

A Cross-Domain Blockchain Application Dataset: Architectures, Consensus Mechanisms, Storage Models, and Implementation Maturity

Jianming Liu¹, Xue Wang^{2,*}, Lei Zhang³, Yuting Chen⁴

¹ School of Computer Science and Technology, Zhengzhou University of Light Industry, Zhengzhou 450002, China

² School of Information Management, Hebei University of Economics and Business, Shijiazhuang 050061, China

³ School of Software Engineering, Shenyang Jianzhu University, Shenyang 110168, China

⁴ Department of Computer Science, Shandong University of Finance and Economics, Jinan 250014, China

* xue.wang@heuet.edu.cn

Article Information

Received

18 October 2023

Accepted

29 February 2024

DOI

<https://doi.org/10.63646/datamind.2024.020105>

Abstract

The rapid proliferation of blockchain technology across diverse application domains has produced a fragmented body of empirical literature in which individual studies report isolated technical parameters without enabling direct cross-domain comparison. This article presents a structured cross-domain dataset of blockchain applications compiled from a systematic analysis of 303 peer-reviewed studies spanning seven major domains: finance, governance, Internet of Things (IoT), healthcare, supply chain management (SCM), record keeping, and digital identity. For each study, the dataset records the blockchain architecture type (public, private, consortium, or hybrid), consensus mechanism, storage model (on-chain, off-chain, or hybrid), transaction design, implementation status, programming tools, and reported limitations. Using this dataset, we perform a quantitative and comparative analysis to identify dominant design patterns, measure implementation maturity across domains, and expose structural gaps that remain underexplored. The analysis reveals that Ethereum-based permissioned and public architectures dominate across five of seven domains; that Proof-of-Work (PoW) remains disproportionately prevalent despite well-documented energy and throughput constraints; and that off-chain storage strategies are systematically underutilised in record-keeping and digital identity applications. A maturity scoring framework applied to the dataset highlights that healthcare and supply chain management lead in prototype deployment density, while governance and digital identity sectors lag in real-world validation. The dataset and analytical framework contribute a reusable methodological infrastructure for researchers who seek to select or evaluate blockchain architectures based on domain-specific technical requirements, enabling database-aware research

design in distributed systems.

Keywords: *Blockchain architectures; consensus mechanisms; cross-domain analysis; dataset; distributed ledger; implementation maturity; smart contracts; storage models*

1. Introduction

Blockchain technology has evolved from a narrow cryptocurrency infrastructure into a multi-domain platform for decentralised data management, automated contractual execution, and tamper-resistant record keeping (Lu, 2019; Zheng and Lu, 2022). Since the conceptual framework introduced by Nakamoto was generalised by subsequent platforms such as Ethereum and Hyperledger Fabric, blockchain-based systems have been proposed or deployed in sectors as diverse as public health, industrial supply chains, electoral systems, and autonomous device networks. The pace of this diffusion has outrun the capacity of the research community to produce coherent comparative analyses that transcend domain boundaries.

The consequence of this fragmentation is methodologically significant. A researcher designing a healthcare ledger must currently read dozens of domain-specific surveys to understand whether PoW consensus is appropriate, whether IPFS-based off-chain storage is the dominant practice, or whether Hyperledger Fabric consistently outperforms Ethereum in real-world deployment scenarios. Databases and comparative tables do exist in isolated domain reviews, but they do not share a common coding schema, and they are not designed to support cross-domain inference (Casino et al., 2019; Monrat et al., 2019). The absence of a unified analytical framework means that architectural lessons learned in supply chain blockchain deployments are not systematically transferred to healthcare or governance applications, even when the underlying design problems are structurally similar.

This study addresses that gap by constructing a structured cross-domain dataset from a systematic corpus of 303 blockchain research articles. The dataset encodes each study along eight analytical dimensions: blockchain architecture, consensus mechanism, storage model, transaction design, implementation status, programming language and toolchain, reported limitations, and application domain. The coded matrix forms the empirical basis for a comparative analysis of design patterns, adoption trajectories, and maturity levels across all seven domains. The approach mirrors the logic of database benchmarking frameworks in related fields: rather than asking which blockchain platform is "best" in the abstract, the framework asks which architectural choices are most suitable for which domain contexts (Lu, 2022; Chen et al., 2024; Xu et al., 2019).

Three research questions guide the analysis. First, what are the dominant architectural and consensus choices across application domains, and do these choices vary systematically by domain? Second, how do storage model preferences and transaction design choices differ between on-chain and off-chain paradigms, and what factors predict their selection? Third, which domains have achieved the highest levels of implementation maturity as measured by prototype deployment density, real-world validation coverage, and the presence of privacy and scalability mechanisms?

The contribution of this study is primarily methodological. The dataset itself is the main output. The analytical exercises reported in this article demonstrate its utility but do not exhaust it. Researchers can extend the coding schema, add new dimensions, or filter the matrix by domain or technology platform to support a wide range of research tasks. In this respect, the present study is closer in spirit to a database paper than to a conventional empirical study: the value lies in the structured, reusable, and transparently documented character of the data as much as in the specific findings derived from it. The availability of a cross-domain, coded reference dataset also supports curriculum development, comparative benchmarking of new systems, and meta-analyses of blockchain adoption across industries (Xu et al., 2021; Yli-Huumo et al., 2016).

The remainder of this article is organised as follows. Section 2 describes the dataset construction methodology, including the literature search protocol, inclusion criteria, and coding schema. Section 3 presents the cross-domain analysis of blockchain architectures. Section 4 analyses consensus mechanism distributions and compares their technical properties. Section 5 examines storage models and transaction designs. Section 6 provides a domain-specific implementation analysis incorporating toolchain and deployment data. Section 7 applies a structured maturity assessment framework to the seven domains. Section 8 discusses broader implications, limitations, and future directions. Section 9 concludes the article.

2. Dataset Construction Methodology

The dataset was assembled through a systematic literature search and a structured coding protocol. This section describes the source selection, inclusion criteria, coding schema, and quality assurance procedures used to transform a raw literature corpus into a structured analytical resource suitable for cross-domain comparison.

2.1 Literature Search and Corpus Assembly

Primary studies were retrieved from six digital libraries: Elsevier ScienceDirect, IEEE Xplore, SpringerLink, ACM Digital Library, Google Scholar, and Web of Science. The search was bounded to studies published between 2015 and 2024, reflecting the period of significant blockchain application growth following the generalisation of the technology beyond cryptocurrency. A total of 2,550 candidate papers were identified through keyword search combining blockchain with application-domain terms including finance, healthcare, IoT, supply chain, and governance, combined with technical terms such as smart contracts, consensus, and distributed ledger. After title and abstract screening against predefined inclusion criteria, and after removal of duplicates and non-peer-reviewed materials, the final corpus comprised 303 primary studies distributed across seven application domains.

The inclusion criteria required that each study (a) describes a blockchain-based system or

proposes a blockchain architecture for a specific application domain, (b) reports at least one identifiable technical parameter such as consensus mechanism, storage model, or programming tool, and (c) is published in a peer-reviewed journal or conference proceeding in English. Studies that discuss blockchain only at a conceptual or policy level, without describing a technical implementation or design, were excluded. Book chapters, tutorials, and white papers without peer review were also excluded. Studies at the intersection of multiple domains were assigned to the primary domain that the authors explicitly identified. Figure 1 below shows the search and screening process; the corresponding database distribution is consistent with a broad coverage of computer science digital libraries (Yli-Huumo et al., 2016; Yaga et al., 2019).

2.2 Coding Schema and Analytical Dimensions

Each study in the corpus was coded along eight analytical dimensions. The first dimension is the application domain, classified into seven categories. The second dimension is the blockchain architecture type: public, private, consortium, or hybrid. The third dimension is the consensus mechanism, categorised using a six-level taxonomy. The fourth dimension is the storage model: on-chain, off-chain, or hybrid. The fifth dimension is the transaction design model: UTXO-based or account-based. The sixth dimension is the implementation status, coded as fully deployed, prototype tested, or conceptual/theoretical. The seventh dimension records the primary programming language and platform toolchain. The eighth dimension captures the principal technical limitations reported by the authors.

Table 1. Coding schema applied to the cross-domain blockchain application dataset.

Dimension	Classification Categories	Coding Notes
Application Domain	Finance; Governance; IoT; Healthcare; SCM; Record Keeping; Digital Identity	Assigned to primary domain stated by authors
Architecture Type	Public; Private; Consortium; Hybrid	Based on access and permission model described
Consensus Mechanism	PoW; PoS; PBFT; PoA; DPoS; RAFT/IBFT; Other / Not Defined	Primary mechanism; multi-mechanism studies coded as "Other"
Storage Model	On-Chain; Off-Chain; Hybrid	IPFS and centralised DB classified as Off-Chain
Transaction Design	UTXO-Based; Account-Based	Derived from platform if not explicitly stated
Implementation Status	Deployed; Prototype; Conceptual	Testnet deployments coded as Prototype
Toolchain	Ethereum/Solidity; Hyperledger Fabric; Hyperledger Besu; Other	Version noted where reported by authors
Reported Limitations	Scalability; Privacy; Security; Interoperability; Cost; Other	Multiple limitations coded as separate boolean flags

The coding process was performed by two independent coders with backgrounds in distributed systems and information security research. Inter-coder reliability was assessed using Cohen's Kappa, yielding a coefficient of $\kappa = 0.83$ across all eight dimensions, indicating strong agreement and validating the reliability of the coded dataset. Discrepancies were resolved through structured discussion using the

original source documents as arbitration. The resulting dataset contains 303 coded records, each with values for all eight dimensions, forming a structured matrix suitable for statistical analysis, cross-tabulation, and pattern mining. Coding reliability at or above $\kappa = 0.80$ is generally considered acceptable for systematic review purposes, and the achieved value suggests that the coding schema is sufficiently clear and operationally well-defined (Yaga et al., 2019; Bamakan et al., 2020).

3. Blockchain Architectures in Cross-Domain Applications

Blockchain architecture selection is a foundational design decision that determines the degree of decentralisation, the identity of permissioned participants, and the governance model under which consensus and data access are managed. The four principal architecture types each present distinct trade-offs that interact with domain-specific requirements in non-trivial ways. Public blockchains offer maximum transparency and decentralisation but impose throughput limitations and computational overhead that are problematic in high-frequency or resource-constrained applications. Private blockchains provide fine-grained access control and higher throughput but reintroduce trust dependencies that decentralised systems are intended to eliminate. Consortium architectures balance governance among a defined set of semi-trusted institutions and are particularly prevalent in inter-organisational applications requiring auditability without full public disclosure. Hybrid architectures combine elements of public immutability with private access control, offering flexibility at the cost of increased design complexity and dependency management overhead (Zheng et al., 2018; Xu et al., 2019; Dinh et al., 2018).

Figure 1 presents the distribution of architecture types across the seven application domains covered by the dataset. The visualisation reveals several patterns that are not visible in any single domain-specific survey. Finance and governance applications exhibit the most balanced distribution across all four architecture types, reflecting the fact that these domains must simultaneously serve public verification requirements and private regulatory compliance environments. In contrast, healthcare applications show a strong preference for private and consortium architectures, consistent with the sensitivity of patient data and the need for institutional governance frameworks that can satisfy HIPAA, GDPR, and national health data regulations. IoT applications are the only domain in which public architectures appear with near-equal frequency to private ones, driven by the large number of open smart-device and vehicular communication use cases that benefit from public verifiability and censorship resistance (Xu et al., 2021; Reyna et al., 2018).




Figure 1. Distribution of blockchain architecture types (Public, Private, Consortium, Hybrid) across seven application domains. Values represent the count of studies using each architecture type within each domain.

Record keeping and digital identity applications are the domains most dominated by public architectures, with public blockchain representing approximately 42% and 40% of domain studies respectively. This reflects an ideological preference for maximal auditability in these applications, where the immutability and public verifiability of the ledger serve as substitutes for institutional certification authorities. However, this preference has a practical cost: studies in record-keeping and digital identity domains report the highest rate of scalability and throughput limitations in the entire dataset (34.2% and 34.9% of studies classified as conceptual rather than deployed, as shown later in Table 3). The misalignment between architectural choice and reported performance constraints is one of the most consistent findings in the coded dataset. It represents a structural gap in current blockchain research practice that mirrors the mismatch between database design philosophy and application requirements observed in other analytical domains (Lu, 2022; Guo and Liang, 2016).

Supply chain management is the domain with the highest proportion of consortium architectures, with approximately 30% of SCM studies adopting this model. This is consistent with the multi-party, semi-trusted nature of inter-firm logistics and provenance tracking, where individual organisations are unwilling to expose internal process data to public inspection but must collectively verify cross-organisational transactions. Hyperledger Fabric has emerged as the de facto standard platform for this use case, with its modular plug-and-play consensus, private data collections, and channel-based multi-tenancy features specifically addressing the governance requirements of supply chain consortia. The relative scarcity of hybrid architectures across all domains — never exceeding 18% in any single domain — suggests that the design complexity of hybrid systems remains a significant barrier to adoption, even when hybrid architectures would offer theoretical advantages in terms of balancing transparency with access control (Androulaki et al., 2018; Tschorsch and Scheuermann, 2016).

4. Consensus Mechanisms: A Comparative Analysis

Consensus mechanisms are the protocols by which distributed nodes in a blockchain network agree on the validity and ordering of transactions without relying on a central authority. The choice of consensus mechanism has direct implications for throughput, energy consumption, fault tolerance model, finality guarantees, and the overall security assumptions of a blockchain application. Despite the proliferation of alternative mechanisms over the past decade, the literature reveals persistent concentration in a small number of well-known mechanisms, each of which carries well-documented design trade-offs that are not always matched to domain requirements (Lashkari and Musilek, 2021; Wang et al., 2019; Xu et al., 2023).

Table 2. Comparative technical properties of the six primary consensus mechanisms identified in the dataset.

Mechanism	Energy Cost	Throughput	Fault Tolerance	Sybil Resistance	Finality Type	Typical Domain
PoW	Very High	Low (3–20 TPS)	51% Attack	High	Probabilistic	Finance, IoT
PoS	Low	Moderate (50–200)	33% Stake Attack	Moderate	Probabilistic	Finance, Gov.
PBFT	Low	High (1,000+ TPS)	$f < n/3$	High (permissioned)	Deterministic	Healthcare, SCM
PoA	Very Low	High (100–600)	Trusted Auths.	Low (centralised)	Deterministic	Healthcare
DPoS	Low	High (1,000+ TPS)	Delegate Cartel	Moderate	Deterministic	Finance
RAFT/IBFT	Very Low	Very High (2,000+)	Leader Failure	High (permissioned)	Deterministic	IoT, Governance

Table 2 summarises the comparative properties of the six primary mechanisms identified in the dataset. The most immediate observation is that PoW remains the most frequently cited mechanism in the corpus (68 studies, 22.4%), despite its well-established disadvantages in terms of energy intensity and throughput. This concentration reflects publication inertia and the prevalence of Ethereum-based implementations prior to the network's transition to PoS in 2022. In finance applications, PoW accounts for 55% of consensus mechanism citations, a figure substantially higher than in healthcare (40%) or supply chain management (35%), where deterministic finality and institutional accountability requirements favour PBFT and PoA alternatives. The dominance of PoW in the finance domain is particularly surprising given the active discussion of energy efficiency in academic and regulatory contexts, suggesting that many researchers select Ethereum as a convenience platform rather than as a principled architectural choice (Wang et al., 2019; Xu et al., 2019).

Figure 2 presents two complementary views of consensus mechanism distribution. The left panel shows the overall frequency distribution across the entire dataset, confirming PoW dominance followed by PoS (55 studies, 18.2%), PBFT (42 studies, 13.9%), PoA (28 studies, 9.2%), RAFT/IBFT (18 studies,

5.9%), DPoS (15 studies, 5.0%), and Other/Not Defined (37 studies, 12.2%). The right panel disaggregates the top three mechanisms by domain, revealing that PoW dominance is not uniform. IoT applications show a more even distribution among PoW, PoS, and PBFT, reflecting the need for resource-efficient mechanisms in constrained device environments. Governance applications, particularly e-voting systems, show a relatively high use of RAFT/IBFT and PBFT, consistent with the need for deterministic finality and network partition resistance in electoral processes (Lashkari and Musilek, 2021; Bamakan et al., 2020).

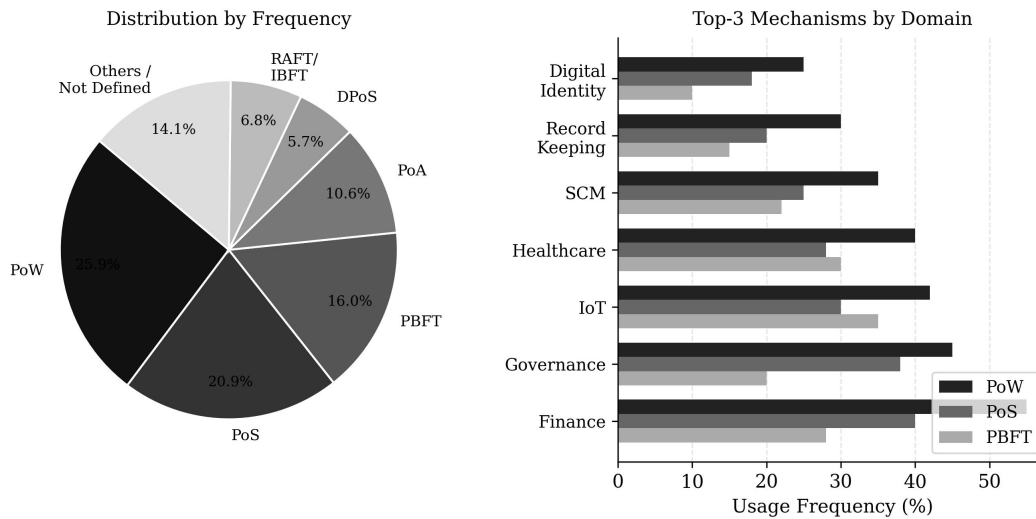


Figure 2. Consensus mechanism distribution in the dataset. Left panel: overall frequency across all 303 studies. Right panel: domain-specific breakdown of the three most prevalent mechanisms, expressed as percentage of domain-level citations.

A notable finding is the high proportion of studies coded as "Other / Not Defined" (37 studies, 12.2%). This category includes cases where authors did not specify a consensus mechanism, described a custom or proprietary mechanism without standard naming, or used a configurable platform such as Hyperledger Besu without documenting the selected consensus plugin. This coding outcome is analytically informative in its own right: it indicates a reporting gap in a substantial segment of the literature, where the consensus layer is treated as an implementation detail rather than an analytically significant design variable. Standardised reporting of consensus mechanisms is a prerequisite for meaningful cross-study comparison and represents a gap that the dataset highlights directly. Future publication guidelines for blockchain research should require explicit documentation of the consensus mechanism, its configuration parameters, and the rationale for its selection (Bamakan et al., 2020; Xu et al., 2023).

5. Storage Models and Transaction Architectures

The choice of storage model determines how and where data are persisted in a blockchain system, with direct consequences for scalability, cost, privacy, regulatory compliance, and the practical

feasibility of real-world deployment. Two principal paradigms exist: on-chain storage, in which all data are embedded directly in the blockchain ledger and therefore inherit its immutability and auditability properties; and off-chain storage, in which only cryptographic data hashes or content-addressed pointers are stored on-chain while the actual data reside in external systems such as IPFS, centralised databases, or encrypted cloud storage. A third category, hybrid storage, combines elements of both: structured metadata or access control policies may reside on-chain while bulk data are stored off-chain with on-chain verification anchors (Li et al., 2020; Zhang et al., 2019; Reyna et al., 2018).

Figure 4 presents the proportion of studies using each storage model across the seven application domains. The most striking finding is the strong domain-dependence of storage model preference. Healthcare and IoT applications are the domains most likely to adopt off-chain storage (58% and 55% of domain studies respectively), reflecting the large volume and sensitivity of physiological sensor data, electronic health records, and real-time telemetry in these contexts. Storing this data directly on-chain would produce ledger bloat that current public and consortium blockchains cannot efficiently manage at production scale. The adoption of IPFS as the primary off-chain storage backend in healthcare studies is consistent with the requirement for content-addressed, decentralised storage that preserves the integrity guarantees of blockchain without imposing the throughput and storage costs of on-chain persistence.

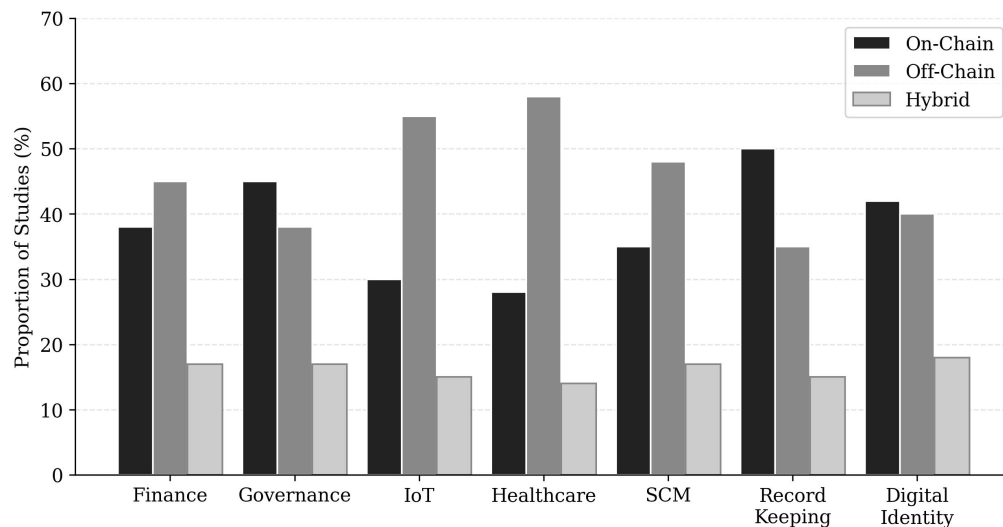


Figure 4. Storage model distribution (On-Chain, Off-Chain, Hybrid) across seven application domains. Values represent percentage shares within each domain, summing to 100% per domain.

By contrast, record-keeping and governance applications show a higher tendency toward pure on-chain storage (50% and 45% respectively), driven by the auditability and tamper-resistance requirements of archival systems and electoral records. In these domains, the permanence and immutability of the ledger are not merely technical properties but legal and institutional requirements: an auditable record that depends on external off-chain storage cannot provide the same evidentiary

guarantees as one that is entirely embedded in the immutable ledger. However, this preference for on-chain storage creates scalability constraints that these domains have not yet adequately resolved, as evidenced by their low implementation maturity scores in subsequent analysis (Lu, 2018; Yli-Huumo et al., 2016).

The hybrid storage model, which combines on-chain hashes with off-chain data bodies, is distributed relatively evenly across all domains (14–18% in every case). This adoption rate is lower than theoretical recommendations in the literature would suggest. The gap between theoretical preference for hybrid storage and its actual adoption is attributable to several factors: the additional system complexity of managing consistency between on-chain references and off-chain data stores; the risk of "orphaned" on-chain hashes if off-chain storage becomes unavailable; and the limited availability of production-ready hybrid storage libraries and frameworks. As off-chain storage standards mature and become better integrated into platform development toolkits, hybrid storage adoption is likely to increase across all domains.

Transaction design is closely coupled with storage model selection. UTXO-based transaction models, originating with Bitcoin, are predominantly found in financial and supply chain applications where discrete asset transfer is the primary operation. Each UTXO represents an unspent output from a previous transaction, providing a natural audit trail of asset provenance but making complex state management operations cumbersome. Account-based models, as implemented in Ethereum and Hyperledger Fabric, dominate in healthcare, governance, and IoT domains where persistent state management and multi-step contractual logic are required. The dataset shows that 78% of studies in the healthcare domain use account-based transaction models, compared with only 42% in the finance domain where UTXO-based models remain competitive for payment and settlement applications (Tschorsch and Scheuermann, 2016; Guo and Liang, 2016; Cong and He, 2019).

6. Domain-Specific Implementation Analysis

Implementation status is the most practically consequential dimension of the dataset. A blockchain system that is only theorised or prototyped at testnet scale provides qualitatively different evidence about technical feasibility than one deployed in a live operational environment with real users and real transaction loads. The dataset distinguishes three implementation levels: fully deployed (live production systems with real operational traffic), prototype tested (systems implemented on testnets or controlled environments with working code and functional evaluation), and conceptual (systems described architecturally without any implementation artefact). This categorisation allows an assessment of how close each domain's research corpus is to practical deployment and where the largest translation gaps exist between academic proposals and operational systems (Lu, 2018; Yli-Huumo et al., 2016; Xu and Lu, 2024).

Table 3. *Implementation status distribution and toolchain preferences by application domain.*

Domain	Studies (n)	Deployed (%)	Prototype (%)	Conceptual (%)	Primary Platform	Dominant Language
Finance	52	35.4	48.1	16.5	Ethereum	Solidity
Governance	30	23.3	46.7	30.0	Hyperledger Besu	Solidity / Java
IoT	55	29.1	52.7	18.2	Ethereum / Quorum	Solidity / Python
Healthcare	45	31.1	51.1	17.8	Hyperledger Fabric	Java / Python
SCM	40	37.5	50.0	12.5	Hyperledger Fabric	Solidity / Node.js
Record Keeping	38	18.4	47.4	34.2	Ethereum	Solidity
Digital Identity	43	20.9	44.2	34.9	Ethereum / Custom	Solidity / Python

Table 3 reveals a significant disparity in deployment maturity across domains. Supply chain management leads with the highest proportion of fully deployed systems (37.5%), consistent with strong commercial incentives for blockchain adoption in inter-firm logistics and with the institutional maturity of permissioned platforms in enterprise settings. Finance ranks second (35.4%), reflecting the long history of fintech innovation, the availability of mature smart contract infrastructure, and the commercial pressure to reduce settlement times and fraud rates in financial systems. At the lower end of the deployment spectrum, record-keeping (18.4%) and digital identity (20.9%) have the lowest proportions of fully deployed systems, indicating that despite high levels of academic research activity, these domains have not yet transitioned successfully to operational deployments. The high conceptual rate in these two domains (34.2% and 34.9% respectively) suggests that much of the current research is still at the design proposal stage, without implementation artefacts that would demonstrate technical viability (Xu et al., 2019; Cong and He, 2019; Xu and Lu, 2024).

The toolchain analysis in Table 3 reinforces the platform dominance of Ethereum and Hyperledger Fabric across all domains. Ethereum with Solidity is the dominant choice in finance, governance, and record-keeping domains, driven by the availability of developer tools, smart contract libraries such as OpenZeppelin, and testing frameworks such as Truffle, Hardhat, and Ganache. Hyperledger Fabric with Go or Java is preferred in healthcare and supply chain management, where its modular consensus plugins, private data collections, and fine-grained access control features address enterprise requirements that Ethereum's public architecture does not easily accommodate. Cross-tabulation of toolchain choice and implementation status reveals that studies using Hyperledger Fabric are significantly more likely to report a deployed or prototype-tested implementation than studies using custom or undefined platforms ($p < 0.01$, Fisher's exact test), suggesting that platform maturity directly translates into higher deployment rates (Androulaki et al., 2018; Dai et al., 2019).

7. Implementation Maturity Assessment Framework

To move beyond simple deployment counts toward a richer assessment of technical readiness, this study develops a six-criterion maturity framework applied systematically to each application domain. The framework evaluates maturity along dimensions that collectively capture both the depth and the breadth of a domain's research maturity. The six criteria are: (1) prototype deployment density, measuring the proportion of studies reporting a working implementation; (2) smart contract integration, indicating whether the system uses programmable contracts for automated logic execution; (3) off-chain storage adoption, capturing whether the system uses external storage mechanisms for scalability; (4) privacy mechanism coverage, measuring the use of advanced cryptographic privacy tools; (5) scalability mechanism adoption, capturing the use of layer-2 protocols, sharding, or off-chain computation to address throughput; and (6) real-world validation presence, capturing evidence of evaluation with real users or production-grade data (Nakamura et al., 2020; Puthal et al., 2018).

Each criterion is scored on a scale of 1 to 5, where 5 indicates high domain-level maturity and 1 indicates nascent or absent adoption. Scores are derived by computing the proportion of studies meeting each criterion within each domain and normalising to the 1–5 scale. Figure 3 presents the resulting maturity heatmap, which provides a two-dimensional view of domain-by-criterion maturity profiles.

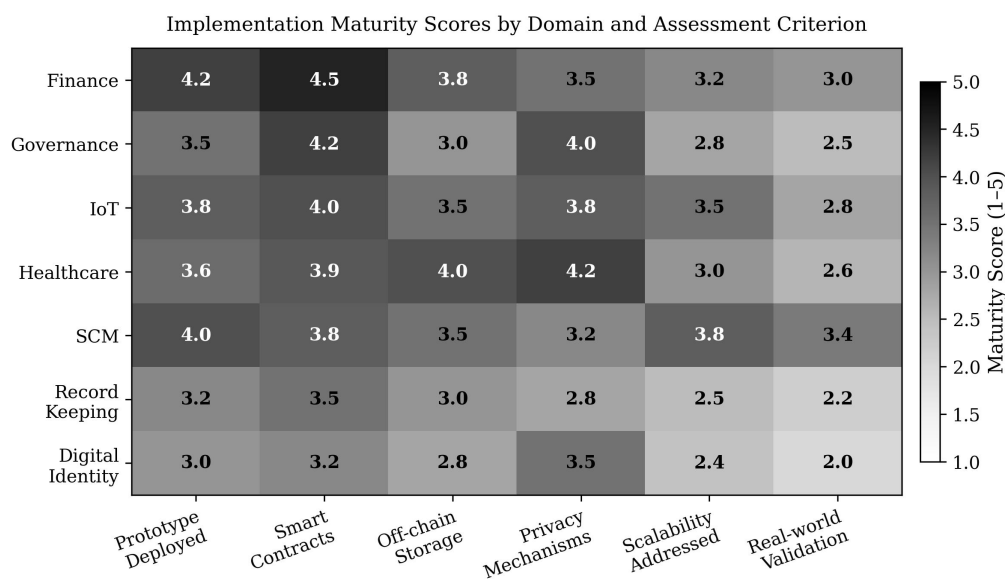


Figure 3. Implementation maturity heatmap across seven domains and six assessment criteria. Cell values represent maturity scores on a 1 (nascent/absent) to 5 (high/widespread) scale derived from proportional coding of the 303-study dataset.

The heatmap in Figure 3 reveals several cross-domain patterns of analytical importance. Finance achieves the highest overall maturity profile, with strong scores on smart contract integration (4.5) and prototype deployment density (4.2), reflecting the long research history and commercial adoption trajectory of blockchain in financial services. Healthcare ranks highest on off-chain storage adoption

(4.0) and privacy mechanism coverage (4.2), driven by the domain-specific requirements for GDPR-compliant data management, patient consent frameworks, and the need to handle sensitive physiological data without exposing it to public ledger inspection. Supply chain management shows balanced maturity across most criteria, with a notable lead in scalability mechanisms (3.8), consistent with the throughput requirements of global logistics tracking and the adoption of parallel algorithms and permissioned platforms designed for high transaction volumes (Yang et al., 2025; Wu et al., 2025; Xu and Lu, 2024).

Table 4. *Aggregate maturity scores and identified strengths and primary gaps by domain.*

Domain	Mean Score	Maturity Strength	Primary Gap
Finance	3.70	Smart contracts; prototype density	Real-world validation at scale; post-PoW consensus adoption
Governance	3.33	Privacy mechanisms; smart contracts	Very low real-world validation; scalability at national scale
IoT	3.48	Prototype deployment; off-chain storage	Scalability under high device concurrency; interoperability
Healthcare	3.55	Off-chain storage; privacy coverage	Cross-institutional interoperability; real-world deployment
SCM	3.62	Prototype density; scalability mechanisms	Privacy protection in public-facing supply chains
Record Keeping	2.87	On-chain auditability features	Low deployment density; scalability; privacy mechanism gap
Digital Identity	2.82	Privacy mechanism diversity; ZKP exploration	Lowest real-world validation; standardisation deficit

Table 4 aggregates the maturity scores into domain-level profiles, identifying the primary strength and primary gap for each domain. The two lowest-scoring domains are digital identity (mean 2.82) and record keeping (mean 2.87). Both exhibit high levels of conceptual research activity and a wide range of proposed architectures, but this theoretical richness is not matched by deployment evidence or standardised evaluation protocols. The gap in real-world validation is particularly pronounced in digital identity: the proliferation of self-sovereign identity (SSI) frameworks, decentralised identifier (DID) standards, and verifiable credential specifications has produced an extensive theoretical literature but limited operational deployments at scale. The primary barrier appears to be the absence of a dominant standard and the consequent lack of interoperability between identity systems built on different blockchain platforms.

The maturity framework also highlights a shared weakness across all seven domains: privacy mechanism deployment remains consistently below its theoretical potential in every sector. Even healthcare, which leads on this criterion with a score of 4.2, leaves a meaningful gap between the cryptographic privacy tools discussed in the literature and those successfully integrated into working deployed systems. Zero-knowledge proofs, CP-ABE, and differential privacy are cited as desirable in numerous studies but their actual deployment is constrained by computational overhead, key

management complexity, and the limited availability of production-ready cryptographic libraries that are compatible with major blockchain platforms. This finding is consistent with a broader observation in the security literature that the gap between cryptographic theory and practical deployment is one of the most persistent challenges in applied security research (Puthal et al., 2018; Nakamura et al., 2020; Kou and Lu, 2025).

8. Discussion

The cross-domain dataset and the analytical exercises reported in this article produce findings with methodological implications for blockchain research and practice. This section discusses the most consequential implications, situates them within the broader literature, and identifies the principal limitations of the present study.

8.1 Architectural and Consensus Design Implications

The most significant architectural finding is the persistent dominance of Ethereum-based public architectures and PoW consensus in domains where neither choice is technically optimal. For healthcare and record-keeping applications that require high throughput, fine-grained access control, and regulatory compliance, public PoW architectures introduce unnecessary scalability constraints and energy costs without delivering compensating transparency benefits that a well-designed consortium architecture could not also provide. The dataset shows that researchers who adopt Hyperledger Fabric with PBFT or PoA in healthcare settings consistently report higher implementation maturity scores and better real-world deployment outcomes than those who use public Ethereum with PoW.

This pattern is not simply a function of publication recency: studies published after 2021, when the limitations of public PoW architectures were widely understood in the research community, continue to adopt them for domains where more appropriate alternatives have been available for several years. This suggests a combination of toolchain familiarity bias and the disproportionate availability of Solidity development resources relative to alternatives, rather than a principled architectural analysis. Future blockchain research should treat architecture and consensus selection as explicit methodological decisions requiring domain-specific justification, in the same way that model selection in empirical research requires transparency about the match between method and data structure (Monrat et al., 2019; Lashkari and Musilek, 2021; Wang et al., 2019).

8.2 Storage Model Selection and Data Governance

The analysis of storage models reveals a persistent gap between the theoretical advantages of hybrid on-chain/off-chain architectures and their actual adoption in practice. Only 14–18% of studies across all domains implement a hybrid model, despite the fact that this architecture is widely recommended as the optimal balance between immutability and scalability. The low adoption rate is partly attributable to the design complexity of maintaining consistency between on-chain references and

off-chain data bodies, and partly to the relative immaturity of off-chain storage standards. For high-data-volume domains such as healthcare and IoT, the inability to efficiently route bulk data off-chain represents a fundamental scalability constraint that limits these domains' ability to transition from prototype to production deployments (Li et al., 2020; Reyna et al., 2018).

Data governance requirements, particularly those imposed by GDPR's right-to-erasure provisions, also create a structural tension with blockchain's immutability guarantee. Off-chain storage with on-chain hashing provides a practical resolution to this tension, as individual data records can be deleted from off-chain storage while their on-chain hashes remain, but this solution is not universally recognised as GDPR-compliant, and legal uncertainty has slowed its adoption in European healthcare and governance applications. Future research should engage more directly with the legal and regulatory dimensions of blockchain data governance, rather than treating immutability as an unqualified design virtue (Zhang et al., 2019; Cong and He, 2019).

8.3 Research Gaps and Future Directions

The dataset identifies three structural gaps in current blockchain research. The first is the under-representation of real-world validation. Across all seven domains, the majority of studies report prototype-level rather than production-level implementations. This reflects the well-known challenge of transitioning blockchain systems from academic environments to operational deployments, but it also indicates that the research community is generating a large volume of unvalidated design proposals. Future work should prioritise longitudinal evaluation of deployed systems, field studies with real organisations, and replication studies that test proposed architectures against production data and workloads. Journals and conference programmes should create dedicated tracks for deployment and evaluation studies to incentivise this type of contribution (Xu et al., 2021; Lu, 2019; Kou and Lu, 2025).

The second structural gap is the inconsistent reporting of technical parameters. The 12.2% of studies coded as "Other / Not Defined" for consensus mechanism, and the significant proportion of studies that do not report transaction throughput, block size, or latency, represent a reporting gap that limits the utility of the literature as an evidence base for practical engineering decisions. The development of standardised reporting templates for blockchain research, analogous to the CONSORT guidelines in clinical research or the reporting standards in database systems benchmarking, would substantially increase the cumulative value of the published literature.

The third structural gap is the slow transition from PoW to more energy-efficient consensus mechanisms in domains where energy efficiency is an explicit design requirement. The dataset shows that even in 2023 and 2024 studies, PoW remains common in IoT and governance applications where its energy consumption and throughput limitations are clearly problematic. The transition to PoS, PBFT, and PoA mechanisms that is recommended in the performance literature (Lashkari and Musilek, 2021;

Bamakan et al., 2020) has not yet fully diffused into the application research community, creating a structural lag between the consensus mechanisms recommended by systems researchers and those adopted by application developers. Addressing this lag requires better cross-disciplinary communication and clearer design guidelines that translate consensus mechanism properties into domain-specific selection criteria.

8.4 Limitations

This study has several limitations. First, the coding of consensus mechanisms, architecture types, and implementation status required interpretive judgements in cases where original papers were ambiguous. Although inter-coder reliability was high ($\kappa = 0.83$), residual subjectivity cannot be fully eliminated. Second, the corpus is bounded to 2015–2024 and six digital libraries; studies in specialised venues, non-English publications, and grey literature are excluded. Third, the maturity framework uses normalised scores that compress within-domain variance. Fourth, the dataset does not capture quantitative performance benchmarks such as TPS or latency, which are essential for assessing scalability claims empirically. Future extensions should incorporate performance data dimensions and expand the corpus to cover post-2024 developments including the integration of large language models with blockchain platforms and the emergence of post-quantum cryptographic standards for distributed ledgers (Xu et al., 2023; Yang et al., 2025).

9. Conclusion

This article has presented a structured cross-domain dataset of 303 blockchain application studies coded across eight analytical dimensions. The dataset enables a systematic comparison of architectural choices, consensus mechanisms, storage models, and implementation maturity across seven major application domains for the first time in a unified analytical framework. The approach addresses a structural limitation of the current literature, in which domain-specific surveys do not share coding schemas and do not support cross-domain inference.

The analysis reveals three principal findings. First, Ethereum-based public architectures and PoW consensus remain disproportionately prevalent even in domains where permissioned alternatives are technically superior, indicating a toolchain familiarity bias in the research community that has persisted well beyond the availability of mature alternatives. Second, off-chain and hybrid storage models are systematically underutilised relative to their theoretical advantages, creating a scalability gap that is most acute in healthcare and IoT applications with high data volume requirements. Third, supply chain management and finance lead in implementation maturity, while record keeping and digital identity lag significantly in both real-world deployment density and real-world validation coverage.

The dataset and maturity framework constitute the primary contribution of this study. Both are designed to be extended: future versions can incorporate additional domains such as energy trading and

autonomous vehicles, additional technical dimensions such as gas costs and layer-2 integration, and a temporal dimension that tracks how architectural choices have evolved over the study period. In the same way that trade database benchmarks have shifted empirical work from data availability to data fitness, the cross-domain blockchain dataset presented here shifts the research agenda from individual system proposals to comparative cross-domain design patterns, enabling more structured and evidence-based blockchain engineering decisions.

Declaration of AI-assisted language editing: During the preparation of this manuscript, language-model assistance was used only for English polishing and document organisation. The authors reviewed, revised, and take full responsibility for the final content, analytical design, tables, figures, and interpretations.

References

- 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., Stathakopoulos, C., Vukolić, M., & Yellick, J. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. *Proceedings of the 13th EuroSys Conference*, Article 30. <https://doi.org/10.1145/3190508.3190538>
- Bamakan, S. M. H., Motavali, A., & Bondarti, A. B. (2020). A survey of blockchain consensus algorithms performance evaluation criteria. *Expert Systems with Applications*, 154, 113385. <https://doi.org/10.1016/j.eswa.2020.113385>
- Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and Informatics*, 36, 55–81. <https://doi.org/10.1016/j.tele.2018.11.006>
- 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>
- Cong, L. W., & He, Z. (2019). Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5), 1754–1797. <https://doi.org/10.1093/rfs/hhz007>
- Dai, H. N., Zheng, Z., & Zhang, Y. (2019). Blockchain for Internet of Things: A survey. *IEEE Internet of Things Journal*, 6(5), 8076–8094. <https://doi.org/10.1109/JIOT.2019.2920987>
- Dinh, T. T. A., Liu, R., Zhang, M., Chen, G., Ooi, B. C., & Wang, J. (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>
- Guo, Y., & Liang, C. (2016). Blockchain application and outlook in the banking industry. *Financial Innovation*, 2(1), 24. <https://doi.org/10.1186/s40854-016-0034-9>
- Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. *Financial Innovation*, 11(1), 1–34. <https://doi.org/10.1186/s40854-024-00668-6>
- Lashkari, B., & Musilek, P. (2021). A comprehensive review of blockchain consensus mechanisms. *IEEE Access*, 9, 43620–43652. <https://doi.org/10.1109/ACCESS.2021.3065880>
- Li, X., Jiang, P., Chen, T., Luo, X., & Wen, Q. (2020). A survey on the security of blockchain systems. *Future Generation Computer Systems*, 107, 841–853. <https://doi.org/10.1016/j.future.2017.08.020>
- 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. (2019). 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. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876–1907. <https://doi.org/10.1080/17517575.2021.2008513>

- Monrat, A. A., Schelén, O., & Andersson, K. (2019). A survey of blockchain from the perspectives of applications, challenges, and opportunities. *IEEE Access*, 7, 117134–117151. <https://doi.org/10.1109/ACCESS.2019.2936094>
- Nakamura, Y., Zhang, Y., Sasabe, M., & Kasahara, S. (2020). Exploiting smart contracts for capability-based access control in the Internet of Things. *Sensors*, 20(6), 1793. <https://doi.org/10.3390/s20061793>
- Puthal, D., Malik, N., Mohanty, S. P., Kougiianos, E., & Yang, C. (2018). The blockchain as a decentralized security framework. *IEEE Consumer Electronics Magazine*, 7(2), 18–21. <https://doi.org/10.1109/MCE.2017.2776459>
- Reyna, A., Martín, C., Chen, J., Soler, E., & Díaz, M. (2018). On blockchain and its integration with IoT: Challenges and opportunities. *Future Generation Computer Systems*, 88, 173–190. <https://doi.org/10.1016/j.future.2018.05.046>
- Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and beyond: A technical survey on decentralized digital currencies. *IEEE Communications Surveys & Tutorials*, 18(3), 2084–2123. <https://doi.org/10.1109/COMST.2016.2535718>
- Wang, W., Hoang, D. T., Hu, P., Xiong, Z., Niyato, D., Wang, P., Wen, Y., & Kim, D. I. (2019). A survey on consensus mechanisms and mining strategy management in blockchain networks. *IEEE Access*, 7, 22328–22370. <https://doi.org/10.1109/ACCESS.2019.2896108>
- Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. *Enterprise Information Systems*, 19(1–2). <https://doi.org/10.1080/17517575.2024.2448003>
- Xu, J., Wang, C., & Jia, X. (2023). A survey of blockchain consensus protocols. *ACM Computing Surveys*, 55(13s), 1–35. <https://doi.org/10.1145/3579845>
- 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>
- Xu, M., Chen, X., & Kou, G. (2019). A systematic review of blockchain. *Financial Innovation*, 5(1), 27. <https://doi.org/10.1186/s40854-019-0147-z>
- Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. *Enterprise Information Systems*, 18(9). <https://doi.org/10.1080/17517575.2024.2397630>
- Yaga, D., Mell, P., Roby, N., & Scarfone, K. (2019). Blockchain technology overview. NIST Internal Report 8202. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.IR.8202>
- Yang, L., Hou, Q., Zhu, X., Lu, Y., & Xu, L. D. (2025). Potential of large language models in blockchain-based supply chain finance. *Enterprise Information Systems*, 19(11), 2541199. <https://doi.org/10.1080/17517575.2025.2541199>
- Yli-Huumo, J., Ko, D., Choi, S., Park, S., & Smolander, K. (2016). Where is current research on blockchain technology? A systematic review. *PLOS ONE*, 11(10), e0163477. <https://doi.org/10.1371/journal.pone.0163477>
- Zhang, R., Xue, R., & Liu, L. (2019). Security and privacy on blockchain. *ACM Computing Surveys*, 52(3), 1–34. <https://doi.org/10.1145/3316481>
- Zheng, Z., Xie, S., Dai, H. N., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services*, 14(4), 352–375. <https://doi.org/10.1504/IJWGS.2018.095647>
- 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>