

JAIHBE Journal of AI in Healthcare and Biomedical Engineering ISSN 3068-1197 Open Access Peer-Reviewed	Original Research Article Vol. 1, No. 4 (2023), pp. 1-17 DOI: 10.63646/jaihbe.2023.010401
--	---

AI and Blockchain for Patient-Centric Genomic Data Governance: A Privacy-Preserving Healthcare Engineering Framework

Farah Nabilah Ahmad¹, Lim Wei Han², Siti Aina Mohd Noor^{3,*}

¹ Faculty of Computing, Universiti Malaysia Pahang Al-Sultan Abdullah, Pekan 26600, Pahang, Malaysia. Email: farah.nabilah@umpsa.edu.my

² Faculty of Information and Communication Technology, Universiti Teknikal Malaysia Melaka, Durian Tunggal 76100, Melaka, Malaysia. Email: limweihaan@utem.edu.my

³ College of Computing, Informatics and Mathematics, Universiti Teknologi MARA, Shah Alam 40450, Selangor, Malaysia. Email: sitiaina.mnoor@uitm.edu.my

* **Corresponding Author. Email:** sitiaina.mnoor@uitm.edu.my

ARTICLE INFO	ABSTRACT
<p>Received 10 August 2023</p> <p>Revised 18 October 2023</p> <p>Accepted 20 November 2023</p> <p>Available Online 30 December 2023</p> <p>DOI 10.63646/jaihbe.2023.010401</p> <p>License CC BY 4.0</p> <p>Publisher INATGI, United States of America</p> <p>Journal JAIHBE – ISSN 3068-1197</p>	<p><i>Genomic data have become central to precision medicine, risk stratification, pharmacogenomics, and population health research, yet their governance remains difficult because sequence data are sensitive, persistent, identifiable, and valuable across multiple future clinical contexts. Conventional electronic health record architectures protect data through institutional access control, but they often provide limited patient agency, weak cross-institutional provenance, and fragmented consent management. This paper develops a patient-centric healthcare engineering framework that integrates artificial intelligence and permissioned blockchain for privacy-preserving genomic data governance. The framework separates encrypted off-chain genomic storage from on-chain metadata, consent events, audit records, and smart-contract access rules. AI modules provide privacy-risk scoring, anomalous access detection, policy recommendation, and data-quality assessment, while blockchain components provide tamper-resistant logs, decentralized identity, consent execution, and verifiable provenance. Drawing on literature in healthcare blockchain, genomic privacy, trustworthy AI, federated learning, and interoperability standards, the paper proposes a layered architecture and conducts a scenario-based engineering evaluation comparing centralized EHR governance, blockchain-only governance, and an AI-blockchain hybrid. Results suggest that the hybrid model improves normalized governance scores for integrity, consent automation, privacy protection, audit completeness, and interoperability, while reducing avoidable access latency through AI triage and off-chain storage. The study contributes a deployable design logic for genomic data stewardship that balances patient autonomy, research utility, regulatory compliance, and engineering scalability.</i></p> <p>Keywords: Artificial Intelligence; Blockchain; Genomic Data Governance; Patient-Centric; Healthcare; Privacy Preservation; Smart Contracts</p>

1. INTRODUCTION

Genomic data have moved from specialized research repositories into routine clinical, public health, and biomedical engineering workflows. Whole-genome sequencing, polygenic risk scores, pharmacogenomic testing, rare-disease diagnosis, cancer profiling, and large biobank initiatives all depend on the ability to collect, store, link, interpret, and share genomic data across institutional boundaries. Precision medicine therefore requires more than computational biology; it requires an engineering infrastructure that treats genomic information as a long-lived, sensitive, and multi-use asset. The clinical value of such data is widely recognized in precision-medicine

scholarship, where genotype-phenotype linkage, population-scale cohorts, and variant interpretation databases are foundational for diagnosis and risk prediction (Collins and Varmus, 2015; Aronson and Rehm, 2015; Ashley, 2016; Kohane, 2015; Khera et al., 2018; Landrum et al., 2018; Karczewski et al., 2020).

The governance problem is difficult because genomic data are not ordinary medical records. They are persistent across a person's lifetime, partially shared with biological relatives, useful for unanticipated future analyses, and potentially re-identifiable even after removal of conventional identifiers. Prior work on genomic privacy has shown that genetic markers can create distinctive disclosure risks, including surname inference, mixture membership inference, familial leakage, and phenotype reconstruction (Homer et al., 2008; Gymrek et al., 2013; Erlich and Narayanan, 2014; Naveed et al., 2015; Humbert et al., 2015). For patient-centric healthcare engineering, this means that merely encrypting a database is insufficient. Governance must also address consent, provenance, auditability, privacy risk, interoperability, model safety, and accountable reuse.

Blockchain has been proposed as a candidate infrastructure for healthcare data governance because it can maintain tamper-resistant logs, encode access rules through smart contracts, and distribute trust across institutions rather than relying on a single data custodian. Early healthcare blockchain studies demonstrated how distributed ledgers can record data access permissions, support patient-driven interoperability, and provide shared audit trails for biomedical collaboration (Azaria et al., 2016; Benchoufi and Ravaud, 2017; Kuo et al., 2017; Gordon and Catalini, 2018; Esposito et al., 2018; Siyal et al., 2019; Hasselgren et al., 2020; Vazirani et al., 2020). However, blockchain alone does not solve the governance problem. Public ledgers may create privacy risks; on-chain storage is unsuitable for large genomic files; smart contracts may be inflexible; and governance events require contextual interpretation rather than mechanical approval.

Artificial intelligence offers complementary capabilities. AI can evaluate privacy risk, identify anomalous access patterns, recommend minimum-necessary data scopes, evaluate consent conflicts, detect data-quality defects, and support clinical decision workflows. Yet AI also introduces new risks: data leakage, membership inference, model inversion, bias, overreliance, and opaque decision logic. Responsible AI in healthcare therefore requires model validation, transparency, reporting standards, audit trails, and secure distributed learning mechanisms (Topol, 2019; Rajkomar et al., 2019; Miotto et al., 2018; Esteva et al., 2019; Shickel et al., 2018; Char et al., 2018; Vayena et al., 2018; Wiens et al., 2019; Liu et al., 2020; Vasey et al., 2022).

This article develops a privacy-preserving healthcare engineering framework that integrates AI and blockchain for patient-centric genomic data governance. Unlike studies that view blockchain as an isolated storage or audit mechanism, the proposed framework separates responsibilities across layers: encrypted off-chain genomic storage, on-chain hash anchoring, decentralized identity, smart-contract consent execution, AI-driven privacy-risk analytics, EHR interoperability, and human governance oversight. This layered view is consistent with broader research on blockchain systems, IoT security, management analytics, and industrial information integration, where technical innovation must be embedded within system architecture and decision processes rather than deployed as a single tool (Lu, 2018; Lu, 2019a; Lu, 2019b; Lu, 2021; Lu, 2022; Lu, 2025; Chen et al., 2024; Xu et al., 2021; Zhang and Lu, 2021; Zheng and Lu, 2022).

The paper addresses three research questions. First, how can AI and blockchain be integrated

into a patient-centric governance architecture for genomic data? Second, which engineering components are necessary to balance privacy protection, data utility, consent enforcement, and auditability? Third, how does the proposed AI-blockchain hybrid compare with centralized EHR governance and blockchain-only governance in a scenario-based evaluation? The remainder of the paper is organized as follows. The next section reviews the related literature. The third section describes the framework and design principles. The fourth section explains the scenario-based evaluation. The fifth section presents results. The sixth section discusses engineering implications, limitations, and future research directions.

2. BACKGROUND AND RELATED WORK

The literature relevant to this study can be organized into five streams: healthcare blockchain, genomic privacy, AI-enabled healthcare analytics, privacy-preserving machine learning, and interoperability standards. Healthcare blockchain research provides the foundation for tamper-resistant data access, patient-mediated permission management, and multi-institutional auditability. Blockchain architectures and permissioned ledgers have been examined from data-processing, security, and smart-contract perspectives, showing that consensus design, transaction throughput, and governance roles strongly influence deployment feasibility (Dinh et al., 2018; Androulaki et al., 2018; Gervais et al., 2016; Christidis and Devetsikiotis, 2016; Dorri et al., 2017; Novo, 2018; Ouaddah et al., 2017).

In healthcare contexts, blockchain has been evaluated as a mechanism for patient-mediated data exchange, medical record auditability, research integrity, and clinical data sharing. MedRec illustrated how blockchain can coordinate access and permission management across medical data custodians, while subsequent health informatics research emphasized patient-driven interoperability and organizational trust (Azaria et al., 2016; Kuo et al., 2017; Gordon and Catalini, 2018). Broader reviews and conceptual studies have identified benefits such as provenance, auditability, transparency, and multi-party coordination, but also warn about privacy, scalability, governance, and workflow integration (Benchoufi and Ravaud, 2017; Esposito et al., 2018; Siyal et al., 2019; Vazirani et al., 2020; Hasselgren et al., 2020).

Genomic privacy research shows why a specialized governance model is necessary. Membership inference from aggregate genomic data, surname linkage, and statistical re-identification demonstrate that genomic data cannot be treated as fully de-identified merely by removing names or hospital identifiers (Homer et al., 2008; Gymrek et al., 2013; Erlich and Narayanan, 2014). Surveys of genomic privacy further show that risk varies with population structure, dataset size, linked phenotypes, and adversarial background knowledge (Naveed et al., 2015). Privacy-preserving genomic computation has therefore explored cryptographic techniques, multiparty computation, homomorphic encryption, and differential privacy as ways to maintain analytical utility while limiting disclosure (Humbert et al., 2015; Raisaro et al., 2018; Cho et al., 2018; Dwork, 2006; Abadi et al., 2016).

A second foundation is zero-knowledge and privacy-preserving smart-contract research. In patient-centric genomic governance, it is sometimes necessary to prove that a data requester satisfies an eligibility condition or that a genomic attribute meets a study criterion without revealing the underlying sensitive data. Zero-knowledge proofs and privacy-preserving smart contracts offer a design vocabulary for this requirement (Goldwasser et al., 1989; Miers et al., 2013; Ben-Sasson et al., 2014; Kosba et al., 2016; Bünz et al., 2018). These techniques are not a complete healthcare system by themselves, but they are useful components for selective verification, confidential policy execution, and auditable proof of compliance.

AI-enabled healthcare analytics provides the decision-intelligence layer of the proposed framework. Deep learning and machine learning have been widely studied for electronic health record analysis, medical prediction, medical imaging, and clinical decision support (Miotto et al., 2018; Shickel et al., 2018; Rajkomar et al., 2019; Esteva et al., 2019; Topol, 2019). Yet responsible deployment requires careful attention to clinical impact, explainability, bias, evaluation, and reporting (Char et al., 2018; Vayena et al., 2018; Kelly et al., 2019; Wiens et al., 2019; Liu et al., 2020; Vasey et al., 2022). In this paper, AI is not used to make autonomous clinical decisions; instead, it supports governance functions such as privacy-risk scoring, anomalous access detection, data-quality evaluation, and policy recommendation.

Privacy-preserving machine learning is especially relevant because genomic datasets are often distributed across hospitals, biobanks, and research networks. Federated learning, secure aggregation, and decentralized learning make it possible to train models without directly pooling sensitive data (Bonawitz et al., 2017; Kairouz et al., 2021; Yang et al., 2019; Rieke et al., 2020; Sheller et al., 2020; Kaissis et al., 2020; Warnat-Herresthal et al., 2021). However, distributed learning is not automatically private. Model inversion, membership inference, and information leakage from collaborative models remain serious threats (Shokri and Shmatikov, 2015; Fredrikson et al., 2015; Shokri et al., 2017; Hitaj et al., 2017; Beaulieu-Jones et al., 2019). The proposed framework therefore treats AI models as governed assets whose training, access, and outputs must be logged and risk-scored.

Interoperability standards complete the background. Genomic governance requires linkage with clinical observations, medication data, phenotypes, and research metadata. EHR phenotyping, data-quality frameworks, terminology systems, and FHIR-based application platforms show how semantic and technical interoperability can support secondary use (Bodenreider, 2004; Hripcsak and Albers, 2013; Kahn et al., 2016; Mandel et al., 2016; Mandl and Kohane, 2016; Wilkinson et al., 2016). Large-scale genomic and biobank resources demonstrate the research value of structured cohorts, while precision medicine literature clarifies why clinical interpretation and long-term stewardship are essential (Sudlow et al., 2015; Bycroft et al., 2018; Collins and Varmus, 2015; Ashley, 2016; Torkamani et al., 2018; Kalia et al., 2017; Landrum et al., 2018; Karczewski et al., 2020).

3. RESEARCH DESIGN AND FRAMEWORK

The study uses design science logic rather than a clinical trial. The objective is to build and evaluate an engineering framework that transforms fragmented genomic data stewardship into a patient-centric governance process. The design problem is derived from the uploaded manuscript's central insight: blockchain can strengthen data integrity, traceability, and access control for bioinformatics data, but practical implementation requires off-chain storage, smart contracts, privacy-preserving verification, and interoperability with healthcare infrastructure. The present paper extends this idea by adding AI modules that make governance adaptive, risk-aware, and clinically contextual.

Figure 1 presents the proposed patient-centric AI-blockchain genomic data governance framework. The framework starts with the patient or data owner, who sets consent preferences, data-sharing boundaries, revocation rules, and notification requirements. The clinical genomics laboratory contributes sequencing outputs, annotations, quality-control results, and phenotype linkages. The AI governance engine evaluates privacy risk, data sensitivity, access anomaly likelihood, consent-policy conflicts, and data-quality status. Smart contracts translate consent and institutional policies into executable access rules. A permissioned blockchain records metadata

hashes, identity assertions, access events, and audit triggers. Encrypted genomic files remain in off-chain vaults or federated repositories, while clinicians and researchers receive only verified and minimum-necessary access.

Patient-Centric AI-Blockchain Genomic Data Governance Framework

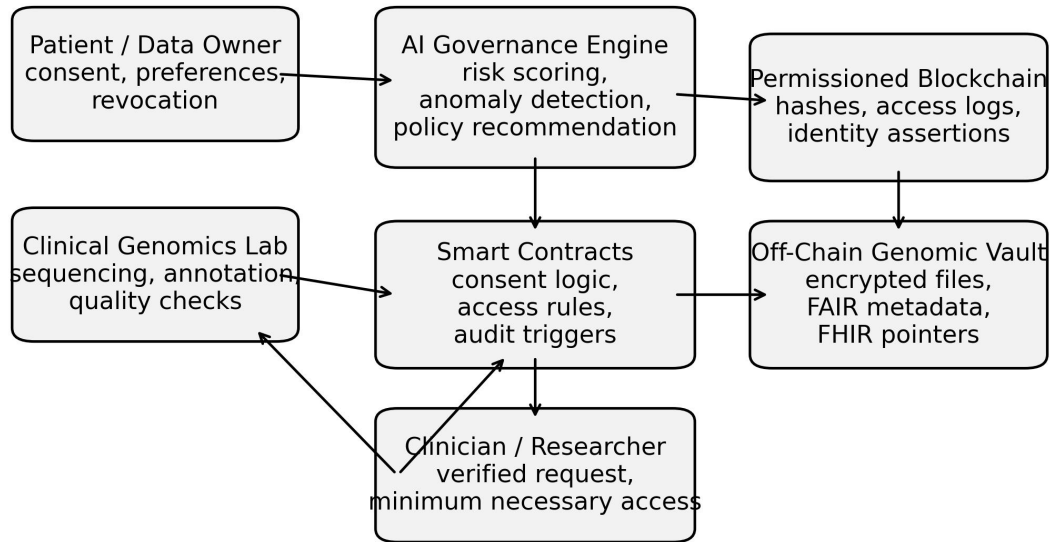


Figure 1. Patient-centric AI-blockchain genomic data governance framework.

The framework follows five design principles. First, patient agency must be computationally represented. Consent is not a static form stored as a PDF; it becomes a machine-readable policy that can be executed, revoked, restricted by purpose, and audited. Second, genomic data should be usable without being copied unnecessarily. Third, data access must be explainable to patients and accountable to regulators. Fourth, AI governance outputs should assist rather than replace human ethics committees, data access committees, and clinicians. Fifth, the architecture must support interoperability with EHR systems, research repositories, and biobank workflows through structured identifiers and standard metadata.

Table 1 summarizes the design requirements and corresponding engineering responses. The table is intentionally framed in operational language because genomic governance often fails not at the conceptual level but at the interface between clinical workflow, technical infrastructure, and institutional accountability. The framework is designed for hospitals and research networks that want to support genomic data sharing while avoiding uncontrolled data duplication, opaque research access, and weak patient feedback mechanisms.

Table 1. Design requirements and engineering responses for patient-centric genomic governance.

Requirement	Governance Need	Engineering Response
Patient agency	Dynamic consent, revocation, notification, purpose limitation	Consent tokens; patient dashboard; smart-contract policy state
Data integrity	Proof that genomic files and derived records have not been silently changed	On-chain hashes; versioned metadata; signed laboratory events
Privacy protection	Limit re-identification, leakage, and excessive disclosure	Encrypted storage; AI risk scoring; zero-knowledge verification
Clinical interoperability	Link variants to EHR phenotypes and care pathways	FHIR pointers; terminology mapping; quality checks

Research utility	Enable legitimate secondary use without uncontrolled copying	Federated queries; access scopes; auditable data-use agreements
Regulatory accountability	Produce evidence for inspection, dispute resolution, and compliance	Immutable audit logs; policy versioning; exception workflow

The evaluation uses a scenario-based engineering analysis with three governance models. The first model is centralized EHR governance, where genomic data are managed through institutional authentication, role-based access control, and local audit logs. The second model is blockchain-only governance, where access records and consent events are logged on a permissioned ledger, but governance decisions are rule-based. The third model is the proposed AI-blockchain hybrid, where smart contracts execute consent and AI modules score privacy risk, data quality, and access anomalies before data retrieval. The goal is not to claim clinical superiority but to examine architectural performance across governance dimensions.


Five normalized metrics are used: integrity verification, consent automation, privacy protection, audit completeness, and interoperability. Integrity verification measures whether data access can be traced to immutable evidence. Consent automation measures whether patient preferences can be applied without manual reinterpretation. Privacy protection measures the degree to which the architecture limits disclosure and detects risky access. Audit completeness measures whether data events produce regulator-ready evidence. Interoperability measures whether the model can exchange metadata with EHR and research systems. Table 2 compares the three governance models conceptually before numerical scenario scoring.

Table 2. Conceptual comparison of genomic data governance models.

Model	Main Logic	Strength	Limitation
Centralized EHR governance	Institutional databases, role-based access, local logs	Low complexity; familiar workflow	Weak cross-institution provenance; limited patient visibility
Blockchain-only governance	Permissioned ledger, smart contracts, on-chain access logs	Strong auditability and consent execution	Limited contextual reasoning; scaling and privacy risks
AI-blockchain hybrid	AI risk scoring plus smart contracts and off-chain vaults	Adaptive privacy protection; improved triage; stronger audit logic	Requires model validation and multi-stakeholder governance

Figure 2 shows the closed-loop consent and access workflow. The sequence begins when a patient registers consent preferences or updates data-sharing restrictions. A decentralized identity module verifies the requester and links the request to a permitted purpose. Smart contracts check whether the request is consistent with active consent, institutional policy, and data-use conditions. The AI engine then assigns a privacy-risk score based on requester history, data sensitivity, cohort size, and purpose. If the request passes, encrypted off-chain data are retrieved or analyzed in place. Finally, an audit event is written back to the ledger, and the patient may receive a notification or summary. This cycle turns consent from a one-time document into an active governance process.

Closed-loop consent and access workflow



The loop converts patient preference into computable access policy, then returns auditable feedback after each data use.

Figure 2. Closed-loop patient consent and genomic data access workflow.

The numerical scenario results in Figure 3 suggest that the AI-blockchain hybrid has the strongest overall governance profile. Centralized EHR governance performs reasonably in basic authentication, but it is weak in patient-facing auditability and cross-institutional provenance. Blockchain-only governance improves integrity and audit completeness, yet it remains limited in privacy-risk interpretation and interoperability because smart contracts cannot understand all clinical and research context. The hybrid model performs better because AI triage reduces inappropriate access, flags anomalous requests, and recommends the minimum data scope needed for a valid purpose.

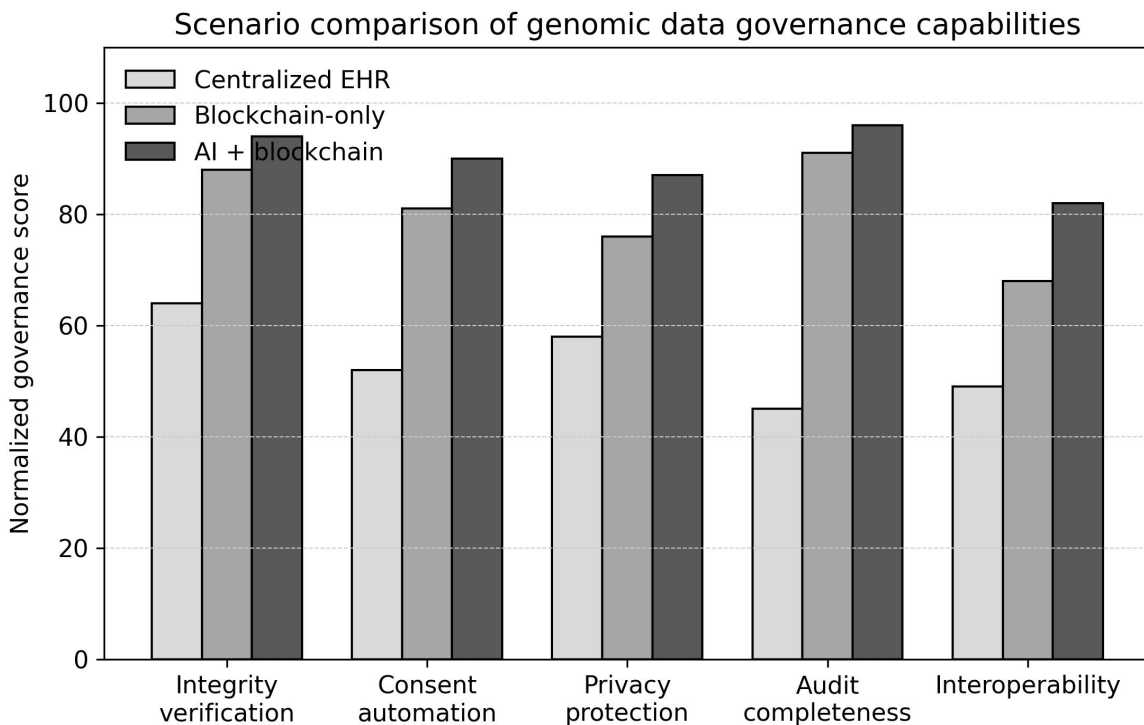


Figure 3. Scenario comparison of centralized EHR, blockchain-only, and AI-blockchain governance.

The trade-off analysis in Figure 4 highlights why off-chain storage is essential. Direct on-chain genomic storage has high latency and poor throughput because blockchains are not designed for large biomedical files. Off-chain vaults with on-chain hashes improve latency and scalability, while AI triage further reduces unnecessary retrieval by rejecting or narrowing high-risk requests before data movement. Federated data spaces offer strong collaboration potential, although they require robust identity, metadata harmonization, and operational governance. The result supports the framework's design decision: blockchain should govern evidence and authorization, not act as a bulk genomic file system.

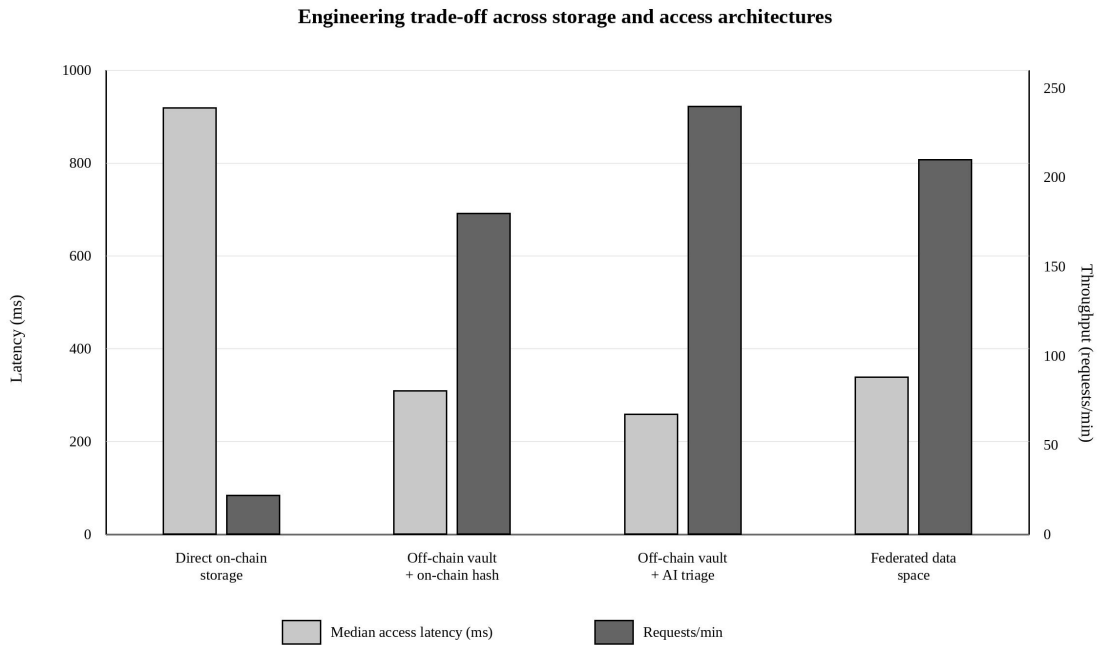


Figure 4. Latency and throughput trade-off across genomic data storage and access architectures.

Table 3 reports the scenario scores. The hybrid model achieves the highest normalized score in all five dimensions. The largest improvements over centralized EHR governance occur in audit completeness and consent automation, while improvements over blockchain-only governance are most visible in privacy protection and interoperability. These results are consistent with the conceptual complementarity between AI and blockchain: blockchain provides evidence, execution, and shared state; AI provides contextual analysis, pattern recognition, and decision support.

Table 3. Scenario-based evaluation of genomic data governance models.

Governance Model	Integrity	Consent	Privacy	Audit	Interop.	Average
Centralized EHR	64	52	58	45	49	53.6
Blockchain-only	88	81	76	91	68	80.8
AI + blockchain	94	90	87	96	82	89.8

4. DISCUSSION

The results have three implications for patient-centric healthcare engineering. First, patient

control should be engineered as a dynamic workflow rather than a symbolic governance statement. Many consent systems give patients limited practical control because institutional forms are difficult to interpret computationally. By contrast, smart contracts can encode purpose limitation, expiration, revocation, and notification rules. AI can assist by detecting when a request is technically permitted but contextually risky. This is important for genomic data because future reuse is difficult to predict at the time of collection.

Second, auditability should cover both data and models. As AI becomes more involved in genomic interpretation and research selection, governance must log not only who accessed data but also which model versions, feature sets, risk scores, and output summaries were produced. Model access can leak information even when raw genomic files are protected, as shown by privacy attacks against machine learning systems (Shokri and Shmatikov, 2015; Fredrikson et al., 2015; Shokri et al., 2017; Hitaj et al., 2017). A blockchain ledger can preserve accountable traces, but ledger entries must be designed carefully to avoid exposing sensitive metadata.

Third, interoperability is not merely a technical standard; it is a governance precondition. If genomic files, clinical phenotypes, consent terms, and audit events are stored in incompatible formats, then patient-centric governance becomes fragmented. FHIR mapping, biomedical terminology alignment, and data-quality assessment must therefore be part of the governance architecture, not post-hoc documentation. The proposed framework treats metadata quality as a safety mechanism because incorrect linkage between genome, phenotype, and consent can lead to inappropriate access decisions (Bodenreider, 2004; Hripcsak and Albers, 2013; Kahn et al., 2016; Mandel et al., 2016; Wilkinson et al., 2016).

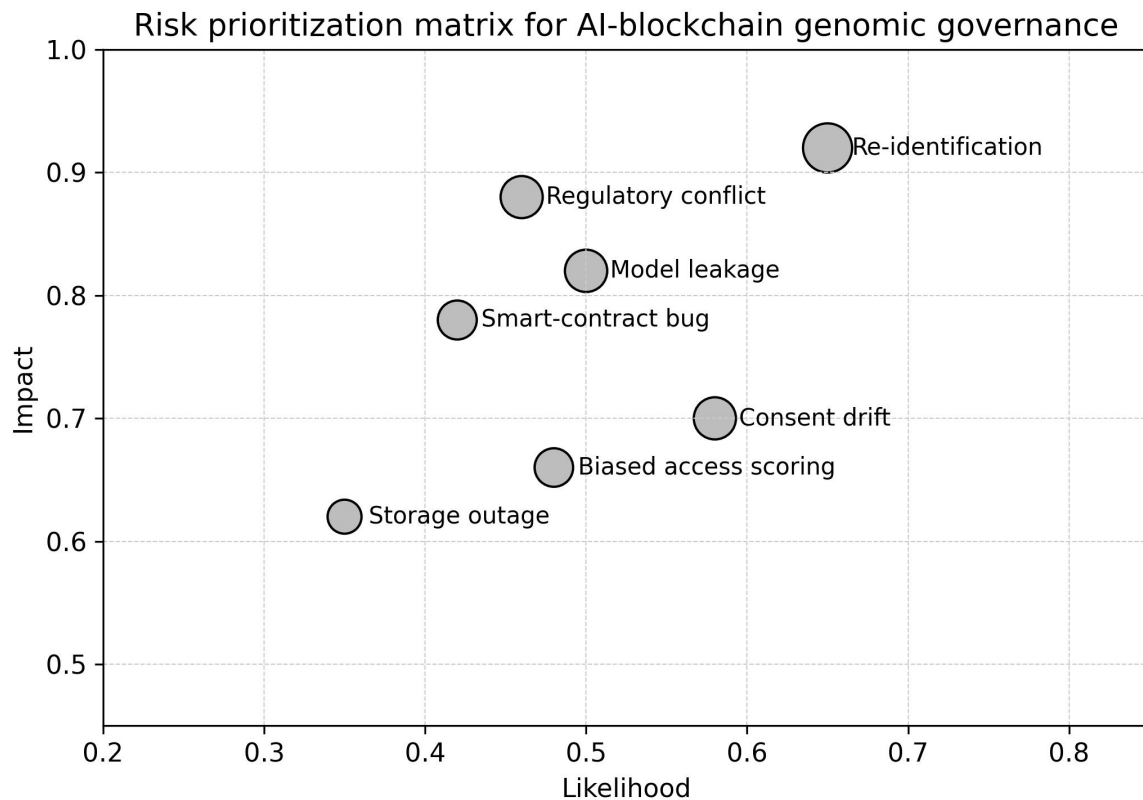


Figure 5. Risk prioritization matrix for AI-blockchain genomic data governance.

The risk matrix in Figure 5 identifies re-identification, regulatory conflict, model leakage, and

consent drift as high-priority risks. Re-identification has high impact because genomic data are inherently distinctive and family-linked. Regulatory conflict is also severe because immutable audit logs may appear to conflict with rights to correction or erasure. Model leakage becomes important when downstream AI tools learn from sensitive cohorts. Consent drift occurs when a patient's preferences change but secondary data users continue to rely on older permissions. The framework addresses these risks through revocable consent tokens, off-chain storage controls, minimum-necessary AI triage, and cryptographic proof mechanisms.


5. HEALTHCARE ENGINEERING IMPLICATIONS

Table 4 translates the framework into healthcare engineering controls. These controls are intended for hospitals, biobank networks, genomic laboratories, and research data platforms. They include governance controls such as data access committees, technical controls such as encryption and secure aggregation, and operational controls such as incident response and patient notification. The key principle is layered accountability: no single mechanism, whether blockchain or AI, should be treated as sufficient. Patient-centric governance emerges only when consent, identity, storage, analytics, and audit functions reinforce one another.

Table 4. Healthcare engineering controls for deployment.

Control Area	Objective	Implementation Mechanism
Consent governance	Purpose limitation, revocation, expiry, and patient notification	Smart-contract templates and patient dashboard
Identity and access	Verify patient, clinician, researcher, and institution roles	Decentralized identity, multi-factor authentication, credential registry
Data protection	Protect raw sequence files and variant annotations	Off-chain encrypted vaults, hash anchoring, key rotation
AI risk analytics	Score access risk, detect anomalies, and recommend minimum scope	Validated risk model, explanation interface, bias review
Interoperability	Connect genomic data with EHR and research metadata	FHIR mapping, controlled terminology, data-quality framework
Audit and compliance	Support evidence generation and regulator inspection	Immutable logs, exportable audit report, incident response workflow

Figure 6 provides an implementation roadmap. Phase 1 establishes the governance charter and determines who may operate validating nodes in the permissioned network. Phase 2 designs metadata models and maps genomic records to EHR identifiers and FHIR-compatible structures. Phase 3 implements consent smart contracts and tests revocation, expiration, and purpose limitation. Phase 4 deploys AI privacy analytics for anomalous access detection and risk scoring. Phase 5 validates the system through audit simulations, bias checks, model leakage tests, and regulatory review. This staged approach reduces the risk of building a technically impressive system that cannot be trusted by patients, clinicians, or regulators.



Each phase produces an auditable artifact: governance policy, interoperable metadata, executable consent, privacy-risk model, and deployment validation evidence.

Figure 6. Implementation roadmap for healthcare organizations and genomic research networks.

The framework also has implications for research collaboration. Multi-institutional genomic studies often require data linkage across hospitals and biobanks, but direct pooling may be legally or ethically constrained. Federated learning and secure aggregation allow models to learn across institutions without centralizing all raw data, while blockchain can record model-training permissions, node participation, and governance events (Bonawitz et al., 2017; Kairouz et al., 2021; Yang et al., 2019; Rieke et al., 2020; Sheller et al., 2020; Kaissis et al., 2020; Warnat-Herresthal et al., 2021). This combination is particularly suitable for rare-disease genomics, where sample sizes are small and collaboration is necessary.

For patient engagement, the system should expose meaningful explanations rather than raw technical logs. A patient dashboard should summarize who requested data, for what purpose, which policy allowed access, whether AI flagged risk, and what data category was released. Such explanations do not require revealing cryptographic details; they require translating governance evidence into understandable patient-facing language. This is consistent with responsible AI guidance that emphasizes clinical impact, explainability, and human oversight rather than opaque automation (Kelly et al., 2019; Wiens et al., 2019; Liu et al., 2020; Vasey et al., 2022).

Table 5. Evaluation metrics for future prototype validation.

Metric	Definition	Interpretive Value
Access latency	Median time between verified request and data availability	Lower is better; direct on-chain storage should be avoided
Revocation propagation	Time required for consent changes to take effect across nodes	Critical for patient agency and legal responsiveness
Audit completeness	Share of data actions with valid identity, policy, time, and hash evidence	Core indicator for trust and inspection readiness
Privacy-risk alert rate	Share of requests flagged for additional review	Must be balanced against false-positive burden
Interoperability coverage	Share of genomic records mapped to clinical metadata standards	Determines downstream clinical and research utility

6. LIMITATIONS AND FUTURE RESEARCH

This study has limitations. First, the evaluation is scenario-based and does not use patient-level operational data. The numerical scores are intended to compare engineering logic, not to estimate clinical outcomes. Second, the framework assumes a permissioned consortium, which may be easier to govern than an open public blockchain but still requires institutional coordination. Third, AI modules may introduce bias if privacy-risk scoring is trained on narrow or historically distorted access logs. Fourth, smart contracts can encode rules precisely, but ethical judgment is often contextual. Human review remains necessary for high-risk access decisions, secondary uses, and requests involving vulnerable populations.

Future research should develop a prototype with synthetic genomic records, FHIR-compatible metadata, and simulated patient consent preferences. Such a prototype should measure transaction latency, revocation propagation time, false-positive anomaly rates, and usability of patient dashboards. Additional work should examine how zero-knowledge proofs can verify eligibility criteria without revealing sensitive genetic attributes, how regulatory rights can be reconciled with immutable audit records, and how privacy attacks against governance AI can be detected. Comparative field studies across hospitals and biobanks would be especially valuable.

7. CONCLUSION

This paper proposed an AI-blockchain framework for patient-centric genomic data governance. The central argument is that genomic data stewardship requires more than secure storage: it needs executable consent, auditable provenance, privacy-risk analytics, interoperable metadata, and meaningful patient agency. Blockchain provides tamper-resistant evidence, shared governance state, and smart-contract execution. AI contributes contextual risk scoring, anomaly detection, data-quality assessment, and policy recommendation. Off-chain storage prevents blockchain bottlenecks and reduces exposure of sensitive genomic content. Scenario-based evaluation suggests that the hybrid model performs better than centralized EHR governance and blockchain-only governance across integrity, consent, privacy, auditability, and interoperability dimensions.

The contribution of the study is a healthcare engineering design that places the patient at the center of genomic data use while preserving scientific utility. By linking AI, blockchain, federated learning, privacy-preserving computation, and EHR interoperability, the framework offers a practical direction for hospitals and research networks seeking to build trustworthy genomic data ecosystems. Its value lies not in treating blockchain as a universal solution or AI as a replacement for governance, but in showing how both technologies can be integrated into a layered, auditable, and privacy-preserving architecture.

Acknowledgement

The authors acknowledge the broader research community working on genomic data protection, healthcare AI, blockchain governance, and patient-centric digital health systems. The manuscript was prepared as a healthcare engineering framework and does not report experiments involving human participants or identifiable patient records.

Funding

The authors received no financial support for the research, authorship, or publication of this article.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability

No patient-level data were used in this study. The scenario values and diagrams are constructed for healthcare engineering analysis and may be reproduced from the information provided in the manuscript.

Author Contributions

Conceptualization, F.N.A. and S.A.M.N.; methodology, F.N.A. and L.W.H.; framework design, S.A.M.N.; data-analysis scenario design, L.W.H.; writing-original draft, F.N.A.; writing-review and editing, all authors; supervision, S.A.M.N.

Use of AI Tools

The manuscript does not use AI-generated patient data, clinical records, or genomic records. Language editing and diagram drafting assistance may be documented by the editorial office if required by journal policy.

Reference

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308-318). <https://doi.org/10.1145/2976749.2978318>
- Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., Manevich, Y., Muralidharan, S., Murthy, C., Nguyen, B., Sethi, M., Singh, G., Smith, K., Sorniotti, A., Stathakopoulou, C., Vukolic, M., ... Yellick, J. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference*. <https://doi.org/10.1145/3190508.3190538>
- Aronson, S. J., & Rehm, H. L. (2015). Building the foundation for genomics in precision medicine. *Nature*, *526*(7573), 336-342. <https://doi.org/10.1038/nature15816>
- Ashley, E. A. (2016). Towards precision medicine. *Nature Reviews Genetics*, *17*(9), 507-522. <https://doi.org/10.1038/nrg.2016.86>
- Azaria, A., Ekblaw, A., Vieira, T., & Lippman, A. (2016). MedRec: Using blockchain for medical data access and permission management. In *2016 2nd International Conference on Open and Big Data (OBD)* (pp. 25-30). <https://doi.org/10.1109/OBD.2016.11>
- Beaulieu-Jones, B. K., Wu, Z. S., Williams, C., Lee, R., Bhavnani, S. P., Byrd, J. B., & Greene, C. S. (2019). Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, *12*(7), e005122. <https://doi.org/10.1161/CIRCOUTCOMES.118.005122>
- Benchoufi, M., & Ravaud, P. (2017). Blockchain technology for improving clinical research quality. *Trials*, *18*(1), 335. <https://doi.org/10.1186/s13063-017-2035-z>
- Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. In *2014 IEEE Symposium on Security and Privacy* (pp. 459-474). <https://doi.org/10.1109/SP.2014.36>
- Bodenreider, O. (2004). The Unified Medical Language System (UMLS): Integrating biomedical terminology. *Nucleic Acids Research*, *32*(Database issue), D267-D270. <https://doi.org/10.1093/nar/gkh061>
- 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. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1175-1191). <https://doi.org/10.1145/3133956.3133982>
- Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., & Maxwell, G. (2018). Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy* (pp. 315-334). <https://doi.org/10.1109/SP.2018.00020>
- Bycroft, C., Freeman, C., Petkova, D., Band, G., Elliott, L. T., Sharp, K., Motyer, A., Vukcevic, D., Delaneau, O., O'Connell, J., Cortes, A., Welsh, S., Young, A., Effingham, M., McVean, G., Leslie, S., Allen, N., Donnelly, P., &

- Marchini, J. (2018). The UK Biobank resource with deep phenotyping and genomic data. *Nature*, 562(7726), 203-209. <https://doi.org/10.1038/s41586-018-0579-z>
- Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing machine learning in health care-addressing ethical challenges. *New England Journal of Medicine*, 378(11), 981-983. <https://doi.org/10.1056/NEJMp1714229>
- Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. *Information Systems Frontiers*, 26(5), 1715-1729. <https://doi.org/10.1007/s10796-022-10248-7>
- Cho, H., Wu, D. J., & Berger, B. (2018). Secure genome-wide association analysis using multiparty computation. *Nature Biotechnology*, 36(6), 547-551. <https://doi.org/10.1038/nbt.4108>
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. *IEEE Access*, 4, 2292-2303. <https://doi.org/10.1109/ACCESS.2016.2566339>
- Collins, F. S., & Varmus, H. (2015). A new initiative on precision medicine. *New England Journal of Medicine*, 372(9), 793-795. <https://doi.org/10.1056/NEJMp1500523>
- Dinh, T. T. A., Wang, J., Chen, G., Liu, R., Ooi, B. C., & Tan, K.-L. (2018). Untangling blockchain: A data processing view of blockchain systems. *IEEE Transactions on Knowledge and Data Engineering*, 30(7), 1366-1385. <https://doi.org/10.1109/TKDE.2017.2781227>
- Dorri, A., Kanhere, S. S., Jurdak, R., & Gauravaram, P. (2017). Blockchain for IoT security and privacy: The case study of a smart home. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops* (pp. 618-623). <https://doi.org/10.1109/PERCOMW.2017.7917634>
- Dwork, C. (2006). Differential privacy. In *Automata, Languages and Programming* (pp. 1-12). Springer. https://doi.org/10.1007/11787006_1
- Erlach, Y., & Narayanan, A. (2014). Routes for breaching and protecting genetic privacy. *Nature Reviews Genetics*, 15(6), 409-421. <https://doi.org/10.1038/nrg3723>
- Espósito, C., De Santis, A., Tortora, G., Chang, H., & Choo, K.-K. R. (2018). Blockchain: A panacea for healthcare cloud-based data security and privacy? *IEEE Cloud Computing*, 5(1), 31-37. <https://doi.org/10.1109/MCC.2018.011791712>
- Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24-29. <https://doi.org/10.1038/s41591-018-0316-z>
- Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (pp. 1322-1333). <https://doi.org/10.1145/2810103.2813677>
- Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 3-16). <https://doi.org/10.1145/2976749.2978341>
- Goldwasser, S., Micali, S., & Rackoff, C. (1989). The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1), 186-208. <https://doi.org/10.1137/0218012>
- Gordon, W. J., & Catalini, C. (2018). Blockchain technology for healthcare: Facilitating the transition to patient-driven interoperability. *Computational and Structural Biotechnology Journal*, 16, 224-230. <https://doi.org/10.1016/j.csbj.2018.06.003>
- Gymrek, M., McGuire, A. L., Golan, D., Halperin, E., & Erlich, Y. (2013). Identifying personal genomes by surname inference. *Science*, 339(6117), 321-324. <https://doi.org/10.1126/science.1229566>
- Hasselgren, A., Kravetska, K., Gligoroski, D., Pedersen, S. A., & Faxvaag, A. (2020). Blockchain in healthcare and health sciences-A scoping review. *International Journal of Medical Informatics*, 134, 104040. <https://doi.org/10.1016/j.ijmedinf.2020.104040>
- Hitaj, B., Ateniese, G., & Perez-Cruz, F. (2017). Deep models under the GAN: Information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 603-618). <https://doi.org/10.1145/3133956.3134012>
- Homer, N., Szlinger, S., Redman, M., Duggan, D., Tembe, W., Muehling, J., Pearson, J. V., Stephan, D. A., Nelson, S. F., & Craig, D. W. (2008). Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genetics*, 4(8), e1000167. <https://doi.org/10.1371/journal.pgen.1000167>
- Hripesak, G., & Albers, D. J. (2013). Next-generation phenotyping of electronic health records. *Journal of the American Medical Informatics Association*, 20(1), 117-121. <https://doi.org/10.1136/amiajnl-2012-001145>
- Humbert, M., Ayday, E., Hubaux, J.-P., & Telenti, A. (2015). Reconciling utility with privacy in genomics. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society* (pp. 11-20). <https://doi.org/10.1145/2808138.2808145>
- Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L.-W. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L.

- A., & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3, 160035. <https://doi.org/10.1038/sdata.2016.35>
- Kahn, M. G., Callahan, T. J., Barnard, J., Bauck, A. E., Brown, J., Davidson, B. N., Estiri, H., Goerg, C., Holve, E., Johnson, S. G., Liaw, S.-T., Hamilton-Lopez, M., Meeker, D., Ong, T. C., Ryan, P., Shang, N., Weiskopf, N. G., Weng, C., Zozus, M. N., & Schilling, L. (2016). A harmonized data quality assessment terminology and framework for the secondary use of electronic health record data. *eGEMs*, 4(1), 1244. <https://doi.org/10.13063/2327-9214.1244>
- 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>
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1-2), 1-210. <https://doi.org/10.1561/22000000083>
- Kalia, S. S., Adelman, K., Bale, S. J., Chung, W. K., Eng, C., Evans, J. P., Herman, G. E., Hufnagel, S. B., Klein, T. E., Korf, B. R., McKelvey, K. D., Ormond, K. E., Richards, C. S., Vlangos, C. N., Watson, M., Martin, C. L., & Miller, D. T. (2017). Recommendations for reporting of secondary findings in clinical exome and genome sequencing, 2016 update. *Genetics in Medicine*, 19(2), 249-255. <https://doi.org/10.1038/gim.2016.190>
- 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>
- Karczewski, K. J., Francioli, L. C., Tiao, G., Cummings, B. B., Alfoldi, J., Wang, Q., Collins, R. L., Laricchia, K. M., Ganna, A., Birnbaum, D. P., Gauthier, L. D., Brand, H., Solomonson, M., Watts, N. A., Rhodes, D., Singer-Berk, M., England, E. M., Seaby, E. G., Kosmicki, J. A., ... MacArthur, D. G. (2020). The mutational constraint spectrum quantified from variation in 141,456 humans. *Nature*, 581(7809), 434-443. <https://doi.org/10.1038/s41586-020-2308-7>
- Khera, A. V., Chaffin, M., Aragam, K. G., Haas, M. E., Roselli, C., Choi, S. H., Natarajan, P., Lander, E. S., Lubitz, S. A., Ellinor, P. T., & Kathiresan, S. (2018). Genome-wide polygenic scores for common diseases identify individuals with risk equivalent to monogenic mutations. *Nature Genetics*, 50(9), 1219-1224. <https://doi.org/10.1038/s41588-018-0183-z>
- Kohane, I. S. (2015). Ten things we have to do to achieve precision medicine. *Science*, 349(6243), 37-38. <https://doi.org/10.1126/science.aab1328>
- Kosba, A., Miller, A., Shi, E., Wen, Z., & Papamanthou, C. (2016). Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In 2016 IEEE Symposium on Security and Privacy (pp. 839-858). <https://doi.org/10.1109/SP.2016.55>
- Kuo, T.-T., Kim, H.-E., & Ohno-Machado, L. (2017). Blockchain distributed ledger technologies for biomedical and health care applications. *Journal of the American Medical Informatics Association*, 24(6), 1211-1220. <https://doi.org/10.1093/jamia/ocx068>
- Landrum, M. J., Lee, J. M., Benson, M., Brown, G. R., Chao, C., Chitipiralla, S., Gu, B., Hart, J., Hoffman, D., Hoover, J., Jang, W., Katz, K., Ovetsky, M., Riley, G., Sethi, A., Tully, R., Villamarin-Salomon, R., Rubinstein, W., & Maglott, D. R. (2018). ClinVar: Improving access to variant interpretations and supporting evidence. *Nucleic Acids Research*, 46(D1), D1062-D1067. <https://doi.org/10.1093/nar/gkx1153>
- Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. *Journal of Management Analytics*, 5(4), 231-255. <https://doi.org/10.1080/23270012.2018.1516523>
- Lu, Y. (2019a). Artificial intelligence: A survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1-29. <https://doi.org/10.1080/23270012.2019.1570365>
- Lu, Y. (2019b). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80-90. <https://doi.org/10.1016/j.jii.2019.04.002>
- Lu, Y. (2021). Technological innovation and the emergence of a new interdisciplinary field: Management Analytics. *Nanotechnologies in Construction*, 13(3), 181-192. <https://doi.org/10.15828/2075-8545-2021-13-3-181-192>
- Lu, Y. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876-1907. <https://doi.org/10.1080/17517575.2021.2008513>
- Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. *Information Systems Frontiers*, 27(1), 215-234. <https://doi.org/10.1007/s10796-021-10221-w>
- 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>
- Mandl, K. D., & Kohane, I. S. (2016). No small change for the health information economy. *New England Journal of Medicine*, 375(6), 680-681. <https://doi.org/10.1056/NEJMp1601381>

- Miers, I., Garman, C., Green, M., & Rubin, A. D. (2013). Zerocoin: Anonymous distributed E-cash from Bitcoin. In 2013 IEEE Symposium on Security and Privacy (pp. 397-411). <https://doi.org/10.1109/SP.2013.34>
- Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2018). Deep learning for healthcare: Review, opportunities and challenges. *Briefings in Bioinformatics*, 19(6), 1236-1246. <https://doi.org/10.1093/bib/bbx044>
- Naveed, M., Ayday, E., Clayton, E. W., Fellay, J., Gunter, C. A., Hubaux, J.-P., Malin, B. A., & Wang, X. (2015). Privacy in the genomic era. *ACM Computing Surveys*, 48(1), 1-44. <https://doi.org/10.1145/2767007>
- Novo, O. (2018). Blockchain meets IoT: An architecture for scalable access management in IoT. *IEEE Internet of Things Journal*, 5(2), 1184-1195. <https://doi.org/10.1109/JIOT.2018.2812239>
- Ouaddah, A., Abou Elkalam, A., & Ait Ouahman, A. (2017). Access control in the Internet of Things: Big challenges and new opportunities. *Computer Networks*, 112, 237-262. <https://doi.org/10.1016/j.comnet.2016.11.007>
- Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347-1358. <https://doi.org/10.1056/NEJMra1814259>
- Raisaro, J. L., Tramèr, F., Ji, Z., Bu, D., Zhao, Y., Carey, K., Lloyd, D., Sofia, H., Baker, D., Flicek, P., Shringarpure, S., Bustamante, C. D., Wang, S., Jiang, X., & Hubaux, J.-P. (2018). Protecting privacy and security of genomic data in i2b2 with homomorphic encryption and differential privacy. *BMC Medical Genomics*, 11(Suppl 4), 30. <https://doi.org/10.1186/s12920-018-0395-1>
- Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3, 119. <https://doi.org/10.1038/s41746-020-00323-1>
- Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R. R., & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10, 12598. <https://doi.org/10.1038/s41598-020-69250-1>
- Shickel, B., Tighe, P. J., Bihorac, A., & Rashidi, P. (2018). Deep EHR: A survey of recent advances in deep learning techniques for electronic health record analysis. *IEEE Journal of Biomedical and Health Informatics*, 22(5), 1589-1604. <https://doi.org/10.1109/JBHI.2017.2767063>
- Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (pp. 1310-1321). <https://doi.org/10.1145/2810103.2813687>
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy* (pp. 3-18). <https://doi.org/10.1109/SP.2017.41>
- Siyal, A. A., Junejo, A. Z., Zawish, M., Ahmed, K., Khalil, A., & Soursou, G. (2019). Applications of blockchain technology in medicine and healthcare: Challenges and future perspectives. *Cryptography*, 3(1), 3. <https://doi.org/10.3390/cryptography3010003>
- Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., & Collins, R. (2015). UK Biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Medicine*, 12(3), e1001779. <https://doi.org/10.1371/journal.pmed.1001779>
- Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44-56. <https://doi.org/10.1038/s41591-018-0300-7>
- Torkamani, A., Wineinger, N. E., & Topol, E. J. (2018). The personal and clinical utility of polygenic risk scores. *Nature Reviews Genetics*, 19(9), 581-590. <https://doi.org/10.1038/s41576-018-0018-x>
- Vazirani, A. A., O'Donoghue, O., Brindley, D., & Meinert, E. (2020). Implementing blockchains for efficient health care: Systematic review. *Journal of Medical Internet Research*, 22(2), e12439. <https://doi.org/10.2196/12439>
- Vasey, B., Nagendran, M., Campbell, B., Clifton, D. A., Collins, G. S., Denaxas, S., Denniston, A. K., Faes, L., Geerts, B. F., Ibrahim, M., Liu, X., Mateen, B. A., Mathur, P., McCradden, M. D., Morgan, L., Ordish, J., Rogers, C., Saria, S., Ting, D. S. W., ... McCulloch, P. (2022). Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *Nature Medicine*, 28(5), 924-933. <https://doi.org/10.1038/s41591-022-01772-6>
- Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine*, 15(11), e1002689. <https://doi.org/10.1371/journal.pmed.1002689>
- Warnat-Herresthal, S., Schultze, H., Shastry, K. L., Manamohan, S., Mukherjee, S., Garg, V., Sarveswara, R., Händler, K., Pickkers, P., Aziz, N. A., Ktena, S., Tran, F., Bitzer, M., Ossowski, S., Casadei, N., Herr, C., Petersheim, D., Behr, J., Kern, F., ... Schultze, J. L. (2021). Swarm Learning for decentralized and confidential clinical machine learning. *Nature*, 594(7862), 265-270. <https://doi.org/10.1038/s41586-021-03583-3>
- Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., Ossorio, P. N., Thadaneys-Israni, S., & Goldenberg, A. (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>
- Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., ... Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, 3, 160018. <https://doi.org/10.1038/sdata.2016.18>
- Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452-10473. <https://doi.org/10.1109/JIOT.2021.3060508>
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1-19. <https://doi.org/10.1145/3298981>
- Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>
- Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. *Enterprise Information Systems*, 16(12), 1939895. <https://doi.org/10.1080/17517575.2021.1939895>
- Liu, X., Rivera, S. C., Moher, D., Calvert, M. J., & Denniston, A. K. (2020). Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. *BMJ*, 370, m3164. <https://doi.org/10.1136/bmj.m3164>