



Two Decades of Internet of Things Scholarship: A Thematic Synthesis of Research Trends, Gaps, and Publication Patterns

Eliza B. Ayo¹, Alvin D. Catalo², Hydie Cruz³, Reggie Mar B. de Castro⁴, Shiela Marie G. Garcia⁵, June S. Rabi⁶
^{1,2,3,4,5,6}La Consolacion Philippines, Malolos, Bulacan, Philippines

ABSTRACT: The Internet of Things (IoT) has evolved over nearly two decades from a niche concept centered on radio-frequency identification into a pervasive technological substrate spanning healthcare, manufacturing, transportation, energy, and consumer environments. Despite accelerating publication volume, the field's literature has not been systematically synthesized across this two-decade horizon to identify thematic evolution, gap patterns, and structural features of publishing. This study addresses that need through a quantitative thematic synthesis of 235 IoT studies published between 2009 and 2026. Using a controlled vocabulary applied to study keywords, key findings, and identified gaps, the corpus was coded across 18 themes, 10 finding-type categories, and 12 gap categories, with cross-tabulations against publication era, indexing platform, and authorship configuration. Three findings stand out. First, architecture and framework proposals (45.5%) and security and privacy (35.7%) dominate IoT scholarship, while artificial intelligence and machine learning integration has risen from absent before 2018 to 29% of 2024–2026 studies—the steepest trend in the corpus. Second, a structural tension exists between gap recognition and gap closure: 33.6% of studies flag the absence of real-world empirical validation as a gap, but only 7.7% deliver such validation. Third, the corpus is dispersed across 169 outlets with no dominant venue, and indexing tier shows only weak association with the depth of reported findings. The study concludes that the most defensible openings for future IoT research lie in long-horizon empirical validation, vendor-neutral interoperability, and governance-oriented studies.

KEYWORDS: Internet of Things, systematic review, thematic synthesis, research gaps, bibliometric analysis, edge computing, artificial intelligence integration.

INTRODUCTION

The Internet of Things (IoT) refers to the networked interconnection of physical objects equipped with sensing, computing, and communication capabilities (Atzori et al., 2010; Gubbi et al., 2013). Since its early conceptualization in the late 2000s, the IoT has grown from a research curiosity built around radio-frequency identification (RFID) into a pervasive technological layer supporting smart cities, connected healthcare, industrial automation, precision agriculture, and consumer devices. Estimates place the number of connected IoT devices in the tens of billions by 2025–2026, and the field's scholarly output has scaled accordingly: published research on the IoT has grown by roughly an order of magnitude over the past decade.

This rapid growth has produced both a strength and a problem. The strength is evidence breadth: substantial published work now exists on nearly every layer of IoT systems, from sensor hardware and communication protocols through edge and cloud architectures to applications and societal effects. The problem is fragmentation. The literature spans hundreds of outlets, dozens of overlapping subdomains, and multiple methodological traditions, making it difficult for new entrants, practitioners, and funders to understand the shape of the field as a whole. Existing systematic reviews of IoT scholarship tend to focus on

individual subdomains, security, edge computing, or healthcare IoT, and few have attempted a cross-cutting synthesis covering the field's full two-decade arc.

The present study offers such a synthesis. It examines 235 IoT studies published between 2009 and 2026 with the goal of characterizing the field's thematic structure, its evolution over time, its dominant types of findings and recurring gaps, its publication patterns, and the structural features of authorship and collaboration. The study is descriptive and exploratory rather than inferential: its purpose is to map, not to prove. The map matters because the IoT is now at a transitional moment. The recent absorption of artificial intelligence (AI) into IoT systems, the tightening of regulatory regimes around data and connected devices, and the increasing maturity of edge–cloud–fog architectures collectively mean that the literature of the next five years will look quite different from that of the previous fifteen. Understanding where the existing literature has been strong, where it has been thin, and where it has consistently failed to deliver on its own stated gaps is a prerequisite for shaping the work that follows.

1.1 Research Questions

This study addresses ten research questions:

RQ1. What are the research themes in IoT scholarship based on the selected studies?

- RQ2.** How have research themes and focus areas in IoT scholarship evolved over time?
- RQ3.** What are the most common types of key findings reported in IoT studies?
- RQ4.** What are the most frequently reported research gaps in the IoT literature?
- RQ5.** How can the identified research gaps be categorized?
- RQ6.** What future research directions can be proposed based on the identified gaps and trends?
- RQ7.** To what extent do existing studies address previously identified gaps in IoT research?
- RQ8.** What are the publication patterns in terms of journals, indexing, and author contributions in IoT research?
- RQ9.** Is there a relationship between journal indexing and the type or depth of research findings?
- RQ10.** What patterns exist in authorship and collaboration within IoT research?

1.2 Contribution

The contribution of the study is threefold. First, it provides a cross-domain thematic synthesis of two decades of IoT scholarship using a transparent controlled-vocabulary coding scheme, allowing the analysis to be replicated and extended. Second, it introduces a gap-coverage matrix that directly compares the frequency with which gaps are cited against the frequency with which they are actually closed in the same corpus, surfacing a structural tension in the field between proposal and proof. Third, it offers an evidence-grounded set of future research directions that intersect frequently cited gaps with rising thematic trends, allowing scholars and funders to prioritise work that is both needed and aligned with where the field is moving.

2. LITERATURE REVIEW

The IoT literature has matured in three overlapping phases, each leaving an identifiable imprint on the present corpus. Understanding these phases is necessary to contextualize the distributions of themes, findings, and gaps reported in Section 4.

2.1 Foundational Vision and Architecture

Early IoT scholarship was concerned primarily with definition and architectural vision. Atzori et al. (2010) outlined the IoT as a paradigm in which uniquely addressable objects interact through standardised communication protocols, and Gubbi et al. (2013) extended this view by introducing a cloud-centric architectural vision in which IoT data are processed at scale through internet-based services. These early contributions set a research agenda dominated by reference architectures, communication protocols, and integration with enterprise computing infrastructures. Their thematic legacy is visible in the present corpus, in which architecture and framework studies remain the single most common category.

2.2 Security, Privacy, and Communication Maturity

As deployments scaled, the literature shifted toward security and privacy. The constrained nature of IoT endpoints—limited memory, intermittent connectivity, and heterogeneous

protocols—rendered conventional cryptographic and access-control approaches partially inadequate, prompting a sustained body of work on lightweight cryptography, authentication for constrained devices, and privacy-preserving data flows. Parallel to this, communication-layer research consolidated around long-range and low-power protocols like LoRa, NB-IoT, 5G, with later studies extending into 6G horizons. The relative share of dedicated communications work has since contracted as these protocols have been absorbed into broader architectural discussions.

2.3 Convergence with Edge Computing and Artificial Intelligence

More recent literature reflects two convergences. First, edge, fog, and cloud computing have moved from competing paradigms to complementary tiers in hybrid IoT architectures, motivated by latency, bandwidth, and data-sovereignty constraints. Second, AI and machine learning (ML) have moved from peripheral application to central design constraint, with federated learning, on-device inference, and AI-assisted orchestration emerging as dominant subthemes. These convergences both expand IoT capability and create new gaps particularly in privacy-preserving analytics, energy-aware inference, and reliability of AI components on resource-constrained devices.

2.4 Reviews of the Literature and the Need for Cross-Cutting Synthesis

Numerous prior reviews have synthesised slices of the IoT literature security-focused surveys, edge-computing reviews, application-domain syntheses in healthcare or industry, and bibliometric analyses of citation patterns. What remains comparatively scarce is a cross-cutting synthesis that traces thematic evolution, gap patterns, and publication structure across the field's full two-decade horizon. The present study addresses this gap by applying a single coding frame to 235 studies spanning 2009–2026 and by introducing a gap-coverage matrix to surface tensions between what the field claims to lack and what it actually delivers.




Fig.1. Conceptual Framework

3. METHODOLOGY

3.1 Research Design

The study employs a systematic literature review with quantitative thematic synthesis. It is descriptive, cross-sectional,

and retrospective: the unit of analysis is the individual published study, and the goal is second-order characterisation of how IoT scholarship has been conducted, rather than testing hypotheses about IoT phenomena themselves. The design draws on three established traditions: PRISMA-style systematic review for corpus assembly (Page et al., 2021), thematic content analysis with a predefined coding frame (Braun & Clarke, 2006), and descriptive bibliometrics for publication, indexing, and authorship patterns. The procedure also follows the conventions for evidence-based systematic reviews articulated by Kitchenham and Charters (2007).

3.2 Data Source and Corpus

The corpus consists of 235 studies spanning 2009–2026, compiled in a single structured spreadsheet containing thirteen fields per study (title, journal name, journal type, indexing platform, DOI, paper keywords, key findings, identified gaps, publication date, author affiliations, and source URL). Studies were retrieved through keyword searches across academic databases including Scopus, Web of Science, IEEE Xplore,

ScienceDirect, SpringerLink, Google Scholar, MDPI, and DOAJ, and were screened for relevance prior to extraction. All 235 records were retained for analysis, as each contained the minimum required fields.

3.3 Coding Scheme

A controlled vocabulary was developed across three coding dimensions: 18 theme buckets, 10 finding-type buckets, and 12 gap categories. Each bucket was operationalised as a set of lexical search terms applied to the relevant free-text fields (keywords, title, key findings, and gaps). Coding was non-mutually-exclusive a study could be assigned to multiple themes, finding types, or gap categories because IoT research is typically multifaceted. This approach is consistent with thematic content analysis in the tradition described by Braun and Clarke (2006), adapted for computational application.

3.4 Statistical Treatment per Research Question

Table 1 summarizes the analytical treatment applied to produce the result for each of the ten research questions

Table 1: Analytical Treatment per Research Question

RQ	Question	Analytical treatment
RQ1	Research themes	Frequency analysis using controlled-vocabulary keyword matching across keywords, title, and key-findings fields. Raw counts and percentages over the 235-study corpus, ranked in descending order. Non-mutually-exclusive coding.
RQ2	Evolution of themes	Temporal binning into five era brackets (≤ 2017 , 2018–2019, 2020–2021, 2022–2023, 2024–2026) following regex-based year extraction. Cross-tabulation of theme by era, expressed as raw counts and era-relative percentages to control for varying publication volume.
RQ3	Finding types	Frequency analysis with controlled-vocabulary pattern matching against the Key Findings field. Ten finding-type categories ranked by frequency and percentage.
RQ4	Reported gaps	Frequency analysis with controlled-vocabulary pattern matching against the Gaps field. Twelve gap categories ranked by frequency.
RQ5	Gap categorization	Higher-order categorical aggregation: the twelve atomic gap categories from RQ4 were grouped into five super categories (technical, theoretical, experimental, scalability/architectural, governance/sociotechnical) by conceptual similarity. Aggregated mention counts reported.
RQ6	Future directions	Cross-referential synthesis: intersection of gap frequencies (RQ4–RQ5) with temporal trends (RQ2). Directions derived where a gap was both high-frequency and aligned with a rising thematic trend. Qualitative interpretation grounded in quantitative ranks.
RQ7	Gap-coverage matrix	Direct ratio comparison between gap frequency (RQ4) and matched finding frequency (RQ3). Coverage status assigned by inspection of the gap-to-finding ratio as “largely addressed,” “partially addressed,” or “largely unaddressed.”
RQ8	Publication patterns	Three separate frequency analyses: (a) outlet ranking from journal-name counts, (b) categorical classification of journal type, and (c) multi-count string matching for indexing platforms. Descriptive statistics only.
RQ9	Indexing vs. depth	Group comparison between high-tier-indexed studies (Scopus and/or Web of Science) and others. Proxies for depth: mean character length of Key Findings, proportion proposing frameworks, and proportion reporting empirical validation. Differences reported as proportions and means.

RQ	Question	Analytical treatment
RQ10	Authorship and collaboration	Heuristic author count from heterogeneous text formats using multi-separator parsing. Bucketed distribution (single, 2, 3, 4–5, 6+); central tendency reported as mean and median. Country detection via gazetteer-based string matching against author-affiliation text.

3.5 Rationale for the Statistical Treatment

The analysis is descriptive and exploratory rather than inferential. No null-hypothesis significance testing was conducted. The rationale is methodological fit: the corpus is the entire selected population (N = 235) rather than a probability sample drawn from a defined universe, and the coding categories are non-mutually-exclusive, which violates the assumptions of standard chi-square tests of independence. Frequency analysis with descriptive statistics is the appropriate treatment for this design and is consistent with comparable systematic reviews and bibliometric studies in the IoT and related fields.

Within this framework, several specific treatments warrant comment. For RQ2, both raw counts and era-relative percentages are reported because the eras differ markedly in size (24 studies in ≤2017 versus 89 in 2024–2026); reporting only raw counts would conflate trend with publication-volume growth. For RQ7, the gap-coverage matrix is the central analytical contribution of the review: it operationalises the question “is the field closing its own gaps?” by directly comparing two frequency distributions (gaps cited and findings delivered) from the same corpus. For RQ9, mean character length of the Key Findings field is used as a rough proxy for descriptive depth. This proxy is acknowledged to be imperfect but defensible, given the absence of an objective rubric for “depth” applicable to free-text findings. For RQ10, the author-count heuristic is conservative and may under-count true team size in entries where authors are separated by formatting that the parser cannot detect; this is reported as a known limitation.

3.6 Methodological Limitations

Several limitations should be acknowledged. First, the controlled vocabulary is bounded—themes, finding types, and

gaps that do not match any term in the lexicon are not captured. Second, the corpus is restricted to studies already screened and entered in the source dataset; selection bias from that earlier screening propagates through the analysis. Third, country detection underestimates international collaboration because author-affiliation text is heterogeneously formatted. Fourth, year normalization from mixed-format date fields could not extract a usable year for a small number of entries with non-standard formats example “fourth quarter,” and conference-date strings, these were handled case-by-case. Three additions would strengthen the methodology if extended in future work: (a) inter-coder reliability with a second coder applying the same vocabulary to a 10–20% sample of studies, reported via Cohen's kappa; (b) co-occurrence network analysis to substantiate claims about themes that travel together; and (c) citation-network triangulation against a larger, automatically extracted corpus.

4. RESULTS

4.1 Research Themes (RQ1)

Eighteen themes were identified across the corpus (Table 2). Architecture and framework design appears in 45.5% of studies, followed by security and privacy (35.7%) and edge, fog, and cloud computing (32.3%). A second tier of themes—network protocols, AI/ML integration, industrial IoT, and data management ranges from 20% to 28% of the corpus. A third tier (energy, scalability, interoperability, block chain, healthcare, and smart cities) appears in 6%–12% of studies. Application-specific themes such as agriculture, vehicular IoT, environmental monitoring, and wearables remain under-represented at less than 3% each.

Table 2: Research Themes Ranked by Frequency

Rank	Theme	Studies (n)	% of corpus
1	Architecture and frameworks	107	45.5
2	Security and privacy	84	35.7
3	Edge / fog / cloud computing	76	32.3
4	Network protocols and communication	65	27.7
5	Machine learning and AI integration	54	23.0
6	Industrial IoT and Industry 4.0	53	22.6
7	Data management and big data analytics	48	20.4
8	Energy efficiency and green IoT	28	11.9
9	Scalability and resource management	27	11.5

Rank	Theme	Studies (n)	% of corpus
10	Interoperability and standardisation	21	8.9
11	Healthcare and medical IoT	18	7.7
12	Blockchain integration	16	6.8
13	Smart cities and smart homes	15	6.4
14	Supply chain and logistics	9	3.8
15	Transportation and vehicular IoT	6	2.6
16	Agriculture and smart farming	5	2.1
17	Environmental monitoring	2	0.9
18	Wearables and consumer IoT	2	0.9

Note. N = 235. Coding is non-mutually-exclusive; studies may appear in multiple themes.




Fig.2. Research Themes Ranked by Frequency

4.2 Evolution of Themes Over Time (RQ2)

The corpus spans 2009–2026 and exhibits marked acceleration after 2019: 89 studies (38% of the corpus) appear in the 2024–2026 era alone. The sustained growth of security and privacy research—more than doubling in relative prevalence from 21% to 42% between the earliest era and 2020–2021—reflects the escalating threat landscape documented by Sicari et al. (2015) and Frustaci et al. (2018), who catalogued systemic vulnerabilities that constrain large-scale IoT deployment. Architecture and framework design rose steadily to 56% of 2024–2026 studies, corroborating Al-Fuqaha et al.’s (2015) projection that protocol and architectural standardisation would become a primary research focus as device heterogeneity increased; Ray (2018) similarly identified architectural diversity as the defining challenge of IoT maturation. AI/ML integration shows the most dramatic trajectory—absent before 2018, it

accounted for 29% of studies in 2024–2026, the steepest trend observed in the corpus—a trajectory foreshadowed by Mahdavinejad et al. (2018), who demonstrated the potential of machine-learning pipelines to transform IoT-generated data into actionable intelligence. Edge and fog computing maintained a steady presence (25%–45% across eras), consistent with the paradigm articulated by Shi et al. (2016) and the earlier fog-computing model of Bonomi et al. (2012), both of which positioned latency and bandwidth constraints as structural drivers for distributed computation. Communications-layer work (5G/6G, MQTT, LoRa) declined relatively from 29%–31% of pre-2020 studies to 12% in the most recent era; Chettri and Bera (2020) contextualise this shift, noting that 5G connectivity is increasingly treated as an infrastructure assumption rather than an independent research problem. Table 3 reports the era-by-era detail.

Table 3: Theme Prevalence by Era (Count and Era-Relative Percentage)

Theme	≤2017	2018–2019	2020–2021	2022–2023	2024–2026
Security and privacy	5 (21%)	6 (17%)	20 (42%)	18 (47%)	35 (39%)
Architecture and frameworks	7 (29%)	11 (31%)	19 (40%)	19 (50%)	50 (56%)
Edge / fog / cloud	8 (33%)	9 (25%)	14 (29%)	17 (45%)	28 (31%)

Theme	≤2017	2018–2019	2020–2021	2022–2023	2024–2026
AI / ML integration	0 (0%)	6 (17%)	6 (12%)	6 (16%)	26 (29%)
Industrial IoT and Industry 4.0	2 (8%)	10 (28%)	9 (19%)	9 (24%)	22 (25%)
5G / 6G and communication	7 (29%)	11 (31%)	13 (27%)	7 (18%)	11 (12%)
Energy / green IoT	2 (8%)	3 (8%)	3 (6%)	4 (11%)	14 (16%)
Blockchain	0 (0%)	1 (3%)	4 (8%)	4 (11%)	7 (8%)
Healthcare IoT	0 (0%)	5 (14%)	0 (0%)	4 (11%)	8 (9%)
Smart cities	0 (0%)	4 (11%)	2 (4%)	3 (8%)	3 (3%)
Total studies in era	24	36	48	38	89

Note. Percentages are era-relative, calculated as the proportion of studies in each era that exhibit the theme. AI = artificial intelligence; ML = machine learning.




Fig.3. Theme Prevalence by Era

4.3 Common Finding Types (RQ3)

Ten finding-type categories were identified (Table 4). Performance and efficiency improvements appear in 64.3% of studies and integration with other technologies in 60.4%; the two finding types co-occur frequently, suggesting that the dominant style of contribution is to combine the IoT with another technology and report efficiency gains—a pattern consistent with the integrative research agenda mapped by Atzori et al. (2010) and the cross-technology enabling-technology survey of Al-Fuqaha et al. (2015). Proposed frameworks and architectures appear in 44.3% of studies, reflecting the sustained demand for reference models

documented by Ray (2018). Security and privacy enhancements, ranked fourth at 27.7%, constitute the applied counterpart to the theoretical threat analyses of Frustaci et al. (2018) and Sicari et al. (2015). Conspicuously rare are explicit empirical validation (7.7%) and formal comparative benchmarking (8.5%): conceptual proposal volume is outpacing rigorous evaluation, a concern raised by Whitmore et al. (2015), who noted that the IoT literature was accumulating architectural propositions faster than it was subjecting them to reproducible field tests. Adoption, investment, or business-impact findings (3.4%) are similarly thin, indicating that the literature continues to lean technical rather than managerial or economic

Table 4: Key Finding Types Ranked by Frequency

Rank	Finding type	Studies (n)	% of corpus
1	Performance / efficiency improvements	151	64.3
2	Integration with other technologies	142	60.4
3	Proposed framework, model, or architecture	104	44.3
4	Security / privacy enhancements	65	27.7
5	Application demonstration or case study	56	23.8
6	Challenges or problems identified	33	14.0

Rank	Finding type	Studies (n)	% of corpus
7	Comparative analysis	20	8.5
8	Empirical / experimental validation	18	7.7
9	Literature review / bibliometric summary	12	5.1
10	Adoption, investment, or business insights	8	3.4

Note. N = 235. Coding is non-mutually-exclusive.




Fig. 4. Key Finding Types Ranked by Frequency

4.4 Frequently Reported Gaps (RQ4)

Twelve gap categories were identified (Table 5). Five gaps stand out as systemic: security and privacy (37.0%), real-world deployment and empirical validation (33.6%), technical implementation (32.3%), interoperability and standardisation (31.1%), and scalability (29.8%). The persistent prominence of security and privacy as the leading gap is consistent with the road-ahead analysis of Sicari et al. (2015) and the critical evaluation of Frustaci et al. (2018), both of which framed unresolved security concerns as the principal barrier to mainstream IoT adoption. The interoperability gap, flagged by 31.1% of studies, mirrors the taxonomic analysis of Noura et al. (2019), who documented that the absence of vendor-

neutral reference profiles remained the single largest obstacle to cross-platform IoT deployment at scale. Scalability concerns (29.8%) align with the edge-computing vision of Shi et al. (2016), who argued that centralised architectures are structurally unable to accommodate the traffic volumes of trillion-device IoT ecosystems. Each of the top five gaps is flagged by roughly one in three studies, and the categories overlap heavily within individual papers. Governance and human-factor gaps, though appearing in only 6.0% and 0.9% of studies respectively, are likely under-flagged given their real-world importance; Weber (2010) identified regulatory and ethical gaps as foundational challenges for IoT governance as early as 2010, suggesting these issues are not new but remain systematically under-researched

Table 5: Most Frequently Reported Research Gaps

Rank	Gap category	Studies (n)	% of corpus
1	Security and privacy	87	37.0
2	Real-world deployment / empirical validation	79	33.6
3	Technical / implementation	76	32.3
4	Interoperability and standardization	73	31.1
5	Scalability	70	29.8
6	Theoretical / conceptual	46	19.6
7	Energy / resource constraints	35	14.9
8	Regulatory / ethical / governance	14	6.0
9	Long-term sustainability / longitudinal	8	3.4

Rank	Gap category	Studies (n)	% of corpus
10	Data quality and management	8	3.4
11	Cross-domain / multi-sector generalisation	2	0.9
12	User / human factors	2	0.9

Note. N = 235. Coding is non-mutually-exclusive.




Fig.5. Most Frequently Reported Research Gaps

4.5 Higher-Order Gap Categorisation (RQ5)

The twelve atomic gap categories aggregate into five higher-order categories (Table 6). The technical category dominates (approximately 206 mentions), encompassing security, implementation, energy, and data-quality concerns; this clustering aligns with the critical security evaluation by Frustaci et al. (2018), who demonstrated that implementation-level vulnerabilities, energy-constrained attack surfaces, and poor data governance share common architectural roots. Scalability and architectural gaps form the next largest cluster (approximately 145 mentions), corroborating the structural analysis of Shi et al. (2016) and Noura et al.'s (2019) finding that interoperability and cross-domain transfer problems are facets

of a single architectural fragmentation problem. Experimental and empirical gaps form a third cluster of approximately 87 mentions; these are not gaps in what to build but in how rigorously the field tests what it has built, a deficit highlighted by Whitmore et al. (2015) in their systematic mapping of IoT research trajectories. Theoretical and conceptual gaps (approximately 46 mentions) are smaller, and governance and sociotechnical gaps (approximately 24 mentions) are the smallest. The latter, however, is likely an under-representation rather than evidence of resolution, as Weber (2010) argued that regulatory frameworks for IoT lag well behind technical deployment realities.

Table 6: Higher-Order Classification of Research Gaps

Higher-order category	Underlying gap clusters	Approx. mentions
Technical	Security and privacy; implementation; energy; data quality	~206
Scalability and architectural	Scalability; interoperability; cross-domain transfer	~145
Experimental / empirical	Real-world deployment; longitudinal; benchmarking	~87
Theoretical / conceptual	Theoretical frameworks; reference models	~46
Governance / sociotechnical	Regulatory and ethical; user and human factors; adoption and ROI	~24

Note. ROI = return on investment. Mentions are non-mutually-exclusive.




Fig.6. Higher-Order Classification of Research Gaps

4.6 Future Research Directions (RQ6)

Future directions were derived by intersecting gap frequencies with rising thematic trends. Table 7 lists ten directions ranked by their combined evidentiary support. The highest-priority direction—lightweight, post-quantum, and zero-trust security for resource-constrained devices—emerges directly from the top-ranked gap (37.0%) and the critical vulnerability analysis of Frustaci et al. (2018), who specifically recommended trust-model innovations for low-power endpoints. The second direction, large-scale empirical validation, responds to the most under-addressed gap in the corpus (33.6% gap, 7.7% coverage) and echoes the longitudinal testbed call of Whitmore et al. (2015). Interoperability via vendor-neutral profiles (Direction 3) is grounded in the open-challenge taxonomy of Noura et al. (2019). AI-driven, edge-native IoT (Direction 4) synthesizes the

accelerating AI/ML trend (29% of 2024–2026 studies) with the edge-computing paradigm of Shi et al. (2016); the recommendation for federated learning draws on the communication-efficient distributed learning framework of McMahan et al. (2017) and the IoT-ML survey of Mahdavinejad et al. (2018). The hybrid block chainIoT direction (Direction 10) is supported by Reyna et al. (2018), who identified trust, provenance, and decentralised identity management as the most tractable near-term applications of block chain for IoT ecosystems. Domain-grounded deployments in under-represented sectors such as agriculture and supply chain (Direction 8) align with the Industry 4.0 agenda mapped by Xu et al. (2018).

Table 7: Proposed Future Research Directions

#	Future research direction
1	Lightweight, post-quantum, zero-trust security for resource-constrained IoT
2	Large-scale, long-horizon empirical validation (real testbeds, multi-year field studies)
3	Interoperability via vendor-neutral reference profiles and shared schemas
4	AI-driven, edge-native IoT (federated learning, on-device inference, AI-assisted orchestration)
5	Scalable orchestration for massive heterogeneous deployments
6	Energy-aware and sustainable IoT (green IoT, energy-harvesting, low-power AI)
7	Privacy-preserving analytics (homomorphic encryption, differential privacy, federated approaches)
8	Domain-grounded deployments in under-represented sectors (agriculture, environment, wearables, supply chain)
9	Governance, ethics, and adoption-economics studies
10	Hybrid block chain IoT for trust and provenance at scale

Note. Directions are ranked by combined evidentiary support derived from gap frequency (RQ4) and rising thematic trends (RQ2).

4.7 Gap-Coverage Matrix (RQ7)

Comparing the frequency of stated gaps (RQ4) against the frequency of matched findings (RQ3) produces a map of which

gaps the field is closing versus restating (Table 8). Three patterns emerge. First, the most-flagged gap—real-world empirical validation, at 33.6%—is also the most under-

addressed: only 7.7% of studies actually deliver such validation, a mismatch that Whitmore et al. (2015) attributed to the incentive structure of IoT research, where architectural novelty is more readily publishable than longitudinal field evaluation. Second, technical and integration gaps are largely addressed, with performance findings (64.3%) and integration findings (60.4%) comfortably exceeding the technical gap rate (32.3%), consistent with the performance-oriented contribution style observed by Al-Fuqaha et al. (2015) and Atzori et al. (2010). Third, governance, longitudinal, cross-domain, and human-

factor gaps are functionally untouched; Weber (2010) and Sicari et al. (2015) both cautioned that these gaps, if left unaddressed, risk producing technically sophisticated but socially undeployable IoT systems. Security and privacy (37.0% gap, 27.7% coverage) remains partially addressed, reflecting ongoing progress documented by Frustaci et al. (2018) but also the difficulty of fully closing a gap that expands with each new device category.

Table 8: Gap-Coverage Matrix

Gap category	Gap frequency	Finding-side coverage	Coverage status
Security and privacy	37.0%	27.7%	Partially addressed
Real-world validation	33.6%	7.7%	Largely unaddressed
Technical / implementation	32.3%	60.4% (integration)	Largely addressed
Interoperability and standardization	31.1%	Few primary contributions	Partially addressed
Scalability	29.8%	8.5% benchmarked	Partially addressed
Theoretical / conceptual	19.6%	44.3% propose frameworks	Partially addressed
Energy / resource constraints	14.9%	11.9% address as theme	Partially addressed
Regulatory / ethical / governance	6.0%	Few studies as primary focus	Largely unaddressed
Long-term sustainability	3.4%	Negligible	Largely unaddressed
Data quality / management	3.4%	20.4% as theme	Partially addressed
Cross-domain transfer	0.9%	Negligible	Largely unaddressed
User / human factors	0.9%	Negligible	Largely unaddressed

Note. Gap frequency from RQ4; finding-side coverage from RQ3. Coverage status assigned by inspection of the gap-to-finding ratio.




Fig.7. Gap-Coverage Matrix

4.8 Publication Patterns (RQ8)

The corpus is highly dispersed: the 235 studies are distributed across 169 distinct journals or conferences (Table 9). Only three outlets contribute more than seven studies each (Table 9a). The dominance of IEEE Access and MDPI-family journals among top outlets reflects the broader open-access shift in engineering publishing; Donthu et al. (2021) observed that high-volume open-access venues disproportionately absorb applied technology literature because their rapid peer-review cycles and absence of subscription barriers lower the cost of early dissemination. Outlet type is dominated by

peer-reviewed journals (50.2%) and open-access peer-reviewed journals (17.9%), with conferences and proceedings at 13.2% (Table 9b). With respect to indexing, 57.9% of studies are indexed in Scopus and 39.6% in Web of Science, with substantial overlap between the two (Table 9c); Donthu et al. (2021) note that dual Scopus–Web of Science indexing is the most commonly applied citation-quality filter in systematic and bibliometric IoT reviews.

Table 9a: Top Publishing Outlets

Outlet	Studies (n)	% of corpus
IEEE Access	13	5.5
Sensors (MDPI)	12	5.1
Future Generation Computer Systems	8	3.4
IEEE Internet of Things Journal	7	3.0
Electronics (MDPI)	5	2.1
Sensors	5	2.1
Future Internet (MDPI)	5	2.1
Applied Sciences (MDPI)	3	1.3
IEEE Communications Surveys & Tutorials	3	1.3
Journal of Network and Computer Applications	3	1.3

Note. Only outlets with three or more studies are shown. The full corpus is dispersed across 169 distinct outlets.




Fig. 8. Top Publishing Outlets

Table 9b: Outlet Type Distribution

Outlet type	Studies (n)	% of corpus
Peer-reviewed journal (general)	118	50.2
Open-access peer-reviewed journal	42	17.9
Conference / proceedings	31	13.2
Other / unspecified type	43	18.3
Book / chapter	1	0.4

Note. N = 235.




Fig.9. Outlet Type Distribution

Table 9c: Indexing Platforms (Multi-Counted)

Index / platform	Studies (n)	% of corpus
Scopus	136	57.9
Web of Science	93	39.6
Google Scholar (implied)	47	20.0
IEEE Xplore	42	17.9
ScienceDirect	14	6.0
SpringerLink	12	5.1
PubMed	10	4.3
DOAJ	9	3.8
ACM Digital Library	7	3.0
INSPEC / EI Compendex / Crossref	≤7 each	≤3.0

Note. Studies may be indexed on multiple platforms; percentages do not sum to 100%. DOAJ = Directory of Open Access Journals.




Fig.10. Indexing Platforms

4.9 Indexing Tier and Depth of Findings (RQ9)

Indexing tier is a weak predictor of the depth or type of findings in this corpus (Table 10). Studies indexed in Scopus or Web of Science report findings in slightly more compact form (mean = 221 characters vs. 286 for non-high-tier outlets), but the underlying composition of contributions is similar. Framework proposals appear in 32% of high-tier-indexed studies versus 39% of others; explicit empirical validation appears in 5% versus 7%, respectively. The differences are too small to support

a meaningful quality gap based on indexing alone, a finding consistent with Donthu et al.'s (2021) observation that indexing tier correlates with citation uptake but not necessarily with methodological rigour or depth of empirical contribution. The IEEE Xplore subset shows slightly higher technical detail and marginally higher empirical-validation rates, reflecting the engineering rigour associated with IEEE peer-review standards (Al-Fuqaha et al., 2015).

Table 10: Indexing Tier and Character of Findings

Indexing group	Studies	Avg. findings length (chars)	% with framework proposals	% with empirical validation
High-tier indexed (Scopus and/or WoS)	139 (59.1%)	221	32	5
Not high-tier indexed	96 (40.9%)	286	39	7
IEEE Xplore subset	42 (17.9%)	273	Higher technical detail	Slightly higher

Note. WoS = Web of Science. Findings length is measured as mean character count of the Key Findings field.




Fig.11. Indexing Tier and Character of Findings

4.10 Authorship and Collaboration Patterns (RQ10)

Authorship is moderately collaborative (Table 11a). Single-author studies remain a noticeable share (21.7%), but the most common configuration is teams of four to five authors (29.8%).

The mean of 3.13 authors per paper and the median of 3 are consistent with engineering and computer-science norms documented in systematic bibliometric analyses (Donthu et al., 2021). Geographic contribution is broad but uneven (Table

11b): the United States leads (12 studies with detectable U.S. affiliations), followed by India (8), with a cluster that includes China, Italy, the United Kingdom, Australia, Pakistan, and the United Arab Emirates. This geographic distribution is broadly consistent with patterns visible in foundational IoT surveys—Atzori et al. (2010) drew heavily on European and Italian-institutional networks, while Al-Fuqaha et al. (2015) represented a North American–Gulf Cooperation Council

collaboration—yet both works drew on a globally diverse citation base, suggesting that IoT research has always been internationally distributed even when formal co-authorship is limited. Cross-border collaboration appears modest at 6.4% of studies, although this is a conservative lower bound given that the gazetteer-based affiliation detection method used here systematically under-counts multi-national team.

Table 11a: Authorship Size Distribution

Authorship bucket	Studies / value	% of corpus
Single author	51	21.7
Two authors	53	22.6
Three authors	40	17.0
Four to five authors	70	29.8
Six or more authors	21	8.9
Mean authors per paper	3.13	—
Median authors per paper	3.0	—

Note. N = 235.




Fig.12. Authorship Size Distribution

Table 11b: Top Geographic Contributors

Country	Studies (detected)
United States	12
India	8
China	5
Italy	5
United Kingdom	4
Australia	3
Pakistan	3
United Arab Emirates	3
Brazil, Indonesia, Ireland, Nigeria, Romania	2 each

Note. Country detection is gazetteer-based and underestimates international collaboration; figures should be read as conservative lower bounds.




Fig.13. *Top Geographic Contributors*

5. DISCUSSION

The synthesis surfaces several findings that together reshape how the IoT literature should be understood. This section discusses six cross-cutting implications.

5.1 The Proposal–Proof Imbalance

The most striking pattern in the corpus is the gap between what the field says it lacks and what it delivers. Real-world empirical validation is identified as a gap in 33.6% of studies but is present as a contribution in only 7.7%. By contrast, framework and architecture proposals are delivered in 44.3% of studies despite being identified as theoretical or conceptual gaps in only 19.6%. The field is producing far more proposals than evaluations of those proposals, and it is doing so with awareness of the problem: studies routinely conclude that real-world deployment validation is needed and then propose another conceptual framework without supplying it. If uncorrected, this pattern leads to an accumulation of unvalidated frameworks and a thinning of the field's empirical base. The most defensible single intervention available to new IoT scholarship is therefore to invert this ratio—to undertake the long-horizon, multi-site, instrumented field studies that the existing literature consistently calls for but rarely delivers.

5.2 The AI–IoT Convergence as Centre of Gravity

Machine learning and AI integration was absent before 2018, accounted for 12%–17% of studies in 2018–2023, and rose to 29% of 2024–2026 studies—the steepest trend in the corpus. This is not a parallel subtopic developing alongside traditional IoT research; it is a centre-of-gravity shift. Federated learning, on-device inference, AI-assisted orchestration, and AI-driven anomaly detection now intersect with security, scalability, energy, and architecture work simultaneously. Studies that fail to engage with the AI–IoT convergence are likely to appear increasingly dated. Conversely, studies that engage with it must also engage with the privacy, energy, and reliability constraints that AI introduces on resource-constrained devices.

5.3 Persistent Gaps as Structural Problems

Security and privacy (37.0%), interoperability and standardisation (31.1%), and scalability (29.8%) are each cited as gaps in roughly one in three studies. These gaps have

remained dominant across all five eras examined, despite substantial investment in security and architecture research. Their persistence suggests structural rather than incremental problems: lightweight cryptography for constrained devices, vendor-neutral interoperability profiles, and scaling pilots into production are problems that do not yield to additional isolated framework proposals. They require coordinated multi-vendor effort, standards-body engagement, and longer evaluation timeframes than typical academic studies provide. The implication is that some of the most important IoT problems may not be solvable through the current dominant research format and may require different institutional vehicles—consortia, multi-year programmes, regulator–academic partnerships—rather than additional individual papers.

5.4 Under-Representation of Governance and Human Factors

Regulatory, ethical, and governance gaps appear in only 6.0% of studies, and user and human factors in only 0.9%. Given the contemporary regulatory environment—including the European Union AI Act, evolving cybersecurity legislation, and sector-specific data regulations in healthcare and finance—and given that IoT systems are ultimately consumed by humans in workplaces and homes, this under-representation is unlikely to reflect resolution of these issues. It more plausibly reflects that the IoT literature is still substantially produced by engineering-discipline researchers whose training, incentives, and review communities favour technical contributions. Closing this gap will require either domain-trained scholars (legal, sociological, organisational) engaging with the IoT, or interdisciplinary teams that include them. Either path is consistent with one of the few clear openings the corpus shows for high-impact differentiation.

5.5 Indexing as a Weak Quality Proxy

The RQ9 finding that high-tier indexing (Scopus, Web of Science) shows only weak association with the depth of reported findings has implications for how readers should evaluate the IoT literature. Indexing is a useful first filter but is not a substitute for explicit screening on methodological criteria—particularly empirical validation, comparative benchmarking, and reproducibility. Reviewers, funders, and

meta-researchers should therefore be wary of treating indexing tier as a quality proxy in this field.

5.6 Implications for Researchers, Funders, and Practitioners

For researchers, the practical implication is to prioritise studies that close gaps the field itself has flagged rather than restate them: long-horizon empirical validation, vendor-neutral interoperability profiles, governance and adoption research, and AI–IoT integration with explicit attention to security, energy, and reliability constraints. For funders, the implication is that some of the most important IoT problems are unlikely to be solved through individual short-cycle papers and may require coordinated, multi-year vehicles. For practitioners, the implication is to read the IoT literature critically—to discount unvalidated frameworks and to assign greater weight to studies that include real-world deployment evidence.

5.7 Limitations and Avenues for Triangulation

The findings rest on a corpus of 235 studies coded with a controlled vocabulary. Although both the corpus size and the coding approach are consistent with comparable systematic reviews, the analysis is descriptive and the categorical assignments are bounded by the lexicon used. Triangulation through a second independent coder, citation-network analysis, or comparison with a larger automatically extracted corpus would strengthen the conclusions. Within these limits, the patterns are sufficiently striking and consistent—particularly the validation gap and the AI/edge trend—to warrant action by the field.

6. CONCLUSION

This study synthesized 235 IoT studies published between 2009 and 2026 across 169 outlets in order to characterize the field's themes, gaps, publication patterns, and authorship structure.

REFERENCES

1. Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A survey on enabling technologies, protocols, and applications. *IEEE Communications Surveys & Tutorials*, 17(4), 2347–2376. <https://doi.org/10.1109/COMST.2015.2444095>
2. Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, 54(15), 2787–2805. <https://doi.org/10.1016/j.comnet.2010.05.010>
3. Bonomi, F., Milito, R., Zhu, J., & Addepalli, S. (2012). Fog computing and its role in the Internet of Things. In *Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing* (pp. 13–16). ACM. <https://doi.org/10.1145/2342509.2342513>
4. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
5. Chettri, L., & Bera, R. (2020). A comprehensive survey on Internet of Things (IoT) toward 5G wireless systems. *IEEE Internet of Things Journal*, 7(1), 16–32. <https://doi.org/10.1109/JIOT.2019.2948888>
6. Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines? *Journal of Business Research*, 133, 285–296. <https://doi.org/10.1016/j.jbusres.2021.04.070>
7. Frustaci, M., Pace, P., Aloï, G., & Fortino, G. (2018). Evaluating critical security issues of the IoT world: Present and future challenges. *IEEE Internet of Things Journal*, 5(4), 2483–2495. <https://doi.org/10.1109/JIOT.2017.2767291>

Three findings deserve emphasis.

First, the IoT literature is dominated by architecture and security proposals, with an unmistakable recent shift toward AI and edge computing. The thematic structure of the field has consolidated around a small number of dominant topics—architecture and frameworks, security and privacy, and edge/fog/cloud computing—while AI and machine learning integration has emerged as the steepest rising trend of the past five years.

Second, the field is structurally over-supplied with conceptual proposals and structurally under-supplied with empirical validation. The gap-coverage matrix introduced in this paper makes this tension explicit: the most-cited gap (real-world empirical validation, 33.6%) is the least closed (delivered in only 7.7% of studies). This is the most defensible single opening for new research, and closing it would address a problem the field has been documenting for over a decade.

Third, indexing tier is a weak proxy for the depth or quality of findings, suggesting that filtering by index alone is inadequate when assessing IoT scholarship. Combined with the highly dispersed publication pattern (169 outlets for 235 studies), this finding has implications for how the field's evidence base should be curated and reviewed going forward. The map provided by this paper is intended to be useful rather than final. As the IoT–AI convergence accelerates and regulatory regimes mature, the thematic mix of the literature will continue to shift, and the gap-coverage matrix introduced here can be re-applied periodically to track whether the field is closing its own self-identified gaps or simply restating them. Until that ratio improves, the central recommendation of this study stands: the next two decades of IoT scholarship will be more valuable to the extent that they deliver the empirical, interoperable, and human-centred evidence that the first two decades have consistently identified as missing.

8. Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645–1660. <https://doi.org/10.1016/j.future.2013.01.010>
9. Kitchenham, B., & Charters, S. (2007). Guidelines for performing systematic literature reviews in software engineering (EBSE Technical Report EBSE-2007-01). Keele University and Durham University Joint Report.
10. Mahdavinnejad, M. S., Rezvan, M., Barekatain, M., Adibi, P., Barnaghi, P., & Sheth, A. P. (2018). Machine learning for Internet of Things data analysis: A survey. *Digital Communications and Networks*, 4(3), 161–175. <https://doi.org/10.1016/j.dcan.2017.10.002>
11. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273–1282). PMLR.
12. Noura, M., Atiquzzaman, M., & Gaedke, M. (2019). Interoperability in Internet of Things: Taxonomies and open challenges. *Mobile Networks and Applications*, 24(3), 796–809. <https://doi.org/10.1007/s11036-018-1089-9>
13. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
14. Ray, P. P. (2018). A survey on Internet of Things architectures. *Journal of King Saud University – Computer and Information Sciences*, 30(3), 291–319. <https://doi.org/10.1016/j.jksuci.2016.10.003>
15. Reyna, A., Martín, C., Chen, J., Soler, E., & Díaz, M. (2018). On block chain 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>
16. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646. <https://doi.org/10.1109/JIOT.2016.2579198>
17. Sicari, S., Rizzardi, A., Grieco, L. A., & Coen-Porisini, A. (2015). Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, 76, 146–164. <https://doi.org/10.1016/j.comnet.2014.11.008>
18. Weber, R. H. (2010). Internet of Things – New security and privacy challenges. *Computer Law & Security Review*, 26(1), 23–30. <https://doi.org/10.1016/j.clsr.2009.11.008>
19. Whitmore, A., Agarwal, A., & Da Xu, L. (2015). The Internet of Things—A survey of topics and trends. *Information Systems Frontiers*, 17(2), 261–274. <https://doi.org/10.1007/s10796-014-9489-4>
20. Xu, L. D., Xu, E. L., & Li, L. (2018). Industry 4.0: State of the art and future trends. *International Journal of Production Research*, 56(8), 2941–2962. <https://doi.org/10.1080/00207543.2018.1444806>