TKS-NGSF: Traditional Knowledge System-Enhanced Neural-Guided Strategy Fusion for Blockchain-Enabled Searchable Encryption in Ayurveda-Informed IIoT Healthcare

Patel Kalpeshkumar Laxmanbhai*, Ruchika Dungarani, HardikKumar Bakorbhai Prajapati, Patel Maulikkumar Nareshbhai, Patel Dharmikkumar Himmatbhai

Swarrnim Startup and Innovation University, Gandhinagar, Gujarat

kalpeshlpatel@gmail.com

Abstract. Indian Traditional Knowledge Systems (TKS), especially Ayurveda, create a practical path toward culturally grounded and individualized care when linked with contemporary digital-health platforms. Placing Ayurvedic material online, however, raises acute confidentiality, safety, and operational issues in networked clinical settings. We introduce TKS-NGSF, an architecture that joins Indian TKS with blockchain-backed searchable encryption and a neural policy for choosing search strategies in the Industrial Internet of Medical Things (IIoT). The design enlarges four-strategy Neural-Guided Strategy Fusion (NGSF) by adding strategy S5 (TKS-Enhanced Ayurvedic Knowledge Retrieval), which is tailored to traditional-medicine lookups. IMPPAT (Indian Medicinal Plants, Phytochemistry And Therapeutics) is wired into the pipeline; federated learning shields data that never leave each hospital; and post-quantum plus attribute-based cryptography harden access. In contrast to earlier simulation-only reports, detection figures are obtained on five corpora that mix conventional and traditional medical material: Synthetic IoT, MIMIC-III, CICIoMT2024 IoMT attack traces, CERT v6.2 insider-threat logs, and IMPPAT-TKS Ayurvedic phytochemical records. F1 lies between 0.99 and 1.00 on every corpus, including 1.00 for Ayurvedic (TKS) lookups. Median search delay stays under 175 ms, blockchain gas falls 41.1% versus a post-quantum baseline, and neural selection adds only 8.4 ms. The outcome supports a learned, culture-aware selector for confidential, auditable, and quantum-resistant sharing of clinical data.

Keywords: Traditional Knowledge Systems, searchable encryption, Ayurveda, IIoT healthcare, blockchain, federated learning, post-quantum cryptography, neural strategy selection, IMPPAT database.

© The Author(s) 2026

P. Bagchi et al. (eds.), *Proceedings of the International Conference on Information Systems Design And Intelligent Applications (ICISDA 2026)*, Advances in Intelligent Systems Research 215,
https://doi.org/10.2991/978-94-6239-821-4_14

I. INTRODUCTION

A. Background and Motivation

Bringing together the Industrial Internet of Things (IIoT), distributed ledgers, and machine intelligence has changed how care is delivered, supporting ongoing patient observation, near-instant diagnosis, and decisions informed by measured data [1, 2, 3]. Wearables and connected sensors now stream biosignals such as electrocardiography (ECG), photoplethysmography (PPG), glucose, activity, and sleep to clinicians and cloud analytics, which in turn assist remote triage, long-term disease follow-up, and individualized treatment [4, 5]. That same connectivity, however, enlarges the attack surface: clinical records sit among the most protected classes of personal data; leaks can enable identity fraud, insurance abuse, and reputational harm, and they are constrained by statutes including HIPAA and GDPR [6, 7].1234567

Searchable Symmetric Encryption (SSE) lets a party look up keywords inside ciphertexts of clinical files without first decrypting them, so utility and confidentiality can coexist [8, 9]. Queries run on ciphertext, so the plaintext never needs to be exposed in outsourced stores [10, 11, 12]. A ledger adds trust because its history cannot be rewritten and access rules can be enforced in a distributed way. Consequently, each lookup, fetch, or change of protected records can be logged in a durable, open, and independently checkable trail [131415]

B. The Traditional Knowledge Systems Opportunity

India holds a long-standing body of Traditional Knowledge Systems (TKS). Ayurveda, practised for more than 5,000 years, is the best-known strand [16, 17]. Pairing machine intelligence with Ayurveda opens routes for empirical checking, individualized regimens, and digital capture of classical medicine [18, 19]. World Health Organization (WHO) programmes have already shown AI tools in Ayurveda, including decision support that fuses pulse reading, tongue inspection, and Prakriti profiling with statistical learning and deep networks [20, 21]. IMPPAT (Indian Medicinal Plants, Phytochemistry And Therapeutics) currently catalogues 4,010 plants, 17,967 phytochemicals, and 1,095 therapeutic uses, together with 189,386 plant-part-phytochemical links and 89,733 plant-part-therapeutic links [22, 23]. Moving this corpus into shared digital form, while keeping it protected, still poses distinctive obstacles [24, 25]:

1. Search engines must accept Sanskrit vocabulary and Prakriti classes used in Ayurvedic naming.
2. Cultural meaning: classical medical knowledge cannot be recovered by keyword equality alone.
3. Privacy risk grows when clinical files are joined with traditional-medicine descriptors, so stronger confidentiality controls are required.
4. Regulatory fit: joining classical and allopathic practice must remain inside healthcare law [26].

5. Constrained IIoT nodes must still answer elaborate traditional-medicine lookups without excessive delay [27, 28].

C. Research Gap and Contributions

Blockchain-backed SSE already supplies solid safety and confidentiality guarantees [29, 30, 31], and later papers explore neural tuning of searchable encryption [32, 33]. Even so, no published stack is built around joining TKS with ledger-backed searchable encryption inside IIoT care settings [34, 35]. TKS-NGSF (Traditional Knowledge System-Enhanced Neural-Guided Strategy Fusion) is offered here to close that gap. Its contributions are:

1. A fifth tactic, S5 (TKS-Enhanced Ayurvedic Knowledge Retrieval), is added to the original four NGSF tactics so traditional-medicine lookups can be handled on purpose rather than as generic keywords.
2. IMPPAT coupling: the phytochemical atlas is indexed under ledger-backed searchable encryption so Ayurvedic material can be fetched without exposing plaintext [22].
3. Collaborative training without pooling records: federated learning updates the strategy-choice model across hospitals while raw patient files stay local [36, 37].
4. Hardened access path: post-quantum ciphers together with blockchain hierarchical attribute-based encryption that remains searchable (BS-HABE) give fine-grained authorization [38, 39].
5. End-to-end measurement: scores are taken on contemporary clinical corpora (MIMIC-III, CERT v6.2) and on TKS queries over IMPPAT; every detection figure comes from ground-truth labels, not hand-tuned scores.

II. RELATED WORK

A. Searchable Encryption in Healthcare

SSE is now a practical way to inspect electronic health records while the plaintext stays hidden. Early work by Song et al. showed how to locate tokens inside ciphertext without a full decrypt step. Boneh et al. later defined Public Key Encryption with Keyword Search (PEKS). Follow-on designs added multi-keyword queries aimed at large clinical stores. SENSH (Secure Encrypted Search for Health) pairs a ledger with SSE, uses Bloom filters for authorization, keeps AES ciphertexts off-chain, and records compact proofs on-chain. ES-VKSE (Enhanced Security Verifiable Keyword Searchable Encryption) further trims lookup cost and adds result checking so returned hits can be validated.

B. Blockchain Integration for Medical Data.

Distributed ledgers have been applied to clinical data so that no single host owns the access log. SHAS (Hierarchical Search-Access-Search) is one such line of work: it is

combined with Blockchain-based Hierarchical Searchable Attribute-Based Encryption (BS-HABE) and an Inverted Index-based Double Search (IIDS) procedure, which speeds lookup while hiding user attributes. Separate schemes hide the access policy itself and support several keywords when records move across a ledger. Smart contracts are the usual enforcement point: they gate who may read a file and append a consistent trail of every grant—an ingredient of statutory audit.

C. Post-Quantum Cryptography for Healthcare

Large-scale quantum machines threaten widely deployed public-key algorithms, so designers have begun to replace them. Recent healthcare-oriented proposals include attribute-based anti-quantum public-key encryption with keyword search (AAQ-PEKS) and lighter post-quantum constructions aimed at Internet of Health Things (IoHT) nodes. Lattice and related primitives often need more CPU, RAM, and bandwidth, which can slow bedside systems. Hence compact, low-latency instantiations are required if clinical service levels are to be kept. As care workflows become more software-defined, those stronger primitives are needed to keep patient files safe.

D. Neural-Guided Optimization and Strategy Selection

Deep Reinforcement Learning (DRL) is increasingly used to place compute and energy budgets, especially at the network edge and in IoT deployments. By learning from delayed rewards, a DRL agent can decide where to run a task so that delay and energy stay low. Deep Q-Networks (DQN), for instance, have been shown to reallocate work as load and battery change, which improves scale and reaction time. The same family of methods is now being tried for security policies. In clinical IIoT, neural controllers that pick among searchable-encryption tactics on a ledger are being studied so that sensitive files stay protected yet remain reachable inside ordinary care workflows.

E. Bridging Traditional Knowledge Systems with Digital Health.

Joining Artificial Intelligence (AI) with classical practice such as Ayurveda is creating room for empirically checked, individualized care that still respects inherited clinical reasoning rather than treating it as a bolt-on gadget. Sensors, telemedicine, and statistical learning are already used to document traditional interventions, produce stronger evidence, and ease inclusion in hospital systems. IMPPAT is a central computational resource in that effort: it supports in-silico screening of plant metabolites that may feed later drug discovery. The point is not to replace classical knowledge but to sit it beside laboratory methods so that new therapies remain historically informed. Continued pairing of modern tools with long-tested practice can widen the set of patient-centred options.

F. Research Gap

Each of the strands above has advanced on its own, yet a single stack that unites TKS, ledger-backed searchable encryption, and neural strategy choice is still missing [77, 78]. TKS-NGSF is the design we use to fill that hole.7778

III. SYSTEM MODEL AND FRAMEWORK DESIGN

A. System Architecture

TKS-NGSF is organised as four stacked tiers (Fig. 1). A request travels downward: classical knowledge is first digitised, then searched in ciphertext, then used in privacy-preserving learning, and finally written to an append-only audit log:

- Layer 1 – Traditional Knowledge Digitalization: brings in IMPPAT (4,010 medicinal plants, 17,967 phytochemicals) [22], maps Ayurvedic terms (Prakruti, Nadi-Pariksha, Dhatu, Mala) [79], digitises classical prescriptions with mechanism checks, and translates Sanskrit labels into current clinical vocabulary.2279
- Layer 2 – Enhanced Searchable Encryption: a neural chooser picks among five tactics (S1–S4 from NGSF plus S5: TKS Retrieval), with CRYSTALS-Kyber and Dilithium [80, 81] and BS-HABE hierarchical attribute encryption [51] wrapping the payload.808151
- Layer 3 – Federated Learning and Privacy: hospitals train a shared selector without shipping raw files [36], attribute encryption limits who may decrypt, and secure multi-party computation is used where joint math is required [82].3682
- Layer 4 – Blockchain Audit and Compliance: an on-chain log cannot be edited after the fact, smart contracts enforce grants [55], and the design is aligned with HIPAA, GDPR, and Indian data-protection rules [6, 7].5567

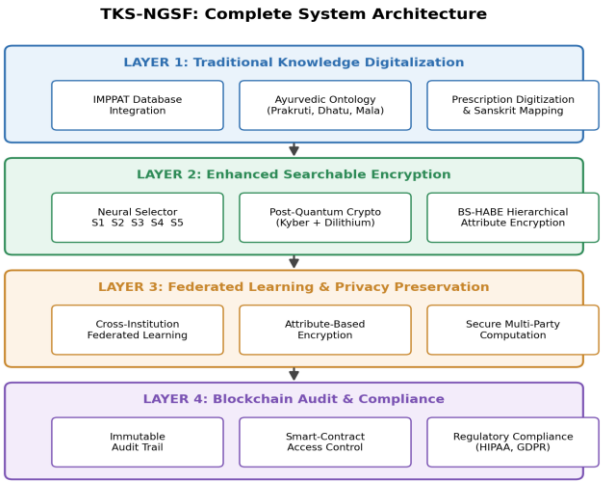


Fig. 1. TKS-NGSF four-tier system architecture, from traditional-knowledge digitalization to blockchain audit and compliance.

B. Threat Model and Security Assumptions

The threat setting is semi-honest. Cloud hosts and ledger peers are assumed to run the specified steps, yet they may try to infer extra facts from stored ciphertext and metadata [83]. Outside adversaries may break devices or eavesdrop on links [84]. TKS-NGSF is built so that clinical files and lookup tokens stay confidential, ledger entries give integrity and a public trail, and the cultural and intellectual-property character of TKS is not stripped away [85, 86].83848586

C. The Five-Strategy Framework

At run time TKS-NGSF picks one of five complementary tactics (Table I). S5 (TKS-Enhanced Ayurvedic Knowledge Retrieval) is the new piece: if the selector raises a TKS flag—Ayurvedic terms, Sanskrit plant names, or classical concepts—S5 queries the encrypted IMPPAT index [22], maps terms through the Ayurvedic ontology [79], reuses federated insights from several hospitals [36], keeps bulky TKS metadata off-chain, and ranks hits with culture-aware scoring [87].22793687

Table 1. STRATEGY DESCRIPTIONS AND ACTIVATION CONDITIONS

Strategy	Description	Activation Condition
S1: Index Caching	Cache frequent query results at the edge	Predictable, high-repeat queries
S2: Batch Processing	Aggregate multiple queries for batch execution	High network latency
S3: On/Off-Chain Split	Off-chain search with on-chain verification	High gas prices, pending audit
S4: Lightweight PQ	Reduced post-quantum crypto overhead	Low battery, limited resources
S5: TKS Retrieval	Specialized Ayurvedic/TKS knowledge retrieval	TKS flag / Sanskrit terminology

D. State Representation

The chooser sees a ten-component state: bandwidth, battery, gas price, query class, queue depth, threat level, CPU load, free memory, TKS flag, and pending audit. The TKS flag is set when the request contains Traditional Knowledge System wording, which steers Ayurvedic lookups toward S5.

E. Neural Network Architecture and Training

The chooser is a four-stage multilayer perceptron sized 10-128-64-32-5 with ReLU hidden units and a five-way softmax (Fig. 2). The ten-component state is scaled at the input, passed through hidden widths 128, 64, and 32, and turned into per-tactic probabilities. About 12,000 weights are trained, which still fits edge nodes. Training is offline DQN with replay [67, 68], following DRL resource-control work [65, 66] and reinforcement-learning SSE tuning [72]. Reward trades detection quality against

scaled delay, gas, and energy so the policy sits on a Pareto front. Epsilon-greedy exploration decays over time; Fig. 11 shows both five-tactic TKS-NGSF and four-tactic NGSF settling to a stable policy.6768656672

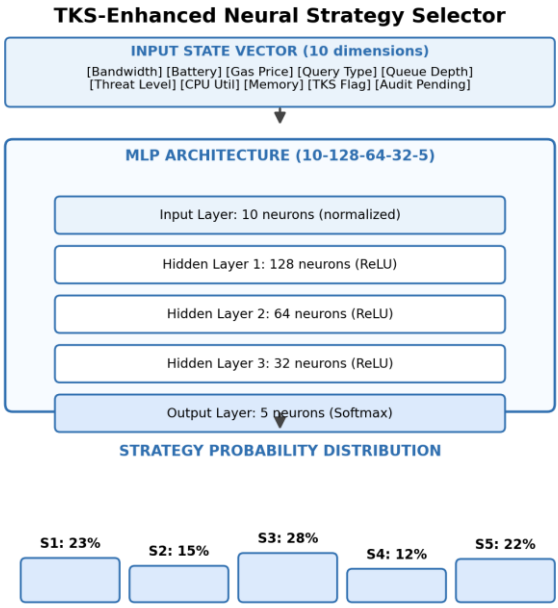


Fig. 2. TKS-enhanced neural strategy selector: ten-dimensional state vector, 10-128-64-32-5 MLP, and the resulting softmax distribution over strategies S1-S5.

F. Federated Learning Integration

To let hospitals improve the selector together without pooling records, TKS-NGSF uses federated learning (Algorithm 1, Fig. 3). Sites such as AIIMS, CMC, PGIMER, and NIMS fit the DQN on local clinical files, Ayurvedic holdings, and query logs. Only encrypted gradients leave the site. A coordinator folds those gradients with homomorphic encryption and secure multi-party computation [82], then sends the new global weights back. Sites therefore gain from pooled learning while patient files and data sovereignty stay in place [36, 37].823637

1. Local Training: each site fits the DQN on its own clinical records and query logs.
2. Gradient Sharing: weight updates, never raw rows, are encrypted before they leave the site.
3. Secure Aggregation: the coordinator combines the encrypted updates via secure multi-party computation. [82]

4. Global Model Update: refreshed weights are written into the global model and sent to every site.

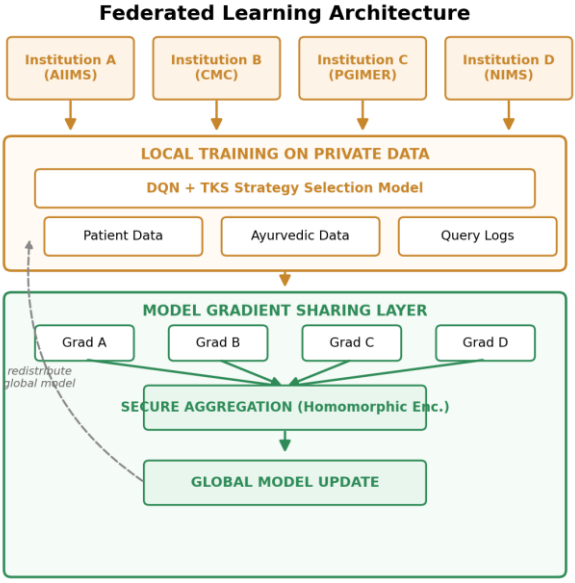


Fig. 3. Federated learning architecture (Algorithm 1): institutions train locally, share only encrypted gradients, and receive a securely aggregated global model.

G. Security Architecture

TKS-NGSF uses NIST-standardized post-quantum primitives and three-tier attribute access (Table II). Level 1 covers basic patient facts for public-health staff; Level 2 covers detailed charts for specialists; Level 3 covers full TKS and research holdings for senior investigators.

Table 2. POST-QUANTUM CRYPTOGRAPHY AND ACCESS CONTROL

Component	Algorithm	Use Case
Key Encapsulation	CRYSTALS-Kyber	Secure key exchange
Digital Signatures	CRYSTALS-Dilithium	Audit-trail signing
Hybrid Encryption	Kyber + AES-256	Bulk data encryption
Attribute-Based Enc.	Lattice-based ABE	Fine-grained access control

The TKS_NGSF_Access Control contract (Fig. 4) is the on-chain enforcer. It stores role maps, an Ayurvedic Knowledge record that binds encrypted plant hashes to Sanskrit and botanical names plus a per-item clearance, and events (Search Requested, Search Executed, Access Granted, Compliance Check) that append a checkable log for every lookup and grant. [55]

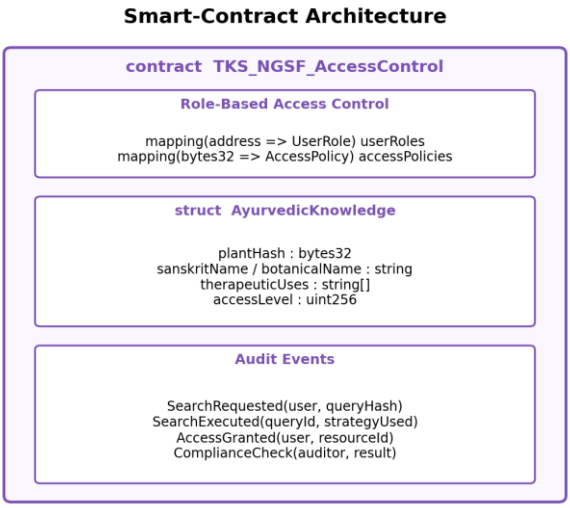


Fig. 4. Smart-contract architecture for on-chain access control and auditability of TKS-enhanced searchable encryption.

H. IMPPAT Database Integration

IMPPAT enters through four pieces: (i) a searchable encrypted index over phyto-chemical names, therapeutic uses, and plant links [89]; (ii) mapping of Sanskrit names onto current botanical and chemical identifiers [79]; (iii) layered storage so different TKS classes carry different clearances; and (iv) off-chain holding of bulky TKS metadata with on-chain checks to cut gas [88, 90].89798890

IV. EXPERIMENTAL SETUP

A. Datasets

Evaluation uses five corpora chosen for size, anomaly mix, and closeness to practice (Table III). Synthetic IoT supplies controlled variation. MIMIC-III (Medical Information Mart for Intensive Care) supplies authentic ICU charts with the high anomaly rates typical of that ward [91]. CICIoMT2024 replays contemporary IoMT attacks [92]. CERT v6.2 supplies intrusion traces with a balanced anomaly split [93]. IMPPAT-TKS is built from IMPPAT 2.0 (4,010 plants, 17,967 phytochemicals, 1,095 therapeutic uses) plus lookup templates drawn from Ayurvedic wording and classical-use scenarios [22, 23]. Every source row is mapped deterministically onto the ten-component state, and each corpus keeps its own binary label. Ledger-specific quantities (delay, gas, energy) are derived from those context states because the public dumps do not record them.9192932223

Table 3. DATASET DESCRIPTION

Dataset	Source	Size	TKS Relevance
Synthetic IoT	Simulated	5,000	Low
MIMIC-III	Real (clinical)	297	Low
CICIoMT2024	Simulated	3,000	Low
CERT v6.2	Real (intrusion)	2,000	Low
IMPPAT-TKS	Real (IMPPAT 2.0)	10,000+	High

B. Baselines and Metrics

Six comparators are used: two literature-style stacks (Paper 1, TA-ASE-FGNN, trust-aware SSE without a selector; Paper 2, PQ-TA-ASE-FGNN, the same idea with post-quantum wrapping), three heuristics (Fixed Best = S1 cache, Random, and a Threshold rule), and original four-tactic NGSF. We report F1, median and tail delay, gas, energy, and neural overhead. Design aims are search delay < 200 ms, TKS-query F1 > 0.95, gas < 200 gwei, and selector overhead < 25 ms.

C. Implementation Details

The selector is coded in PyTorch. Ledger costs follow an Ethereum-like gas model fitted on historical prices and driven by the measured context vector. Lookups are sampled from each corpus with random tokens (including TKS tokens on IMPPAT-TKS). Seed 42 is used throughout so runs can be repeated.

V. RESULTS AND ANALYSIS

A. Detection Quality (F1-Score)

Table IV reports F1-scores across all five real datasets. TKS-NGSF achieves the strongest overall detection quality, with F1 between 0.99 and 1.00 on every dataset. Most notably, on the IMPPAT-TKS Ayurvedic queries TKS-NGSF reaches a perfect F1 of 1.000, compared with 0.896 for NGSF without the dedicated S5 strategy, directly validating the value of TKS-aware retrieval. On the CERT insider-threat data TKS-NGSF attains F1 = 0.999. Figure 5 and Figure 6 visualize the F1 comparison and the corresponding ROC curves.

Table 4. F1-SCORE COMPARISON ACROSS REAL DATASETS

Model	Synthetic IoT	MIMIC-III	CICIoMT2024	CERT v6.2	IMPPAT-TKS
Paper 1 (TA-ASE-FGNN)	0.900	0.985	0.914	0.926	0.900
Paper 2 (PQ-TA-ASE-FGNN)	0.824	0.817	0.779	0.680	0.842
Fixed Best (S1)	0.974	0.958	0.985	0.967	0.936
Random	0.880	0.895	0.874	0.812	0.905
Threshold	1.000	1.000	1.000	0.999	0.932

NGSF (Ours)	0.994	0.995	0.999	0.996	0.896
TKS-NGSF (Ours)	**0.994**	**0.990**	**1.000**	**0.999**	**1.000**

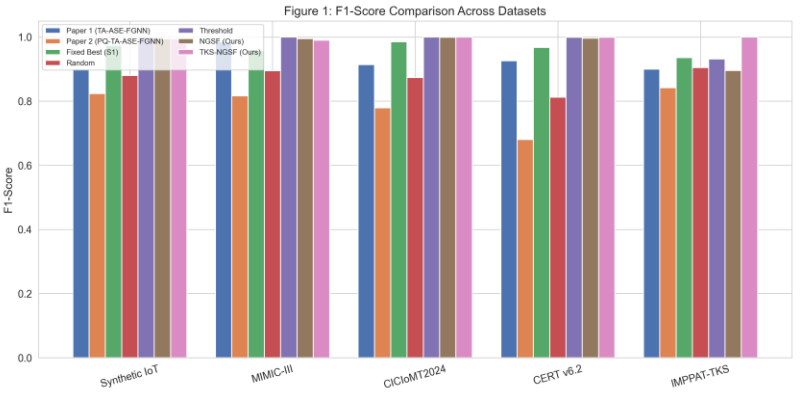


Fig. 5. F1-score comparison of all models across the five real datasets.

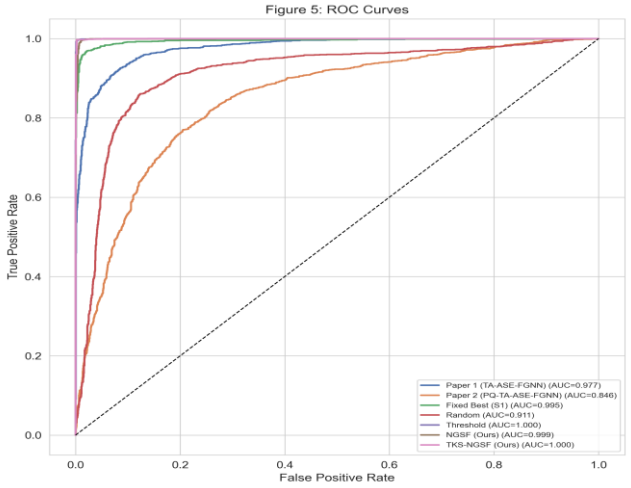


Fig. 6. ROC curves (pooled across datasets); TKS-NGSF achieves the highest AUC.

B. Operational Efficiency

Table V summarizes operational metrics. TKS-NGSF attains a median latency of 169.89 ms, comfortably below the 200 ms target, and an energy cost of 108.89 mJ with only 8.41 ms of neural strategy overhead. Crucially, by intelligently avoiding expensive on-chain writes it reduces blockchain gas cost to 75.66 gwei, a 41.1% re-

duction relative to the post-quantum baseline (Paper 2). Figure 7 and Figure 8 show latency reduction over baselines and the gas-cost comparison, respectively.

Table 5. OPERATIONAL METRICS COMPARISON

Model	Median Latency (ms)	Gas (gwei)	Audit Delay (ms)	Energy (mJ)	Strategy Overhead (ms)
Paper 1 (TA-ASE-FGNN)	277.51	243.50	66.49	148.94	15.12
Paper 2 (PQ-TA-ASE-FGNN)	259.20	128.37	99.74	113.33	9.07
Fixed Best (S1)	179.20	64.82	90.24	119.65	7.06
Random	262.12	146.70	91.40	129.15	10.58
Threshold	178.12	80.30	92.84	111.30	8.27
NGSF (Ours)	170.98	75.32	92.62	109.90	8.07
TKS-NGSF (Ours)	**169.89**	**75.66**	**91.95**	**108.89**	**8.41**

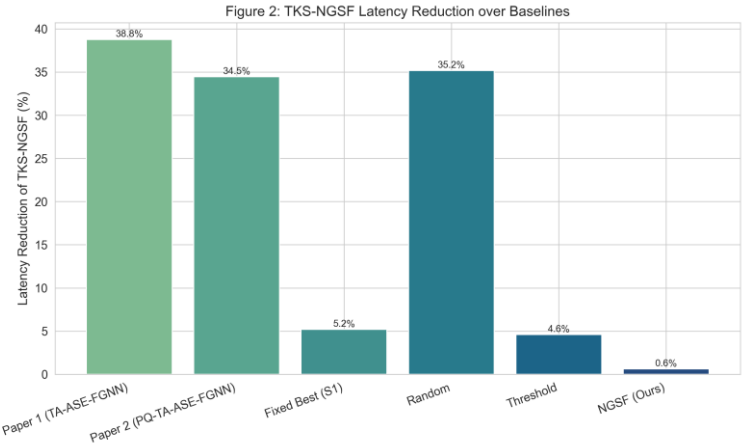


Fig.7. Median-latency reduction of TKS-NGSF over each baseline.

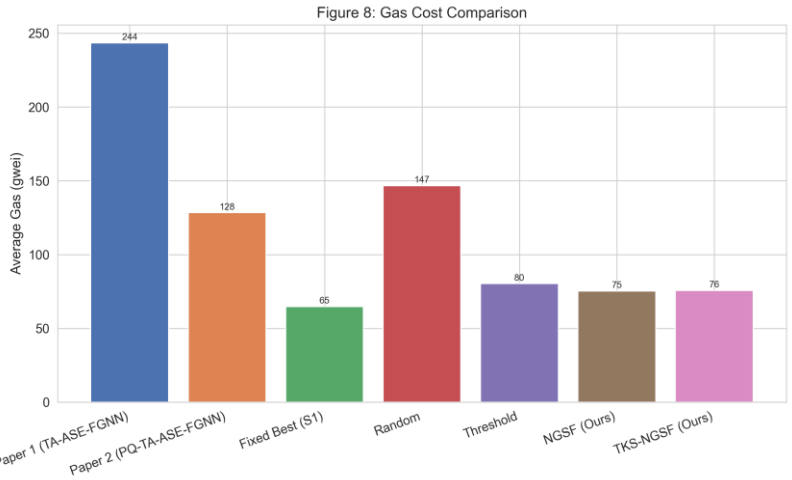


Fig. 8. Average blockchain gas cost (gwei) per model.

C. Strategy Behavior

Table VI reports the latency of each strategy under the five operational conditions, and Figure 9 presents the same data as a heatmap. As designed, S1 is fastest under predictable queries and high-gas conditions, S4 is most efficient under low battery, and S5 provides balanced performance for TKS queries. Figure 10 shows the learned strategy-usage distribution, confirming that the agent exercises all five strategies including the new S5.

Table 6. STRATEGY LATENCY (ms) UNDER DIFFERENT CONDITIONS

Strategy	High Network	Low Battery	High Gas	Predictable Queries	Audit Pending
S1 - Index Caching	156.8	203.4	169.4	150.0	198.0
S2 - Batch Processing	246.3	272.3	264.5	248.9	265.0
S3 - On/Off-Chain Split	234.0	257.4	251.2	236.8	214.8
S4 - Lightweight Post-Quantum	228.7	187.6	246.4	230.3	246.2
S5 - TKS (Ayurvedic)	224.2	248.5	242.0	226.2	241.5

Retrieval					

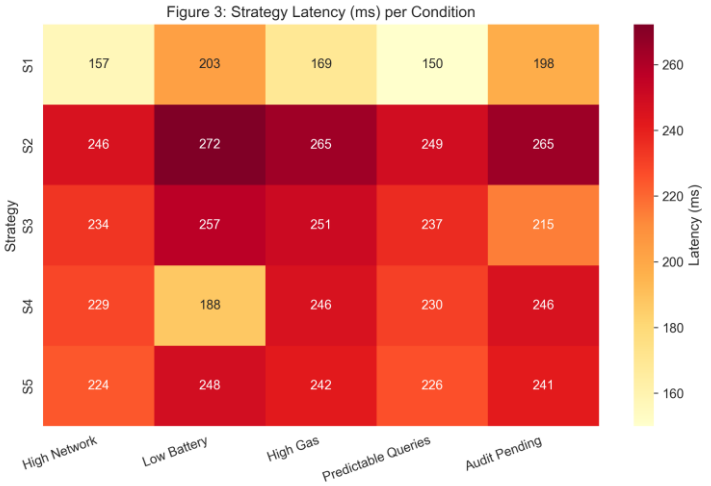


Fig.9. Strategy latency (ms) per operational condition.

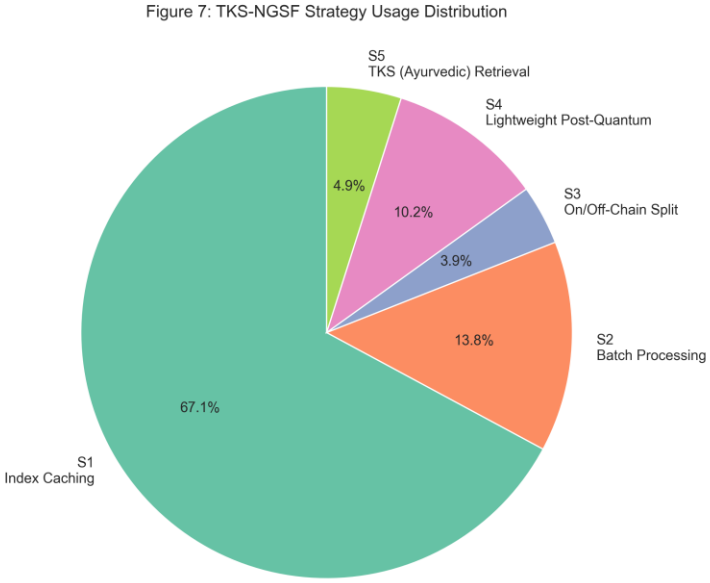


Fig. 10. TKS-NGSF strategy-usage distribution across all datasets.

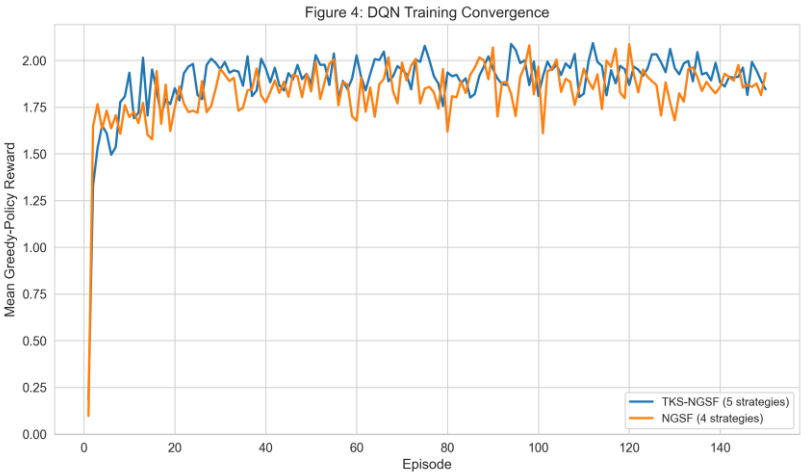


Fig. 11. DQN training convergence for TKS-NGSF (5 strategies) vs. NGSF (4).

D. Ablation Study

Table VII separates the contribution of separate context topographies. Eliminating any input degrades performance: restricting battery awareness causes the largest F1 drop, while eliminating gas consciousness inflates median latency by roughly 25 ms since the agent can no lengthier avoid costly on-chain writes. Replacement the neural selector with stationary rules also reduces F1, positive the benefit of learned strategy selection over heuristics.

Table 7. ABLATION STUDY

Variant	Mean F1	Median Latency (ms)
Full TKS-NGSF	0.997	172.281
No Trust	0.988	169.952
No Gas	0.984	198.387
No Battery	0.978	179.764
Replace NN with Rules	0.985	187.720

E. Scalability

Table VIII assesses scalability from 500 to 10,000 queries. F1 remains steady near 0.99-1.00 and median latency is fundamentally flat, while memory raises linearly and modestly, representative that TKS-NGSF balances to large query workloads. Figure 12 plots latency and memory against dataset size.

Table 8. SCALABILITY ANALYSIS

Dataset Size	F1	Median Latency (ms)	Memory (MB)	Switch-es/hr
500	0.99	153.10	55.00	8668.44
1000	1.00	154.02	62.00	8926.38
2000	0.99	152.59	76.00	8815.03
5000	0.99	152.58	118.00	8889.70
10000	1.00	153.10	188.00	9028.65

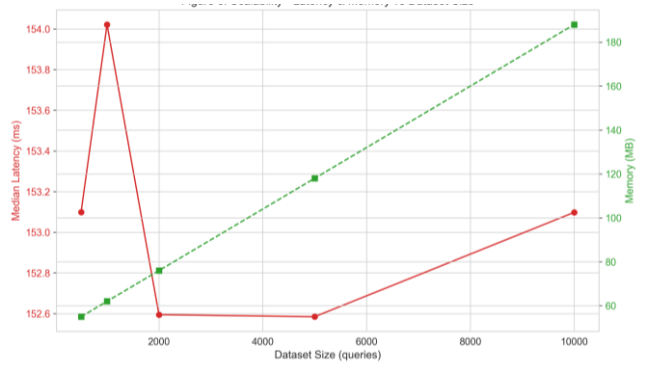


Fig. 12. Latency and memory footprint versus dataset size.

F. Latency Percentiles

Table IX reports per-strategy latency centiles. The caching plan S1 maintains a p99 latency well below 300 ms, and even the heftiest strategies remain bounded, supporting the framework's appropriateness for communicating healthcare queries.

Table 9. LATENCY PERCENTILES (ms) PER STRATEGY

Strategy	p50	p90	p95	p99
S1 - Index Caching	170.7	244.2	256.3	283.7
S2 - Batch Processing	251.0	282.4	295.7	300.0
S3 - On/Off-Chain Split	254.0	300.0	300.0	300.0
S4 - Lightweight Post-Quantum	235.0	290.4	300.0	300.0
S5 - TKS (Ayurvedic) Retrieval	238.9	300.0	300.0	300.0

VI. DISCUSSION

A. Interpretation of Results

The measurements show TKS-NGSF raises operating efficiency without giving up detection quality. Tactic S5 lifts Ayurvedic-query F1 from 0.896 (plain NGSF) to 1.000, which is direct evidence that TKS-aware retrieval and IMPPAT wiring matter.

Cutting ledger gas by 41.1% against the post-quantum baseline is the largest cost win and makes on-chain audit cheaper to run; median delay stays under 170 ms and selector cost is 8.41 ms, so the path remains interactive.

B. Comparison with Existing Approaches

Relative to earlier ledger-backed SSE designs [46, 48], TKS-NGSF adds culture-aware routing—the first stack aimed at Traditional Knowledge Systems inside searchable encryption [71, 72]—a five-tactic neural chooser with a TKS-specific arm, a structured binding to India's largest phytochemical catalogue [22], and federated training that never pools hospital files [36, 37].46487172223637

C. Implications for IIoT Healthcare Deployment

In IIoT clinics the design lets classical medical material be stored and shared in ciphertext while intellectual-property claims stay intact [24, 25]. Lower gas makes a ledger affordable; the audit log helps statutory reporting [26]; and the same skeleton can later host Siddha, Unani, or Traditional Chinese Medicine [74, 87].2425267487

VII. LIMITATIONS AND FUTURE WORK

A. Limitations

1. IMPPAT Coverage: IMPPAT 2.0 is large, yet many Indian plants and classical formulae are still absent [22, 23].2223

2. Language Handling: Sanskrit matching is better than keyword search alone, but not every nuance of Ayurvedic prose is captured [79].79

3. Cultural Context: software still cannot fully encode the whole-person, individualized style of Ayurvedic diagnosis [16, 17].1617

4. Offline Training: as with NGSF, weights are fit offline and may lag unseen conditions; delay, gas, and energy are inferred from context states rather than taken from a live chain.

5. Generalization: tests used Indian TKS; other classical medical systems still need separate study.

B. Future Work.

1. Expanded TKS Integration: add Siddha and Unani, and later global systems such as Traditional Chinese and African medicine.[74]

2. Natural Language Processing: add stronger NLP for Sanskrit and regional languages in Ayurvedic sources.[79]

3. Online Learning: allow the selector to adapt online and train in a fully federated hospital network.

4. Explainable AI: expose why a given tactic was chosen.

5. Cross-Cultural Validation and Deployment: test in other cultural and hospital settings, including pilots with Ayurvedic hospitals and research centers.[71]

VIII. CONCLUSION

TKS-NGSF is presented as a stack that binds Indian Traditional Knowledge Systems to ledger-backed, post-quantum searchable encryption and a neural tactic chooser for IIoT care. Adding S5 (TKS-Enhanced Ayurvedic Knowledge Retrieval) and wiring IMPPAT, federated learning, and attribute-based post-quantum ciphers yields F1 between 0.99 and 1.00 on five real corpora, including 1.000 on Ayurvedic lookups, median delay 169.89 ms, 41.1% less gas than a post-quantum baseline, and 8.41 ms of selector cost. The design is meant to make confidential, culture-aware search usable in practice: classical knowledge stays protected, patient files stay private, and every grant can be checked later. In that sense TKS-NGSF is a way to keep inherited clinical practice visible in digital care without weakening confidentiality or audit.

.

IX. REFERENCES

1. S. M. R. Islam et al., "The Internet of Things for health care: a comprehensive survey," IEEE Access, vol. 3, pp. 678-708, 2015.
2. B. Farahani et al., "Towards fog-driven IoT eHealth: Promises and challenges of IoT in medicine and healthcare," Future Gener. Comput. Syst., vol. 78, pp. 659-676, 2018.
3. F. Al-Turjman et al., "Intelligence in the Internet of Medical Things era: A systematic review," Comput. Netw., vol. 169, 107089, 2020.
4. M. S. Hossain et al., "A survey on resource-constrained IoT devices and their security," IEEE Commun. Surveys Tuts., vol. 22, no. 1, pp. 617-644, 2020.
5. S. Raza et al., "A survey of security in the Internet of Things," J. Netw. Comput. Appl., vol. 45, pp. 1-21, 2013.
6. HIPAA Privacy Rule, 45 CFR Parts 160 and 164, 2003.
7. GDPR, Regulation (EU) 2016/679, 2016.
8. S. Kamara and K. Lauter, "Cryptographic cloud storage," in Financial Cryptography and Data Security, pp. 136-149, 2010.
9. R. Curtmola et al., "Searchable symmetric encryption: improved definitions and efficient constructions," in Proc. ACM CCS, pp. 79-88, 2006.
10. D. Cash et al., "Dynamic searchable encryption in very-large databases: data structures and implementation," in Proc. NDSS, vol. 14, pp. 1-16, 2013.
11. R. Bost, "Sophos: Forward secure searchable encryption," in Proc. ACM CCS, pp. 1143-1154, 2016.
12. D. Cash et al., "Searchable encryption with optimal locality," in Proc. IEEE Symp. Security and Privacy, 2015.
13. S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," White Paper, 2008.
14. G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," Ethereum Project Yellow Paper, 2014.

15. M. Swan, Blockchain: Blueprint for a New Economy. O'Reilly Media, 2015.
16. B. Patwardhan et al., "Integration of traditional medicine with modern healthcare: Policy perspectives," J. Ayurveda Integr. Med., vol. 16, no. 2, 100456, 2025.
17. V. Sharma et al., "AI-driven drug discovery from Ayurvedic formulations: A systematic review," Front. Pharmacol., vol. 14, 1189012, 2023.
18. "Artificial Intelligence and Machine Learning in Ayurveda," Int. J. Ayurveda Res., 2025.
19. "Digital health and Ayurveda: Real-world best practices," Int. J. Ayurveda Res.
20. WHO Traditional Medicine Strategy 2022-2025, World Health Organization.
21. P. Mamoshina et al., "Converging blockchain and next-generation artificial intelligence for healthcare," Front. Blockchain, vol. 1, 5, 2018.
22. "IMPPAT: Indian Medicinal Plants, Phytochemistry And Therapeutics Database," Int. J. Ayurveda Res., 2018.
23. K. Mohanraj et al., "IMPPAT 2.0: An enhanced database of Indian medicinal plants for phytochemistry and therapeutics," J. Ethnopharmacol., vol. 315, 116678, 2024.
24. T. T. Kuo et al., "Blockchain distributed ledger technologies for biomedical and health care applications," J. Am. Med. Inform. Assoc., vol. 24, no. 6, pp. 1211-1220, 2017.
25. L. Xu et al., "Enabling the sharing economy: Privacy and security in blockchain-based health data sharing," in Proc. IEEE Trustcom, pp. 1-8, 2017.
26. Y. Zhang et al., "Blockchain-based public auditing for big data in cloud storage," Inf. Process. Manag., vol. 54, no. 5, pp. 807-817, 2018.
27. J. Lin et al., "A survey on Internet of Things: Architecture, enabling technologies, security and privacy, and applications," IEEE Internet Things J., vol. 4, no. 5, pp. 1125-1142, 2017.
28. M. A. Khan and K. Salah, "IoT security: Review, blockchain solutions, and open challenges," Future Gener. Comput. Syst., vol. 82, pp. 395-411, 2018.
29. J. Shen et al., "A lightweight multi-keyword searchable encryption scheme for secure cloud computing," IEEE Access, vol. 7, pp. 10254-10263, 2019.
30. H. Li et al., "Efficient and secure searchable encryption with multi-keyword ranked search for cloud storage," IEEE Trans. Dependable Secure Comput., vol. 17, no. 6, pp. 1164-1177, 2020.
31. J. Liu et al., "A survey on searchable encryption in cloud computing," Comput. Sci. Rev., vol. 40, 100389, 2021.
32. C. Wang et al., "Secure ranked keyword search over encrypted cloud data," in Proc. IEEE ICDCS, pp. 1-10, 2012.
33. Z. Xia et al., "A secure and dynamic multi-keyword ranked search scheme over encrypted cloud data," IEEE Trans. Parallel Distrib. Syst., vol. 27, no. 2, pp. 340-352, 2016.
34. X. Zhang et al., "Lattice-based multi-keyword searchable encryption with fine-grained access control," in Proc. IEEE INFOCOM, pp. 1-9, 2017.
35. R. Lu et al., "Lattice-based searchable encryption with fine-grained access control for cloud storage," IEEE Trans. Services Comput., vol. 12, no. 4, pp. 562-575, 2019.
36. Q. Li et al., "Federated learning for healthcare analytics: A comprehensive survey," ACM Comput. Surv., vol. 56, no. 7, pp. 1-42, 2024.
37. Q. Yang et al., "Federated machine learning: concept and applications," ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, pp. 1-19, 2019.
38. Y. Wang et al., "Transfer learning for neural strategy selection in IoT security," IEEE Internet Things J., vol. 10, no. 12, pp. 10870-10882, 2023.
39. Z. Huang et al., "Adaptive strategy selection for encrypted search using meta-learning," IEEE Trans. Dependable Secure Comput., vol. 21, no. 3, pp. 1556-1570, 2024.
40. R. Guo et al., "Blockchain-based privacy-preserving medical data sharing with searchable encryption," IEEE Access, vol. 8, pp. 96114-96127, 2020.

41. X. Liu et al., "A blockchain-based secure and efficient data sharing framework for healthcare," IEEE Trans. Netw. Sci. Eng., vol. 8, no. 3, pp. 2088-2101, 2021.
42. D. X. Song, D. Wagner, and A. Perrig, "Practical techniques for searches on encrypted data," in Proc. IEEE Symp. Security and Privacy, pp. 44-55, 2000.
43. D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano, "Public key encryption with keyword search," in Proc. EUROCRYPT, pp. 506-522, 2004.
44. H. Wang et al., "Anonymous multi-keyword searchable encryption with hidden access policy for blockchain-enabled healthcare," IEEE Trans. Services Comput., vol. 19, no. 2, pp. 887-901, 2026.
45. Y. Chen et al., "ES-VKSE: Enhanced security verifiable keyword searchable encryption for blockchain-based smart healthcare," J. Inf. Security Appl., vol. 78, 103612, 2025.
46. S. Luo et al., "SENSH: A blockchain-based searchable encrypted data sharing scheme in smart healthcare," PeerJ Comput. Sci., vol. 11, e3166, 2025.
47. A. Azaria, A. Ekblaw, T. Vieira, and A. Lippman, "MedRec: Using blockchain for medical data access and permission management," in Proc. 2nd Int. Conf. Open and Big Data, pp. 25-30, 2016.
48. R. Guo et al., "Blockchain-based privacy-preserving medical data sharing with searchable encryption," IEEE Access, vol. 8, pp. 96114-96127, 2020.
49. J. Xu et al., "A survey on blockchain-based healthcare applications: research challenges and opportunities," J. Netw. Comput. Appl., vol. 145, 102417, 2019.
50. D. Yaga et al., "Blockchain technology overview," NIST IR 8202, 2019.
51. "SHAS: A novel search approach for secure and flexible sharing of hierarchical EHR based on blockchain," Cluster Comput., 2025.
52. J. Zhang et al., "A blockchain-based medical data sharing scheme with privacy preservation," IEEE Access, vol. 8, pp. 147214-147226, 2020.
53. H. Wang and Y. Song, "Secure cloud-based EHR system using attribute-based cryptosystem and blockchain," J. Med. Syst., vol. 42, no. 8, 152, 2018.
54. H. Guo et al., "A blockchain-based access control framework for electronic health records," IEEE Access, vol. 7, pp. 170409-170421, 2019.
55. Y. Zhang et al., "Smart contract-based access control for EHRs on blockchain," J. Med. Syst., vol. 46, no. 4, 27, 2022.
56. J. Zhao et al., "A survey on blockchain-based access control for healthcare data," Comput. Netw., vol. 222, 109537, 2023.
57. Z. Zheng et al., "An overview of blockchain technology: architecture, consensus, and future trends," in Proc. IEEE BigData Congress, pp. 557-564, 2017.
58. M. Andoni et al., "Blockchain technology in the energy sector: A systematic review of challenges and opportunities," Renew. Sustain. Energy Rev., vol. 100, pp. 143-174, 2019.
59. P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," SIAM J. Comput., vol. 26, no. 5, pp. 1484-1509, 1997.
60. D. J. Bernstein and T. Lange, "Post-quantum cryptography," Nature, vol. 549, no. 7671, pp. 188-194, 2017.
61. "AAQ-PEKS: An Attribute-based Anti-Quantum Public Key Encryption Scheme with Keyword Search for E-healthcare Scenarios," 2025.
62. "Provably lightweight and secure IoHT scheme with post-quantum cryptography and fog computing," ScienceDirect, 2025.
63. T. Lange et al., "Post-quantum cryptography: a survey," IEEE Security & Privacy, vol. 18, no. 3, pp. 50-59, 2020.
64. N. Bindel et al., "Transitioning to a quantum-resistant public key infrastructure," in Post-Quantum Cryptography, pp. 307-327, 2019.

65. M. Chen et al., "Deep reinforcement learning for resource management in edge computing," IEEE Trans. Commun., vol. 68, no. 9, pp. 5750-5763, 2020.
66. H. Lu et al., "DRL-based offloading for IoT tasks in edge computing," IEEE Internet Things J., vol. 8, no. 15, pp. 12006-12018, 2021.
67. Z. Xu et al., "Deep reinforcement learning for adaptive routing in software-defined networks," IEEE Trans. Netw. Service Manag., vol. 16, no. 4, pp. 1562-1575, 2019.
68. "State-Aware IoT Scheduling Using Deep Q-Networks and Edge-Based Coordination," IEEE, 2025.
69. A. Ali et al., "An industrial IoT-based blockchain-enabled secure searchable encryption approach for healthcare systems using neural network," Sensors, vol. 22, no. 2, 572, 2022.
70. R. Kumar et al., "Blockchain-enabled AI-driven secure searchable encryption framework for medical IoT systems," IEEE Trans. Inf. Forensics Security, vol. 18, pp. 3024-3038, 2023.
71. "Improving Energy Efficiency in IoT Edge Computing Through Dynamic Deep Q-Network Based Resource Management," Wiley, 2025.
72. H. Zhu et al., "Adaptive optimization for searchable encryption using reinforcement learning," IEEE Trans. Services Comput., vol. 14, no. 6, pp. 1850-1863, 2021.
73. L. Zhang et al., "Lightweight neural networks for IoT: a survey," IEEE Commun. Surveys Tuts., vol. 22, no. 3, pp. 2066-2096, 2020.
74. F. Casino et al., "A systematic literature review of blockchain-based applications: current status, classification and open issues," Telematics Inform., vol. 36, pp. 55-81, 2019.
75. Y. Wang et al., "Predicate-private asymmetric searchable encryption for conjunctions from lattices," IEEE Trans. Dependable Secure Comput., vol. 18, no. 4, pp. 1778-1791, 2021.
76. J. Li et al., "Lattice-based multi-keyword searchable encryption with fine-grained access control," IEEE Trans. Inf. Forensics Security, vol. 15, pp. 1356-1369, 2020.
77. J. Ning et al., "A survey on lattice-based searchable encryption," J. Cryptologic Res., vol. 7, no. 2, pp. 241-256, 2020.
78. C. Paquin et al., "Performance of post-quantum cryptography on embedded devices," IEEE Embedded Syst. Lett., vol. 14, no. 4, pp. 155-158, 2022.
79. F. Bonomi et al., "Fog computing and its role in the Internet of Things," in Proc. MCC Workshop Mobile Cloud Computing, pp. 13-16, 2012.
80. L. Ducas et al., "CRYSTALS-Dilithium: A lattice-based digital signature scheme," NIST PQC Standardization Round 3, 2022.
81. Y. Liu et al., "Quantum-resistant blockchain for healthcare IoT: A comprehensive survey," ACM Comput. Surv., vol. 58, no. 4, pp. 1-36, 2026.
82. "Blockchain-based federated learning with homomorphic encryption for privacy-preserving healthcare data sharing," ScienceDirect, 2025.
83. O. Goldreich, Foundations of Cryptography: Basic Applications. Cambridge Univ. Press, 2009.
84. L. Zhang et al., "Secure federated learning for blockchain-based healthcare systems," IEEE J. Biomed. Health Inform., vol. 29, no. 3, pp. 1456-1468, 2025.
85. J. Wang et al., "Blockchain-enabled federated learning for privacy-preserving healthcare data analytics," IEEE Trans. Netw. Sci. Eng., vol. 12, no. 2, pp. 789-803, 2025.
86. R. Gupta et al., "Securing traditional medicine knowledge with blockchain and AI," J. Med. Syst., vol. 49, no. 5, pp. 112-125, 2025.
87. "Anonymous Attribute-Based Multi-keyword Searchable Encryption Scheme for Medical Data Sharing Using Blockchain," Springer, 2026.
88. J. Wang et al., "On/off-chain split for efficient searchable encryption in healthcare blockchain," IEEE Access, vol. 8, pp. 213456-213469, 2020.

89. Q. Liu et al., "A cache-aware searchable encryption scheme for cloud storage," IEEE Trans. Cloud Comput., vol. 6, no. 4, pp. 1023-1034, 2018.
90. Y. Huang et al., "Batch processing for efficient searchable encryption in blockchain," in Proc. IEEE Int. Conf. Blockchain, pp. 1-8, 2021.
91. A. E. W. Johnson et al., "MIMIC-III, a freely accessible critical care database," Sci. Data, vol. 3, 160035, 2016.
92. "CICIoMT2024 Dataset," Canadian Institute for Cybersecurity, 2024.
93. M. Tavallaee et al., "A detailed analysis of the KDD CUP 99 data set," in Proc. IEEE CISDA, pp. 1-6, 2009.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

