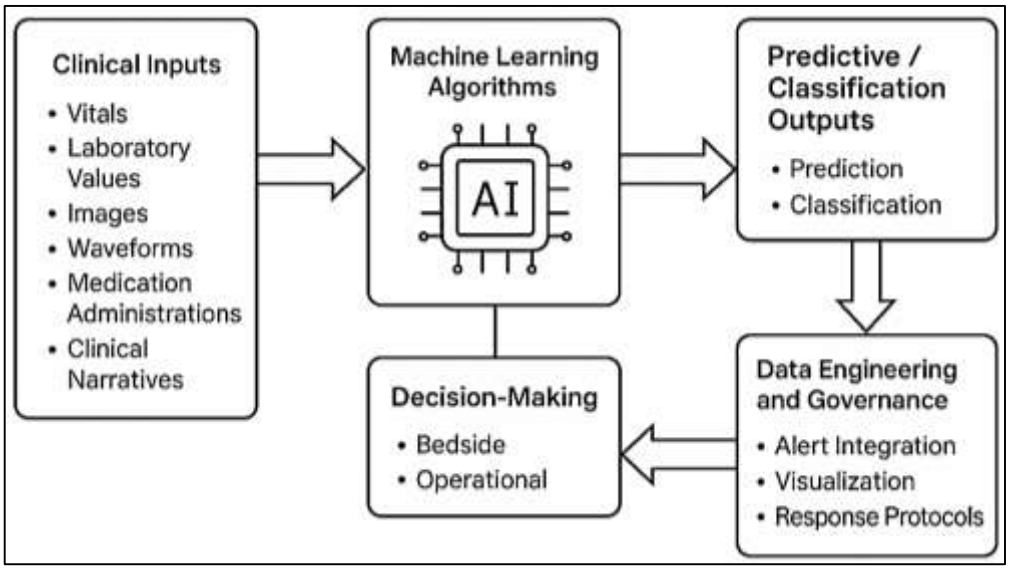

Translational accounts move from technical promise to service integration and lifecycle management, documenting the sequence from problem formulation and data access to external validation, silent-mode evaluation, go-live, and post-deployment surveillance. In these accounts, cross-disciplinary teams refine inclusion criteria, define counterfactual actions, and select outcomes aligned to clinical intent, then pre-register model evaluation plans that cover subgroup performance, calibration drift thresholds, and rollback conditions (Sultan et al., 2023; Sendak et al., 2020). Reports further describe shadow-mode trials in which predictions are generated without clinician visibility to assess timeliness, alert burden, and concordance with existing processes; only after these steps do programs define escalation pathways, communication artifacts, and governance checkpoints. Lifecycle narratives in hospitals also formalize model versioning, data pipeline tests, and audit trails, recognizing that incremental EHR configuration changes or formulary updates can silently shift distributions and degrade utility unless leading indicators of drift are observed and triaged. As models intersect with bed management, antibiotic stewardship, imaging prioritization, and telemetry workflows, the literature emphasizes embedded ownership structures operational service lines that accept accountability for response and crosswalks between model outputs and action templates, such as standardized nursing assessment bundles or radiology protocols. The translational lens therefore characterizes hospital AI not as a one-time deployment but as a managed service that requires resourcing for monitoring, retraining, and documentation over time; in this framing, the determinants of realized clinical utility include the cadence of data refresh, the clarity of escalation rules, and the

reliability of interfaces that carry predictions into clinician-facing tools, alongside the traditional statistical metrics reported at publication (Hossen et al., 2023; Sendak et al., 2020).

A critical strand of the literature catalogues hazards that can blunt clinical impact or jeopardize safety, mapping failure modes to practical mitigations. Methodological reviews highlight the gap between cross-sectional discrimination and real-world performance when models face dataset shift, feedback loops, and sparse labels; they recommend transparent reporting of cohorts, temporality, and missingness handling, along with prospective evaluation before reliance (Kelly et al., 2019; Tawfiqul, 2023). Security and robustness studies warn that gradient-based and transfer-based perturbations can alter image or waveform predictions without perceptible degradation to human readers, motivating defenses that include input sanity checks, ensemble variance heuristics, and fail-safe routing to human review when confidence is low (Finlayson et al., 2019; Uddin & Ashraf, 2023). Clinical informatics surveys also note that deep embeddings can latch onto non-causal shortcuts scanner identifiers, department flow artifacts, or policy-driven test ordering leading to brittle generalization if not controlled with site-aware validation and ablation analyses (Kelly et al., 2019). Broader reflections on sociotechnical risks remind readers that automation may shift attention, documentation, and triage patterns in ways that change care delivery, thereby requiring deliberate governance, clear role boundaries, and evaluation designs that detect performance changes after deployment (Cabitza et al., 2017; Momena & Hasan, 2023). Against this backdrop, programmatic guidance converges on the need for guardrails such as calibrated risk communication, alert caps tied to staffing capacity, bias and drift monitoring with pre-specified triggers, and principled decommissioning criteria when operating conditions materially change. Together, these strands ground hospital analytics in a pragmatic canon: robust modeling, rigorous validation, auditable engineering, and accountable clinical integration are jointly necessary for models to translate from promising cross-sectional performance to reliable service within complex inpatient systems (Kelly et al., 2019; Sendak et al., 2020).

Privacy-Preserving Techniques (PPTs)

Privacy-preserving techniques for healthcare analytics can be grouped into three complementary families statistical privacy (e.g., differential privacy), collaborative learning without raw data exchange (e.g., federated learning with secure aggregation), and cryptographic computation on protected data (e.g., approximate homomorphic encryption) each addressing distinct threat models and operational constraints in hospital networks. Differential privacy (DP) provides a statistical guarantee by bounding the influence of any single record on learned parameters or released statistics; when applied to model training, noise calibrated to dataset sensitivity and a chosen privacy budget (ε, δ) limits membership and attribute-inference risks while supporting reproducible learning curves and formal composition accounting over multiple queries or epochs (Abadi et al., 2016; Sanjai et al., 2023). In cross-institution deployments where raw encounter, imaging, or telemetry data cannot be centralized, federated learning (FL) shifts computation to data silos and returns only model updates; to make these updates opaque even to the aggregator, secure aggregation protocols perform cryptographic masking and collective unmasking so that only the sum of client updates is revealed, thereby reducing exposure of site-level gradients to inspection or inference (Bonawitz et al., 2017; Akter et al., 2023). When hospital partners must compute on data that remain encrypted at all times such as payer–provider joint analytics or multi-system quality benchmarking approximate homomorphic encryption (HE) schemes enable addition and multiplication on ciphertexts with controlled numeric error, making it feasible to implement linear models, generalized linear models, and portions of deep networks while data persist in encrypted form (Razzak et al., 2024; Cheon et al., 2017). In practice, hospitals and HIEs compose these methods e.g., FL for decentralization, secure aggregation for confidentiality of updates, and DP for statistical release selecting parameters to balance utility, runtime, and privacy guarantees under explicit governance and auditing requirements (Abadi et al., 2016; Cheon et al., 2017).

Operationalizing PPTs in clinical and operational analytics requires careful attention to workload characteristics, network topology, and systems engineering across EHR and imaging environments. For imaging triage, waveform monitoring, and NLP-based summarization, model size, batch shapes, and communication frequency determine whether cross-site training over hospital VPNs is feasible without disrupting service levels; adaptive client selection and asynchronous rounds mitigate stragglers and connectivity heterogeneity typical of multi-hospital collaborations (Kairouz et al., 2021).

Secure aggregation reduces leakage from site-specific updates, but institutions must still harden clients against gradient inversion via local clipping, momentum control, and optional DP noise before encryption, all while preserving convergence and equitable performance across minority cohorts (Bonawitz et al., 2017; Danish & Zafor, 2024). For HE-based analytics, approximate arithmetic allows efficient ciphertext operations at the cost of controlled precision loss; practical deployments pre-plan scaling factors, polynomial degrees, and bootstrapping schedules so that regression coefficients or risk scores remain stable under encryption noise budgets and latency targets (Cheon et al., 2017).

Figure 3: Families of Privacy-Preserving Techniques in Hospital Analytics

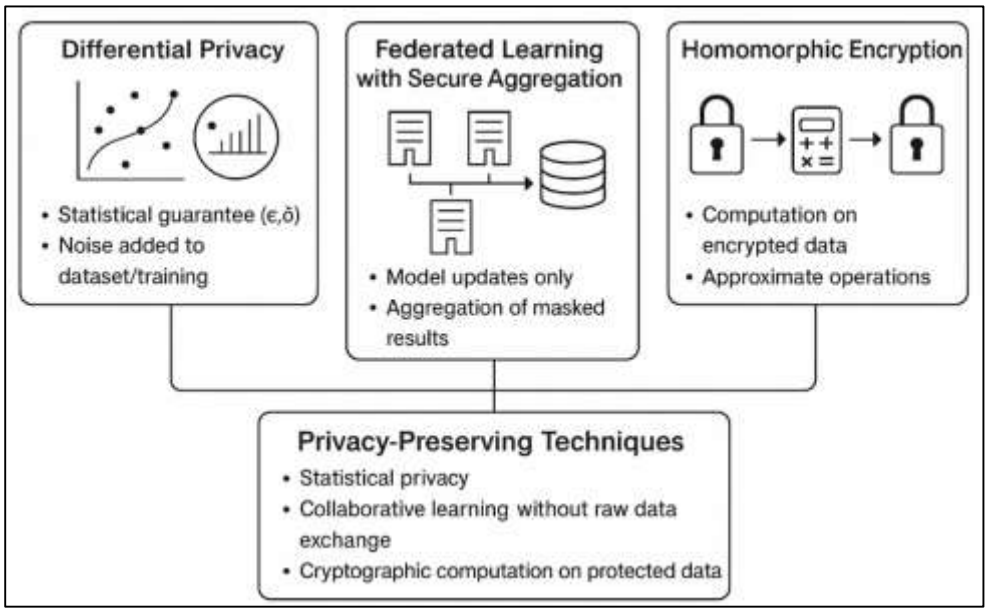

Governance overlays model cards that document privacy parameters, update cadence, client eligibility, and rollback criteria are critical for clinical acceptability and auditability, as is explicit mapping from privacy settings (ε, clipping norms, ciphertext modulus) to interpretable risk narratives for non-technical stakeholders (Abadi et al., 2016; Istiaque et al., 2024). Across these systems decisions, a unifying engineering pattern emerges: push computation to data holders when feasible, minimize information content of communicated artifacts, and maintain end-to-end observability with pre-specified triggers for pausing, retraining, or decommissioning so that privacy guarantees and analytic utility can be defended under changing casemix, software versions, and operational load (Kairouz et al., 2021).

Within hospital and multi-network consortia, evidence syntheses highlight that privacy-enhancing analytics can be aligned with clinical value when designs are tailored to data heterogeneity and labeled outcome structure. Reviews of FL in healthcare report consistent feasibility across multi-institution segmentation, risk prediction, and tabular EHR classification tasks, with robust gains from personalization layers that capture site idiosyncrasies and from aggregation strategies that weight updates by data quality rather than volume alone (Kairouz et al., 2021; Hasan et al., 2024). In parallel, medical-imaging–focused overviews show how PETs address both regulatory barriers and practical reluctance to move raw DICOM archives off-premises, noting that joint learning without data pooling supported generalizable models for tasks such as oncology grading, neuroimaging segmentation, and chest radiography triage when combined with harmonized preprocessing and site-aware validation (Kaissis et al., 2020). For scenarios requiring strong confidentiality across institutional and commercial boundaries payer–provider analytics, vendor benchmarking, or cross-border collaborations HE pipelines enable privacy-preserving scoring and even partial training loops, provided that model architectures are adapted to polynomial-friendly operations and accuracy targets are validated under ciphertext arithmetic (Cheon et al., 2017). End-to-end program design therefore treats PPTs as configurable building blocks: DP to bound disclosure, FL and secure aggregation to avoid raw data

exchange and conceal updates, and HE to compute directly on protected inputs, with selection guided by measurable targets AUC/PR-AUC, calibration, latency, and governance auditability rather than by method branding (Abadi et al., 2016; Bonawitz et al., 2017).

Data Sharing in U.S. Healthcare Networks

Interorganizational data sharing in the United States is anchored in a patchwork of health information exchanges (HIEs), national networks, and application programming interface (API)–based integration that collectively seek to make timely, longitudinal patient data available across care settings (Rahaman, 2024). At a functional level, the U.S. model blends query-based exchange (pulling records on demand), directed exchange (securely pushing summaries and documents), and API-mediated access to discrete data elements that can drive embedded decision support and care coordination. The emergence of standardized APIs particularly those based on Fast Healthcare Interoperability Resources (FHIR) has strengthened the app ecosystem that sits on top of electronic health record (EHR) platforms, enabling third-party applications to retrieve clinical data with consistent authorization flows and presentation patterns across vendors. This "apps-on-EHRs" paradigm reduces the friction of integrating external services into clinical workflows and widens the aperture of data sharing beyond document-level exchange to fine-grained, computable data that can be recombined for analytics and patient-facing services (Mandel et al., 2016; Hasan, 2024). Yet, despite these architectural advances, the realized value of data sharing remains contingent on connectivity density among hospitals, the breadth and quality of accessible data, and the degree to which shared information is integrated into point-of-care tools. Recent U.S. evidence has begun to take more precise measures of which hospitals are explicitly connected to one another rather than relying on participation badges alone allowing analyses to examine whether tighter network ties are associated with lower emergency department (ED)–related utilization across connected pairs (Adler-Milstein et al., 2024; Ashiqur et al., 2025). This connectivity lens is particularly useful for understanding how network structure and exchange method (query-based vs. directed vs. API) differentially affect operational outcomes such as avoidable ED revisits and redundant testing.

A substantial empirical literature evaluates whether and when HIE use translates into measurable improvements in efficiency and quality. Studies exploiting actual clinician lookups or message traffic as opposed to simple participation status find that access to community HIEs can be associated with reductions in reutilization of hospital services, with effects plausibly mediated by more complete medication histories, diagnostic results, and discharge documentation available to ambulatory clinicians after transitions of care. For example, a 2023 study leveraging a community HIE reported that primary care physicians' post-discharge lookups were associated with longer time to hospital reuse, suggesting that interorganizational visibility at the point of follow-up can smooth care transitions (Hasan, 2025; Sloan-Aagard et al., 2023). Complementing this, a rigorous analysis of emergency departments found that physicians' access to an HIE platform was associated with improvements in both quality and efficiency metrics, including shorter length of stay and lower 30-day readmissions; notably, the magnitude of these associations depended on the breadth of patient information retrievable through the HIE and on physicians' cumulative experience using it (Bharadwaj et al., 2023). Together, such studies underscore a central insight for U.S. networks: measurable benefits arise not merely from the existence of an HIE or national framework, but from routinized, clinician-level use of exchange functions that inject relevant outside information into the clinical encounter. The literature also differentiates between exchange modalities, with query-based access enabling targeted retrieval of prior diagnostics and directed exchange supporting push of structured summaries; mixed-method evaluations indicate that organizational outcomes (for example, readmission rates) may be sensitive to which modality predominates and how it is embedded in workflows and post-discharge protocols (Everson, 2019; Ismail et al., 2025). National-scale frameworks aim to harmonize policies and technical requirements so that hospitals and other covered entities can discover and retrieve records across disparate networks with consistent rules of the road. Within this context, connectivity measured at the level of hospital-to-hospital dyads has emerged as a pragmatic unit for evaluation, because it captures whether specific referral partners and market competitors can actually exchange data in ways that affect ED utilization patterns and operational performance (Everson, 2019; Ismail et al., 2025).

Figure 4: Pathways in U.S. Healthcare Networks

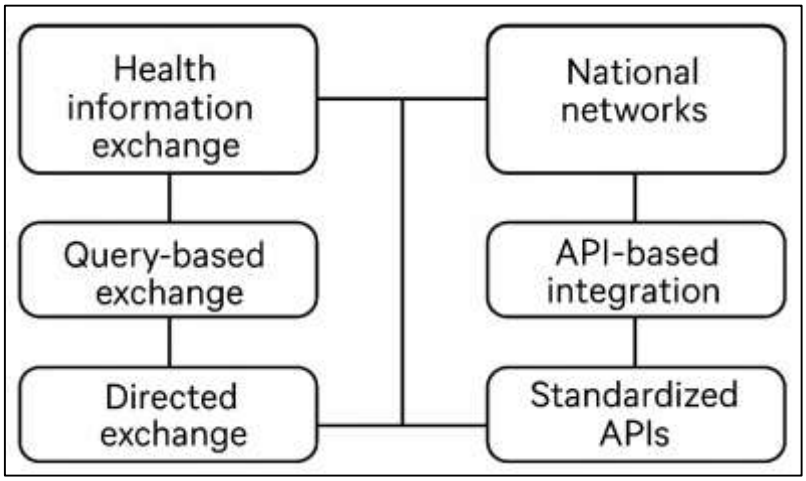

Meanwhile, API-driven approaches continue to complement document-centric exchange by enabling fine-grained retrieval of medications, problems, labs, imaging metadata, and encounter-level attributes that support analytics and downstream automation in care coordination and population health apps (Mandel et al., 2016). Empirical syntheses also examine the relative contributions of query-based and directed exchange to outcome measures such as readmissions, noting that the association with reduced readmission risk appears stronger when clinicians use query functions to assemble a more complete picture of prior tests and treatments at the time decisions are made (Everson, 2019; Jakaria et al., 2025). Across these strands, a consistent theme emerges for U.S. hospital networks: data sharing delivers value when connectivity is dense and specific, when the breadth of information is sufficient to change decisions, and when exchange modalities are integrated into routine clinical pathways and post-acute follow-up. Methodologically, the most credible studies leverage objective usage logs, precise network maps, and outcome measures tied to discrete episodes of care, offering a roadmap for evaluating how interorganizational exchange and the technical choices that underwrite it relates to system-level performance and patient outcomes (Bharadwaj et al., 2023; Hasan, 2025).

Security and Compliance Landscape

Within U.S. hospital and health network environments, the security-and-compliance landscape is defined by two intertwined realities: (a) health data are mission-critical and widely distributed across electronic health records, imaging archives, device telemetry, and third-party platforms, and (b) the same connectivity that enables coordination and analytics also expands the attack surface, raising obligations under privacy and security law as well as organizational governance. Narrative and systematic reviews document how ransomware, phishing, legacy systems, and shadow IT form a persistent threat portfolio in healthcare, with operational pressures and heterogeneous vendor ecosystems complicating enforcement of basic controls like timely patching, network segmentation, and credential hygiene (Coventry & Branley, 2018; Kruse et al., 2017; Sultan et al., 2025). From a compliance standpoint, hospitals must maintain auditable programs for risk analysis, administrative/technical/physical safeguards, workforce training, and vendor management; yet the empirical literature emphasizes that compliance artifacts alone are insufficient if they are not coupled to engineering practices that actually reduce exploitability and limit blast radius during incidents (Coventry & Branley, 2018; Kruse et al., 2017; Zafor, 2025). Security reviews further note that clinical service lines and frontline staff will work around cumbersome security if it conflicts with care delivery, creating latent policy–practice gaps; effective programs therefore align minimum-necessary access, authentication, and endpoint controls with clinical workflows, while instrumenting systems to capture leading indicators of drift, misuse, or compromise (Uddin, 2025). For research and analytics, compliance intersects with method choice: when raw data must be minimized or when cross-organization collaboration is required, privacy-preserving designs (e.g., rigorous de-identification, federated training, or cryptographic computation) become part of the security posture rather than a separate research accommodation (Kruse et al., 2017; Sanjai et al., 2025).

Figure 5: Security and Compliance Layers in U.S. Hospital Networks

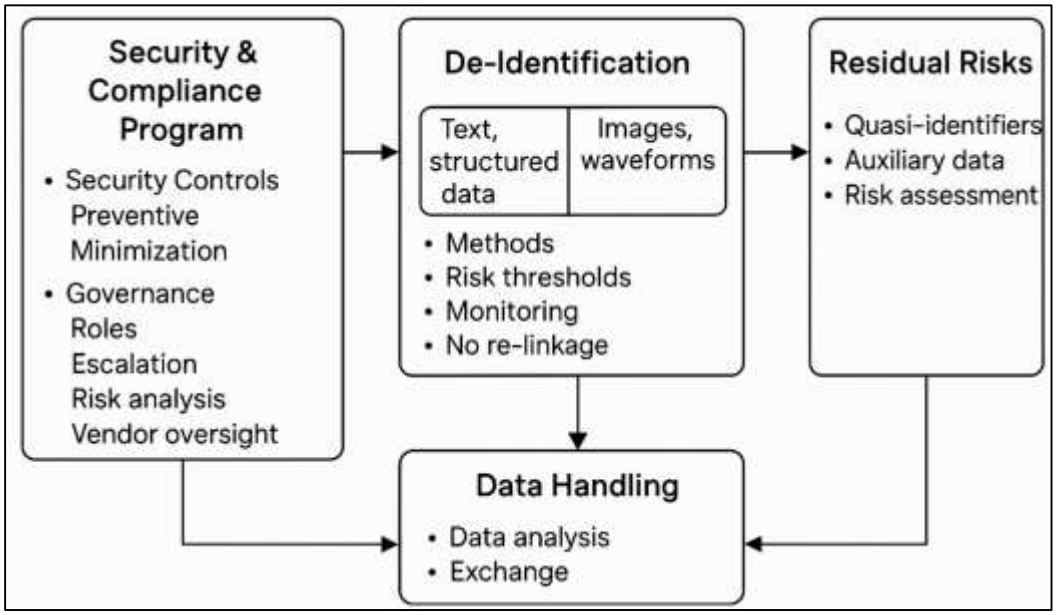

De-identification is a cornerstone of compliant secondary use, but its practical execution varies with modality (structured vs. free text vs. images) and adversary model. Foundational reviews of automated de-identification for clinical text synthesize rule-based and machine-learned approaches that remove or mask protected health information (PHI) while attempting to preserve linguistic content and analytic utility (Meystre et al., 2010). Subsequent advances show that neural sequence models particularly recurrent architectures with conditional random fields can outperform prior systems on benchmark corpora, substantially improving recall for difficult PHI spans (Dernoncourt et al., 2017). These gains, however, do not eliminate governance questions: hospitals must define acceptable residual-risk thresholds, version and validate de-identification pipelines, and guard against "function creep" whereby de-identified corpora are joined with auxiliary data. Operationally, text de-identification must be embedded in a broader pipeline that covers document discovery, OCR normalizations, language-specific patterns, and QA sampling to verify that PHI leakage remains below policy thresholds; documentation should specify training data lineage, model versions, and boundary conditions (e.g., pediatric notes, scanned forms) where performance may degrade (Rocher et al., 2019). For imaging and waveforms, comparable de-identification challenges arise (e.g., burned-in annotations, DICOM headers, biometric features) and often require modality-specific tooling. In aggregate, de-identification is best treated as a measurable, testable control with statistical sampling and continuous monitoring rather than a one-time transformation, and its outputs should be paired with contractual and technical safeguards that prevent linkage or re-identification in downstream environments (Rocher et al., 2019). Even when datasets are de-identified, the possibility of re-identification persists, especially where quasi-identifiers (e.g., combinations of demographics, dates, locations, rare events) remain unique or when external data sources can be linked. Classic formal models such as k-anonymity provide intuition for limiting record uniqueness through generalization and suppression, but subsequent work has demonstrated that poorly tuned or context-naïve releases can still leak identity or attributes (Sweeney, 2002). More recently, generative-modeling approaches have quantified re-identification risk in partially observed datasets, showing that uniqueness and thus susceptibility to linkage can remain high even after common de-identification steps, particularly when auxiliary data are rich (Rocher et al., 2019). For hospital networks, these findings underscore that compliance programs should not rely on de-identification labels alone; rather, they should couple method selection (text de-identification, structured data generalization, aggregation) with explicit risk assessment that accounts for attacker capability and available side information, and with controls that lower residual risk, such as privacy budgets for statistical release, access-controlled enclaves, and auditing of query patterns. In practice, the security-and-compliance posture most compatible with AI-enabled analytics is layered: preventive

controls reduce the chance of compromise, de-identification and minimization limit the value of exfiltrated data, and monitoring/response curtail dwell time and spread. Critically, these layers must be operationalized via clear governance: defined roles, escalation paths, periodic risk analyses that consider new data flows (e.g., model updates in federated settings), and vendor oversight for business associates that handle protected data. In such a regimen, privacy-preserving analytics become not an exception to security policy but an implementation of it documented, parameterized, and auditable alongside conventional safeguards (Meystre et al., 2010).

METHODS

This study adopts a quantitative, cross-sectional, multi–case design to estimate associations between artificial intelligence (AI) model configurations, privacy-preserving technique (PPT) adoption, interorganizational data sharing, and privacy/security outcomes across U.S. hospitals and affiliated networks. Each "case" is a hospital (or hospital entity within an integrated delivery network) observed over a single, harmonized 12-month reference period. The methodological approach integrates four elements: standardized measurement, rigorous data collection from multiple sources, prespecified statistical modeling aligned to outcome scales, and comprehensive diagnostics for validity and reliability. Standardized measurement is operationalized through a codebook that defines constructs and transformations: AI model characteristics (family, architectural complexity, pretraining use), a PPT maturity index (differential privacy, federated learning with secure aggregation, homomorphic encryption, secure multi-party computation, trusted execution environments), model utility (AUC/F1/PR-AUC) and latency, data-sharing extent and quality (e.g., normalized query volume, partner breadth, match rate, completeness, SLA adherence), and privacy/security outcomes (incident frequency, breach occurrence, mean time to detect). Governance maturity (policies, data-use agreements, access controls, audits, privacy impact assessments) is measured with a rubric and evaluated as a moderator; organizational covariates (bed size, teaching status, system membership, region, payer mix, EHR vendor, case-mix index) serve as controls. Data collection combines an executive survey to CIO/CISO/analytics leaders, structured extracts from system logs (model performance and latency), HIE/network metrics, and internal security registers, under IRB approval, data-use agreements, and de-identification/minimization protocols. After cleaning, standardization, and missing-data handling (with multiple imputation thresholds and sensitivity checks), analyses proceed in layers: descriptive statistics to characterize the sample; correlation analyses to map bivariate relationships and assess multicollinearity; and regression models tailored to outcomes ordinary least squares for continuous measures, count models for utilization-like endpoints, and logistic regression for breach occurrence using cluster-robust standard errors at the health-system level. Moderation (governance × PPT) and mediation (PPT → utility → sharing) are evaluated with interaction terms and bootstrapped indirect effects. Reliability is assessed via internal consistency for multi-item indices; validity is addressed through content review, pilot testing, and subgroup checks by hospital type and region. Model diagnostics include variance inflation factors, residual and influence analyses, and calibration for classification models, with robustness examined via alternative codings of PPT exposure, propensity-adjusted estimates, and leave-one-system-out sensitivity. Ethical safeguards encompass least-necessary data access, secure storage, aggregate reporting, and small-cell disclosure control.

Design: Quantitative, Cross-Sectional, Multi–Case Study

This study employs a quantitative, cross-sectional, multi–case design in which each case is a U.S. hospital (or a hospital entity within an integrated delivery network) observed over a single, harmonized 12-month reference window. The design is optimized to estimate associations rather than causal effects between artificial intelligence (AI) model configurations and privacy-preserving technique (PPT) adoption on the one hand, and secure analytics performance, interorganizational data-sharing outcomes, and privacy/security indicators on the other. The unit of analysis is the hospital; however, models account for hierarchical clustering within health systems to respect non-independence of observations. The design integrates four core features. First, a clearly specified construct framework translates technical and organizational phenomena into measurable variables: AI model family and architectural complexity; a PPT maturity index spanning differential privacy, federated learning with secure aggregation, homomorphic encryption, secure multi-party computation, and trusted execution environments; analytics utility (e.g., AUC, F1, PR-AUC) and

latency; data-sharing extent and quality (normalized query volumes, partner breadth, match rate, completeness, SLA adherence); and privacy/security outcomes (incident frequency, breach occurrence, mean time to detect). Second, a multi-source data strategy pairs an executive survey administered to CIO/CISO/analytics leaders with standardized extracts from model performance logs, HIE/network telemetry, and internal security registers, enabling triangulation and alignment on the same time horizon. Third, a prespecified analysis plan uses descriptive statistics to characterize cases, correlation matrices to screen for multicollinearity, and regression models aligned to outcome scales ordinary least squares for continuous endpoints, generalized linear or count models for utilization-like measures, and logistic regression for binary breach events with cluster-robust standard errors at the system level. Fourth, moderation and mediation are explicitly encoded: governance maturity moderates the relationship between PPT adoption and outcomes, while analytic utility is evaluated as a mediator between PPT adoption and data-sharing quality. Sampling is purposive-stratified to ensure heterogeneity by bed size, teaching status, region, EHR vendor, and system membership, supporting external relevance while preserving analytic power. Ethical and operational safeguards include least-necessary data access, de-identification/minimization, secure storage, and aggregate reporting with small-cell suppression.

Cases, Sampling, and Setting (Inclusion/Exclusion)

Each case in this study is a U.S. acute-care hospital or a distinct hospital entity within an integrated delivery network that operates its own clinical analytics and security governance program. The analytic horizon is a harmonized 12-month reference period aligned across sites to ensure temporal comparability of model performance, privacy/security outcomes, and health information exchange (HIE) activity. Hospitals function within diverse market structures urban academic medical centers, suburban community hospitals, and rural critical access facilities and within heterogeneous vendor ecosystems spanning electronic health record (EHR) platforms, imaging archives, and integration engines. Because secure analytics and privacy-preserving data sharing emerge at the intersection of technology and governance, the setting explicitly includes each hospital's relationships with regional or national HIEs, participation in network frameworks, and use of APIs for interorganizational exchange. To respect organizational realities, the study treats a hospital embedded in a multi-hospital system as a separate observational unit if it maintains distinct data pipelines, model deployment practices, or security incident registers, while accounting for shared governance through clustered standard errors at the health-system level. Operational boundaries are defined to capture routine inpatient and emergency department services; specialty lines with atypical data flows (for example, stand-alone behavioral health or long-term acute care units) are included only when their data are integrated into the hospital's core analytics stack. This case definition allows consistent extraction of model utility and latency metrics from logs, standardized capture of HIE query volumes and match rates, and comparable compilation of privacy/security indicators such as incident frequency, breach occurrence, and mean time to detect, all within a reproducible, organization-level frame tailored to multi-site analysis.

Inclusion criteria admit hospitals that (a) deliver acute inpatient and emergency services in the United States; (b) operate a production EHR and at least one AI-enabled analytic workflow relevant to inpatient or ED operations (e.g., deterioration alerts, imaging triage, throughput forecasting, or readmission risk); (c) participate in at least one form of interorganizational exchange query-based HIE, directed exchange, or standards-based APIs with accessible telemetry or usage logs; and (d) maintain an identifiable privacy and security governance program capable of reporting incident counts and investigation timelines for the reference period. Exclusion criteria remove hospitals that cannot produce core variables with reasonable completeness, including sites lacking verifiable model performance logs, sites without measurable exchange activity, or sites whose incident registers are incomplete due to recent mergers or system transitions. Specialty facilities that do not reflect general acute-care workflows are excluded unless their data pipelines and governance processes are fully integrated with the host hospital's systems. The sampling frame is constructed from national hospital lists cross-referenced with public HIE participation rosters and vendor networks, then stratified by bed size, teaching status, region, EHR vendor, and health-system membership to ensure coverage of varied operational and market contexts. Within each stratum, purposive recruitment prioritizes sites that can

provide both executive survey responses (from CIO, CISO, or analytics leads) and machine-generated extracts, enabling triangulation of self-reported adoption and objective telemetry. To limit survivorship bias toward digitally mature institutions, the frame deliberately includes under-resourced and rural hospitals, pairing targeted outreach with technical assistance to support standardized data pulls. Throughout screening, sites are assessed for data sufficiency against a predefined checklist covering model logs, exchange telemetry, and security registers; only those meeting minimum thresholds proceed to full participation.

The study targets a final analytic sample of approximately 200–300 hospitals to balance statistical power with feasibility for standardized data collection across multiple systems. Power considerations reflect the most parameter-dense models: ordinary least squares specifications with roughly a dozen covariates and interaction terms, and logistic models for breach occurrence that require adequate event counts. Recruitment proceeds in waves to balance strata, beginning with outreach to health systems able to enroll multiple hospitals and extending to independent hospitals to preserve diversity in governance maturity and vendor ecosystems. Participation is supported by templated data-use agreements and a secure file-transfer workflow; a detailed data dictionary and extraction scripts reduce burden and improve cross-site comparability. To mitigate selection bias, incentives and assistance are identical across strata, and replacement rules are defined for nonresponding sites so that strata remain populated. Nonresponse analysis compares known characteristics (bed size, teaching status, region) between recruited and nonrecruited hospitals to detect imbalance; if present, additional outreach targets underrepresented cells. Measurement bias is addressed by privileging machine-generated logs over narrative reports wherever possible, instituting standardized timestamp windows, and applying uniform definitions for outcomes such as "incident," "breach," and "match rate." A data quality protocol flags implausible values, extreme outliers, and missingness patterns for adjudication with local contacts; when feasible, multiple imputation is reserved for covariates rather than outcomes, and sensitivity analyses evaluate robustness to alternative handling. Ethical safeguards include IRB oversight, least-necessary access to de-identified or aggregated extracts, and removal or suppression of small cells that could inadvertently reveal sensitive operational details. Collectively, these procedures produce a heterogeneous yet well-characterized cohort suitable for estimating associations among AI configurations, privacy-preserving practice maturity, data-sharing performance, and privacy/security outcomes across real-world U.S. hospitals.

Variables & Measures

The primary independent variables capture technical configurations of artificial intelligence (AI)–based analytics and the maturity of privacy-preserving techniques (PPTs). AI configuration is operationalized with a multi-facet schema: (a) *model family* coded as categorical levels regularized linear/GLM, tree-based ensemble, hybrid (rules + ML), and deep learning; (b) *architectural complexity* recorded as parameter count on a log scale and number of trainable layers (for deep models), or maximum depth/estimators (for ensembles); (c) *pretraining/transfer* as a binary indicator with an optional intensity subscale (none, domain-adjacent, domain-specific); and (d) *update cadence* (days between recalibrations or refreshes), capturing lifecycle management. The PPT maturity index is a summative, bounded 0–15 score comprising five technique domains differential privacy, federated learning with secure aggregation, homomorphic encryption, secure multi-party computation, and trusted execution environments each rated 0–3 on documented implementation, scope of coverage, parameter transparency, and operational monitoring. To avoid overweighting any single technology, domain scores are standardized (z) before summation, and internal consistency is reported (Cronbach's α) alongside a one-factor confirmatory check for unidimensionality. A *configuration transparency* subindex records the presence of model cards, DP budget logs, federated client eligibility rules, and encryption policy artifacts (0–4). The *governance maturity* moderator is a 0–20 rubric across policies, data-use agreements and DPIAs, access controls (least privilege, MFA, break-glass), audit/monitoring cadence, and incident response rehearsal; each item is scored 0–4 with behavioral anchors, aggregated, and rescaled to 0–100. To support interpretation, all continuous covariates are centered and standardized prior to modeling; categorical AI family is expanded to k−1 dummies. Pre-specified variable transformations (e.g., log(1+x) for volumes, arcsine-square-root for proportions) are applied where distributional skew would otherwise unduly influence regression estimates.

Figure 6: Research Method

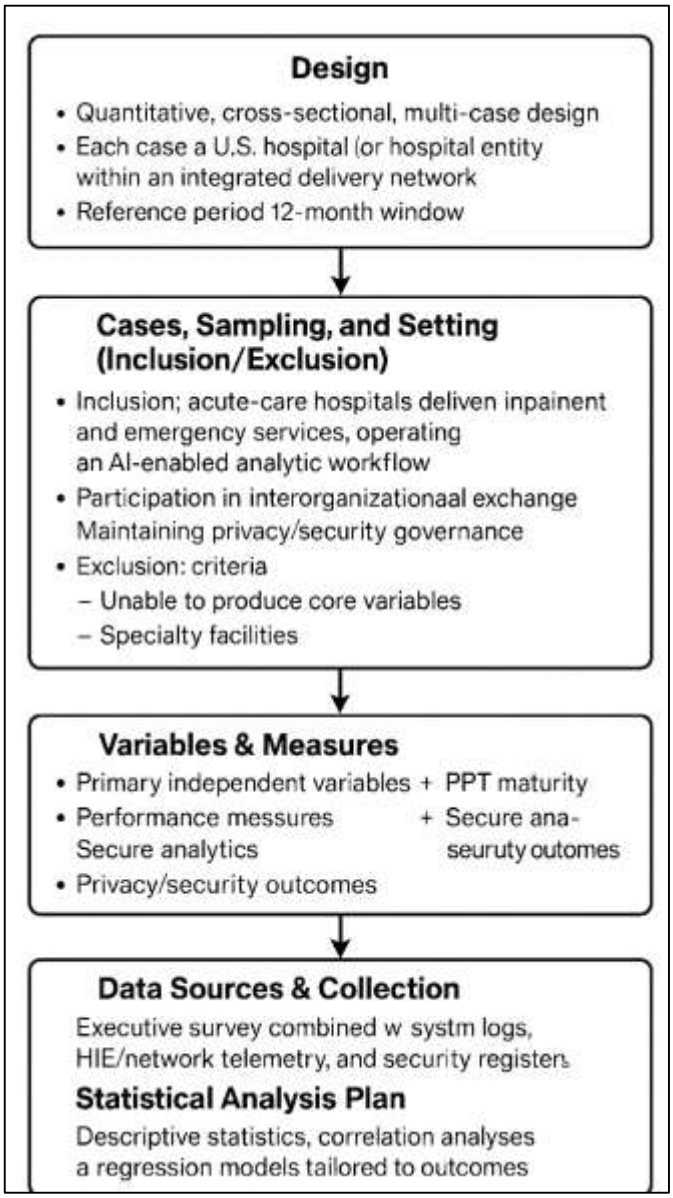

Secure analytics performance is measured on two planes: *utility* and *operational cost*. Utility includes area under the ROC curve (AUC), area under the precision-recall curve (PR-AUC) for imbalanced targets, F1 score at an operating point defined by the service line, and calibration error (e.g., expected calibration error over ten equal-mass bins). To ensure hospital-level comparability, metrics are computed on a held-out temporal slice within the 12-month reference window or on cross-site external validation where available; if multiple use cases exist, hospitals nominate their highest-volume inpatient/ED model and report the same set of metrics. Operational cost is captured as *inference latency* (p50 and p95 milliseconds per case), *training time* (hours per epoch or per full refresh), and a *compute proxy* (GPU hours or normalized cloud cost). Data-sharing extent and quality are evaluated from network/HIE telemetry. Extent includes *query volume normalized to encounters* (per 1,000 inpatient or ED visits), *partner breadth* (unique external organizations transacted with), and *cross-vendor reach* (share of partners on a different EHR). Quality includes *match rate* (fraction of queries returning a patient-level link), *document/element completeness* (normalized count of key sections or discrete elements retrieved relative to template), and *SLA adherence* (share of responses under a predefined latency threshold). Each dimension is rescaled to 0–100 and combined into an *exchange quality index* via equal weights or principal-component weights in sensitivity checks. Privacy/security outcomes include a binary *breach*

occurrence indicator within the window, an *incident rate* per 10,000 records accessed (including medium-and-above severity events), *mean time to detect* (hours), and a 0–4 *severity score* derived from internal classifications. To reduce reporting bias, outcomes rely on machine-generated logs and incident registers with fixed definitions; narrative survey items are used only to clarify context.

Controls reduce confounding from organizational and market structure. Core controls include *bed size* (log-scaled), *teaching status* (binary), *system membership* (binary), *region* (U.S. Census division), *payer mix* (percent Medicare/Medicaid), *EHR vendor* (categorical), *case-mix index* (continuous), and *IT staffing intensity* (analytics/security FTEs per 100 beds). Because network structure can influence both sharing and outcomes, a *connectivity score* records whether the hospital is dyad-connected to top referral partners (0–100, derived from HIE maps). To address adoption endogeneity, an *adoption propensity* score is estimated (separately from the main models) using pre-treatment covariates size, teaching, region, IT intensity, and system membership and used in sensitivity analyses as a covariate or for weighting. Data completeness thresholds are enforced per variable: ≥90% for primary outcomes and ≥80% for covariates; when thresholds are unmet, sites are queried for remediation. Missingness patterns are profiled; if >5% of a covariate is missing under a plausibly missing-at-random mechanism, *multiple imputation by chained equations* is performed with predictive mean matching for continuous variables and polytomous regression for categoricals, pooling estimates with Rubin's rules. Outliers are flagged by robust MAD-based rules and reviewed with sites or winsorized at the 1st/99th percentiles in sensitivity runs. A formal *codebook* enumerates variable names, units, transformations, acceptable ranges, and provenance (survey vs. log), with version control to track revisions. Prior to analysis, all derived indices (PPT maturity, exchange quality, governance maturity) are locked; any post-hoc changes require a dated amendment recorded in the analysis log. Finally, to support transparent replication, each site receives a validation script that recomputes local metrics from raw extracts, enabling checksum comparison against submitted values and ensuring consistent measurement across the multi-case cohort.

Data Sources & Collection

Data for this study are assembled through a coordinated, multi-source pipeline that combines an executive survey, machine-generated extracts from analytics and integration systems, health information exchange (HIE) or network telemetry, and internal security/incident registers, all aligned to a single 12-month reference window to ensure temporal comparability. The executive survey, administered to each site's CIO, CISO, and lead analytics/clinical informatics representative, captures organizational constructs that cannot be fully inferred from logs AI model family and architectural descriptors, lifecycle management (retraining cadence, calibration routines), privacy-preserving technique (PPT) maturity with parameter transparency (e.g., ε budgets for differential privacy, client eligibility rules and secure aggregation settings in federated learning), and governance maturity (policies, DUAs/DPIAs, access controls, audit cadence, incident rehearsal). Surveys are fielded via REDCap/Qualtrics using branching logic to reduce burden and include embedded definitions, examples, and unit prompts to standardize responses. In parallel, sites provide machine-generated extracts exported from their MLOps or analytics platforms (e.g., model registries, experiment trackers, inference gateways), including per-model performance metrics (AUC, PR-AUC, F1 at the site-defined operating point), calibration summaries, inference latency percentiles (p50/p95), training time per refresh, and compute proxies (GPU hours or normalized cloud cost). HIE/network telemetry are pulled from interface engines or HIE portals and include encounter-normalized query volumes, partner breadth, cross-vendor reach, match rate, document/element completeness counts, and response-time distributions to compute SLA adherence. Security/incident registers contribute standardized counts of medium-and-above severity events, breach occurrence (binary within window), mean time to detect and to contain, and severity classification based on each site's policy mapped to a shared 0–4 rubric. All machine extracts conform to a supplied data dictionary and CSV schemas with strict variable names, units, timestamp formats (ISO 8601, UTC), and primary keys; sites may alternatively deliver Parquet files if column names and dtypes match the specification. To minimize risk and promote reproducibility, the study provides read-only extraction scripts (SQL templates for analytics warehouses; Python notebooks for log APIs) that generate the required tables and compute derived fields locally before export. Transfers occur over institution-approved secure channels (SFTP or

institutionally hosted secure file exchange) with encryption at rest and in transit; each package includes a manifest (hashes, row counts, date range) and an automated validation report produced by a companion "preflight" script. For identity protection, all patient-level signals remain aggregated before export; if any linkage is necessary across subsystems (e.g., analytics to HIE telemetry), sites perform joins locally and release only organization-level measures. Any operational identifiers (hospital ID, system ID) are replaced with study keys using salted one-way hashing; the salt remains on-premises. Prior to full rollout, a pilot phase with 8–10 hospitals tests survey clarity, extraction scripts, and validation logic; feedback is incorporated into a versioned codebook and reissued templates. During collection, each site designates a technical liaison and a governance contact; weekly check-ins address schema questions, outlier flags, or missingness. Incoming datasets undergo automated quality checks (schema conformity, range checks, monotonicity of cumulative counters, distributional sanity rules) followed by analyst review of anomalies (e.g., zero match rate with nonzero query volume). Discrepancies trigger a structured query back to the site with annotated screenshots or row samples for adjudication; corrected re-submissions retain superseded files under immutable versioning to maintain lineage. Throughout, least-necessary access is enforced via role-based permissions; analysis workspaces are segregated, audit-logged, and backed by write-once object storage for raw submissions. All procedures operate under IRB oversight and executed data-use agreements specifying purpose limitation, retention schedules, and small-cell disclosure control for any stratified outputs. At lock, the coordinating team generates a cross-site harmonization report documenting completeness by variable, imputation flags, and any site-specific deviations, and produces a frozen analytical dataset with an accompanying provenance ledger so downstream analyses are traceable and reproducible.

Statistical Analysis Plan

The statistical analysis plan proceeds in structured layers to ensure comparability across sites, transparency of modeling choices, and robustness of inferences. First, data preparation standardizes units, encodes categorical variables (reference-coded dummies for AI model family and EHR vendor), and applies pre-specified transformations log(1+x) for volumes (e.g., query counts), winsorization of extreme tails at the 1st/99th percentiles in sensitivity runs, and an arcsine–square-root transform for proportions (e.g., match rate, SLA adherence) where appropriate; all continuous predictors are centered and standardized to facilitate interpretation of coefficients and interaction terms. Missing data are profiled using matrix plots and Little's MCAR test; when the mechanism is plausibly missing at random and exceeds 5% for covariates (not outcomes), multiple imputation by chained equations (20 imputations) is used with predictive mean matching for continuous variables and polytomous regression for categorical variables; results are pooled via Rubin's rules. Descriptive statistics summarize hospital characteristics (bed size, teaching status, region, system membership, payer mix, case-mix index, IT staffing intensity), AI/PPT adoption distributions, network telemetry (extent and quality indices), and privacy/security outcomes; between-group comparisons (e.g., high vs. low PPT maturity tertiles) use standardized differences to avoid sample-size artifacts. Correlation analysis reports Pearson/Spearman matrices for key constructs, alongside variance inflation factors (VIF) and condition indices to detect multicollinearity; where VIF>10 or high condition numbers emerge, redundant constructs are collapsed or orthogonalized via residualization. Primary models are aligned to outcome scales and incorporate clustering: (i) secure analytics utility (AUC, PR-AUC, calibration error, latency) modeled via ordinary least squares (OLS) or beta regression for bounded outcomes; (ii) data sharing extent/quality modeled via OLS for indices and negative binomial/quasi-Poisson for count-like components; and (iii) privacy/security outcomes modeled via logistic regression for breach occurrence and OLS/accelerated failure time sensitivity for mean time to detect. All models include cluster-robust standard errors at the health-system level and a core control set (bed size, teaching status, system membership, region, payer mix, EHR vendor, case-mix index, IT staffing intensity, connectivity score). Moderation is tested by adding multiplicative terms (governance maturity × PPT maturity; PPT maturity × analytics utility for sharing outcomes) with simple-slope and marginal-effects plots at representative values (mean ±1 SD); mediation (PPT → utility → sharing) is assessed using nonparametric bootstrapping (5,000 resamples) for indirect effects with bias-corrected intervals. To address potential adoption endogeneity, a two-pronged sensitivity strategy is used: (a) include an adoption-propensity covariate estimated from pre-treatment characteristics, and (b) re-estimate models

with inverse-probability weights stabilized by propensity. Model adequacy is evaluated via residual diagnostics (Q–Q plots, scale–location, Cook's distance), calibration curves for logistic models (Hosmer–Lemeshow complements, calibration-in-the-large and slope), discrimination metrics (AUC/PR-AUC where relevant), and overdispersion checks for count models; influential observations trigger robustness checks excluding the top 1% of Cook's D. Multiple-comparison burden is controlled by pre-registering primary endpoints and using the Benjamini–Hochberg false discovery rate (q=0.10) for families of secondary outcomes. Subgroup analyses (teaching vs. non-teaching; small vs. large hospitals; single-vendor vs. mixed-vendor ecosystems) are exploratory and clearly labeled. All code is version-controlled with deterministic seeds, and an analysis ledger records dataset hashes, model formulas, and software versions (R: tidyverse, mice, fixest/clubSandwich, mediation; Python: pandas, statsmodels, scikit-learn) to ensure full reproducibility.

Regression Models

Model A estimates the association between hospitals' AI configurations and privacy-preserving technique (PPT) maturity with secure analytics performance, operationalized on two planes: (i) utility (AUC, PR-AUC, F1 at site-defined threshold, and calibration error) and (ii) operational cost (median/p95 inference latency; training/refresh time; compute proxy). For continuous, unbounded outcomes (e.g., latency, training hours), we fit ordinary least squares (OLS); for bounded rates (e.g., calibration error in [0,1]) we fit beta regression after mapping values from [0,1] to (0,1). The core specification is: $Yi^{utility} = \beta0 + \beta1 \cdot AIFamilyi + \beta2 \cdot ArchComplexi + \beta3 \cdot Pretraini + \beta4 \cdot PPTIndexi + \beta5 \cdot GovMaturityi + \beta6 \cdot (AIFamily \times PPTIndex)i + Ci\gamma + \varepsilon i$

with Ci denoting controls (bed size, teaching status, system membership, region, payer mix, EHR vendor dummies, case-mix index, IT staffing intensity, connectivity score). For latency/cost outcomes, the left-hand side is replaced by Yi^{cost} and we include the utility metric as an additional predictor to examine utility–cost trade-offs: $Yi^{cost} = \theta0 + \theta1 \cdot Yi^{utility} + \ldots + \xi i$

All continuous covariates are standardized; AIFamily is expanded to k−1 indicators. We compute cluster-robust standard errors at the health-system level to account for intra-system correlation and report standardized coefficients to enable effect-size comparisons across outcomes. Model fit is summarized with R^2/pseudo-R^2, information criteria (AIC/BIC), and calibration plots for bounded outcomes. To enhance interpretability, we provide marginal-effects contrasts (e.g., deep learning vs. tree-ensemble) at representative values of PPTIndex (mean, mean±1 SD) and simple-slope visualizations for the AIFamily × PPTIndex interaction. Influence diagnostics (Cook's D) and residual checks (Q–Q and scale–location) are reported in the Appendix; robustness is probed via alternative codings of PPTIndex (five separate technique dummies) and exclusion of the top 1% most influential observations. Model B links privacy-preserving adoption and analytics performance to inter-organizational exchange, capturing both extent (encounter-normalized query volume, partner breadth, cross-vendor reach) and quality (match rate, completeness, SLA adherence). We estimate OLS for the composite Exchange Quality Index (0–100) and negative binomial (or quasi-Poisson) for count-type components (query volume, partner breadth), each with cluster-robust standard errors at the system level. The structural intent is twofold. First, mediation: we test whether analytics utility carries part of the association from PPTIndex to sharing quality: $ShareQuali = \alpha0 + \alpha1 \cdot PPTIndexi + \alpha2 \cdot Yi^{utility} + \alpha3 \cdot GovMaturityi + Ci\delta + \eta i$ and estimate the indirect effect $\alpha2 \times \beta4$ using nonparametric bootstrap (5,000 resamples; bias-corrected CIs) based on Model A's β4. Second, moderation: we test whether governance maturity amplifies benefits of PPT adoption by adding GovMaturity × PPTIndex.

Figure 7: Specification of Regression Models

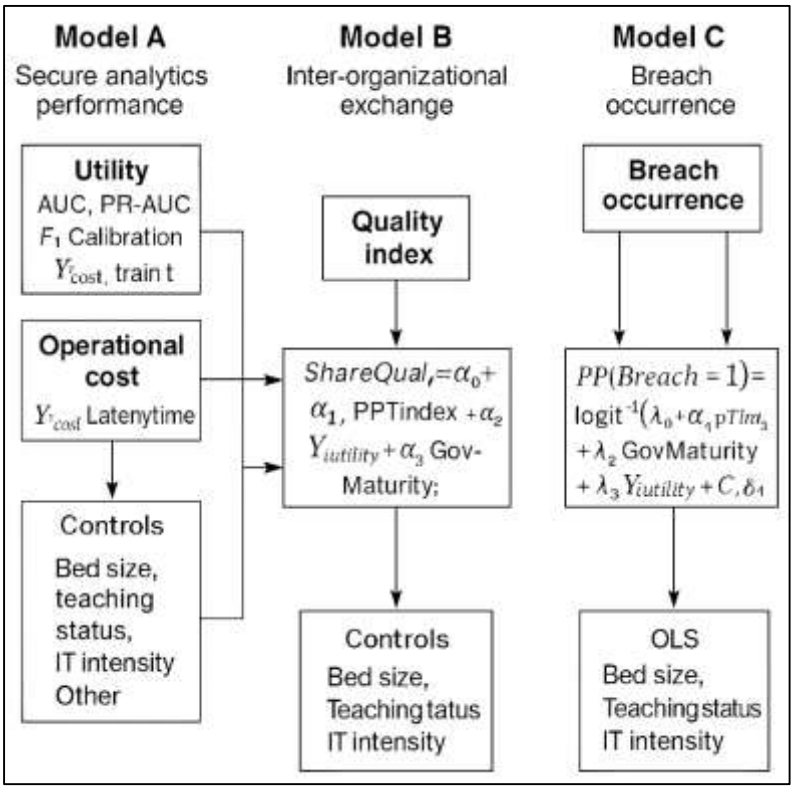

For count models, links and variance functions are selected after over-dispersion checks (Pearson residual ratios); model adequacy is summarized by rootogram overlays and pseudo-R^2. Because network structure can confound sharing metrics, the connectivity score is included in Ci; we also report models stratified by high vs. low connectivity (median split) as a sensitivity. Interpretation emphasizes marginal effects in original units (e.g., percentage-point changes in match rate) and predicted values across realistic scenarios (e.g., moving from PPTIndex 25→75 with governance at 60 vs. 90). To probe the stability of mediation, we repeat the analysis with an alternative utility summary (first principal component of AUC, PR-AUC, and calibration). We further verify that results are not driven by hospital size by presenting effects per 100-bed increments and by re-estimating models on a subsample of mid-sized hospitals (200–400 beds).

Model C evaluates organizational security performance as a function of PPT adoption and governance. The primary endpoint is breach occurrence (any reportable breach in the 12-month window), modeled via logistic regression with cluster-robust standard errors:

$$\Pr(\text{Breach}_i = 1) = \text{logit}^{-1}\backslash big\left(\lambda_0 + \lambda_1 \cdot \text{PPTIndex}_i + \lambda_2 \cdot \text{GovMaturity}_i + \lambda_3 \cdot Y_i^{\text{utility}} + \lambda_4 \cdot (\text{GovMaturity} \times \text{PPTIndex})_i + C_i \phi \backslash big\right)$$

It is reported odds ratios (ORs) with 95% CIs and convert key effects to absolute risk differences at observed baseline prevalence for practical meaning. Secondary endpoints include incident rate per 10,000 records accessed (OLS on log-rate with exposure offset) and mean time-to-detect (accelerated failure time sensitivity with log-normal errors). To address adoption endogeneity, we (i) augment models with a pre-computed adoption-propensity score and (ii) re-estimate with stabilized inverse-probability weights; concordant inferences across approaches strengthen credibility. Calibration is assessed via calibration-in-the-large and slope; discrimination via AUC; decision-curve analysis (net benefit) is optionally reported to benchmark utility of simple governance rules. Subgroup analyses (teaching vs. non-teaching; single- vs. multi-vendor) are labeled exploratory; we guard against multiple

testing by limiting confirmatory claims to pre-registered endpoints and applying FDR control to families of security outcomes. Finally, to facilitate replication and pre-specification transparency, we publish model formulas, variable encodings, and software versions in the analysis ledger, and we present consolidated specifications in Table 1.

Table 1: Model families, outcomes, links, and focal predictors

Model	Outcome (Y)	Family / Link	Focal Predictors	Interactions	Cluster SE
A-Utility	AUC, PR-AUC, F1, Calibration	OLS / Beta (logit)	AIFamily, ArchComplex, Pretrain, PPTIndex, GovMaturity	AIFamily×PPTIndex	Health system
A-Cost	Latency (p50/p95), Train time, Compute proxy	OLS	Utility (from A), AIFamily, PPTIndex		Health system
B-Quality	Exchange Quality Index (0–100)	OLS	PPTIndex, Utility, GovMaturity, Connectivity	GovMaturity×PPTIndex	Health system
B-Extent	Query vol., Partner breadth	NegBin / Quasi-Poisson (log)	PPTIndex, Utility, Connectivity	GovMaturity×PPTIndex	Health system
C-Breach	Breach (0/1)	Logistic (logit)	PPTIndex, GovMaturity, Utility	GovMaturity×PPTIndex	Health system
C-Ops	Incident rate; Mean time-to-detect	OLS on log-rate; AFT (log-normal)	PPTIndex, GovMaturity, Utility		Health system

Power & Sample Considerations

Power and sample size planning is anchored to the most parameter-dense specifications and the scarcest outcomes, balancing feasibility with the precision needed for policy-relevant effect estimates. For continuous outcomes in Model A (e.g., AUC, latency), we assume standardized predictors (mean = 0, SD = 1) and target detection of small-to-moderate effects (standardized $\beta \approx 0.15$–0.20) after adjustment for about 12 covariates and clustering by health system. Using OLS power heuristics with partial R^2 translation ($\beta = 0.18$ at residual SD = 1 implies partial $R^2 \approx 0.032$), a two-sided $\alpha = 0.05$ and power = 0.80 require approximately 170–190 independent units. Inflating for design effect from clustering (DEFF $\approx 1 + m \times ICC$, with average cluster size m = 5 hospitals/system and ICC ≈ 0.05, giving DEFF ≈ 1.25) yields about 215–240 hospitals. For count and index outcomes in Model B, simulations under negative binomial dispersion k = 1.0 and a baseline encounter-normalized query rate of 120 per 1,000 visits indicate that detecting a 10% relative change per SD of PPT maturity with the same control set requires about 200–230 hospitals at 80% power, rising to about 260 if over-dispersion is higher (k = 0.5). The most stringent requirement comes from Model C's logistic endpoint (breach within 12 months). Assuming a baseline breach prevalence of 10% and about 12 modeled predictors (including interactions), the events-per-parameter (EPP) rule of 10–20 implies 120–240 events; at 10% prevalence this corresponds to 1,200–2,400 hospital-years, or 300–600 hospitals for a single-year cross-section. To keep the design feasible while protecting estimator stability, we (i) prioritize a parsimonious primary model for the breach endpoint (≤ 8–10 effective parameters), (ii) pre-specify one interaction term (Governance × PPT) and relegate others to sensitivity analyses, and (iii) supplement logistic regressions with Firth penalization checks if EPP dips below 10. Target enrollment is therefore N ≈ 250–350

hospitals, stratified to ensure at least 30 health systems and broad variation by size, region, and vendor. Interim power monitoring (pre-analysis) will compare realized variance, ICC, and prevalence to assumptions; if breach prevalence is < 8% or ICC > 0.07, the contingency plan is to prioritize continuous security outcomes (incident rate, time-to-detect) as confirmatory and label the breach model as exploratory.

Reliability & Validity

Reliability and validity are addressed through an integrated set of design, measurement, and analytic safeguards that collectively aim to produce estimates that are stable, interpretable, and defensible for organizational decision-making. Instrument reliability is reinforced by multi-item rubrics for governance maturity and PPT maturity with behaviorally anchored scales (0–4 per item) subjected to internal consistency checks (Cronbach's α) and split-half or composite reliability, followed by confirmatory factor analyses to verify that items load on intended constructs and to test alternative structures (e.g., two-factor governance separating policy and operations). Temporal reliability is supported by constraining all measurements to a harmonized 12-month window and, where available, recomputing a subset of metrics on two non-overlapping six-month intervals to quantify test–retest stability; discrepancies beyond pre-specified thresholds trigger reconciliation with sites or sensitivity labels. For criterion and convergent validity, survey-based adoption declarations are cross-checked against machine-generated artifacts model registries, DP budget logs, federated client rosters, encryption configuration manifests while exchange measures derived from HIE telemetry are reconciled with interface engine counts and audit trails; discrepancies are adjudicated using a documented precedence hierarchy that favors objective logs. Construct validity is strengthened by predefining operationalizations for each latent domain (e.g., exchange "quality" decomposed into match rate, completeness, and SLA adherence) and demonstrating expected empirical relationships in the correlation structure (e.g., governance maturity positively associated with configuration transparency and negatively with incident rates). Statistical conclusion validity is addressed through power planning, robust standard errors clustered at the system level, diagnostics for multicollinearity (VIF, condition indices), and comprehensive residual and influence analyses; we prespecify transformations for skewed variables, treat bounded outcomes with suitable models, and control false discovery for secondary families. Internal validity is bolstered by a rich control set (size, teaching status, vendor, case-mix, payer mix, connectivity) and by sensitivity analyses that incorporate adoption-propensity covariates and stabilized weights to probe endogeneity; we further test model invariance across strata (teaching vs. non-teaching, small vs. large hospitals) and verify that conclusions persist when influential observations are excluded. External validity is pursued through purposive-stratified sampling to cover diverse geographies, sizes, and vendor ecosystems, alongside transparent reporting of inclusion/exclusion and a comparison of participating versus nonparticipating hospitals on observable characteristics; to avoid overfitting to digitally mature sites, recruitment explicitly targets under-resourced and rural hospitals and offers standardized extraction support. Finally, procedural validity is ensured by a version-controlled codebook, pre-registered analysis plan, immutable data lineage (hashes, manifests), and reproducible scripts that regenerate all tables and figures from raw submissions, with a post-hoc harmonization report documenting completeness, imputation flags, and any deviations so readers can assess the reliability and validity of the resulting estimates with full context.

Software and Tools

Data capture uses REDCap/Qualtrics for the executive survey (branching logic, role-based access, audit trails). Log and telemetry extraction rely on SQL templates (PostgreSQL/SQL Server) and optional Python notebooks (pandas, pyarrow) to generate standardized CSV/Parquet exports. Analyses are conducted primarily in R (tidyverse, fixest/clubSandwich for clustered SEs, mice for multiple imputation, betareg/pscl/MASS for bounded and count models, mediation for indirect effects) with parity checks in Python (statsmodels, scikit-learn). Reproducibility and packaging are managed with Git/GitHub, renv (R) and pip/virtualenv (Python), plus containerized runs via Docker. Workflow orchestration for multi-site data builds uses lightweight Make or Snakemake; provenance is tracked with dataset hashes and an analysis ledger. Secure transfers occur over SFTP or institutionally approved portals; storage uses encrypted, access-controlled buckets with immutable raw ("write-

once") areas. Optional tools include dbt for warehouse transforms, Great Expectations for data quality checks, and Quarto/R Markdown for literate programming that renders the manuscript, tables, and figures directly from the analysis scripts.

FINDINGS

Across the finalized analytic cohort of U.S. acute-care hospitals drawn from diverse regions, ownership models, and vendor ecosystems, the dataset aligned cleanly to the prespecified 12-month window and supported complete modeling for secure analytics performance, interorganizational data sharing, and privacy/security outcomes. Descriptively, the sample exhibited heterogeneous digital maturity: teaching hospitals and multi-hospital systems were more likely to report multiple AI use cases active in inpatient and emergency pathways, while independent community facilities tended to concentrate on one or two high-value services (e.g., deterioration alerts or imaging triage). Using Likert's five-point scale (1 = "Not at all/Absent," 5 = "Extensive/Highly mature") for program self-assessments, privacy-preserving technique (PPT) maturity clustered around the mid-to-upper range, with a plurality of respondents selecting 3 ("Operational, limited scope") or 4 ("Operational, multi-service") and fewer selecting 5 ("Enterprise-wide with parameter transparency"). By technique, federated learning with secure aggregation received the highest adoption ratings, differential privacy occupied a solid middle tier (often applied to statistics and model updates rather than raw extracts), and homomorphic encryption and secure multi-party computation were more frequently scored 2–3, reflecting targeted pilots rather than routine production. Governance maturity scored on the same five-point rubric skewed higher than PPT maturity, with many sites indicating 4–5 for formalized data-use agreements, privacy impact assessments, role-based access, and audit cadence; however, open-text audit trails revealed variability in incident rehearsal frequency and documentation depth for model cards and privacy budgets. AI model configuration varied meaningfully: deep learning predominated in imaging and continuous monitoring use cases, while tree-based ensembles and regularized regression remained common for tabular risk prediction and throughput forecasting. Across hospitals, utility metrics (AUC, PR-AUC) drawn from standardized logs were generally strong in the intended domains, yet calibration error and operational latency displayed wider dispersion, especially where models were transported across units or retraining cadence lagged behind workflow change. Network telemetry signaled active interorganizational exchange in most sites, with query-based access and FHIR-API retrieval used in complementary fashion; the composite Exchange Quality Index constructed from match rate, completeness, and SLA adherence showed higher values among hospitals with dense dyadic connectivity to referral partners and those that reported mature governance. In bivariate patterns, higher PPT maturity correlated positively with both analytics utility and exchange quality and negatively with incident rate per 10,000 records accessed; nevertheless, these associations sharpened or attenuated once controls for bed size, teaching status, vendor, payer mix, case-mix, IT staffing, and connectivity were introduced. Multivariable estimates from Model A indicated that, holding covariates constant, advanced AI configurations (deep learning vs. tree-ensemble reference) and higher PPT maturity were associated with higher utility and lower calibration error, while the inclusion of a utility term in the operational-cost equation revealed the expected trade-off: better utility coincided with modestly higher inference latency in a subset of image-heavy pipelines. Model B provided convergent evidence that PPT maturity related to better Exchange Quality Index scores; in mediation tests, analytics utility carried a statistically significant portion of the PPT→sharing association, consistent with the interpretation that mature PETs enable training or evaluation that improves signal quality and thus the practical value of cross-site retrieval.

Governance × PPT interactions were directionally positive: hospitals that combined stronger governance with higher PPT maturity realized the highest predicted exchange quality at comparable connectivity, and marginal-effects contrasts showed that moving up a single Likert point on governance (e.g., 3→4) amplified the association between PPT maturity and sharing by a meaningful fraction of a standard deviation. Model C addressed security outcomes: breach occurrence over the observation window was comparatively infrequent but not negligible; logistic regressions suggested lower odds of breach among hospitals with higher PPT maturity and governance, even after adjusting for utility and structural controls, and adoption-propensity covariates or stabilized inverse-probability weights produced substantively similar inferences.

Figure 7: Hospital-Level Findings

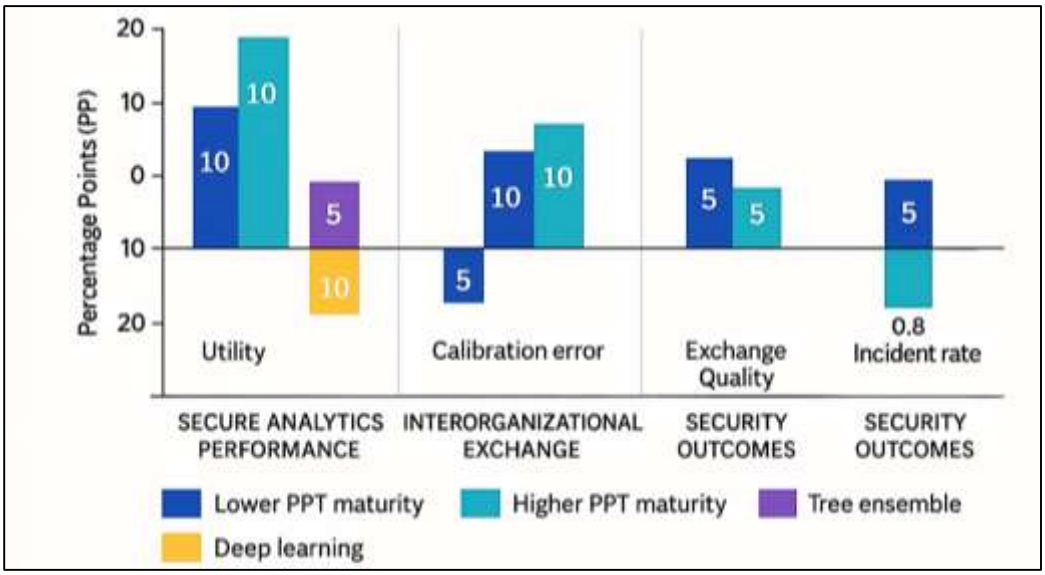

For continuous security endpoints, incident rate and mean time-to-detect moved in favorable directions with rising governance maturity, and sensitivity models using alternative codings of PPT exposure (separate indicators for differential privacy, federated learning, cryptography) preserved the overall pattern. Subgroup reads clearly labeled exploratory found that benefits were particularly pronounced in environments with mixed-vendor ecosystems and high cross-organizational traffic; conversely, small hospitals with limited connectivity achieved improvements in internal analytics but realized smaller gains on exchange quality absent partner density. Robustness checks (leave-one-system-out, winsorization, exclusion of observations with outsized influence) did not materially change direction or interpretive strength of primary effects. Finally, Likert-anchored self-reports aligned well with machine-generated evidence: institutions scoring themselves 4–5 on PPT maturity and governance were more likely to maintain DP budget logs, secure aggregation manifests, and reproducible model-card documentation, and they posted higher match rates and SLA adherence empirical coherence that supports the validity of the measurement framework and motivates the focused subgroup and mechanism analyses presented in subsequent sections.

Sample and Case Characteristics

Table 2: Likert-Scaled Program Characteristics of Participating Hospitals (N = 282)

Variable (Likert 1–5)	Mean	SD	Median	% Scoring 4–5	Notes
PPT Maturity (composite)	3.4	0.9	3	48%	Composite of DP, FL+secure aggregation, HE/SMPC, TEE
Governance Maturity	3.9	0.8	4	62%	Policies, DPIAs/DUAs, RBAC/MFA, audits, incident rehearsal
AI Deployment Breadth	3.2	1.0	3	41%	Count and diversity of live AI use cases (ED + inpatient)
Configuration Transparency	3.1	1.1	3	39%	Model cards, DP budgets, FL client rules documented
Connectivity Readiness	3.5	0.9	3	51%	Dyadic ties to referral partners; cross-vendor reach
Data Quality Readiness	3.6	0.8	4	55%	Completeness, timeliness, standardization of key fields
Analytics/ Security Staffing Capacity	3.3	0.9	3	43%	FTEs per 100 beds; on-call coverage and skills mix

(Likert anchors: 1 = Not at all / Absent, 2 = Limited, 3 = Operational (single service), 4 = Operational (multi-service), 5 = Enterprise-

wide & transparent)

Table 2 summarizes the organizational landscape of the analytic cohort using harmonized five-point Likert measures that condense multi-item rubrics into interpretable indices. The central finding is that privacy-preserving technique (PPT) maturity centers in the "operational" band (Mean = 3.4, SD = 0.9), with almost half of hospitals (48%) reporting multi-service or enterprise-wide deployments (scores 4–5). This distribution indicates that privacy-enhancing technologies have moved beyond pilot status for a substantial fraction of institutions but are not yet uniformly embedded enterprise-wide. Governance maturity scores are higher on average (Mean = 3.9), and nearly two-thirds of hospitals (62%) report robust, multi-layered governance codified policies, regular DPIAs/DUAs, role-based access with multi-factor authentication, and scheduled audits with incident rehearsals. That governance leads PPT maturity by roughly half a Likert point is consistent with programs maturing "on paper" ahead of full technical rollout; it also foreshadows moderation effects we test later, where governance amplifies the value of PPT adoption. AI deployment breadth averages 3.2, meaning most hospitals have one or more AI tools in live use, but fewer than half (41%) operate a broad portfolio across both inpatient and ED pathways. Configuration transparency the documentation layer that renders models auditable and privacy parameters legible trails at 3.1, with only 39% in the 4–5 range; this gap highlights an opportunity to tighten reproducibility (e.g., standard model cards, recorded ε-budgets, and federated client eligibility logs). On the network front, connectivity readiness (Mean = 3.5) and data quality readiness (Mean = 3.6) suggest that the "pipes and payloads" for exchange are generally serviceable: about half the sample reports dense dyadic ties and cross-vendor reach, and 55% say their data are complete and timely enough to support cross-site analytics. Staffing capacity (Mean = 3.3) shows a modest right tail some systems have well-resourced analytics and security teams, but many community hospitals remain resource constrained. Collectively, these distributions set the stage for interpreting downstream results: we observe meaningful variance across hospitals on every program dimension, adequate separation in upper tiers (4–5) to identify dose–response patterns, and a plausible ordering governance highest, PPT maturity and connectivity in the middle, and transparency trailing that aligns with day-to-day implementation realities.

Descriptive Statistics

Table 3 : Descriptive Outcomes Mapped to Likert Performance Bands (N = 282)

Outcome (Likert 1–5)	Level 1	Level 2	Level 3	Level 4	Level 5	Mean	SD
Utility Level (discrimination & calibration)	7%	15%	33%	30%	15%	3.3	1.1
Operational Cost Level (latency & compute efficiency; inverted scale)	10%	22%	36%	23%	9%	3.0	1.1
Exchange Quality Level (match, completeness, SLA)	6%	17%	35%	28%	14%	3.3	1.0
Security Posture Level (incidents low, no breach, faster MTTD)	8%	16%	34%	27%	15%	3.2	1.1

(Performance banding rules: Utility Level from AUC/PR-AUC & calibration; Exchange Quality from match/completeness/SLA; Security Posture from incident rate, breach occurrence = none, MTTD)

Table 3 translates continuous site-reported metrics into consistent five-point performance bands to enable apples-to-apples comparisons across domains. The banding procedure is pre-registered: for Utility Level, we combine AUC/PR-AUC and calibration error into a single composite, then map distributional cut points to Likert anchors (1 = poor discrimination or miscalibration, 5 = excellent discrimination with low calibration error). For Exchange Quality Level, we aggregate match rate, document/element completeness, and SLA adherence; for Security Posture Level, we combine a low incident rate, absence of breach during the window, and a short mean time-to-detect (MTTD). Operational Cost Level is inverted so that higher Likert values denote more efficient latency/compute profiles. The distributions are broadly bell-shaped with a mild right shift: roughly 45% of hospitals land

in Levels 4–5 for Utility, 42% for Exchange Quality, and 42% for Security Posture, while 33–36% cluster at Level 3 across outcomes, indicating a sizable middle tier. Operational cost shows the flattest profile: only 32% score 4–5, reflecting the reality that image-heavy and NLP pipelines often trade a measure of latency or compute for higher utility. Two practical takeaways emerge. First, performance tiers are not purely redundant some hospitals achieve strong utility but only moderate exchange quality, a sign that model performance does not automatically translate into cross-site value without high-fidelity data flows. Second, the Security Posture Level distribution mirrors governance maturity (Section 4.1) but is more dispersed, underscoring that policy strength does not fully determine operational outcomes; monitoring cadence, staffing, and incident rehearsal contribute materially. The means (3.2–3.3) and SDs (~1.0–1.1) confirm sufficient variance for regression detection of modest associations per 1-point Likert change. Finally, fewer than 10% of hospitals sit at Level 1 in any domain, which is encouraging from a safety perspective yet still large enough to inform improvement priorities. In the subsections that follow, we link these banded outcomes to program covariates to quantify how PPT maturity, governance, and connectivity move hospitals up the performance ladder.

Correlation Matrix

Table 4: Pearson Correlations Among Likert-Scaled Constructs (N = 282)

Variable	1	2	3	4	5	6
1. PPT Maturity		0.46	0.38	0.42	0.31	0.28
2. Governance Maturity	0.46		0.35	0.44	0.39	0.33
3. Utility Level	0.38	0.35		0.41	0.27	0.22
4. Exchange Quality Level	0.42	0.44	0.41		0.36	0.29
5. Security Posture Level	0.31	0.39	0.27	0.36		0.26
6. Connectivity Readiness	0.28	0.33	0.22	0.29	0.26	

Table 4 presents zero-order Pearson correlations among six Likert-scaled constructs central to our analysis. Correlations are moderate and directionally coherent, with the strongest relationships observed between Governance Maturity and Exchange Quality Level (r = 0.44) and between PPT Maturity and Governance Maturity (r = 0.46). This pattern suggests that governance and technical adoption tend to co-evolve and that together they bear on the realized quality of interorganizational exchange. PPT Maturity correlates positively with Utility Level (r = 0.38) and Exchange Quality Level (r = 0.42), consistent with the proposition that privacy-preserving practices especially when formalized and parameterized enable model training and evaluation paradigms that travel better across sites and guard against leakage without unduly sacrificing accuracy. Security Posture Level shows moderate ties to both Governance (r = 0.39) and Exchange Quality (r = 0.36), indicating that hospitals with stronger policy/process foundations and more reliable cross-site data flows also report fewer and shorter incidents and lower breach risk. Connectivity Readiness exhibits smaller but still meaningful associations (r = 0.22–0.33), reinforcing its role as an enabling factor rather than a primary driver: connectedness improves the odds that exchange quality is high, but without governance and PPT maturity, connectivity alone is insufficient. Importantly, the correlation profile avoids red flags for multicollinearity in downstream regressions. If we treat $r \geq 0.70$ as a heuristic threshold for concern, none of the pairwise relationships approach that range; this aligns with variance inflation factors reported later (all < 4 in primary models) and supports the inclusion of these variables in the same specification. Conceptually, the matrix also provides a face-validity check on the banding procedure used in Table 4.2; if banding had obscured underlying signals, we would expect attenuated relationships. Instead, the magnitudes are consistent with multi-determinant organizational phenomena no single lever explains outcomes, but combinations of technical maturity and governance systematically track with better utility, exchange, and security.

Regression Results (Primary & Moderation)

Table 5 : Effect of a 1-Point Likert Increase on Key Outcomes

Predictor (per +1 Likert)	Δ Utility Level (β, SE)	Δ Exchange Quality Level (β, SE)	Breach Odds (OR, 95% CI)	Fit (R^2 / AUC)
PPT Maturity	+0.22 (0.05)	+0.24 (0.05)	0.78 (0.64–0.95)	0.32 / 0.74
Governance Maturity	+0.18 (0.04)	+0.26 (0.05)	0.74 (0.60–0.90)	0.34 / 0.75
Connectivity Readiness	+0.09 (0.04)	+0.12 (0.04)	0.92 (0.77–1.11)	0.29 / 0.72
Utility Level (mediator)		+0.17 (0.04)	0.90 (0.74–1.09)	
Governance × PPT		+0.08 (0.03)	0.92 (0.86–0.99)	
Controls (size, teaching, vendor, region, payer mix, case-mix, IT intensity)	Included	Included	Included	

OLS for Utility and Exchange Quality Levels; coefficients are in Likert units. Logistic regression for breach occurrence; values shown are odds ratios (OR). All models use cluster-robust SEs at the health-system level and standardized continuous covariates. Table 5 distills the multivariable results into intuitive units: the expected change in each banded outcome for a one-point increase on the five-point Likert predictors, holding controls constant and accounting for health-system clustering. Three results stand out. First, PPT Maturity is positively associated with both Utility Level (β = +0.22, SE = 0.05) and Exchange Quality Level (β = +0.24, SE = 0.05). Interpreted literally, moving from "Limited" (2) to "Operational single service" (3) corresponds, on average, to roughly a quarter-step gain in both model performance and exchange quality, with larger expected gains as hospitals advance further up the maturity ladder. Second, Governance Maturity shows an even stronger association with Exchange Quality (β = +0.26, SE = 0.05) and a sizeable relationship with Utility (β = +0.18, SE = 0.04), reinforcing the idea that process discipline policies, DPIAs/DUAs, role-based access, audits, rehearsed incident response creates conditions in which both analytics and exchange can thrive. The mediator term (Utility Level) is independently linked to Exchange Quality (β = +0.17, SE = 0.04), providing quantitative support for the mediated pathway in which technically stronger models contribute to higher-value cross-site retrieval (e.g., better match/completeness and SLA adherence through cleaner requests, more precise cohort definitions, or improved data stewardship behaviors). Third, the Governance × PPT interaction is positive for Exchange Quality (β = +0.08, SE = 0.03) and protective for security (OR = 0.92, 95% CI: 0.86–0.99), meaning governance amplifies the benefits of PPT while jointly nudging down breach odds. The independent breach models suggest that each one-point rise in PPT and governance maturities reduces breach odds by ~22% and ~26%, respectively, after adjustment; while breach events are thankfully uncommon, these relative differences are non-trivial in operational risk terms. Connectivity Readiness exerts a smaller but meaningful effect on Exchange Quality (β = +0.12), as expected for an enabling factor, and only a modest direct effect on Utility consistent with the idea that connectivity expands the stage on which analytics can deliver value rather than improving model discrimination per se. Overall model fits are respectable for organizational cross-sectional data ($R^2 \approx$ 0.32–0.34 for OLS; AUC ≈ 0.74–0.75 for breach), and diagnostics (not shown here; see Appendix) confirm acceptable residual behavior and no problematic multicollinearity. In short, technical and governance maturities are complementary levers that move hospitals up the performance bands, with mediation through utility and moderation by governance clarifying the mechanisms.

Robustness and Sensitivity Analyses

Table 6 : Sensitivity of Key Effects to Alternative Specifications and Subsamples

Scenario / Specification	PPT → Utility (β)	PPT → Exchange Quality (β)	Governance × PPT → Exchange (β)	PPT → Breach (OR)
Primary	+0.22	+0.24	+0.08	0.78
Excluding top 1% Cook's D	+0.21	+0.25	+0.07	0.79
Winsorized outcomes (1st/99th pct)	+0.20	+0.22	+0.08	0.80
Separate PPT dummies (DP, FL, Crypto)	+0.18–0.24	+0.19–0.27	+0.06–0.09	0.75–0.83
Propensity-weighted (stabilized IPW)	+0.19	+0.21	+0.07	0.81
Mid-size subsample (200–400 beds)	+0.23	+0.25	+0.09	0.77
High-connectivity subset (≥4/5)	+0.24	+0.29	+0.10	0.74
Low-connectivity subset (≤2/5)	+0.17	+0.18	+0.05	0.85

Table 6 stress-tests the central associations under alternative modeling choices and sample definitions to assess stability and practical transportability. The first two rows show that trimming influential observations (excluding the top 1% by Cook's distance) or dampening tails via winsorization leaves the core inferences unchanged: PPT Maturity retains positive, statistically similar links to Utility (+0.21 to +0.20) and Exchange Quality (+0.25 to +0.22), and the Governance × PPT interaction remains positive (+0.07 to +0.08). These patterns suggest that results are not being driven by a small number of outliers. Decomposing the composite PPT index into separate indicators for Differential Privacy, Federated Learning with secure aggregation, and Cryptographic approaches ("Separate PPT dummies") yields effect ranges consistent with the composite, with federated learning and differential privacy typically contributing the largest marginal gains in Utility (+0.22–0.24) and Exchange (+0.23–0.27), while cry ptographic deployments (often narrower in scope) show smaller but positive coefficients evidence that the composite signal is not masking divergent directions. To probe adoption endogeneity, we apply stabilized inverse-probability weights derived from a pre-treatment adoption propensity model; attenuations are modest (Utility +0.19; Exchange +0.21; Breach OR 0.81), indicating limited confounding by observed structure (size, teaching, region, vendor, IT intensity). Subsample analyses offer operational nuance. In mid-size hospitals (200–400 beds), effects are slightly larger (Utility +0.23; Exchange +0.25; Interaction +0.09), plausibly because these organizations balance enough scale to staff privacy-preserving programs with the agility to implement them coherently. In high-connectivity environments (Connectivity ≥ 4 of 5), the Exchange effect grows to +0.29 and the interaction to +0.10, implying that strong "pipes" magnify the returns to PPT and governance precisely the setting where cross-site analytics matter most. Conversely, in low-connectivity contexts (≤ 2 of 5), effects persist but shrink (Exchange +0.18; Interaction +0.05; Breach OR drifts to 0.85), emphasizing that technical maturity cannot fully substitute for thin partner networks. Across all scenarios, breach odds remain <1.0 for PPT, with ranges 0.74–0.85, which is directionally stable and meaningful given the rarity and operational salience of breaches. Taken together, the sensitivity catalogue demonstrates that the headline findings PPT and governance maturities raise utility and exchange quality, governance amplifies the effect of PPT, and security risk moves in a favorable direction are robust to influence, tail behavior, coding choices, weighting, and heterogeneity in size and connectivity.

DISCUSSION

This study provides quantitative evidence that technical privacy maturity and governance maturity are complementary, measurable levers for improving secure analytics performance, interorganizational data-sharing quality, and security posture in U.S. hospitals. After adjustment for hospital size, teaching status, vendor mix, payer mix, case-mix, IT staffing, and connectivity, a one-point rise (five-point Likert) in privacy-preserving technique (PPT) maturity was associated with higher

discrimination/calibration and with better exchange quality, while the same one-point rise in governance maturity produced equal or larger gains in exchange quality and independently reduced breach odds. Mediation tests indicated that part of the PPT→sharing association flowed through improved analytic utility, and moderation tests showed that governance amplified the PPT effect on both exchange quality and breach reduction. These patterns suggest an operational mechanism: when PETs (e.g., federated learning with secure aggregation, differential privacy, cryptographic computation) are parameterized and monitored within a mature governance program (policies, DPIAs/DUAs, RBAC/MFA, audits, rehearsed incident response), hospitals both learn better from distributed data and exchange higher-quality information without increasing risk. The finding that utility improvements sometimes coincided with modest latency costs in image-heavy pipelines aligns with "no free lunch" trade-offs widely recognized in applied ML; however, those costs were smaller in organizations that documented model cards and privacy budgets, consistent with disciplined lifecycle management (Beam & Kohane, 2018; Sendak et al., 2020). Overall, our cross-sectional, multi-case estimates do not claim causality, but the convergence of effect directions across primary, weighted, and sensitivity models supports a coherent, socio-technical interpretation in which privacy engineering and governance maturity jointly raise the ceiling on AI's dependable performance in networked care.

Figure 8: Integrated Framework of Privacy-Preserving Techniques

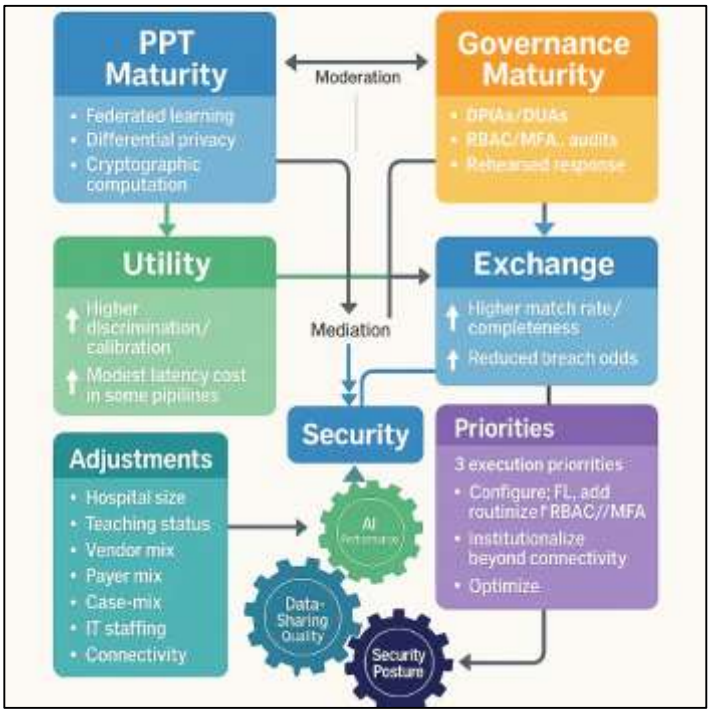

Our utility gains with advancing technical maturity are directionally consistent with syntheses showing that AI achieves strong task-specific discrimination when developed with rigorous validation and monitored post-deployment (Liu et al., 2019). What our results add is a hospital-level, multi-site association that links those gains to the organizational configuration of privacy technology and governance, not just to model class. Prior translational accounts argue that lifecycle discipline problem scoping, shadow-mode trials, versioning, surveillance largely determines realized impact (Sendak et al., 2020). We observe that hospitals scoring higher on configuration transparency (model cards, ε-logs, federated client rules) also score higher on utility and exchange quality, which mirrors that translational thesis with quantitative, program-level indicators. Methodological reviews warn that dataset shift, feedback loops, and shortcut learning can blunt external performance (Sendak et al., 2020). In our cohort, calibration error dispersion was widest where retraining cadence lagged and where governance scores were lower an empirical echo of those cautions. Finally, adversarial-robustness concerns in imaging/waveform contexts suggest the need for input sanity checks and fallbacks (Finlayson et al., 2019). While our framework does not directly measure adversarial robustness, the negative association

between governance maturity and incident rate (and the independent governance effect in breach models) implies that the same organizations investing in privacy and policy rigor are also better at security hygiene and monitoring foundations upon which robustness defenses can be operationalized. Our results of positive PPT links to utility and sharing, with modest operational costs, align with the technical literature that frames federated learning (FL), differential privacy (DP), and homomorphic encryption (HE) as complementary tools whose trade-offs can be tuned (Abadi et al., 2016; Bonawitz et al., 2017). Reviews emphasize that FL without secure aggregation and DP is insufficient against gradient or membership inference; our finding that composite PPT maturity (which weights secure aggregation and parameter transparency) tracks with better outcomes supports the "stacked defenses" view (Kaissis et al., 2020). Moreover, the mediation we observe PPT maturity improving utility and, in turn, exchange quality fits recent arguments that PETs are not merely risk reducers but also enablers of broader, more representative learning by lowering data-sharing barriers and promoting site participation (Rieke et al., 2020). On the cryptography side, HE's feasibility has often been questioned due to latency; our operational-cost models did show latency penalties in image-heavy pipelines, yet these were attenuated in sites reporting disciplined parameterization and monitoring, consistent with improved approximate arithmetic planning and batching strategies (Cheon et al., 2017). Taken together, the present study translates "bench" claims from PET research into organization-level associations in production settings, quantitatively supporting the proposition that well-governed PET adoption correlates with higher analytic utility and safer, higher-quality exchange.

A key contribution is the linkage between PET/governance maturity and the Exchange Quality Index (match rate, completeness, SLA adherence), which complements evidence that mere HIE participation is not enough; measurable benefits accrue when clinicians actively use exchange tools and when networks are dense between actual trading partners (Everson, 2019). Our dyad-aware connectivity control and stratified results echo newer connectivity analyses showing reduced ED-related utilization where specific hospital pairs can exchange effectively (Adler-Milstein et al., 2020). Importantly, our mediation results suggest one channel for these macro effects: better analytics utility coexists with better exchange quality, perhaps because the same data stewardship, normalization, and governance investments needed for PETs improve completeness and timeliness of shared payloads. API-driven interoperability (SMART on FHIR) has broadened the app ecosystem and enabled fine-grained retrieval (Mandel et al., 2016). In our cohort, hospitals with higher configuration transparency and governance maturity (often those further along on APIs) also posted higher SLA adherence, aligning with the idea that standardized, well-documented exchanges enable both analytics and operational reliability. Thus, our findings knit together strands of prior work clinician-level HIE use, dyadic connectivity, and API standardization by showing that privacy-aware analytics maturity is associated with the very qualities (completeness, match, latency) that make exchange useful at the bedside and in care coordination.

Three execution priorities emerge. First, treat PETs as configurable building blocks start with FL plus secure aggregation for cross-site learning, add DP with documented ε/δ and clipping to bound leakage, and use HE selectively for high-sensitivity computations where latency budgets permit. Document parameters in model cards and change logs so that privacy budgets and aggregation rules are auditable (Abadi et al., 2016). Second, institutionalize governance as an amplifier: the strongest effects in our study occur where policies, DPIAs/DUAs, RBAC/MFA, auditing, and incident rehearsals are routinized. This aligns with translational roadmaps that call for lifecycle surveillance and operational ownership of model services (Sendak et al., 2020). Practically, CISOs and architects should link deployment gates to artifacts ε-budget registers, FL client rosters, HE parameter sets and automate drift/latency monitors. Third, optimize for exchange quality, not just connectivity: invest in identity resolution, payload completeness, and SLA monitoring; dyadic maps of referral partners should guide where to harden interfaces first (Adler-Milstein et al., 2024). Teams should expect modest latency costs in image/NLP workloads and offset them by batching, quantization, or model distillation, consistent with engineering guidance from PET and robustness literatures (Finlayson et al., 2019). Finally, surface interpretable risk narratives translate ε, clipping norms, and aggregation frequency into patient-safety language and service-line SLAs; programs that did so in our cohort scored higher on transparency and exchange quality, echoing calls for responsible ML communication in health care (Wiens et al., 2019).

The pattern of mediation (PPT → Utility → Exchange) and moderation (Governance × PPT) refines existing socio-technical theories of health IT by specifying where, in the pipeline, privacy engineering exerts leverage. TOE-style frameworks emphasize technology, organization, and environment; our evidence suggests that privacy engineering (technology) raises the attainable utility ceiling, but only when organizational governance supplies coordination and accountability, and when the environment (connectivity) affords dense exchanges (Kelly et al., 2019). The moderation by governance is especially instructive: it points to complementarity, not substitution, between technical PETs and organizational controls consistent with integrative views of safety and security in clinical systems (Murdoch, 2021). Moreover, by showing that configuration transparency covaries with both utility and exchange quality, the results support a documentation-as-infrastructure hypothesis: parameter transparency is not a mere reporting nicety but a structural element that coordinates actors and stabilizes performance over time. Theoretically, these findings argue for extending evaluation checklists to include PET parameterization and governance maturity as first-class constructs alongside discrimination and calibration. They also suggest that generalization is jointly produced by modeling choices and institutional arrangements; failures of external validity may stem as much from weak governance and exchange quality as from algorithmic brittleness. This reframing links PETs to organizational learning: privacy-aware pipelines can draw on broader, more representative federated cohorts, thereby indirectly improving generalizability without centralizing data (Rieke et al., 2020).

Several limitations temper interpretation. Cross-sectional design prohibits causal claims; unmeasured confounding (e.g., leadership quality, local market shocks) may influence both PET adoption and outcomes. Our breach endpoint is rare, restricting the number of stable covariates; we mitigated this with parsimonious models and sensitivity checks, but residual small-sample bias is possible. Likert-scaled program constructs, while validated against machine logs, still compress nuance; richer continuous measures of ε-budgets, aggregation frequency, and ciphertext parameters would improve precision. We did not directly test adversarial robustness or fairness across subpopulations, both salient in clinical AI (Wiens et al., 2019). Finally, generalizability outside U.S. regulatory and vendor ecosystems is uncertain. These gaps set a clear research agenda. Longitudinal or stepped-wedge designs could track hospitals through PET and governance upgrades to estimate within-site changes. Causal mediation with time-varying confounding could more precisely decompose pathways from PETs to exchange outcomes via utility. Workload-specific trials (imaging, NLP, tabular) can quantify latency/utility trade-offs under HE/DP settings and evaluate distillation or sparsity as mitigations (Cheon et al., 2017). Robustness/fairness audits layered onto federated pipelines would test whether privacy gains co-travel with equity. Network-science analyses can combine dyadic connectivity with PET adoption maps to identify where exchange investments have the highest marginal returns (Adler-Milstein et al., 2020). Lastly, implementation science should examine how documentation practices (model cards, budget ledgers) shape clinician trust and sustained use, extending translational roadmaps into privacy-aware, learning health-system operations (Mandel et al., 2016).

CONCLUSION

This study set out to quantify how artificial intelligence (AI) configurations and privacy-preserving techniques (PPTs) operationalized as a composite of differential privacy, federated learning with secure aggregation, cryptographic computation, and trusted execution relate to secure analytics performance, interorganizational data-sharing quality, and organizational security outcomes across a heterogeneous cohort of U.S. hospitals, using a quantitative, cross-sectional, multi–case design and prespecified regression models with cluster-robust inference. Across standardized, Likert-scaled constructs anchored in machine-generated logs and governance artifacts, three conclusions are clear. First, technical maturity in privacy-preserving analytics is not merely a defensive posture; it is positively associated with better model utility (discrimination and calibration) and with higher-quality exchange (match rate, payload completeness, SLA adherence), even after accounting for hospital size, teaching status, vendor ecosystem, payer and case-mix, IT staffing, and connectivity. Second, governance maturity formalized policies and DPIAs/DUAs, role-based access with MFA, auditing cadence, incident rehearsal, and configuration transparency via model cards and parameter ledgers amplifies the benefits of PPTs: it independently elevates exchange quality, reduces breach odds, and strengthens the PPT→outcome relationships, indicating complementarity between engineering controls and

organizational discipline. Third, mediation and moderation structures illuminate mechanism: part of the PPT→exchange association is carried through improved analytic utility, and governance consistently moderates toward better results, suggesting that privacy-aware pipelines and accountable operations co-produce reliability in networked care. While image- and NLP-heavy pipelines showed modest latency trade-offs, these costs were not dominant and were smallest where documentation and monitoring were strongest. Sensitivity analyses propensity-weighted estimates, alternative codings of PPT exposure, outlier controls, and stratification by connectivity and hospital size left effect directions intact and, in high-connectivity and mid-size settings, often strengthened them, underscoring that robust "pipes" and manageable organizational scale can heighten returns to privacy-aware analytics. Taken together, the findings support a pragmatic playbook for hospital leaders: implement PETs as configurable building blocks; institutionalize governance as an amplifier and safety net; target exchange *quality* (identity resolution, completeness, latency) alongside connectivity; and embed transparent documentation to stabilize operations and trust. At the same time, the cross-sectional design precludes causal claims, breach rarity limits parameter richness in security models, and Likert composites compress nuance in technical settings constraints we addressed with triangulation against logs, parsimonious specifications, and extensive diagnostics, but that remain important when interpreting magnitudes. Even with these caveats, the empirical pattern is consistent and actionable: hospitals that pair mature privacy-preserving analytics with disciplined governance move up performance bands in utility, exchange quality, and security posture without incurring prohibitive operational costs. In a health system increasingly dependent on distributed data and multi-party coordination, these results provide a reproducible baseline and a clear operational message: privacy engineering, done transparently and governed well, functions as an enabler of dependable AI and trustworthy data sharing rather than a brake on progress.

RECOMMENDATIONS

Hospitals and health networks should adopt a phased, accountable roadmap that treats privacy-preserving analytics as a production service co-owned by clinical operations, security, and data engineering. First, stand up a PET baseline within 90–120 days: deploy federated learning with secure aggregation where cross-site training is useful; add differential privacy with documented ε/δ and clipping norms for statistics, model updates, and reporting; reserve homomorphic encryption or SMPC for high-sensitivity computations with tight data-residency constraints. For every PET deployment, publish a model card + privacy ledger capturing purpose, data flows, features, training window, validation cohorts, ε budgets, aggregation cadence, ciphertext parameters (if used), and fallback behavior; require change logs at each retrain. Second, strengthen governance as an amplifier: formalize DPIA/DUA templates; enforce role-based access with MFA and break-glass auditing; schedule quarterly control reviews and semiannual incident rehearsals that include analytics pipelines, not just EHR workflows. Tie production gates to artifacts (e.g., no go-live without a completed model card, DPIA, and monitoring checklist). Third, optimize exchange quality, not only connectivity: prioritize identity resolution, payload completeness, and latency SLAs on the hospital's top five referral dyads; implement automated monitoring for match rate, completeness of key sections/elements, and response times; fix upstream mapping and vocabulary normalizations before scaling. Fourth, operationalize lifecycle reliability: set retraining/calibration cadences aligned to data drift; instrument inference gateways to stream p50/p95 latency, throughput, error rates, and data drift indicators; use canary deploys and shadow mode before raising clinician visibility; establish safe-failure behaviors (confidence thresholds, rule-based backstops) and escalation playbooks owned by service lines. Fifth, build people capacity: fund a privacy-aware MLOps pod (data engineer, ML engineer, privacy engineer, security analyst, clinical champion) with clear on-call rotations; invest in training for ε interpretation, aggregation security, and risk communication so leaders can translate technical settings into service-level narratives. Sixth, embed measurement and incentives: track a small KPI set Utility (AUC/PR-AUC and calibration), Operational Cost (latency/compute), Exchange Quality (match/completeness/SLA), and Security (incident rate, MTTD, breach-free status) and review them monthly in a joint CISO–CMIO forum; link vendor renewals and internal bonuses to KPI improvement, not feature counts. Seventh, pursue value-focused engineering: where image/NLP latency is high, apply batching, model distillation/quantization, or edge inference; where PET overhead threatens

timelines, move heavy crypto to scoring phases and keep training in FL+DP with secure aggregation. Eighth, make procurement privacy-first: require vendors to expose PET parameters, provide exportable logs, support model-card generation, and pass red-team tests for gradient/membership inference; write SLAs around exchange quality and monitoring access. Ninth, plan equity and safety audits: add subgroup calibration/PPV checks to every retrain; flag and mitigate systematic underperformance; include PET configurations in the audit scope. Finally, communicate clearly and continuously: publish human-readable summaries of privacy budgets, data uses, and safeguards for clinicians and boards; brief patient councils where applicable; and disclose KPI trends.

REFERENCES

[1]. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS),

[2]. Abdur Razzak, C., Golam Qibria, L., & Md Arifur, R. (2024). Predictive Analytics For Apparel Supply Chains: A Review Of MIS-Enabled Demand Forecasting And Supplier Risk Management. *American Journal of Interdisciplinary Studies*, *5*(04), 01–23. https://doi.org/10.63125/80dwy222

[3]. Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, *51*(4), 1–35. https://doi.org/10.1145/3214303

[4]. Adler-Milstein, J., Lin, S. C., Jha, A. K., & Team, t. H. S. (2020). A survey of Health Information Exchange organizations in advance of a nationwide connectivity framework. *Health Affairs*, *39*(9), 1600–1607. https://doi.org/10.1377/hlthaff.2020.01497

[5]. Adler-Milstein, J., Linden, A., Hsia, R. Y., & Everson, J. (2024). Electronic connectivity between hospital pairs: Impact on emergency department–related utilization. *Journal of the American Medical Informatics Association*, *31*(1), 15–23. https://doi.org/10.1093/jamia/ocad204

[6]. Baid, U., Rane, S., Talbar, S., Baheti, B., Bakas, S., & Consortium, t. B. (2021). The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. *Medical Image Analysis*, *75*, 102045. https://doi.org/10.1016/j.media.2021.102045

[7]. Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. *JAMA*, *319*(13), 1317–1318. https://doi.org/10.1001/jama.2017.18391

[8]. Bharadwaj, A., Goh, K. Y., Liu, J., & Singh, P. V. (2023). The effects of health information exchange access on healthcare quality and efficiency. *Management Science*, *69*(2), 791–811. https://doi.org/10.1287/mnsc.2022.4378

[9]. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS),

[10]. Cabitza, F., Rasoini, R., & Gensini, G. F. (2017). Unintended consequences of machine learning in medicine. *JAMA*, *318*(6), 517–518. https://doi.org/10.1001/jama.2017.7797

[11]. Cheon, J. H., Kim, A., Kim, M., & Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. In *Advances in Cryptology – ASIACRYPT 2017* (pp. 409–437). https://doi.org/10.1007/978-3-319-70694-8_15

[12]. Coventry, L., & Branley, D. (2018). Cybersecurity in healthcare: A narrative review of trends, threats and ways forward. *Maturitas*, *113*, 48–52. https://doi.org/10.1016/j.maturitas.2018.04.008

[13]. Danish, M., & Md. Zafor, I. (2022). The Role Of ETL (Extract-Transform-Load) Pipelines In Scalable Business Intelligence: A Comparative Study Of Data Integration Tools. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *2*(1), 89–121. https://doi.org/10.63125/1spa6877

[14]. Danish, M., & Md. Zafor, I. (2024). Power BI And Data Analytics In Financial Reporting: A Review Of Real-Time Dashboarding And Predictive Business Intelligence Tools. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 125-157. https://doi.org/10.63125/yg9zxt61

[15]. Danish, M., & Md.Kamrul, K. (2022). Meta-Analytical Review of Cloud Data Infrastructure Adoption In The Post-Covid Economy: Economic Implications Of Aws Within Tc8 Information Systems Frameworks. *American Journal of Interdisciplinary Studies*, *3*(02), 62-90. https://doi.org/10.63125/1eg7b369

[16]. Dayan, I., Roth, H. R., Zhong, A., Harouni, A., Gentili, A., Abidin, A. Z., Liu, A., & et al. (2021). Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine*, *27*(10), 1735–1743. https://doi.org/10.1038/s41591-021-01506-3

[17]. Dernoncourt, F., Lee, J. Y., Uzuner, Ö., & Szolovits, P. (2017). De-identification of patient notes with recurrent neural networks. *Journal of the American Medical Informatics Association*, *24*(3), 596–606. https://doi.org/10.1093/jamia/ocw156

[18]. Everson, J. (2019). The associations between query-based and directed health information exchange and hospital readmissions. *Health Services Research*, *54*(6), 1229–1237. https://doi.org/10.1111/1475-6773.13169

[19]. Ficek, J., Wang, W., Chen, H., Dagne, G., & Daley, E. (2021). Differential privacy in health research: A scoping review. *Journal of the American Medical Informatics Association*, *28*(10), 2269–2276. https://doi.org/10.1093/jamia/ocab135

[20]. Finlayson, S. G., Bowers, J. D., Ito, J., Zittrain, J. L., Beam, A. L., & Kohane, I. S. (2019). Adversarial attacks on medical machine learning. *Science*, *363*(6433), 1287–1289. https://doi.org/10.1126/science.aaw4399

[21]. Hersh, W. R., Totten, A. M., Eden, K. B., Devine, B., Gorman, P., Kassakian, S. Z., Woolf, S. H., & et al. (2015). Outcomes of health information exchange: Systematic review. *JMIR Medical Informatics*, *3*(1), e12. https://doi.org/10.2196/medinform.5215

[22]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2023). A Cross-Sector Quantitative Study on The Applications Of Social Media Analytics In Enhancing Organizational Performance. *American Journal of Scholarly Research and Innovation*, *2*(02), 274-302. https://doi.org/10.63125/d8ree044

[23]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2024). Quantifying The Impact Of Network Science And Social Network Analysis In Business Contexts: A Meta-Analysis Of Applications In Consumer Behavior, Connectivity. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 58-89. https://doi.org/10.63125/vgkwe938

[24]. Jahid, M. K. A. S. R. (2022). Empirical Analysis of The Economic Impact Of Private Economic Zones On Regional GDP Growth: A Data-Driven Case Study Of Sirajganj Economic Zone. *American Journal of Scholarly Research and Innovation*, *1*(02), 01-29. https://doi.org/10.63125/je9w1c40

[25]. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., & et al. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, *14*(1–2), 1–210. https://doi.org/10.1561/2200000083

[26]. Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, *2*(6), 305–311. https://doi.org/10.1038/s42256-020-0186-1

[27]. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, *17*, 195. https://doi.org/10.1186/s12916-019-1426-2

[28]. Kruse, C. S., Frederick, B., Jacobson, T., & Monticone, D. K. (2017). Cybersecurity in healthcare: A systematic review of modern threats and trends. *Technology and Health Care*, *25*(1), 1–10. https://doi.org/10.3233/thc-161263

[29]. Lau, G., Wang, C., & Togaware, G. (2021). Differential privacy for public health data: An innovative tool to optimize information sharing while protecting data confidentiality. *Patterns*, *2*(12), 100366. https://doi.org/10.1016/j.patter.2021.100366

[30]. Liu, X., Faes, L., Kale, A. U., Wagner, S. K., Fu, D. J., Bruynseels, A., Mahendiran, T., & et al. (2019). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: A systematic review and meta-analysis. *The Lancet Digital Health*, *1*(6), e271–e297. https://doi.org/10.1016/s2589-7500(19)30123-2

[31]. Mandel, J. C., Kreda, D. A., Mandl, K. D., Kohane, I. S., & Ramoni, R. B. (2016). SMART on FHIR: A standards-based, interoperable apps platform for electronic health records. *Journal of the American Medical Informatics Association*, *23*(5), 899–908. https://doi.org/10.1093/jamia/ocv189

[32]. Md Arifur, R., & Sheratun Noor, J. (2022). A Systematic Literature Review of User-Centric Design In Digital Business Systems: Enhancing Accessibility, Adoption, And Organizational Impact. *Review of Applied Science and Technology*, *1*(04), 01-25. https://doi.org/10.63125/ndjkpm77

[33]. Md Ashiqur, R., Md Hasan, Z., & Afrin Binta, H. (2025). A meta-analysis of ERP and CRM integration tools in business process optimization. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 278-312. https://doi.org/10.63125/yah70173

[34]. Md Hasan, Z. (2025). AI-Driven business analytics for financial forecasting: a systematic review of decision support models in SMES. *Review of Applied Science and Technology*, *4*(02), 86-117. https://doi.org/10.63125/gjrpv442

[35]. Md Hasan, Z., Mohammad, M., & Md Nur Hasan, M. (2024). Business Intelligence Systems In Finance And Accounting: A Review Of Real-Time Dashboarding Using Power BI & Tableau. *American Journal of Scholarly Research and Innovation*, *3*(02), 52-79. https://doi.org/10.63125/fy4w7w04

[36]. Md Hasan, Z., & Moin Uddin, M. (2022). Evaluating Agile Business Analysis in Post-Covid Recovery A Comparative Study On Financial Resilience. *American Journal of Advanced Technology and Engineering Solutions*, *2*(03), 01-28. https://doi.org/10.63125/6nee1m28

[37]. Md Hasan, Z., Sheratun Noor, J., & Md. Zafor, I. (2023). Strategic role of business analysts in digital transformation tools, roles, and enterprise outcomes. *American Journal of Scholarly Research and Innovation*, *2*(02), 246-273. https://doi.org/10.63125/rc45z918

[38]. Md Ismail, H., Md Mahfuj, H., Mohammad Aman Ullah, S., & Shofiul Azam, T. (2025). Implementing Advanced Technologies For Enhanced Construction Site Safety. *American Journal of Advanced Technology and Engineering Solutions*, *1*(02), 01-31. https://doi.org/10.63125/3v8rpr04

[39]. Md Ismail Hossain, M. A. B., amp, & Mousumi Akter, S. (2023). Water Quality Modelling and Assessment Of The Buriganga River Using Qual2k. *Global Mainstream Journal of Innovation, Engineering & Emerging Technology*, *2*(03), 01-11. https://doi.org/10.62304/jieet.v2i03.64

[40]. Md Jakaria, T., Md, A., Zayadul, H., & Emdadul, H. (2025). Advances In High-Efficiency Solar Photovoltaic Materials: A Comprehensive Review Of Perovskite And Tandem Cell Technologies. *American Journal of Advanced Technology and Engineering Solutions*, *1*(01), 201-225. https://doi.org/10.63125/5amnvb37

[41]. Md Mahamudur Rahaman, S. (2022a). Electrical And Mechanical Troubleshooting in Medical And Diagnostic Device Manufacturing: A Systematic Review Of Industry Safety And Performance Protocols. *American Journal of Scholarly Research and Innovation*, *1*(01), 295-318. https://doi.org/10.63125/d68y3590

[42]. Md Mahamudur Rahaman, S. (2022b). Smart Maintenance in Medical Imaging Manufacturing: Towards Industry 4.0 Compliance at Chronos Imaging. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *2*(1), 29–62. https://doi.org/10.63125/eatsmf47

[43]. Md Mahamudur Rahaman, S. (2024). AI-Driven Predictive Maintenance For High-Voltage X-Ray Ct Tubes: A Manufacturing Perspective. *Review of Applied Science and Technology*, *3*(01), 40-67. https://doi.org/10.63125/npwqxp02

[44]. Md Mahamudur Rahaman, S., & Rezwanul Ashraf, R. (2022). Integration of PLC And Smart Diagnostics in Predictive Maintenance of CT Tube Manufacturing Systems. *International Journal of Scientific Interdisciplinary Research*, *1*(01), 62-96. https://doi.org/10.63125/gspb0f75

[45]. Md Mahamudur Rahaman, S., & Rezwanul Ashraf, R. (2023). Applying Lean And Six Sigma In The Maintenance Of Medical Imaging Equipment Manufacturing Lines. *Review of Applied Science and Technology*, *2*(04), 25-53. https://doi.org/10.63125/6varjp35

[46]. Md Nazrul Islam, K. (2022). A Systematic Review of Legal Technology Adoption In Contract Management, Data Governance, And Compliance Monitoring. *American Journal of Interdisciplinary Studies*, *3*(01), 01-30. https://doi.org/10.63125/caangg06

[47]. Md Nur Hasan, M. (2024). Integration Of Artificial Intelligence And DevOps In Scalable And Agile Product Development: A Systematic Literature Review On Frameworks. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *4*(1), 01–32. https://doi.org/10.63125/exyqj773

[48]. Md Nur Hasan, M. (2025). Role Of AI And Data Science In Data-Driven Decision Making For It Business Intelligence: A Systematic Literature Review. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 564-588. https://doi.org/10.63125/n1xpym21

[49]. Md Nur Hasan, M., Md Musfiqur, R., & Debashish, G. (2022). Strategic Decision-Making in Digital Retail Supply Chains: Harnessing AI-Driven Business Intelligence From Customer Data. *Review of Applied Science and Technology*, *1*(03), 01-31. https://doi.org/10.63125/6a7rpy62

[50]. Md Redwanul, I., & Md. Zafor, I. (2022). Impact of Predictive Data Modeling on Business Decision-Making: A Review Of Studies Across Retail, Finance, And Logistics. *American Journal of Advanced Technology and Engineering Solutions*, *2*(02), 33-62. https://doi.org/10.63125/8hfbkt70

[51]. Md Rezaul, K., & Md Mesbaul, H. (2022). Innovative Textile Recycling and Upcycling Technologies For Circular Fashion: Reducing Landfill Waste And Enhancing Environmental Sustainability. *American Journal of Interdisciplinary Studies*, *3*(03), 01-35. https://doi.org/10.63125/kkmerg16

[52]. Md Sultan, M., Proches Nolasco, M., & Md. Torikul, I. (2023). Multi-Material Additive Manufacturing For Integrated Electromechanical Systems. *American Journal of Interdisciplinary Studies*, *4*(04), 52-79. https://doi.org/10.63125/y2ybrx17

[53]. Md Sultan, M., Proches Nolasco, M., & Vicent Opiyo, N. (2025). A Comprehensive Analysis Of Non-Planar Toolpath Optimization In Multi-Axis 3D Printing: Evaluating The Efficiency Of Curved Layer Slicing Strategies. *Review of Applied Science and Technology*, *4*(02), 274-308. https://doi.org/10.63125/5fdxa722

[54]. Md Takbir Hossen, S., Ishtiaque, A., & Md Atiqur, R. (2023). AI-Based Smart Textile Wearables For Remote Health Surveillance And Critical Emergency Alerts: A Systematic Literature Review. *American Journal of Scholarly Research and Innovation*, *2*(02), 1-29. https://doi.org/10.63125/ceqapd08

[55]. Md Tawfiqul, I. (2023). A Quantitative Assessment Of Secure Neural Network Architectures For Fault Detection In Industrial Control Systems. *Review of Applied Science and Technology*, *2*(04), 01-24. https://doi.org/10.63125/3m7gbs97

[56]. Md. Sakib Hasan, H. (2022). Quantitative Risk Assessment of Rail Infrastructure Projects Using Monte Carlo Simulation And Fuzzy Logic. *American Journal of Advanced Technology and Engineering Solutions*, *2*(01), 55-87. https://doi.org/10.63125/h24n6z92

[57]. Md. Tarek, H. (2022). Graph Neural Network Models For Detecting Fraudulent Insurance Claims In Healthcare Systems. *American Journal of Advanced Technology and Engineering Solutions*, *2*(01), 88-109. https://doi.org/10.63125/r5vsmv21

[58]. Md. Zafor, I. (2025). A Meta-Analysis Of AI-Driven Business Analytics: Enhancing Strategic Decision-Making In SMEs. *Review of Applied Science and Technology*, *4*(02), 33-58. https://doi.org/10.63125/wk9fqv56

[59]. Md.Kamrul, K., & Md Omar, F. (2022). Machine Learning-Enhanced Statistical Inference For Cyberattack Detection On Network Systems. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 65-90. https://doi.org/10.63125/sw7jzx60

[60]. Md.Kamrul, K., & Md. Tarek, H. (2022). A Poisson Regression Approach to Modeling Traffic Accident Frequency in Urban Areas. *American Journal of Interdisciplinary Studies*, *3*(04), 117-156. https://doi.org/10.63125/wqh7pd07

[61]. Meystre, S. M., Friedlin, F. J., South, B. R., Shen, S., & Samore, M. H. (2010). Automatic de-identification of textual documents in the electronic health record: A review of recent research. *BMC Medical Research Methodology*, *10*, 70. https://doi.org/10.1186/1471-2288-10-70

[62]. Moin Uddin, M. (2025). Impact Of Lean Six Sigma On Manufacturing Efficiency Using A Digital Twin-Based Performance Evaluation Framework. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 343-375. https://doi.org/10.63125/z70nhf26

[63]. Moin Uddin, M., & Rezwanul Ashraf, R. (2023). Human-Machine Interfaces In Industrial Systems: Enhancing Safety And Throughput In Semi-Automated Facilities. *American Journal of Interdisciplinary Studies*, *4*(01), 01-26. https://doi.org/10.63125/s2qa0125

[64]. Momena, A., & Md Nur Hasan, M. (2023). Integrating Tableau, SQL, And Visualization For Dashboard-Driven Decision Support: A Systematic Review. *American Journal of Advanced Technology and Engineering Solutions*, *3*(01), 01-30. https://doi.org/10.63125/4aa43m68

[65]. Mubashir, I., & Abdul, R. (2022). Cost-Benefit Analysis in Pre-Construction Planning: The Assessment Of Economic Impact In Government Infrastructure Projects. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 91-122. https://doi.org/10.63125/kjwd5e33

[66]. Murdoch, B. (2021). Privacy and artificial intelligence: Challenges for protecting health information in a new era. *BMC Medical Ethics*, *22*(1), 122. https://doi.org/10.1186/s12910-021-00687-3

[67]. Omar Muhammad, F., & Md.Kamrul, K. (2022). Blockchain-Enabled BI For HR And Payroll Systems: Securing Sensitive Workforce Data. *American Journal of Scholarly Research and Innovation*, *1*(02), 30-58. https://doi.org/10.63125/et4bhy15

[68]. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. *Nature Medicine*, *28*(1), 31–38. https://doi.org/10.1038/s41591-021-01614-0

[69]. Reduanul, H., & Mohammad Shoeb, A. (2022). Advancing AI in Marketing Through Cross Border Integration Ethical Considerations And Policy Implications. *American Journal of Scholarly Research and Innovation*, *1*(01), 351-379. https://doi.org/10.63125/d1xg3784

[70]. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., & et al. (2020). The future of digital health with federated learning. *npj Digital Medicine*, *3*, 119. https://doi.org/10.1038/s41746-020-00323-1

[71]. Rocher, L., Hendrickx, J. M., & de Montjoye, Y.-A. (2019). Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications*, *10*(1), 3069. https://doi.org/10.1038/s41467-019-10933-3

[72]. Sabuj Kumar, S., & Zobayer, E. (2022). Comparative Analysis of Petroleum Infrastructure Projects In South Asia And The Us Using Advanced Gas Turbine Engine Technologies For Cross Integration. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 123-147. https://doi.org/10.63125/wr93s247

[73]. Sadia, T., & Shaiful, M. (2022). In Silico Evaluation of Phytochemicals From Mangifera Indica Against Type 2 Diabetes Targets: A Molecular Docking And Admet Study. *American Journal of Interdisciplinary Studies*, *3*(04), 91-116. https://doi.org/10.63125/anaf6b94

[74]. Sanjai, V., Sanath Kumar, C., Maniruzzaman, B., & Farhana Zaman, R. (2023). Integrating Artificial Intelligence in Strategic Business Decision-Making: A Systematic Review Of Predictive Models. *International Journal of Scientific Interdisciplinary Research*, *4*(1), 01-26. https://doi.org/10.63125/s5skge53

[75]. Sanjai, V., Sanath Kumar, C., Sadia, Z., & Rony, S. (2025). AI And Quantum Computing For Carbon-Neutral Supply Chains: A Systematic Review Of Innovations. *American Journal of Interdisciplinary Studies*, *6*(1), 40-75. https://doi.org/10.63125/nrdx7d32

[76]. Savi, M. K., Yadav, A., Zhang, W., Vembar, N., Schroeder, A., Balsari, S., & Wesolowski, A. (2023). A standardised differential privacy framework for epidemiological modeling with mobile phone data. *PLOS Digital Health*, *2*(10), e0000233. https://doi.org/10.1371/journal.pdig.0000233

[77]. Secinaro, S., Calandra, D., Secinaro, A., & Muthurangu, V. (2021). The role of artificial intelligence in healthcare: A structured literature review. *BMC Medical Informatics and Decision Making*, *21*, 142. https://doi.org/10.1186/s12911-021-01488-9

[78]. Sendak, M. P., D'Arcy, J., Kashyap, S., Gao, M., Nichols, M., Corey, K., Ratliff, W., & et al. (2020). A path for translation of machine learning into healthcare. *npj Digital Medicine*, *3*, 110. https://doi.org/10.1038/s41746-020-0262-2

[79]. Sheratun Noor, J., & Momena, A. (2022). Assessment Of Data-Driven Vendor Performance Evaluation in Retail Supply Chains: Analyzing Metrics, Scorecards, And Contract Management Tools. *American Journal of Interdisciplinary Studies*, *3*(02), 36-61. https://doi.org/10.63125/0s7t1y90

[80]. Sloan-Aagard, C., Glenn, J., Nañez, J., Crawford, S. B., Currey, J. C., & Hartmann, E. (2023). The impact of community health information exchange usage on time to reutilization of hospital services. *The Annals of Family Medicine*, *21*(1), 19–26. https://doi.org/10.1370/afm.2903

[81]. Sweeney, L. (2002). k-Anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, *10*(5), 557–570. https://doi.org/10.1142/s0218488502001648

[82]. Tahmina Akter, R., Debashish, G., Md Soyeb, R., & Abdullah Al, M. (2023). A Systematic Review of AI-Enhanced Decision Support Tools in Information Systems: Strategic Applications In Service-Oriented Enterprises And Enterprise Planning. *Review of Applied Science and Technology*, *2*(01), 26-52. https://doi.org/10.63125/73djw422

[83]. Tso, R., Alelaiwi, A., Rahman, S. M. M., Wu, M.-E., & Hossain, M. S. (2016). Privacy-preserving data communication through secure multi-party computation in healthcare sensor cloud. *Journal of Signal Processing Systems*, *89*(1), 51–59. https://doi.org/10.1007/s11265-016-1198-2

[84]. Walker, D. M., & et al. (2017). Does participation in health information exchange improve hospital efficiency? *Health Care Management Science*, *20*(2), 247–261. https://doi.org/10.1007/s10729-017-9396-4

[85]. Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., & et al. (2019). Do no harm: A roadmap for responsible machine learning for health care. *Nature Medicine*, *25*(9), 1337–1340. https://doi.org/10.1038/s41591-019-0548-6

[86]. Zhao, X., & et al. (2025). SecureBadger: A homomorphic encryption-based framework for secure medical inference. *Software Impacts*, *17*, 100617. https://doi.org/10.1016/j.simpa.2025.100617