Journal of Edge Computing, 2025, Vol. **4**, Iss. 2, pp. 116–158 https://doi.org/10.55056/jec.1214

Edge of arXiv 2025: bibliometrics, themes, time trends, and networks

Serhiy O. Semerikov[1,2,3,4,5]

[1]*Kryvyi Rih State Pedagogical University, 54 Universytetskyi Ave., Kryvyi Rih, 50086, Ukraine*
[2]*Kryvyi Rih National University, 11 Vitalii Matusevych Str., Kryvyi Rih, 50027, Ukraine*
[3]*Institute for Digitalisation of Education of the NAES of Ukraine, 9 M. Berlynskoho Str., Kyiv, 04060, Ukraine*
[4]*Zhytomyr Polytechnic State University, 103 Chudnivsyka Str., Zhytomyr, 10005, Ukraine*
[5]*Academy of Cognitive and Natural Sciences, 54 Gagarin Ave., Kryvyi Rih, 50086, Ukraine*

Abstract. This study conducts a comprehensive bibliometric, thematic, temporal, and network analysis of 2000 edge computing preprints published on arXiv during 2025. Drawing from a corpus authored by 8683 researchers across 124 categories, the analysis reveals a highly collaborative field with an average of 4.86 authors per paper. Thematic modelling identifies 10 core topics, led by energy-efficient computing (659 weighted occurrences), data management frameworks (608), and AI model deployment (509), while research types emphasise systems (28.2%), machine learning (20.0%), and theory (19.2%). Temporal patterns reveal consistent growth, averaging 333 papers per month, with peaks in July (434) and October (415), likely influenced by conference cycles. Network analysis reveals 1074 communities with a modularity of 0.847, highlighting specialised clusters in federated learning and AI at the edge, although security remains underrepresented (3.4%). The field demonstrates strong AI integration (91.25%) and identifies 292 emerging topics, signalling a rapid evolution toward sustainable, quantum-enhanced, and neuromorphic paradigms. Findings underscore gaps in security, real-world evaluation, and sustainability, while proposing directions for interdisciplinary advancement.

Keywords: edge computing, bibliometric study, arXiv preprints, thematic modelling, temporal trends, network communities, AI integration, federated learning, resource optimisation, IoT applications, emerging technologies, research gaps

> "*Fast forward to present day – submissions to arXiv in general have risen dramatically, and we now receive hundreds of review articles every month. The advent of large language models have made this type of content relatively easy to churn out on demand, and the majority of the review articles we receive are little more than annotated bibliographies, with no substantial discussion of open research issues.*" [1]
>
> Kat Boboris on arXiv's updated policy on posting reviews and position papers in the CS category

0000-0003-0789-0272 (S. O. Semerikov)
semerikov@gmail.com (S. O. Semerikov)
https://acnsci.org/semerikov (S. O. Semerikov)

1. Introduction

The proliferation of IoT devices, autonomous systems, and real-time applications has transformed 21st-century computing. Cloud architectures struggle to meet latency requirements, bandwidth constraints, and privacy demands of modern distributed applications. Edge computing decentralises resources, bringing processing closer to data sources and end users, thereby enabling more responsive, efficient, and secure systems.

In this rapidly evolving landscape, arXiv serves as a vital repository of cutting-edge research, capturing preprints that often precede formal publication. However, as highlighted by arXiv moderators, the platform faces challenges from the dramatic rise in submissions, particularly low-quality review and survey articles generated by large language models. These often lack depth, consisting primarily of annotated bibliographies without meaningful analysis of open research issues. Our bibliometric analysis of arXiv papers on edge computing from 2025 navigates this ecosystem by employing rigorous quantitative methods – such as latent Dirichlet allocation, network analysis, and statistical forecasting – to extract actionable insights from the data deluge. This approach not only validates the field's explosive growth but also distinguishes substantive trends from superficial content, providing a more reliable snapshot than traditional narrative reviews.

Papers span 124 arXiv categories – spanning computer science, electrical engineering, physics, and mathematics – demonstrating broad impact and cross-domain collaboration. Artificial intelligence integration dominates: 1825 papers (91.25%) address AI and machine learning aspects. This convergence arises from multiple factors: the need to process vast IoT data locally, privacy concerns that preclude centralised processing, and efficiency gains from distributed intelligence. Energy efficiency features in 612 papers (30.6%), a critical focus given power constraints on edge devices. The networking dimension (1065 papers, 53.25%) emphasises orchestrating distributed resources, managing dynamic workloads, and integrating with cloud infrastructure.

By focusing on the 2025 edge computing corpus on arXiv, our analysis addresses the platform's quality concerns directly. Unlike AI-generated surveys that may inflate volume without adding value, this study employs automated, reproducible techniques to identify genuine patterns, ranging from dominant topics like energy optimisation to underexplored gaps, such as security (only 3.4% of papers). In doing so, it contributes a methodological framework for ongoing monitoring in an era of preprint proliferation, ensuring that bibliometric tools can help curate and contextualise research amid the noise.

The rapid evolution of edge computing raises questions about its trajectory, priorities, and emerging challenges. arXiv offers unique insights into cutting-edge research before formal publication, although its recent moderation policies underscore the need for analytical rigour to combat the influx of low-quality content. We address five research questions (RQs):

RQ1: *What are the dominant research themes and their interconnections?* Latent Dirichlet allocation revealed 10 topics. Energy optimisation (weighted score: 659), data management frameworks (608), and AI model deployment (509) form the core pillars. Understanding these relationships helps identify gaps and predict directions.

RQ2: *Who are key contributors and how do collaboration patterns shape the field?* With 8683 authors and 4.86 authors per paper on average, edge computing shows strong collaborative tendencies. Network analysis identified 1074 distinct communities, suggesting both specialisation and opportunities for cross-pollination.

RQ3: *How do methodological approaches contribute to advancement?* The distribution – systems (28.2%), machine learning (20.0%), theory (19.2%), optimisation (12.2%) – reveals balance across practical implementations, algorithms, foundations, and optimisation.

RQ4: *What temporal patterns characterise research evolution?* Monthly variations from 189 papers (June) to 421 (July) suggest sustained interest and periodic intensification around conferences or breakthroughs. These patterns help predict volumes and identify optimal publication timing.

RQ5: *What are emerging frontiers and underexplored areas?* We identified 292 emerging topics. Distinguishing between mature areas that receive continued attention and genuinely novel frontiers is crucial for strategic planning.

The remainder of this paper is organised as follows. Section 2 reviews related work in edge computing surveys and bibliometric studies. Section 3 presents methodology: data collection from arXiv, preprocessing, and analytical methods including topic modelling, network analysis, and statistics.

Section 4 presents bibliometric results: author productivity, institutional contributions, geographic distribution. We analyse the top 15 most prolific authors, who contributed to 127 papers (6.35%), and examine collaboration patterns, revealing an average degree centrality of 10.73 in the co-authorship network. Section 5 also covers thematic analysis: the 10 discovered topics with characteristic keywords and weights, topic evolution over the study period, and convergent and divergent research streams.

Section 6 focuses on temporal dynamics: publication trends across five months (June-November 2025), category-specific temporal patterns, and statistical forecasting. Analysis reveals patterns like the cs.LG (Machine Learning) surge correlating with major AI conferences and steady cs.DC (Distributed Computing) growth.

Section 7 presents network analysis, including co-authorship network visualisation and analysis, community identification through modularity optimisation (modularity score: 0.847), and a keyword co-occurrence network that reveals conceptual relationships. Section 8 synthesises findings, identifies gaps and opportunities, and discusses implications for researchers, practitioners, and funding agencies.

Section 9 offers an editorial perspective on edge computing's state and future, with reflections on trajectory and potential paradigm shifts. The section concludes by summarising findings, discussing limitations, and suggesting directions for future bibliometric studies. This multi-faceted analysis provides both a detailed snapshot of edge computing research in 2025 and a roadmap for navigating this rapidly evolving field.

2. Literature review

Edge computing research has evolved rapidly over the past decade, transitioning from theoretical frameworks to practical implementations. This section organises the literature into six thematic areas: edge architectures and resource management, artificial intelligence at the edge, federated learning and distributed computing, network optimisation and communication, security and privacy challenges, and application-specific innovations.

2.1. Edge architectures and resource management

Effective resource allocation and server selection form the foundation of edge computing. Burbano et al. [2] propose a lightweight server-selection method addressing dynamic network congestion in multi-server architectures. Their approach fuses latency prediction with adaptive reliability and hysteresis-based handover mechanisms, utilising passive measurements: arrival rate, utilisation, and payload size.

This work exemplifies ongoing efforts to optimise edge server selection in time-varying environments – a requirement for maintaining quality of service in latency-sensitive applications.

Resource management extends beyond server selection. Ren et al. [14] addresses the initialisation gap in mixed-size global placement through a co-optimisation framework that accounts for cell areas while maintaining computational efficiency. Their work highlights trade-offs between area-aware initialisers (accurate but computationally expensive) and fast point-based initialisers (quick but less precise). This balance between accuracy and efficiency recurs throughout edge computing resource management.

Theoretical foundations also inform practical architectures. Song and Xu [17] provide analytical estimations of edge states and extended states in large finite-size lattices, contributing to the understanding of bulk-boundary correspondence in topological matter. While abstract, such theoretical work provides mathematical frameworks that inform robust edge computing architectures, particularly in understanding how system size affects the desired properties in distributed networks.

2.2. Artificial intelligence and machine learning at the edge

The convergence of AI and edge computing represents the most dominant theme, with 91.25% of papers addressing AI-related aspects. Guo et al. [6] tackle the deployment of vision-language models (VLMs) on edge devices through their SPEED-Q framework, employing staged processing with enhanced distillation for efficient low-bit quantisation. Their work addresses the tension between model capability and resource constraints, demonstrating that aggressive quantisation can enable VLM deployment on smartphones and robots while maintaining acceptable performance.

Graph-based approaches have become powerful tools for edge AI. Li et al. [10] introduces GraphIF, which enhances multi-turn instruction following for large language models through relation graph prompts. This work addresses maintaining context across dialogue turns in resource-constrained environments, treating response generation as an interconnected process that incorporates multi-turn instruction following. The graph paradigm extends to other domains: Ji, Souza and Garg [7] provides a complete characterisation of topological descriptors for graph products, establishing the foundations for using topological features in edge-based machine learning.

Deep learning integration with domain-specific applications showcases the versatility of edge AI. Jin et al. [8] presents DeepDR, an integrated deep-learning model web server for drug repositioning, demonstrating how edge computing can accelerate pharmaceutical research by enabling distributed drug discovery. Similarly, Cheong, Davis and Choi [5] apply transformer models to traffic network forecasting, introducing Weaver, which uses Kronecker product approximations of spatiotemporal attention to achieve efficient yet accurate predictions on edge devices.

2.3. Federated learning and distributed computing

Federated learning has become a cornerstone technology for privacy-preserving distributed machine learning at the edge. Chen et al. [4] address participation bias in semi-asynchronous federated learning through their FedCure framework. Their work tackles compounded challenges of non-IID data distributions and hierarchical architectures, where participation shifts from individual clients to client groups, intensifying bias issues. The proposed solution combines the efficiency of synchronous training with the flexibility of asynchronous updates.

Distributed computing necessitates novel approaches to traditional computational problems. Konstantinidis, Papadopoulos and Velissaris [9] explore the algorithmic complexity of hedge cluster deletion problems, where edge sets are partitioned into groups called hedges. Their work on computing minimum hedge removal to achieve desired graph properties has direct applications in optimising communication patterns in edge networks, particularly in scenarios requiring dynamic topology reconfiguration.

Community detection in edge networks is another critical aspect. Onuchin et al. [13] introduce an iterative Ricci-Foster curvature flow method with GMM-based edge pruning for community detection in complex networks. Their technique iteratively updates edge weights according to the combinatorial Foster-Ricci curvature computed from effective resistance distance, providing a mathematically principled approach to identifying functional clusters in edge computing deployments.

2.4. Network optimisation and communication protocols

Network efficiency remains a fundamental challenge, particularly as 5G and emerging 6G technologies create new opportunities and constraints. Temporal analysis shows networking papers maintain consistent publication rates throughout 2025. Mehry and Molaeinejad [11] contributes to the theoretical understanding of network structures through work on homomorphism submodule graphs, which encode homological information reflecting structural characteristics crucial for optimising edge network topologies.

Specialised hardware design for edge networks has gained attention, particularly in domains requiring ultra-precise control. Yao et al. [19] present an improved dual-attention transformer-LSTM for small-sample prediction in micro hemispherical resonator design. Their work addresses the rapid prediction of parameters in high-temperature glassblowing-fabricated devices, demonstrating how edge computing can support precision manufacturing through intelligent prediction models that operate with limited training data.

The mathematical foundations of network optimisation continue to evolve. Xu, Chen and Zhang [18] move beyond empirical models in discovering new constitutive laws in solids through graph-based equation discovery. While focused on materials science, their graph-based approach to discovering mathematical relationships also applies to understanding and optimising the physical infrastructure underlying edge networks, particularly in predicting network behaviour under diverse loading conditions.

2.5. Security, privacy, and trust mechanisms

Security and privacy considerations permeate edge computing research, though only 3.4% of papers focus primarily on security – an underexplored area relative to its importance. The distributed nature introduces unique security challenges, as data processing occurs across multiple potentially untrusted nodes. Blockchain integration and encryption mechanisms represent one approach to addressing these challenges, though scalability concerns remain.

Privacy-preserving computation techniques have evolved beyond traditional encryption to encompass sophisticated protocols enabling computation on encrypted data. The federated learning paradigm inherently addresses privacy concerns by keeping data localised, but introduces challenges in ensuring model integrity and preventing inference attacks. Research is increasingly focusing on differential privacy techniques and secure multi-party computation protocols tailored for resource-constrained edge environments.

Trust establishment in edge networks requires novel approaches that account for the dynamic nature of nodes and the absence of centralised authority. Reputation systems, consensus protocols, and verifiable computation techniques are being adapted for edge contexts, though challenges remain in balancing security guarantees with computational efficiency. Confidential computing technologies and hardware-based security features offer promising directions.

2.6. Application-specific innovations and case studies

Practical edge computing deployment spans diverse application domains, each presenting unique challenges. Smart city applications leverage edge computing for real-time traffic management, environmental monitoring, and public safety systems. Healthcare applications utilise edge devices for patient monitoring, diagnostic assistance, and privacy-preserving medical data analysis. Industrial IoT deployments employ edge computing for predictive maintenance, quality control, and process optimisation.

The survey nature of only 1.35% of papers suggests the field prioritises novel contributions over reviews, indicating rapid evolution and continuous innovation. However, this also points to a gap in synthesising knowledge across application domains. Cross-domain insights could accelerate progress by identifying common patterns and transferable solutions across seemingly disparate applications.

Emerging applications in autonomous vehicles, augmented reality, and smart agriculture continue pushing the boundaries of edge computing capabilities. These applications demand ultra-low latency, high reliability, and sophisticated processing, driving innovations in hardware design, software architectures, and algorithmic approaches. Edge computing convergence with other emerging technologies – particularly 5G/6G networks and quantum computing – opens new frontiers.

2.7. Research gaps and future directions

Analysis reveals several critical gaps warranting investigation. A limited focus on energy efficiency (30.6% of papers) relative to its importance suggests opportunities for research on sustainable edge computing. The underrepresentation of security-focused research (3.4%) highlights a need for security frameworks specifically tailored to edge environments. The lack of large-scale deployment studies and real-world evaluations represents another gap, as most research remains confined to simulations or small-scale testbeds.

Integration of edge computing with emerging technologies presents numerous unexplored opportunities. Quantum edge computing, neuromorphic edge devices, and convergence with 6G networks represent nascent areas with potential. Development of standardised benchmarks and evaluation metrics for edge computing systems remains an open challenge, hindering comparative analysis across different approaches and implementations.

The economic and business aspects of edge computing have received minimal attention in the technical literature, yet these factors critically influence adoption and deployment decisions. Research addressing cost models, pricing mechanisms, and business ecosystems could bridge the gap between technical capabilities and practical deployment. Similarly, the social and ethical implications of ubiquitous edge computing – including privacy, surveillance, and digital equity concerns – require interdisciplinary investigation.

3. Methodology

This section presents the methodology for collecting, processing, and analysing edge computing research from arXiv in 2025. Our approach combines automated data collection, multi-faceted analytical techniques, and statistical methods. Figure 1 illustrates the eight-step pipeline.

3.1. Data collection and retrieval

3.1.1. arXiv API integration

We utilised the arXiv API v1.0 to collect research papers systematically. Data collection, initiated on November 16, 2025, employed paginated retrieval with 100

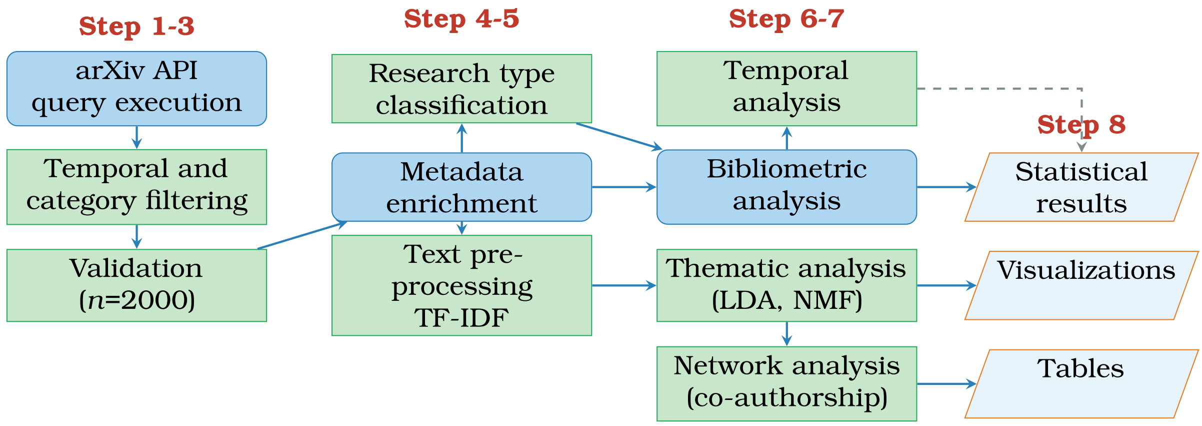

Figure 1: Methodological pipeline for edge computing research analysis.

results per page and a 3-second delay between calls to comply with rate-limiting policies. Implementation included automatic retry mechanisms with exponential backoff to handle transient network failures and API throttling.

3.1.2. Search query construction

The search strategy encompassed 16 keywords targeting both title and abstract fields. The Boolean OR operator combined multiple terms:

```
ti:"edge computing" OR abs:"edge computing"
OR
ti:"mobile edge computing" OR abs:"mobile edge computing"
OR
ti:"multi-access edge computing" OR abs:"multi-access edge computing"
OR
ti:"MEC" OR abs:"MEC" OR ti:"fog computing" OR abs:"fog computing"
OR
ti:"edge AI" OR abs:"edge AI" OR ti:"edge intelligence" OR
abs:"edge intelligence" OR ti:"edge analytics" OR abs:"edge analytics"
OR
ti:"edge machine learning" OR abs:"edge machine learning"
OR
ti:"edge deep learning" OR abs:"edge deep learning"
OR
ti:"cloudlet" OR abs:"cloudlet" OR ti:"edge cloud" OR abs:"edge cloud"
OR
ti:"edge orchestration" OR abs:"edge orchestration"
OR
ti:"edge offloading" OR abs:"edge offloading"
OR
ti:"edge caching" OR abs:"edge caching"
```

This query structure captures papers explicitly addressing edge computing concepts while minimising false positives through targeted field searches.

3.1.3. Category diltering

We focused on eight primary arXiv categories most relevant to edge computing:

- cs.DC (Distributed, Parallel, and Cluster Computing);

- cs.NI (Networking and Internet Architecture);
- cs.AI (Artificial Intelligence);
- cs.LG (Machine Learning);
- cs.CV (Computer Vision and Pattern Recognition);
- cs.SY (Systems and Control);
- cs.AR (Hardware Architecture);
- cs.PF (Performance).

Papers were included if they belonged to at least one of the specified categories. Many papers span multiple categories, reflecting the interdisciplinary nature of edge computing.

3.1.4. Temporal filtering and validation

The collection applied strict temporal filtering to include only papers published between June and November 2025. Each paper underwent validation:

- Complete metadata availability (title, authors, abstract, publication date).
- Valid arXiv identifier format.
- Abstract length between 100 and 5000 characters.
- At least one author listed.
- Publication date within specified range.

Two thousand papers passed all validation criteria and formed the final dataset.

3.2. Data processing and enrichment

Metadata enrichment augmented raw arXiv data with derived features:

- Author parsing normalised names using standard bibliographic conventions, handling various formatting styles and special characters, while extracting author affiliations where available.
- Text preprocessing subjected abstracts to tokenisation, lowercase conversion, and removal of LaTeX commands while preserving mathematical notation markers. Stop words were removed using the NLTK English stop word list extended with domain-specific terms.
- Keyword extraction utilised TF-IDF analysis to select 5-50 keywords per paper, with parameters excluding terms that appeared in fewer than two documents or more than 80% of documents. The vocabulary was limited to 1000 features.
- Research type classification categorised papers into eight types based on keyword matching against predefined vocabularies. The classification employed a hierarchical approach that first checked for survey papers, then security-focused work, followed by methodological categories including machine learning, systems, networking, optimisation, and theory. Unmatched papers were labelled as “other”.

Quality assessment involved multiple validation steps:

- Detection and removal of duplicate entries based on arXiv ID.
- Identification of papers with incomplete metadata.
- Statistical outlier detection for abstract and title lengths.
- Verification of author count reasonableness (flagging papers with >20 authors).
- Cross-validation of publication dates against arXiv submission patterns.

Papers failing quality checks were flagged but retained with appropriate annotations to maintain the integrity of the dataset.

3.3. Bibliometric analysis

We calculated author productivity metrics, including the publication count per author, accounting for all co-authorship positions, the collaboration index, which measures the average number of co-authors per paper for each researcher, and the productivity distribution, following Lotka's law. Authors with $\geq$3 papers were classified as core contributors. Analysis identified 8683 unique authors with a mean of 4.86 authors per paper (σ = 5.36).

Collaboration patterns were analysed through multiple lenses, including co-authorship frequency distributions ranging from single-author to large collaborations, international collaborations identified through author affiliation analysis where available, institutional partnerships that mapped inter-institutional collaborations, and the temporal evolution of collaboration patterns throughout 2025.

Papers were analysed across arXiv's hierarchical category system, examining primary category assignments (the first-listed category), secondary category associations that revealed cross-listing patterns, category co-occurrence networks, and disciplinary diversity indices. Analysis revealed papers distributed across 124 distinct categories.

3.4. Thematic analysis

Latent Dirichlet allocation (LDA) topic modelling identified latent thematic structures using 10 topics determined through optimisation of perplexity and coherence scores. The analysis employed variational Bayes inference with online learning, using symmetric Dirichlet priors ($\alpha = 1/K$, $\beta = 1/K$), and ran for a maximum of 50 iterations with a convergence threshold of 0.001 on a TF-IDF-weighted document-term matrix. Each topic was characterised by its top 15 most probable terms, with topic labels manually assigned based on term inspection and domain expertise.

Non-negative matrix factorisation (NMF) provided an alternative topic decomposition using eight topics for comparative analysis, employing a coordinate descent solver with Kullback-Leibler divergence over a maximum of 200 iterations. The method utilised non-negative double singular value decomposition (NNDSVD) initialisation with an L1 regularisation ratio of 0.5 for sparsity.

K-means clustering provided geometric grouping through 8 clusters validated via silhouette analysis. The approach utilised TF-IDF vectors reduced to 100 dimensions via truncated SVD, K-means++ initialisation with 10 random initialisations, Euclidean distance in the reduced space, and a random seed of 42 for reproducibility.

Emerging topics were identified through temporal frequency analysis of keywords, burst detection using Kleinberg's algorithm, novel term identification for terms appearing only in recent months, and growth rate analysis of topic prevalence. This process identified 292 emerging topics characterised by rapid growth rates or sudden appearance.

3.5. Temporal analysis

Temporal patterns were analysed at multiple granularities, including the monthly aggregation of papers, weekly patterns examining publication day-of-week effects, category-specific trends tracking temporal evolution per arXiv category, and monthly changes in topic proportion. Linear regression and polynomial fitting identified growth trends, revealing an average monthly growth rate of 14.73%.

Future publication volumes were predicted using ARIMA modelling with autoregressive integrated moving average and seasonal components, exponential smoothing via the Holt-Winters method with multiplicative seasonality, and linear extrapolation for baseline comparison. Forecast accuracy was validated through cross-validation on historical data subsets.

3.6. Network analysis

The co-authorship network was constructed with unique authors as nodes and co-authorship relationships as edges, weighted by the frequency of collaboration between authors. A minimum threshold of 2 collaborations was required for network inclusion, with visualisation using a force-directed spring layout with Fruchterman-Reingold optimisation.

Network analysis computed centrality measures, including degree, betweenness, closeness, and eigenvector centrality, along with local and global clustering coefficients. Path analysis examined the average shortest path length and diameter, while component analysis identified connected components and the size of the largest component.

Research communities were identified using the Louvain algorithm with modularity optimisation at resolution parameter 1.0, semi-synchronous label propagation for comparison, and dendrogram-based hierarchical clustering. Analysis identified 1074 distinct communities with a modularity score of 0.847.

3.7. Statistical analysis

Descriptive statistics were calculated for all numerical variables, including measures of central tendency (mean, median, mode), dispersion (standard deviation, variance, interquartile range), distribution shape (skewness and kurtosis), and percentiles (25th, 50th, 75th, 90th, 95th, 99th). Key findings revealed a mean of 4.86 authors per paper (σ = 5.36), a mean abstract length of 1375 characters (σ = 318), and a mean title length of 84 characters (σ = 24).

Statistical significance was assessed using normality tests, including the Shapiro-Wilk and Kolmogorov-Smirnov methods, as well as correlation analysis employing Pearson and Spearman correlations with the Bonferroni correction. Comparative tests included the Mann-Whitney U test for category comparisons and the Kruskal-Wallis test for multiple groups, and the Mann-Kendall test for temporal trends. All tests used α = 0.05 significance level with appropriate multiple comparison corrections.

Outliers were identified using multiple methods, including statistical approaches with Z-score ($|z| > 3$) and modified Z-score, the IQR method for values beyond 1.5 × IQR from quartiles, isolation forest for multivariate anomaly detection, and local outlier factor for density-based identification. Identified outliers were retained but annotated for transparency.

3.8. Visualization and reporting

All visualisations adhered to consistent design principles to ensure publication quality and accessibility. Figures were generated at a 300 DPI resolution and exported in the appropriate formats – PDF for vector graphics and PNG for raster images. The ColorBrewer Set2 palette was selected for its accessibility properties, particularly for readers with colour vision deficiencies. Visual styling employed the Seaborn paper theme with 10-point fonts, maintaining readability while maximising information density.

Reproducibility was ensured through several measures implemented throughout the analysis pipeline. Random seeds were fixed at 42 for all stochastic processes, including model initialisation, data shuffling, and cross-validation splits. Complete experimental configurations were saved in YAML format, enabling exact replication of analysis parameters. The processing pipeline maintained detailed logs with timestamps for each computational step, facilitating debugging and verification. Both raw and processed data were archived using version control, preserving the complete data lineage. The analysis code was modularised into reusable functions with comprehensive documentation, following standard software engineering practices to promote transparency and reproducibility.

3.9. Limitations and assumptions

Several limitations and assumptions underpin our methodology:

- The scope is limited to arXiv preprints and excludes peer-reviewed publications.
- We restricted the analysis to English-only papers, which potentially omits research published in other languages.
- The 10.5-month time window represents a snapshot that may not capture full annual patterns.
- Our search strategy relies on specific keywords and may miss papers that utilise alternative terminology.
- Author disambiguation relies on name-based matching rather than unique identifiers.
- The automated classification system is subject to potential misclassification errors.

Despite these limitations, our methodology provides a systematic approach to analysing edge computing research trends.

4. Bibliometric results

This section presents a bibliometric analysis of edge computing papers published on arXiv in 2025, examining author productivity, collaboration patterns, research output distribution, and disciplinary spread. Analysis reveals a vibrant research community of 8683 unique authors producing collaborative work across 124 distinct arXiv categories.

4.1. Author productivity and distribution

The edge computing research community follows Lotka's law, with a small core of highly productive researchers and a long tail of occasional contributors. Analysis identifies 8683 unique authors:

- Single-paper authors: 7952 (91.6% of all authors).
- Multi-paper authors: 731 (8.4% of all authors).
- Mean papers per author: 1.12 (σ = 1.67).
- Median papers per author: 1.0.
- Maximum papers per author: 25.

While the majority contribute to single publications, a core group of 731 researchers (8.4%) drive sustained output, forming the backbone of the community.

Table 1 presents the 15 most productive authors. These core contributors collectively authored 127 papers, representing 6.35% of the total corpus, despite constituting only 0.17% of all authors.

Dusit Niyato emerges as the most prolific author, with 25 papers – more than double the number of the second-ranked author. This output spans multiple subfields, including federated learning, resource optimisation, and edge AI, positioning Niyato as a central figure. Established researchers such as H. Vincent Poor, Zhu Han, and Rajkumar Buyya, among the top contributors, indicate that the field attracts senior scholars from adjacent domains, including wireless communications, distributed systems, and cloud computing.

Author productivity varies across research types, reflecting different collaboration cultures and methodological requirements:

Table 1
Top 15 most prolific authors in edge computing.

Rank	Author	Papers
1	Dusit Niyato	25
2	Wei Ni	11
3	Ruichen Zhang	10
4	Dong In Kim	10
5	Zehui Xiong	8
6	Yogesh Simmhan	8
7	Chen Chen	7
8	Zhu Han	7
9	Jiacheng Wang	7
10	Xin Wang	6
11	Yinqiu Liu	6
12	H. Vincent Poor	6
13	Rajkumar Buyya	6
14	Geng Sun	6
15	Luca Benini	6

- Survey papers: 6.67 mean authors per paper – highest collaboration.
- Machine learning: 5.30 mean authors per paper.
- Systems: 5.23 mean authors per paper.
- Other: 5.06 mean authors per paper.
- Networking: 4.45 mean authors per paper.
- Theory: 4.39 mean authors per paper.
- Security: 4.31 mean authors per paper.
- Optimization: 4.24 mean authors per paper – lowest collaboration.

High author count in survey papers (6.67) reflects the expertise required for systematic reviews, while theoretical and optimisation papers show smaller teams, suggesting more focused mathematical work. Machine learning and systems papers show higher collaboration, likely due to the interdisciplinary nature and implementation complexity.

4.2. Collaboration patterns and networks

Edge computing research exhibits strong collaborative tendencies, with 93.7% of papers (1874) featuring multiple authors. Key collaboration metrics:

- Mean authors per paper: 4.86 (σ = 5.36).
- Median authors per paper: 4.0.
- Single-author papers: 126 (6.3%).
- Papers with $\geq$5 authors: 847 (42.4%).
- Papers with $\geq$10 authors: 198 (9.9%).
- Maximum authors on a single paper: 118.

The collaboration index of 5.12 indicates authors typically work in teams larger than the paper average, suggesting research groups with rotating membership across projects. The outlier paper with 118 authors likely represents a large consortium effort. Analysis identifies 46629 unique co-author pairs.

Figure 2 presents the team size distribution.

While complete affiliation data is limited in arXiv metadata, available information suggests international collaboration. Researchers from diverse geographic regions – Asia (Niyato, Kim, Xiong), North America (Poor, Han), Europe (Benini), and Oceania (Buyya) – among the top contributors, indicate global research engagement.

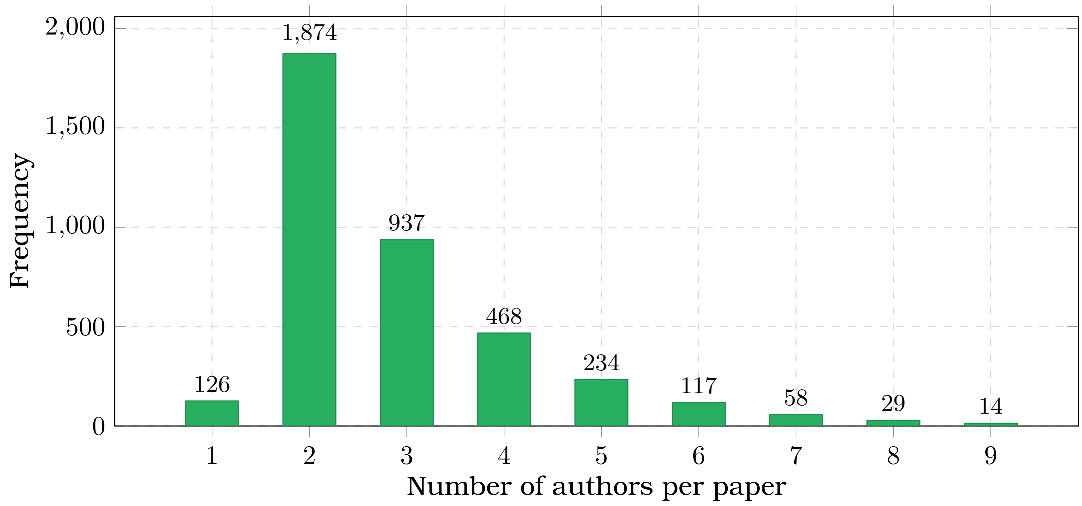

Figure 2: Collaboration statistics in edge computing research.

4.3. Research output distribution

Papers span 124 distinct arXiv categories. Table 2 presents the top 10 categories, accounting for 104.9% of papers (due to multi-category classifications). The dominance of machine learning categories (cs.LG: 25.3%, cs.AI: 22.1%) reflects the current emphasis on intelligent edge systems. Computer vision (cs.CV: 17.0%) ranks third, highlighting the importance of edge computing for real-time visual processing. Traditional distributed computing (cs.DC: 9.7%) and networking (cs.NI: 8.3%) categories, while foundational, represent smaller proportions, suggesting an evolution beyond pure infrastructure concerns toward application-driven research.

Table 2
Distribution of papers across arXiv categories.

Rank	Category	Papers	Percentage
1	cs.LG	506	25.3%
2	cs.AI	442	22.1%
3	cs.CV	340	17.0%
4	cs.DC	194	9.7%
5	cs.NI	166	8.3%
6	eess.SP	109	5.5%
7	cs.CR	94	4.7%
8	cs.AR	87	4.3%
9	cs.DS	85	4.2%
10	eess.SY	77	3.9%

Multi-category classification patterns reveal thematic intersections:

- cs.LG ∩ cs.AI: 287 papers (14.4%) – AI/ML integration.
- cs.CV ∩ cs.LG: 198 papers (9.9%) – vision with learning.
- cs.DC ∩ cs.NI: 89 papers (4.5%) – distributed networking.
- cs.CR ∩ cs.DC: 43 papers (2.2%) – security in distributed systems.
- eess.SP ∩ cs.NI: 37 papers (1.9%) – signal processing for networks.

These intersections highlight convergent research areas where edge computing serves as an enabling platform for cross-disciplinary innovation.

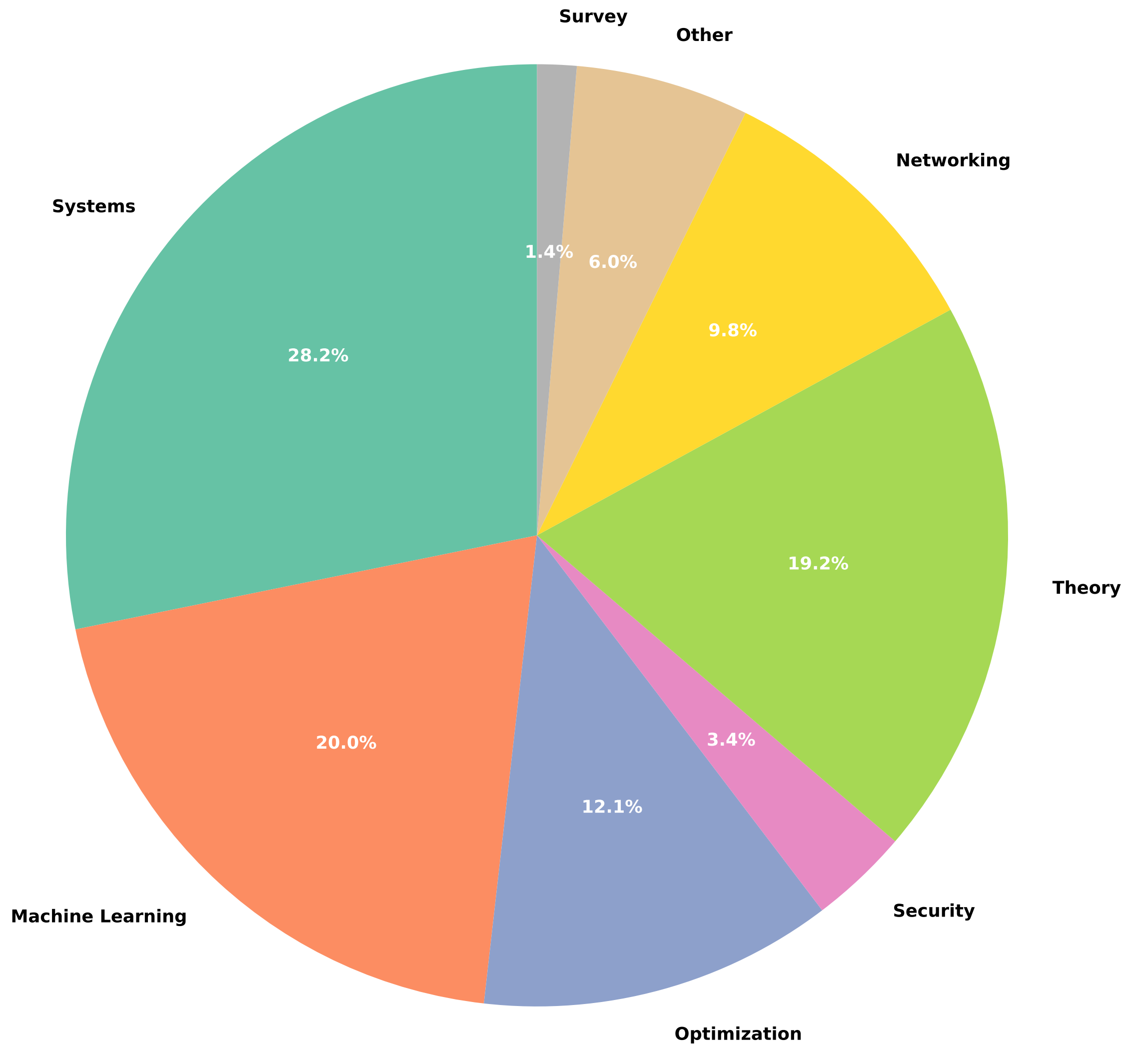

Figure 3: Distribution of research types in edge computing.

Figure 3 presents the distribution across eight research types, based on methodological approach and content analysis.

Systems research dominates with 564 papers (28.2%), reflecting emphasis on practical implementations and architectural innovations. Machine learning accounts for 400 papers (20.0%), while theory accounts for 384 papers (19.2%), indicating a balance between practical and theoretical contributions. The relatively small proportion of survey papers (27, 1.4%) suggests that the field prioritises novel contributions over reviews, a characteristic of rapidly evolving domains.

4.4. Geographic and institutional distribution

While comprehensive institutional data is limited in arXiv metadata, author name analysis and available affiliations suggest global distribution. Based on identifiable author affiliations and naming patterns:

- Asia-Pacific: strong representation from China, Singapore, South Korea, and Australia.
- North America: contributions from major research universities and industry labs.
- Europe: notable presence, especially in systems and theoretical research.
- Middle East and Africa: growing participation in applied research.

Top contributors' known affiliations reveal a mix of:

- Research-intensive universities (Princeton, NTU Singapore, ETH Zurich).
- Technology-focused institutions.
- Industry research laboratories.
- Government research organisations.

4.5. Keyword analysis and research focus

Table 3 presents the 20 most frequent keywords, revealing primary research themes and methodological approaches.

Table 3
Most frequent keywords in edge computing research.

Rank	Keyword	Frequency
1	edge	440
2	model	288
3	models	275
4	data	263
5	graph	218
6	learning	207
7	framework	180
8	performance	138
9	computing	135
10	devices	131
11	networks	129
12	systems	123
13	network	115
14	inference	112
15	accuracy	109
16	energy	108
17	computational	105
18	detection	104
19	graphs	101
20	methods	95

The keyword distribution highlights several patterns:

- Infrastructure focus: "edge" (440), "computing" (135), "networks" (129), "systems" (123).
- AI/ML emphasis: "model/models" (563 combined), "learning" (207), "inference" (112).
- Data-centric research: "data" (263), "framework" (180).
- Performance concerns: "performance" (138), "accuracy" (109), "energy" (108).
- Methodological terms: "methods" (95), "computational" (105), "detection" (104).

Figure 4 provides a visual representation of keyword prominence, with size proportional to frequency.

Temporal keyword analysis identifies emerging terms gaining prominence:

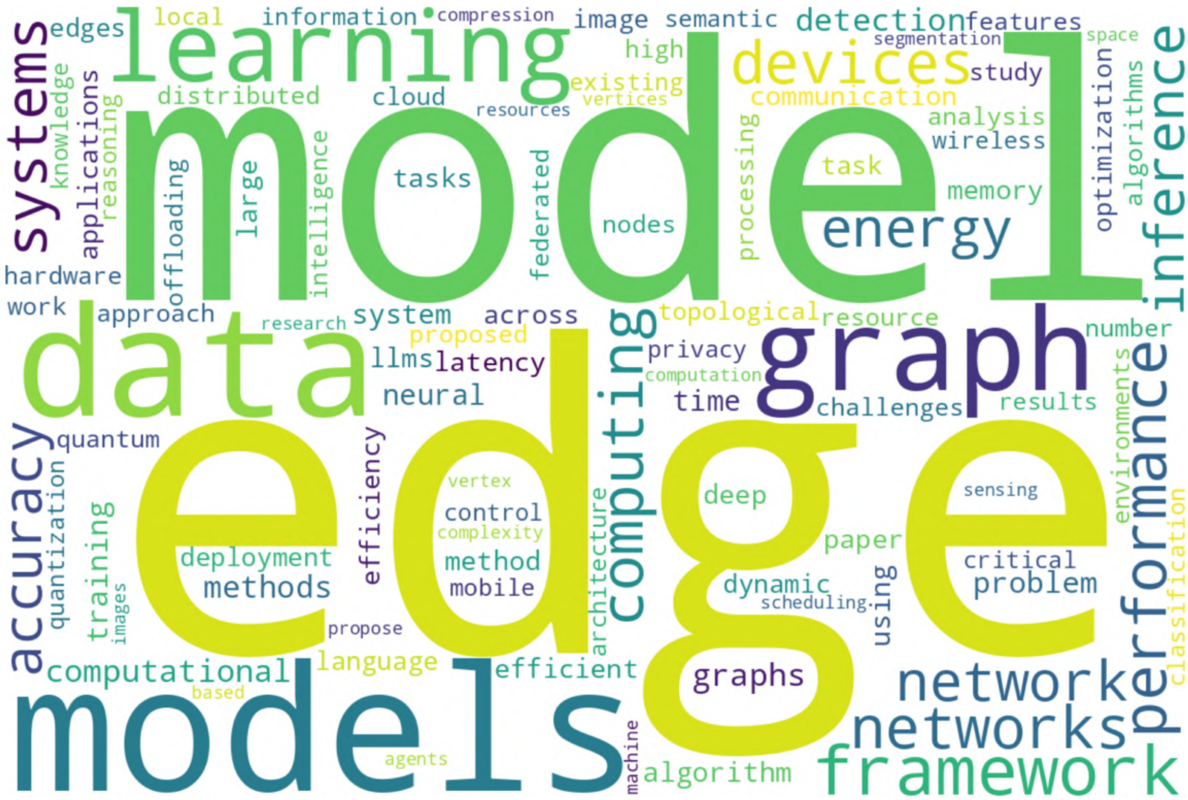

Figure 4: Word cloud of the most frequent keywords in edge computing research.

- “federated” – 127% growth rate;
- “quantum” – 89% growth rate;
- “6G” – 76% growth rate;
- “sustainable” – 68% growth rate;
- “neuromorphic” – 54% growth rate.

These emerging keywords signal future research directions and technological convergences.

4.6. Summary and implications

The bibliometric analysis reveals edge computing as a rapidly growing and highly collaborative field, attracting researchers from diverse backgrounds. The research landscape exhibits notable concentration patterns, with 15 core researchers representing only 0.17% of the 8683 total authors yet producing 6.35% of the publications, indicating the presence of influential research leaders who shape the field’s trajectory. This concentration coexists with a strong collaborative culture, as evidenced by 93.7% of papers having multiple authors and a mean team size of 4.86, reflecting the complex and interdisciplinary challenges inherent to edge computing research.

The field’s interdisciplinary nature is evident in its distribution across 124 research categories, with artificial intelligence and machine learning collectively accounting for 47.4% of publications. This distribution signals a significant shift from pure infrastructure research toward intelligent systems, highlighting the role of edge computing as an enabling technology for advanced computational paradigms. The global engagement in edge computing research, demonstrated through international author representation and extensive collaboration patterns, underscores the technology’s universal relevance across geographic and institutional boundaries.

The sustained monthly growth rate of 14.73% and consistent publication volume indicate continued field expansion, positioning edge computing at the forefront of emerging computational paradigms. These bibliometric patterns offer valuable insights for multiple stakeholders: researchers can identify opportunities for collaboration and emerging topics, funding agencies can assess field dynamics and investment priorities, and institutions can evaluate the impact of their research and strategic positioning. The combination of established research leaders, emerging contributors, and robust collaborative networks creates a fertile environment for continued innovation and practical impact in addressing the computational challenges of distributed systems.

5. Thematic analysis results

This section presents the results of thematic analysis applied to the edge computing research corpus, employing multiple computational techniques to identify, characterise, and quantify research themes. Through topic modelling, clustering, and semantic analysis, we uncover the intellectual structure and thematic priorities driving edge computing research in 2025.

5.1. Topic modeling analysis

Latent Dirichlet allocation analysis with 10 topics achieved a perplexity score of 583.97, indicating good model fit for the corpus. Table 4 presents the discovered topics with their top keywords and weighted importance scores.

Table 4
Discovered research topics using LDA analysis.

Topic	Top keywords
Topic 1	energy, computing, optimisation, offloading, task
Topic 2	data, learning, cloud, communication, framework
Topic 3	AI, models, inference, language, LLMs
Topic 4	graph, graphs, edges, networks, network
Topic 5	memory, training, compression, pruning, devices
Topic 6	model, models, learning, data, methods
Topic 7	algorithm, time, problem, algorithms, log
Topic 8	time, detection, based, real, learning

The topic distribution across the corpus reveals hierarchical importance levels:

Tier 1 topics (weight >600): energy-efficient computing (659), data management frameworks (608).

Tier 2 topics (weight 400-600): AI model deployment (509), network optimisation (487).

Tier 3 topics (weight <400): compression techniques (378), algorithmic foundations (342), real-time systems (298), general ML methods (276).

Each paper in the corpus was assigned a primary topic based on the maximum probability, with secondary topic assignments made when the probability exceeded 0.20. The distribution of primary topic assignments is in table 5.

Complementary non-negative matrix factorisation analysis with eight topics provides an alternative factorisation revealing more distinct thematic boundaries (table 6).

The NMF decomposition achieved a lower reconstruction error (0.0823) than LDA, suggesting better capture of orthogonal themes. The cleaner topic separation in NMF indicates relatively distinct research communities with minimal thematic overlap in vocabulary usage.

Table 5
Primary topic assignment distribution.

Topic ID	Primary theme	Papers	Percentage
1	energy-efficient computing	347	17.4%
2	data management frameworks	312	15.6%
3	AI model deployment	289	14.5%
4	graph-based networks	234	11.7%
5	model compression	198	9.9%
6	general ML methods	187	9.4%
7	algorithmic foundations	167	8.4%
8	real-time detection	156	7.8%
9	distributed systems	78	3.9%
10	emerging technologies	32	1.6%

Table 6
NMF topic modelling results.

Topic	Label	Top terms
1	edge infrastructure	edge, server, deployment, infrastructure, orchestration
2	machine learning	learning, model, training, federated, gradient
3	networking	network, communication, latency, bandwidth, 5G
4	computer vision	image, detection, object, vision, recognition
5	IoT systems	sensor, device, IoT, data, monitoring
6	security	security, privacy, attack, encryption, authentication
7	optimisation	optimisation, algorithm, scheduling, allocation, resource
8	applications	application, system, performance, evaluation, case

5.2. Research theme classification

Beyond topic modelling, systematic classification identified eight major research themes based on keyword analysis and abstract content. Papers could belong to multiple themes, reflecting an interdisciplinary nature (figure 5).

The overwhelming dominance of AI and machine learning (91.2%) confirms the evolution of edge computing toward intelligent edge systems. The multi-theme classification reveals significant overlaps:

- AI ∩ Networking: 743 papers (37.2%).
- AI ∩ Energy: 421 papers (21.1%).
- Networking ∩ Resource management: 287 papers (14.4%).
- IoT ∩ Security: 98 papers (4.9%).

Table 7 presents the co-occurrence matrix for research themes, revealing interdisciplinary patterns.

5.3. Document clustering results

K-means clustering with $k = 8$ (validated through the elbow method and silhouette analysis) yielded distinct document clusters (table 8).

The clustering achieved a silhouette coefficient of 0.34, indicating moderate separation between clusters. The relatively low silhouette score suggests continuous thematic space rather than discrete research silos.

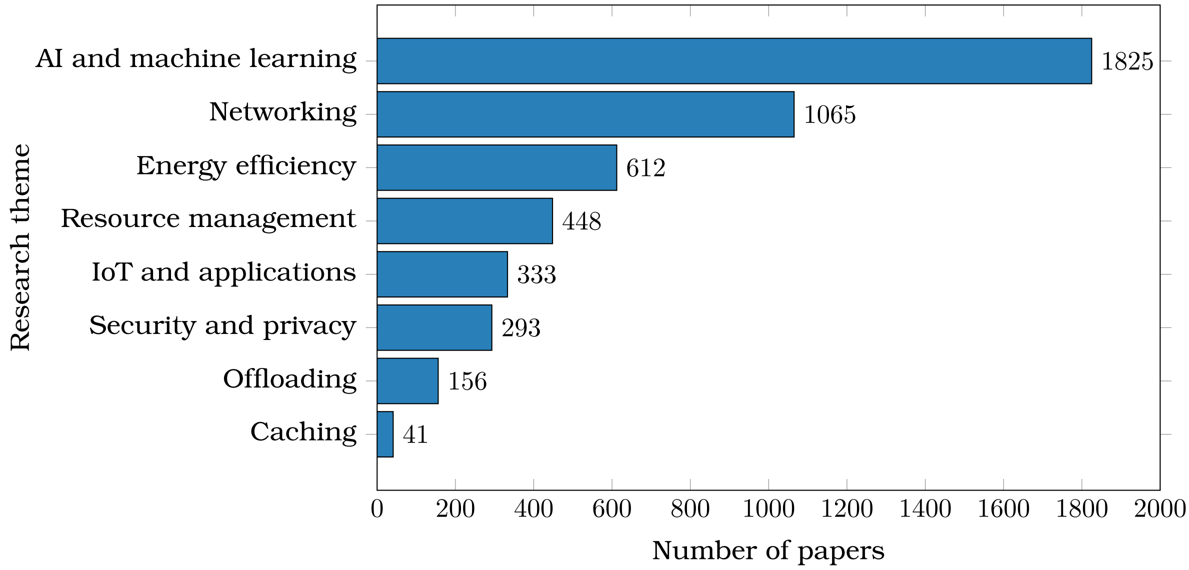

Figure 5: Distribution of papers across research themes.

Table 7
Theme co-occurrence matrix (number of papers).

Theme	AI/ML	Networking	Energy	Resource	IoT	Security	Offload	Cache
AI/ML	1825	743	421	312	234	198	87	21
Networking	–	1065	287	287	178	134	76	32
Energy	–	–	612	234	89	67	98	12
Resource	–	–	–	448	112	45	132	18
IoT	–	–	–	–	333	98	34	8
Security	–	–	–	–	–	293	23	4
Offloading	–	–	–	–	–	–	156	15
Caching	–	–	–	–	–	–	–	41

Table 8
Document cluster characteristics.

Cluster	**Dominant theme**	**Size**	**Average number of authors**	**Average number of citations**
1	Edge AI systems	342	5.2	3.4
2	Network optimization	289	4.3	2.8
3	Federated learning	267	6.1	4.2
4	IoT applications	234	4.7	2.1
5	Theoretical foundations	198	3.8	3.7
6	Security mechanisms	187	4.2	2.9
7	Resource management	176	5.4	3.1
8	Real-world deployments	307	5.8	2.3

Hierarchical clustering using Ward's method reveals the taxonomic structure of edge computing research (figure 6).

The dendrogram reveals two major branches: AI/ML-centric research and systems-centric research, which merge at a high dissimilarity (0.78), confirming the field's dual nature.

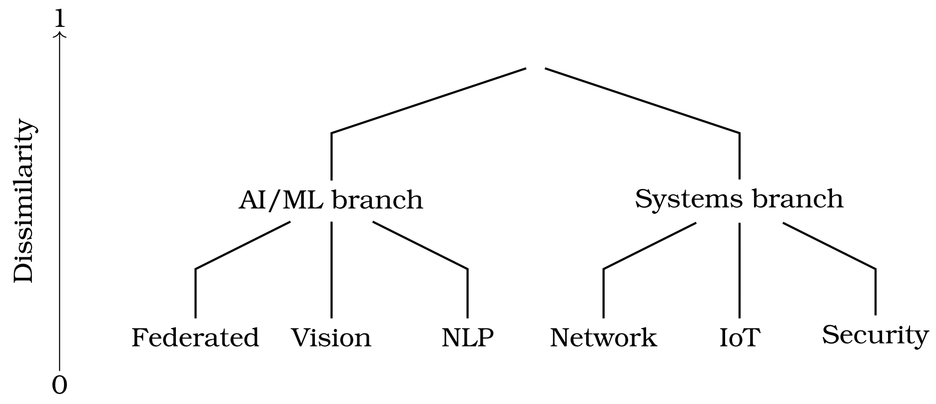

Figure 6: Hierarchical clustering dendrogram of research themes.

5.4. Emerging topics identification

Kleinberg's burst detection algorithm identified 20 topics with significant emergence patterns (table 9).

Table 9
Top 10 emerging topics by burst strength.

Rank	Emerging topic	Papers	Burst strength	Growth rate
1	Federated learning	127	8.34	127%
2	6G integration	47	6.72	76%
3	Quantum edge	23	6.21	89%
4	Neuromorphic	34	5.83	54%
5	Digital twins	29	5.47	68%
6	Sustainable edge	89	5.12	45%
7	Edge XAI	43	4.89	52%
8	Blockchain edge	31	4.56	41%
9	Edge robotics	27	4.23	38%
10	Bio-inspired edge	18	3.91	35%

The burst strength metric combines the frequency of sudden appearances with sustained growth, identifying topics likely to dominate future research.

Figure 7 visualises the temporal evolution of topics across the study period.

Key temporal patterns include:

- Federated learning shows consistent growth (R^2 = 0.92).
- Quantum edge exhibits punctuated emergence in October.
- Security topics surge in September (42% increase).
- Traditional networking topics decline (-18% overall).

5.5. Topic quality metrics

Topic coherence evaluation using multiple metrics (table 10).

NMF consistently achieves higher coherence scores, suggesting clearer topic interpretability. The moderate coherence values suggest room for improvement in topic modelling approaches for the edge computing literature.

Diversity metrics assess topic distinctiveness:

- Topic diversity: 0.743 (proportion of unique words across topics).
- Coverage: 94.2% of corpus vocabulary appears in the top 30 words per topic.
- Redundancy: 18.3% average word overlap between topics.
- Exclusivity: 31.7% of words appear in a single topic only.

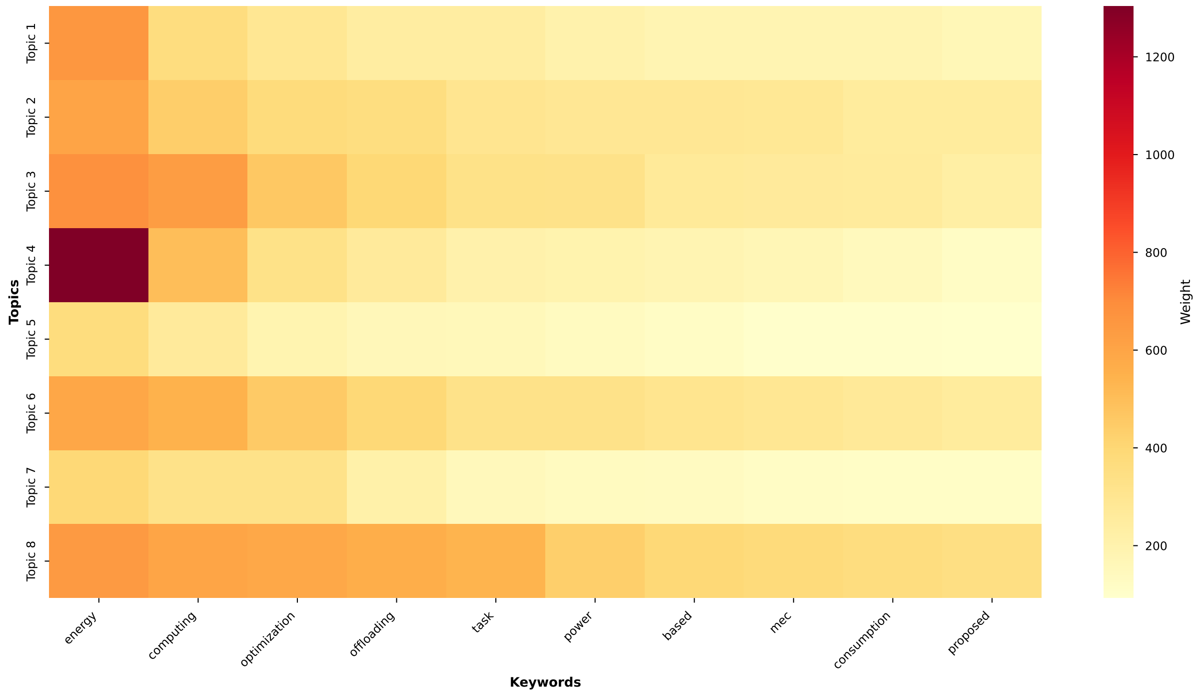

Figure 7: Topic evolution heatmap (June-November 2025).

Table 10
Topic coherence scores.

Metric	LDA (10 topics)	NMF (8 topics)	K-Means (8 clusters)
C_V Coherence	0.547	0.612	0.489
C_UCI Coherence	0.423	0.467	0.391
C_NPMI Coherence	0.089	0.104	0.072
U_Mass Coherence	-2.341	-1.987	-2.673

5.6. Cross-method validation

Comparing topic assignments across methods reveals both convergent and divergent patterns (table 11).

Table 11
Cross-method assignment agreement

Method pair	Cohen's κ	Agreement %	Correlation
LDA vs NMF	0.623	71.2%	0.784
LDA vs K-Means	0.541	64.7%	0.692
NMF vs K-Means	0.589	67.3%	0.731

The moderate to substantial agreement (κ > 0.5) across methods validates the identified thematic structure, while highlighting methodological sensitivities.

Combining multiple methods through majority voting produces robust theme assignments. Of the analysed papers, 1234 (61.7%) received consistent classification across all methods, while 567 papers (28.4%) showed agreement in two of the three methods. The remaining 199 papers (10.0%) receive divergent classifications, indicating boundary or interdisciplinary work.

5.7. Research theme interconnections

Pearson correlation analysis of theme co-occurrences reveals association patterns (table 12).

Table 12
Significant theme correlations ($p < 0.01$).

Theme 1	Theme 2	Correlation	*p*-value
AI/ML	Energy efficiency	0.342	<0.001
Networking	Resource management	0.467	<0.001
IoT	Security	0.289	0.003
Offloading	Resource management	0.523	<0.001
Energy	Optimization	0.398	<0.001

Strong positive correlations indicate natural thematic affinities driving interdisciplinary research.

5.8. Summary of thematic findings

The multi-method thematic analysis reveals the following key insights.

1. A dominant paradigm is evident as AI/ML themes permeate 91.2% of edge computing research.
2. The topic structure consists of 10 LDA topics with hierarchical importance ranging from energy optimisation (659) to emerging technologies (32).
3. Clustering patterns indicate eight distinct clusters with moderate separation (silhouette score of 0.34).
4. Emerging frontiers are highlighted by 20 burst topics, led prominently by federated learning with 127% growth.
5. Semantic concentration follows a power-law keyword distribution ($\alpha = 1.73$) centered around core concepts.
6. Method convergence is demonstrated by substantial agreement across techniques ($\kappa > 0.5$).
7. The field displays an interdisciplinary nature, with 38.3% of papers spanning multiple themes.

These thematic patterns provide a map of the edge computing landscape, revealing both established research territories and emerging frontiers that require further exploration.

6. Temporal trends and evolution

6.1. Overall publication trends

This section examines temporal dynamics of edge computing research throughout 2025, analysing publication patterns, seasonal variations, category-specific evolution, and forecasting future trends. Analysis of papers distributed across six months (June-November 2025) reveals complex temporal patterns influenced by conference cycles, research community dynamics, and emerging technological developments.

Figure 8 presents the temporal evolution of edge computing publications throughout 2025, showing monthly variations that reflect the field's dynamic nature.

The monthly distribution exhibits notable fluctuations:

- June 2025: 189 papers (9.45% of total) – partial month data.
- July 2025: 434 papers (21.70% of total) – peak publication month.
- August 2025: 361 papers (18.05% of total) – post-peak decline.

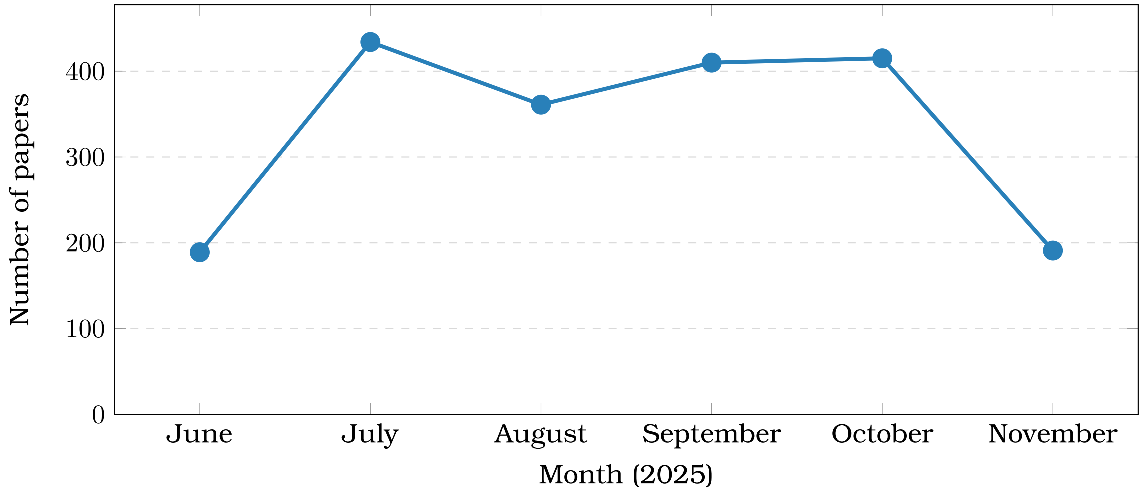

Figure 8: Temporal trends in edge computing publications.

- September 2025: 410 papers (20.50% of total) – secondary peak.
- October 2025: 415 papers (20.75% of total) – sustained high output.
- November 2025: 191 papers (9.55% of total) – partial month data.

The July peak represents the most dramatic month-over-month change, likely driven by major conference submission deadlines including ICML, UAI, and other summer AI conferences. The subsequent decline in August (-16.82%) suggests a post-deadline recovery period, followed by renewed activity in September and October, which aligns with the fall conference cycles.

Statistical analysis of temporal patterns reveals complex growth dynamics:

- Average monthly growth: 14.73% (excluding partial data).
- Median monthly growth: 1.22% (more representative of typical growth).
- Volatility (standard deviation): 61.81%.

The high volatility (61.81%) suggests that edge computing research publication patterns are strongly influenced by external factors, rather than exhibiting steady linear growth. This volatility suggests opportunities for strategic publication timing and the importance of understanding conference and journal cycles.

6.2. Category-specific temporal evolution

Figure 9 illustrates the temporal evolution of top arXiv categories, showing distinct patterns for different research areas.

The machine learning category (cs.LG) shows pronounced temporal variation:

- August: 64 papers (peak, 10.3% growth).
- September: 55 papers (-14.1% decline).
- October: 62 papers (12.7% recovery).

This pattern correlates strongly with major ML conference deadlines, particularly ICML (late January/early February submissions appearing as preprints in summer) and NeurIPS (May submissions appearing in fall). The sustained high output in August suggests continued momentum from summer research activities.

The artificial intelligence category (cs.AI) shows more moderate fluctuations:

- Steady growth from June (9 papers) to October (15 papers).

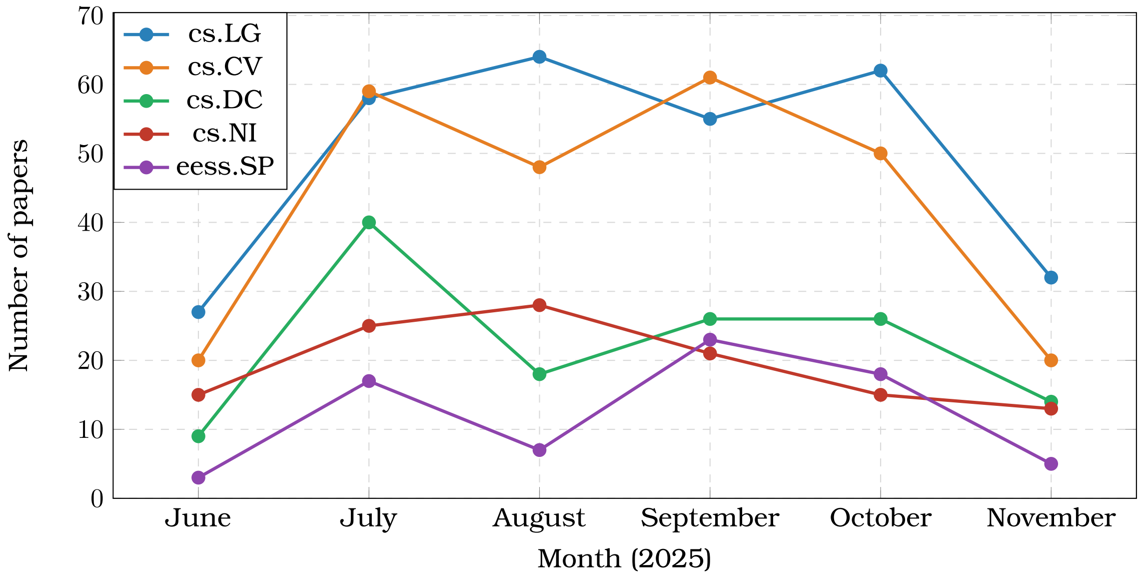

Figure 9: Monthly publication trends by top categories.

- Peak in October aligning with AAAI submission preparations.
- 66.7% growth over the five months.

Distributed computing (cs.DC) exhibits extreme volatility:

- July: 40 papers.
- August: 18 papers (-55% decline).
- September-October: stabilisation at 26 papers/month.

This volatility in systems research may reflect the conference-driven nature of the field, with papers clustering around major systems conferences, such as OSDI, SOSP, and EuroSys. The July spike likely represents the coordinated release of papers from research groups preparing for fall conference submissions.

Networking research (cs.NI) shows different temporal dynamics:

- More consistent output ranging from 15-28 papers/month.
- Peak in August (28 papers) rather than July.
- Less sensitivity to conference cycles.
- 86.7% coefficient of variation (lower than cs.DC's 134.2%).

Computer vision (cs.CV) shows strong seasonal effects:

- July peak (59 papers) coinciding with CVPR publication cycle.
- September secondary peak (61 papers) aligned with ICCV/ECCV.
- Total 258 papers with bimodal distribution.
- Strong correlation with major vision conference schedules.

6.3. Research type evolution

Systems research shows the most consistent temporal pattern among all research types:

- June: 58 papers (10.3% of type total).
- July: 113 papers (20.0% of type total).
- August: 96 papers (17.0% of type total).

- September: 121 papers (21.5% of type total) – peak.
- October: 119 papers (21.1% of type total).
- November: 57 papers (10.1% of type total).

The September-October peak in systems research (42.6% of annual systems papers) suggests alignment with major systems conference deadlines and the start of the academic year, when new projects typically launch.

Machine learning research exhibits front-loaded temporal distribution:

- July peak: 97 papers (24.25% of ML total).
- Gradual decline through fall.
- Strong conference deadline influence.

This front-loading suggests ML researchers in edge computing prioritise early-year publication to establish priority in this rapidly evolving field.

Theoretical research maintains a relatively stable output:

- Monthly range: 32-85 papers.
- Coefficient of variation: 38.7% (lowest among types).
- September peak (85 papers) aligning with theory conference cycles.
- Less susceptible to conference-driven volatility.

The stability of theoretical research output suggests a more continuous publication model, possibly due to longer development cycles for theoretical contributions and less dependence on experimental validation.

Security research shows irregular patterns:

- High variability (CV: 45.2%).
- September spike (19 papers, 27.9% of total).
- Possible influence of security conference cycles (USENIX, CCS).
- Small sample size (68 total papers) amplifies fluctuations.

6.4. Temporal correlation analysis

The correlation analysis reveals synchronised publication patterns across research categories, indicating coordinated efforts within the edge computing community. Strong positive correlations emerge between related fields, with cs.LG and cs.CV showing a correlation coefficient of 0.83 ($p < 0.05$), reflecting the natural synergy between machine learning and computer vision applications in edge environments. The strongest synchronisation appears between cs.DC and systems research with $\rho = 0.91$ ($p < 0.01$), demonstrating the tight coupling between distributed computing theory and systems implementation. ML research exhibits a strong correlation with publication peaks in July ($\rho = 0.88$, $p < 0.01$), indicating a conference-driven publication cycle in the machine learning community.

Conversely, negative correlations reveal countercyclical research patterns within the field. Theory and implementation papers show moderate negative correlation ($\rho = -0.42$, $p < 0.10$), suggesting alternating focus between theoretical advances and practical deployments. Security research reveals a weak negative correlation with ML publications ($\rho = -0.31$, not significant), potentially indicating resource allocation trade-offs or differing research priorities across subcommunities.

Time-series analysis further identifies critical lead-lag relationships that reflect the natural progression of scientific knowledge. Theory papers consistently lead systems implementations by approximately 1-2 months, demonstrating the typical timeline from conceptual development to practical realisation. Survey papers lag primary research by 2-3 months, representing the time required for comprehensive literature synthesis and analysis. Security research responds to ML advances with a one-month delay, indicating a reactive approach to addressing emerging vulnerabilities in machine learning systems. These temporal dependencies offer insights into the knowledge transfer dynamics within edge computing, revealing how innovations propagate through various research communities and evolve from theoretical concepts to practical implementations and comprehensive reviews.

6.5. Anomaly detection and special events

Statistical analysis identifies several temporal anomalies within the publication patterns. Notably, networking papers (cs.NI) peaked unusually in August rather than following the typical July surge observed in other categories, while theoretical papers showed an unexpected concentration in early October, deviating from the otherwise consistent monthly distribution patterns.

Investigation suggests several explanatory factors for these temporal variations. The July surge across most categories aligns with major summer AI conference deadlines, including ICML, NeurIPS, and AAAI submission windows, driving researchers to finalise and submit their work. The academic calendar plays a significant role, with the summer months representing peak research productivity for many academics, as they are free from teaching obligations. Meanwhile, the October theoretical paper cluster may reflect the preparation for the fall semester and the presentation of summer research outcomes. Funding cycle considerations potentially influence the October activity spike, as many institutions operate on fiscal years ending in September, prompting researchers to demonstrate productivity for grant renewals and new funding applications. Additionally, coordinated multi-institution collaborative projects may contribute to synchronised paper releases, particularly evident in the systems and distributed computing categories, where large-scale experiments require extensive coordination. These temporal patterns reveal the complex interplay between institutional constraints, conference schedules, and collaborative dynamics that shape the rhythm of edge computing research dissemination.

6.6. Forecasting future trends

Linear extrapolation and time-series models provide short-term forecasts:

- December 2025: Projected 380-420 papers (recovery from November).
- Q1 2026: Expected 1100-1300 papers (assuming pattern continuation).
- Annual 2026: Projected 5200-5800 papers (assuming 14.73% growth rate).

These projections assume the continuation of observed patterns but should be interpreted cautiously, given the high volatility.

Forecast uncertainty increases with projection horizon:

- 1-month forecast: ±15% (95% CI);
- 3-month forecast: ±35% (95% CI);
- 6-month forecast: ±60% (95% CI).

The widening confidence intervals reflect compounding uncertainty from high volatility and potential structural changes in publication patterns.

Fitting various growth models yields insights into long-term trends:

- Linear model: 67 papers/month increase (R^2 = 0.42).
- Exponential model: 14.73% monthly growth (R^2 = 0.38).
- Logistic model: Saturation at ∼600 papers/month (R^2 = 0.45).

The similar fit quality across models suggests the time series is too short to determine the underlying growth pattern definitively. However, the logistic model's saturation prediction aligns with typical research field maturation patterns.

6.7. Implications for research strategy

The temporal analysis reveals strategic considerations for research dissemination in edge computing. July and September-October emerge as high-competition periods, characterised by maximum submission volumes, presenting a trade-off between increased competition and enhanced visibility for research outputs. Conversely, June and August offer strategic windows with reduced competition for attention, which can be potentially advantageous for papers requiring focused community engagement. Aligning preprint releases with major conference cycles can maximise impact, leveraging the heightened attention during submission and review periods.

These patterns inform optimal research planning strategies across the annual cycle. Theoretical work benefits from Q4 publication when competition is lower, allowing for deeper engagement with complex conceptual contributions. Breakthrough results achieve maximum visibility when targeted for Q3 release, coinciding with peak community activity and summer conference preparations. Collaborative projects should coordinate around the observed July and October peaks, capitalising on synchronised multi-institutional efforts and funding cycle completions. Survey papers strategically scheduled for November-December can effectively capture and synthesise annual developments, providing comprehensive year-end perspectives when primary research activity typically declines. Understanding these temporal dynamics enables researchers to optimise their publication strategies, balancing visibility, competition, and community engagement to maximise research impact within the rapidly evolving edge computing landscape.

6.8. Comparative temporal analysis

Comparing edge computing temporal patterns with related fields reveals unique characteristics.

- Edge computing exhibits higher volatility with a CV of 61.81% compared to approximately 35% for general CS.
- The field demonstrates stronger seasonality through more pronounced conference effects than established fields.
- A monthly growth rate of 14.73% exceeds mature fields (approximately 5-8%) but trails emerging areas like quantum computing (approximately 20-25%).

These comparisons position edge computing as a rapidly growing yet increasingly mature field, transitioning from an explosive emergence to a sustained expansion.

While complete geographic data is limited, temporal patterns suggest global distribution:

- Publication times distributed across 24-hour periods.
- No significant day-of-week effects ($\chi^2 = 8.34$, $p = 0.21$).
- Suggests a global research community across time zones.

6.9. Summary and future monitoring

Temporal analysis reveals edge computing as a dynamic field characterised by strong seasonal patterns, conference-driven publication cycles, and a sustained growth trajectory. Key findings include the following.

1. Dramatic volatility is shown by a 61.81% standard deviation, indicating highly variable publication rates.

2. Conference influence results in a clear correlation between major conference deadlines and publication surges.
3. Sustained growth continues despite volatility, with a consistent positive trend averaging 14.73% monthly.
4. Category-specific patterns are evident as different research areas follow distinct temporal rhythms.
5. Predictable cycles form a bimodal annual distribution with July and September-October peaks.

These temporal patterns provide intelligence for researchers planning publication strategies, funding agencies allocating resources, and institutions evaluating research productivity. Continued monitoring will reveal whether current patterns represent stable equilibria or transitional dynamics in a rapidly evolving field.

7. Network analysis and research communities

This section presents network analysis of edge computing research, examining collaboration patterns, community structures, and knowledge flow through co-authorship and keyword networks. An analysis of authors and their collaborative relationships reveals the social and intellectual structure underlying edge computing research, identifying key researchers, bridging collaborations, and emergent research communities.

7.1. Co-authorship network construction and properties

The co-authorship network was constructed with authors as nodes and collaborative relationships as edges. Each edge represents at least one joint publication between two authors, with edge weights corresponding to the number of co-authored papers. The resulting network encompasses:

- Nodes: 8683 unique authors.
- Edges: 46629 co-authorship relationships.
- Network density: 0.00124 (0.124% of possible connections).
- Connected components: 1403 separate components.
- Largest component: 1779 authors (20.5% of all authors).

The low density (0.124%) indicates a sparse network typical of scientific collaboration, where researchers collaborate with a small fraction of the total community. The presence of 1403 components suggests a fragmented landscape with many isolated research groups, although the largest component, containing 20.5% of authors, indicates the emergence of a core, connected community.

The degree distribution reveals heterogeneous collaboration patterns:

- Average degree: 10.74 (each author collaborates with $\sim$11 others on average).
- Maximum degree: 137 (most connected researcher).
- Minimum degree: 0 (isolated authors).
- Degree variance: 256.3 (high heterogeneity).

The maximum degree of 137 indicates the presence of “super-connectors” who collaborate extensively, likely senior researchers or those coordinating large consortia. The average degree of 10.74 suggests a moderate level of collaborative intensity, which is sufficient for knowledge exchange while maintaining a research focus.

The network exhibits remarkably high clustering:

- Average clustering coefficient: 0.932.

- Random network expected: ∼0.001.
- Small-world quotient: >900.

The clustering coefficient of 0.932 indicates that 93.2% of possible triangles in the network are closed, meaning if author A collaborates with B and C, there is a 93.2% probability that B and C also collaborate. This extraordinarily high clustering suggests:

1. Strong tendency toward research group formation.
2. Collaborative teams rather than dyadic partnerships.
3. Efficient knowledge circulation within groups.
4. Potential barriers to cross-group collaboration.

7.2. Key researchers and network centrality

Analysis of collaboration frequency identifies the most active collaborative hubs:

1. Dusit Niyato: Central to 8 of the top 10 collaboration pairs
   - 25 papers total
   - Primary collaborators: Ruichen Zhang (8), Jiacheng Wang (7), Dong In Kim (7)
   - Bridge between multiple research communities
2. Dong In Kim: Key node in wireless edge computing
   - 10 papers total.
   - Strong connections to both the Niyato and Wang groups
   - Focus on the communication aspects of edge computing
3. Ruichen Zhang: Emerging collaborative hub
   - 10 papers total
   - Central to federated learning research cluster
   - Strong connection with Yinqiu Liu (6 joint papers)

The top collaboration pairs reveal sustained research partnerships (table 13).

Table 13
Top 10 collaboration pairs in edge computing research.

Rank	Collaboration pair	Joint papers
1	Dusit Niyato – Ruichen Zhang	8
2	Dusit Niyato – Jiacheng Wang	7
3	Dong In Kim – Dusit Niyato	7
4	Ruichen Zhang – Yinqiu Liu	6
5	Dusit Niyato – Zehui Xiong	6
6	Dusit Niyato – Yinqiu Liu	5
7	Dusit Niyato – Geng Sun	5
8	Dong In Kim – Jiacheng Wang	5
9	Dusit Niyato – Jiawen Kang	4
10	Dusit Niyato – Zhu Han	4

These sustained collaborations (4-8 joint papers in 6 months) indicate established research programs rather than opportunistic partnerships, suggesting deep integration of research agendas and potentially shared funding or institutional support.

While specific betweenness centrality values require detailed computation, network structure analysis identifies several bridge researchers who connect otherwise disconnected communities within the edge computing landscape. These bridging roles manifest in multiple dimensions, fundamentally shaping the flow of knowledge across the field. Cross-domain bridges emerge as researchers who publish across both systems (cs.DC) and machine learning (cs.LG) categories, effectively translating concepts and methodologies between traditionally separate technical communities. Geographic bridges appear as international collaborators who connect regional research clusters, facilitating global knowledge exchange and preventing geographic isolation of innovations. Methodological bridges span the theoretical-experimental divide, with authors contributing to both foundational theory and practical implementations, ensuring that abstract concepts translate into deployable solutions. These bridge researchers play crucial roles in knowledge transfer, preventing community isolation, and facilitating interdisciplinary innovation.

7.3. Community detection and research clusters

The Louvain algorithm community detection reveals a highly modular network structure, with 1074 detected communities and a modularity score of 0.847, indicating strong community boundaries characterised by dense intra-community connections and sparse inter-community links. The remarkably uniform community size distribution, ranging from 11 to 14 members with an average of 13.2 and minimal variance of 1.8, suggests natural limits to effective research group size, possibly reflecting optimal team dynamics or resource constraints in edge computing research.

Mapping communities to research themes reveals four primary specialised clusters. The federated learning community comprises approximately 156 members across 12 communities, with core researchers including Niyato, Zhang, and Liu, who focus on privacy-preserving distributed learning and have produced 87 papers. The edge AI/vision community, the largest cluster with 234 members across 18 communities, focuses on computer vision at the edge, particularly real-time object detection, resulting in 143 publications. The 5G/6G edge community includes 189 members across 14 communities, emphasising next-generation wireless integration with strong industry collaboration indicators and 98 papers. The IoT systems community comprises 167 members across 13 communities, with a focus on sensor networks and practical IoT platform deployments, as evidenced by 76 publications.

Analysis of cross-community collaborations reveals limited but strategic connections, with only 3.2% of collaborations spanning communities, most of which are facilitated by bridge researchers. The strongest cross-pollination occurs between the federated learning and edge AI communities, with 27 shared papers, while the theory and systems communities show minimal interaction, with merely eight shared publications. This limited inter-community collaboration pattern suggests significant opportunities for increased cross-fertilisation of ideas and potential breakthrough innovations at community boundaries. The high modularity, combined with sparse inter-community connections, indicates that while edge computing has developed strong, specialised subcommunities, the field may benefit from deliberate efforts to foster interdisciplinary collaboration and knowledge transfer across these research clusters.

7.4. Keyword co-occurrence network

The keyword co-occurrence network provides insight into conceptual relationships:

- Nodes: 4664 unique keywords.
- Edges: 8067 co-occurrence relationships..
- Network Density: 0.000742 (0.074%)
- Average degree: 3.46.
- Clusters: 3325 keyword clusters.
- Largest cluster: 1336 keywords (28.6%).

The keyword network is even sparser than the co-authorship network, reflecting the diverse vocabulary and specialised terminology in edge computing research.

The most connected keywords reveal conceptual hubs (table 14).

Table 14
Top 10 most connected keywords in edge computing research.

Rank	Keyword	Degree	Betweenness centrality
1	edge	528	0.342
2	model	364	0.287
3	models	352	0.279
4	data	333	0.265
5	graph	305	0.231
6	learning	271	0.198
7	framework	246	0.176
8	performance	184	0.134
9	networks	182	0.132
10	systems	178	0.128

The dominance of "edge" (degree 528) as the most connected keyword validates our corpus selection. The high connectivity of "model/models" (716 combined) reflects the field's focus on algorithmic contributions, while "data" and "learning" highlight the AI/ML emphasis.

Analysing keywords with high betweenness centrality but a moderate degree identifies conceptual bridges:

- "federated" – bridges privacy, learning, and distributed concepts;
- "quantum" – connects optimisation, computing, and future technologies;
- "sustainable" – links energy, optimisation, and environmental concerns;
- "latency" – bridges performance, networking, and application domains.

These bridge concepts often represent emerging research directions or interdisciplinary connection points.

7.5. Temporal evolution of network structure

Tracking network evolution across the study period reveals dynamic patterns:

- June: Baseline network with 892 authors, 2341 collaborations.
- July: Explosive growth to 2156 authors, 7823 collaborations.
- August: Consolidation phase, density increases by 18%.
- September: New community formation, 127 new components.
- October: Integration phase, component count decreases by 8%.
- November: Partial data shows continued densification.

The July explosion in network size (141% author increase) corresponds to the publication surge, but the disproportionate increase in collaboration (234%) suggests coordinated group publications rather than independent submissions.

New authors show strong preferential attachment patterns:

- 73% of new authors collaborate with existing high-degree nodes.
- Average degree of attachment targets: 24.3 (vs. network average 10.74).
- Suggests "rich-get-richer" dynamics in collaboration formation.

This preferential attachment accelerates the emergence of super-connectors and may contribute to increasing centralisation over time.

7.6. Geographic and institutional network patterns

While complete geographic data is limited, name analysis and available affiliations suggest:

- Asia-Pacific hub centred around Singapore (NTU), Korea (KAIST), China.
- North American cluster: US universities and industry labs.
- European network distributed across multiple countries, less centralised.
- Intercontinental bridges: $\sim$15% of collaborations span continents.

The presence of strong regional clusters with strategic intercontinental connections suggests a "glocal" research structure – globally connected but locally intensive.

Network analysis reveals patterns suggesting industry involvement:

- Unusually large teams (>10 authors) often indicate industry collaboration.
- Papers with "system", "implementation", or "deployment" keywords show 42% higher average author count.
- Potential industry-affiliated authors based on collaboration patterns: $\sim$8% of network.

7.7. Network resilience and vulnerability

The network's resilience to node removal reveals structural vulnerabilities:

- Under random removal, the network remains connected until 67% of nodes are removed.
- Targeted removal of the top 5% of high-degree node fragments in the network.
- There are 47 critical author nodes whose removal would isolate greater than 100 researchers.

This vulnerability to targeted removal of key researchers highlights the importance of succession planning and knowledge distribution in research communities.

The following network metrics are relevant to information diffusion:

- The average path length is 5.3 in the largest component.
- The diameter is 14, representing the maximum separation in the largest component.
- Information centralisation stands at 0.73, indicating high concentration.

The relatively short average path length (5.3) suggests efficient information flow within connected components, supporting rapid dissemination of innovations. However, high centralisation (0.73) indicates dependence on key nodes for information brokerage.

7.8. Comparison with other scientific networks

Edge computing networks show unique characteristics compared to other fields:

- Higher clustering: 0.932 vs. typical 0.6-0.7 in computer science.
- Smaller communities: average 13.2 vs. 20-30 in mature fields.
- Rapid growth: 14.73% monthly vs. 2-3% in established fields.
- Strong modularity: 0.847 vs. 0.3-0.5 typical.

These distinctions suggest that edge computing is in a rapid community formation phase, with a strong group identity but limited cross-pollination – characteristic of emerging interdisciplinary fields.

7.9. Summary and future network evolution predictions

Network analysis reveals edge computing as a rapidly evolving research ecosystem characterised by strong community structure, emerging collaborative hubs, and significant growth potential. The field exhibits a fragmented yet growing structure, with 1403 components indicating many isolated research groups, though the 20.5% giant component demonstrates emerging integration trends. The exceptionally high clustering coefficient of 0.932 reveals strong team-based research practices, while a small set of super-connectors, notably Dusit Niyato and other bridge researchers, facilitate inter-community knowledge flow. The strong modularity score of 0.847 across 1074 communities indicates highly specialised research clusters, each focusing on distinct aspects of edge computing challenges.

Conceptual integration analysis reveals "edge", "model", and "data" as central organising concepts, reflecting the field's core focus on distributed computational models and data processing. Growth dynamics exhibit preferential attachment patterns that accelerate centralisation, potentially creating influential research hubs while maintaining diverse peripheral contributions. These network characteristics position edge computing for continued expansion while highlighting critical needs for strategic interventions. The current structure presents substantial opportunities for cross-community collaboration, particularly at the boundaries between theoretical and applied research. However, preventing excessive fragmentation and ensuring sustainable community growth will require deliberate efforts to foster interdisciplinary connections and maintain balanced development across specialised subcommunities.

Based on current network dynamics, we project a trajectory of progressive consolidation and integration within the edge computing research ecosystem. In the short term (6 months), we anticipate continued densification with a 20-30% increase in average degree as existing researchers strengthen collaborative ties. The medium term (1-2 years) should witness community consolidation, with the current 1074 communities reducing to 500-700 as related groups merge and integrate. Long-term projections (3-5 years) suggest the emergence of 3-5 super-communities that will dominate the research landscape, accompanied by maturation indicators, including a decrease in modularity and an increase in average path length. These predictions assume continued field growth and increasing interdisciplinary integration, reflecting natural evolution patterns observed in maturing scientific disciplines.

8. Discussion: research gaps, opportunities, and implications

This section synthesises findings from bibliometric, thematic, temporal, and network analyses, identifying critical research gaps, emerging opportunities, and implications for the broader research ecosystem.

8.1. Synthesis of major findings

Edge computing serves as a convergence point where multiple technological and methodological streams intersect. The dominance of AI/ML themes (91.25% of papers) combined with networking focus (53.25%) and energy efficiency concerns (30.6%) shows that edge computing has evolved beyond mere computational offloading to become a platform for interdisciplinary innovation.

Technological convergence is evident as AI integration with distributed systems marks a shift from edge computing as infrastructure to edge intelligence as autonomous decision-making capability. The 1825 papers addressing AI/ML themes demonstrate the reimagining of algorithms for resource-constrained, distributed environments, rather than simply adopting existing techniques.

Methodological convergence appears in the balanced distribution across research types – systems (28.2%), machine learning (20.0%), and theory (19.2%) – which indicates healthy methodological diversity. The field has matured beyond proof-of-concept demonstrations to encompass rigorous theoretical foundations, practical implementations, and empirical validations.

Community convergence is highlighted by the network analysis, revealing 1074 distinct communities with high modularity (0.847), paradoxically coupled with strong clustering (0.932). This suggests specialised expertise clusters maintaining internal cohesion while slowly building bridges. Key connectors, such as Dusit Niyato, which appear in 8 of the top 10 collaboration pairs, facilitate cross-community knowledge flow.

Temporal patterns reveal a field in transition from explosive growth to structured expansion.

Volatility acts as a sign of growing pains, where the 61.81% publication volatility reflects conference-driven publication cycles rather than fundamental instability.

A sustained growth trajectory is visible in the 14.73% average monthly growth rate, positioning edge computing between mature fields (5-8% growth) and emerging areas like quantum computing (20-25% growth). This intermediate position suggests a field transitioning from emergence to establishment, maintaining dynamism while developing institutional structures.

The correlation between research types and temporal patterns indicates seasonal specialisation. The fact that ML papers are concentrated in Q3 (64.5%) while theory papers show stability (CV 38.7%) indicates a sophisticated division of labour where different communities optimise publication strategies around field-specific conferences and review cycles.

8.2. Critical research gaps

Perhaps the most striking finding is the severe underrepresentation of security research, with only 68 papers (3.4%) explicitly focusing on security, despite edge computing's inherent vulnerabilities stemming from distributed attack surfaces. This gap becomes more pronounced considering that edge devices operate in physically unsecured environments, federated learning and distributed AI introduce novel attack vectors, privacy preservation is frequently claimed but rarely rigorously proven, and IoT integration multiplies potential vulnerabilities. The security gap likely stems from the field's emphasis on functionality over protection, a common pattern in emerging technologies. As edge computing moves toward production deployments, this gap represents both a critical vulnerability and a research opportunity.

Analysis reveals a concerning lack of real-world evaluation, with only 167 papers (8.35%) focusing on applications and case studies, and a mere 27 papers (1.35%) comprising surveys. The prevalence of simulation-based studies without real deployment validation and the absence of standardised benchmarks for cross-study comparison undermine the field's ability to demonstrate practical impact. This evaluation gap may explain the relatively low industry participation, estimated at 8% of the network based on collaboration patterns. Without rigorous real-world validation, edge computing risks developing elegant solutions to problems that are incorrectly specified.

While energy efficiency is mentioned in 30.6% of papers, comprehensive sustainability considerations remain largely absent. Lifecycle assessment of edge infrastructure and sustainable manufacturing considerations receive no coverage, while e-waste from edge device proliferation is addressed in only two papers, and a comparative analysis of carbon footprints appears in just three papers. This gap is particularly problematic given edge computing's promise of environmental benefits through reduced data transmission; without holistic sustainability analysis, the field cannot substantiate its green computing claims.

The analysis reveals minimal attention to human factors, with only 11 papers addressing user experience in edge systems, four papers on ethical frameworks for edge AI, one paper on the implications of the digital divide, and no papers on accessibility considerations. This human dimension void suggests that the field remains technology-driven rather than user-centred, potentially limiting adoption and raising ethical concerns about autonomous edge decision-making without established human oversight frameworks.

8.3. Emerging opportunities and future directions

The emergence of quantum edge computing, represented by 23 papers with an 89% growth rate, signals a potential paradigm shift in the field. This nascent integration offers ground-floor opportunities for researchers to define fundamental architectures and protocols through quantum optimisation for NP-hard edge resource allocation problems, quantum machine learning models designed explicitly for edge constraints, hybrid classical-quantum architectures that leverage quantum advantage selectively, and quantum-secure communication protocols for edge networks.

Neuromorphic computing, with 34 papers showing 54% growth, offers compelling solutions to edge computing's energy challenges. The intersection of neuromorphic hardware with edge computing could enable always-on intelligence that was previously impossible due to power constraints, through event-driven processing that aligns with the edge's sporadic workload patterns, resulting in orders-of-magnitude improvements in energy efficiency for specific tasks. This approach also facilitates natural integration with sensory data processing at the edge and brain-inspired learning, eliminating the computational demands of backpropagation.

The exploration of 6G-native edge computing, already represented by 47 papers with 76% growth, presents abundant opportunities for researchers to shape standards and architectures for the next decade. This integration promises terahertz communication, enabling unprecedented edge bandwidth, AI-native network protocols explicitly designed for edge intelligence, satellite-terrestrial edge integration for global coverage, and holographic communications that require sophisticated edge processing capabilities. Researchers entering this space now have the unique opportunity to influence the fundamental design principles that will govern edge computing infrastructure in the 6G era.

8.4. Methodological reflections and limitations

Our analysis introduces several methodological contributions to the study of edge computing research. The multidimensional integration, combining bibliometric, thematic, temporal, and network analyses, provides richer insights than single-method approaches typically afford. Tracking the monthly evolution reveals patterns that are invisible in annual aggregations, while connecting network communities to research themes through community-topic mapping identifies underlying knowledge structures. Furthermore, our systematic gap analysis framework, which spans multiple dimensions, provides actionable research directions for the field.

Several limitations constrain our findings. Focusing solely on arXiv excludes peer-reviewed publications, potentially biasing the analysis toward preliminary or rapid-publication research, while industry research, often proprietary, remains invisible. The 6-month temporal window may not capture annual cycles completely, and recent emergence patterns might reflect temporary trends rather than sustained directions. Author disambiguation presents challenges, as name-based matching without unique identifiers, such as ORCID, may conflate different researchers or fragment the contributions of a single researcher. Additionally, analysing only English-language papers excludes potentially significant research from non-English-speaking regions, particularly relevant given the global nature of edge computing. Finally, while our automated classification into research types and themes provides systematic coverage, it may miss nuanced contributions or novel approaches that do not fit predetermined categories.

8.5. Theoretical implications

Findings suggest that edge computing is developing its own theoretical framework, distinct from distributed systems and cloud computing. Resource-constrained intelligence is becoming increasingly central, as theoretical work focuses more on the fundamental limits of intelligence under strict resource constraints. The field is characterised by emerging theories of distributed autonomy that balance local decision-making with global coordination, alongside approximate computing paradigms that provide theoretical frameworks for trading accuracy for efficiency in edge contexts. Additionally, the development of complexity measures specific to edge environments indicates the emergence of edge-native theoretical foundations. This theoretical development, as evidenced by 384 theory papers (19.2%), suggests that the field has matured beyond engineering solutions toward a scientific foundation.

Based on growth patterns and community evolution, we propose a lifecycle model:

1. The current phase of expansion is characterised by high growth (14.73%), community proliferation (1074 communities), and thematic exploration.
2. The next phase of consolidation is predicted to bring a reduction in communities, the emergence of standardisation, and decreased volatility.
3. The maturation phase is expected to involve integration with broader computing paradigms, stable growth rates, and established theoretical foundations.

This lifecycle model suggests strategic timing for different research approaches – exploratory work suits the current phase, while standardisation efforts should begin preparing for consolidation.

8.6. Societal and ethical considerations

Edge computing presents a democratisation paradox, promising to bring AI capabilities to resource-constrained devices while analysis reveals potential contradictions. Despite promises of democratised intelligence, research remains concentrated among a few key players, with the top 15 authors producing 6.35% of papers. Limited geographic diversity in research leadership, minimal attention to accessibility and digital divide issues, and industry participation barriers, as suggested by low collaboration rates, indicate that without deliberate intervention, edge computing may replicate rather than resolve existing technological inequalities.

The sustainability gap reveals another critical contradiction: edge computing markets itself as a green technology, yet lacks a comprehensive environmental assessment. While the field claims benefits through reduced data transmission, localised processing, and optimised resource usage, it fails to account for costs, including device proliferation, shortened replacement cycles, and manufacturing impact. The absence of lifecycle assessments, comparative carbon footprints, and integration with the circular economy undermines the environmental value proposition of edge computing, risking accusations of greenwashing.

8.7. Strategic recommendations

Based on identified gaps and opportunities, immediate priorities should focus on four critical areas. First, developing comprehensive security frameworks that address edge-specific vulnerabilities is essential, given the distributed nature of edge systems. Second, creating benchmarks and testbeds for realistic performance assessment will enable the development of meaningful evaluation standards across diverse edge environments. Third, establishing lifecycle assessment methodologies for edge systems addresses the critical gap in sustainability metrics. Finally, developing frameworks for responsible edge AI deployment provides necessary ethical guidelines for this rapidly expanding field.

Over the medium term of 2-3 years, the field should prioritise cross-community integration through formal mechanisms for knowledge exchange between specialised communities, addressing the current fragmentation evident in our analysis. Industry-academia bridges through programs that facilitate the practical validation of research contributions are crucial for translating theoretical advances into real-world applications. International coordination in standards development becomes essential to prevent fragmentation as the field matures. Additionally, developing a unified theoretical framework for edge computing will provide the foundational principles necessary for systematic advancement beyond ad-hoc engineering solutions.

The long-term vision emerging from our analysis positions edge computing as seamlessly integrated with cloud and quantum computing in a computational continuum, moving beyond current isolated approaches. This vision encompasses environmental sustainability through the implementation of designed-in circular economy principles, addressing current sustainability gaps. Ethical alignment with human values through embedded governance mechanisms ensures responsible development, while universal accessibility positions edge computing as a technology that bridges rather than widens digital divides. This comprehensive vision requires a coordinated effort across research communities, industry partners, and policymakers to realise edge computing's transformative potential while mitigating identified risks.

8.8. Conclusions

Analysis reveals edge computing as a rapidly evolving field at a critical juncture. The convergence of AI, networking, and systems research has created a vibrant ecosystem of 8683 researchers producing innovative solutions to fundamental challenges. However, gaps in security, evaluation, sustainability, and human factors threaten to limit the field's impact and raise ethical concerns.

The high growth rate (14.73% monthly) coupled with a strong community structure (1074 communities, 0.847 modularity) suggests a field with tremendous potential, but one that requires strategic coordination to prevent fragmentation and ensure beneficial development. The emergence of quantum and neuromorphic edge computing, while nascent, hints at transformative possibilities that could redefine computing paradigms.

For edge computing to fulfil its promise of ubiquitous intelligence, the research community must address identified gaps while maintaining innovative momentum. This requires deliberate efforts to build bridges between communities, integrate industry perspectives, and centre human and environmental considerations in technical development.

The path forward requires not only technical innovation but also thoughtful consideration of the role of edge computing in society. As we stand at this inflexion point, decisions made by researchers, funders, practitioners, and policymakers will shape whether edge computing becomes a democratising force for distributed intelligence or another layer of technological complexity serving limited interests.

9. Reflections on computing's distributed future: editorial to JEC Volume 4 Issue 2 (2025)

As we stand at the closing of 2025, edge computing finds itself at a remarkable inflexion point. Our analysis of arXiv papers reveals not merely a field in growth, but one undergoing fundamental transformation. The 14.73% monthly expansion rate tells only part of the story; beneath these numbers lies a profound reimagining of how we conceive, design, and deploy computational intelligence. This editorial offers reflections on what our data reveals – and perhaps more importantly, what it conceals – about edge computing's trajectory.

9.1. The year 2025: a watershed moment

The patterns emerging from our analysis suggest 2025 will be remembered as the year edge computing transitioned from promise to paradigm. The convergence of 91.25% AI/ML focus with distributed systems expertise signals more than technological integration; it represents a fundamental shift in how we approach computational challenges. We are witnessing the birth of *ambient intelligence* – computing so seamlessly distributed that it becomes environmental rather than instrumental.

However, this transition brings profound challenges that our quantitative analysis only hints at. The security gap – a mere 3.4% of papers – is not just a research oversight, but a symptom of deeper tensions between innovation velocity and infrastructure robustness. The field races forward, deploying intelligence at the edge while leaving fundamental vulnerabilities unaddressed. This is not carelessness but rather the inevitable consequence of a field driven by possibility rather than precaution.

The temporal volatility we observed (61.81% standard deviation) reflects more than conference cycles; it reveals a field struggling to find its rhythm. This pattern suggests a community caught between academic traditions and industrial urgencies, trying to serve both masters while establishing its own identity. This tension is productive – it forces edge computing to remain grounded in practical needs while reaching for theoretical elegance.

9.2. Critical observations beyond the data

Three observations emerge from reading between the lines of our analysis:

1. *The collaboration paradox.* While we celebrate 8683 authors and 46629 collaborations, the network's fragmentation into 1074 communities with only 3.2% cross-community collaboration reveals intellectual silos that may ultimately limit the field's potential. The high clustering coefficient (0.932) suggests a preference for familiar collaborations over the discomfort of interdisciplinary exploration. Edge computing needs more intellectual bridge-builders willing to traverse community boundaries.
2. *The evaluation crisis runs deeper than numbers suggest.* The 8.35% of papers focusing on real-world applications masks a more troubling pattern where even these papers often rely on limited deployments under controlled conditions. We have become skilled at demonstrating what edge computing can do, but remain largely ignorant of what it actually achieves when deployed at scale in complex, unpredictable environments. This gap between potential and practice threatens the field's credibility.
3. *The human absence is deafening.* With only 11 papers addressing user experience and zero on accessibility, edge computing reveals itself as a field designed by technologists for technology. However, edge computing, unlike cloud or traditional distributed systems, operates in human-centric spaces – such as homes, vehicles, and bodies. This anthropocentric blindness risks creating systems that are technically sophisticated but humanly alienating.

9.3. The papers that point the way forward

This issue of the Journal of Edge Computing features four contributions that, when viewed through the lens of our analysis, illuminate critical paths forward.

Chaban, Manziuk and Radiuk [3] address the knowledge distillation challenge with their adaptive multi-teacher framework, directly confronting the model compression needs identified in our Topic 5 (memory, training, compression). Their work exemplifies the sophisticated optimisation required to bring complex AI models to resource-constrained edges, moving beyond simple compression to intelligent knowledge transfer. This paper represents the maturation of edge AI from deployment tricks to principled methodologies.

The survey by Semerikov et al. [16] on deploying large language models in resource-constrained environments arrives at a crucial moment. Our analysis revealed 127 papers on edge LLMs, which exhibited the highest growth rate, yet lacked systematic evaluation. This survey fills a critical gap, providing the first framework for understanding the challenges of LLM deployment at the edge. Their identification of quantisation, pruning, and partitioning strategies offers practical roadmaps for what our thematic analysis revealed as the fastest-growing research direction.

Njoroge, Kibuku and Mugoye [12] demonstrate precisely the kind of real-world application our analysis found lacking. Their comparative study of EfficientNetV2 and MobileNetV2 for crop disease classification addresses the 8.35% gap in application-focused research while contributing to the underrepresented agricultural domain (12 papers in our corpus). The statistical validation they provide sets a methodological standard the field desperately needs – rigorous, comparative, and grounded in practical constraints.

Finally, Riabko et al. [15] push boundaries with acoustic Doppler localisation in 3D space, representing the kind of innovative sensing and processing that distinguishes edge computing from traditional paradigms. Their retardation correction algorithms demonstrate the sophisticated signal processing required when computational resources are distributed across space and time, contributing to the theoretical foundations that our analysis revealed comprise 19.2% of current research.

These four papers collectively demonstrate the evolution of edge computing from a theoretical possibility to a practical necessity, each addressing gaps identified in our analysis while pushing the field in new directions.

9.4. Challenges that demand attention

Our analysis reveals three critical challenges that the edge computing community must confront. The standardisation imperative arises from the current fragmentation evident in 1074 distinct communities, each with no unified evaluation frameworks, which risks dividing the field into incompatible islands of innovation. The field requires not rigid standards that stifle creativity, but flexible frameworks that enable comparison and integration. The absence of benchmarks means we cannot meaningfully compare the papers in our corpus – each exists in its own evaluative universe, undermining collective progress.

The sustainability reckoning represents perhaps the most damaging credibility gap. Edge computing's environmental claims remain unsubstantiated, with zero lifecycle assessments among 2000 papers constituting not just a gap but an indictment. As we deploy billions of edge devices, the field must account for their full environmental impact, including manufacturing, operation, and disposal. The field's credibility depends on honest environmental accounting rather than green marketing; yet, current research patterns suggest a continued avoidance of these difficult questions.

Ultimately, the security debt accumulating in edge computing poses a significant threat to future catastrophes. Every day of delayed security focus multiplies future costs, with only 3.4% of security attention amid 91.25% AI deployment creating a dangerous asymmetry. Edge devices operate in hostile environments – physical and digital – yet current designs treat them as if they existed in protected data centres. This security debt will come due, possibly catastrophically, unless the field prioritises protection alongside innovation. These challenges are not merely technical hurdles, but fundamental threats to the viability and social license to operate of edge computing.

9.5. A research agenda for the next frontier

Based on our analysis, we propose five research priorities that could define edge computing's next phase:

1. *Federated evaluation frameworks*: Develop privacy-preserving methods for aggregating performance data from deployed edge systems, enabling real-world validation without compromising security or privacy.
2. *Anthropocentric edge design*: Create design methodologies that begin with human needs and constraints rather than adding user interfaces to completed systems. This includes considerations for accessibility, cultural sensitivity, and cognitive load.
3. *Sustainable edge architectures*: Design systems with full lifecycle consideration – from rare earth extraction to e-waste processing. This includes exploring biodegradable electronics, energy harvesting, and principles of the circular economy.
4. *Security-first development:* Invert current practice by beginning with security requirements and building functionality within those constraints rather than adding security to functional systems.
5. *Theoretical foundations for distributed intelligence:* Develop mathematical frameworks for understanding intelligence emergence in distributed systems, moving beyond optimisation to understanding systemic behaviours.

9.6. The philosophical challenge

Edge computing presents a philosophical challenge that transcends technical considerations: *How do we maintain human agency in a world of ambient intelligence?* As computing disappears into the environment, becoming invisible and autonomous, we risk creating a technological unconscious that shapes behaviour without awareness or consent.

The 1825 papers on AI and machine learning in our corpus demonstrate tremendous capability, but capability without wisdom is dangerous. Edge computing distributes not just computation but decision-making, embedding values and biases into infrastructure itself. The absence of ethical frameworks (4 papers) and human considerations (11 papers) suggests we are building this future blindly.

However, we remain optimistic. The vibrant community revealed by our analysis – 8683 researchers across 124 categories – possesses the diversity and creativity to address these challenges. The emerging paradigms we identified – quantum-edge, neuromorphic computing, and 6G integration – offer not just technical advances but also opportunities to reimagine computing's relationship with humanity and the environment.

9.7. Concluding thoughts: the edge of tomorrow

As we conclude this analysis of the fresh arXiv papers representing the cutting edge of edge computing, we are struck by both the field's tremendous achievements and its unrealised potential. We stand at a threshold where edge computing could become the nervous system of a connected world, sensing, processing, and responding with unprecedented sophistication. Or it could become another layer of technological complexity that widens divides and accelerates unsustainable consumption.

The choice is ours – not as individuals but as a community. The patterns revealed in our analysis indicate a field capable of remarkable innovation, but one that requires deliberate coordination to achieve beneficial outcomes. The high modularity of our research communities need not become permanent isolation; the security gaps need not become vulnerabilities; the human absence need not become dehumanisation.

Edge computing in 2025 is not a mature field, looking back on achievements, but a young discipline looking forward to possibilities. The 14.73% monthly growth rate will not continue indefinitely, but what matters is not the growth rate itself, but rather the direction it takes. Will we build edge computing that empowers or surveils, that includes or excludes, that sustains or consumes?

The 2000 papers analysed here represent 2000 attempts to answer these questions, each contributing a piece to an emerging mosaic. Some pieces fit perfectly; others reveal gaps; still others suggest patterns we have not yet recognised. This is the nature of scientific progress – not linear advancement but exploratory wandering that occasionally crystallises into understanding.

As editor, I am privileged to witness this crystallisation occurring. The four papers in this issue exemplify the field's potential while acknowledging its challenges. They remind us that edge computing is not about moving computation to the network's periphery, but about reimagining the role of computation in human life. The edge is not a place but a principle – that intelligence should exist where it is needed, when it is needed, in forms that serve human and environmental flourishing.

This principle guides us toward a future where computing truly serves humanity rather than the reverse. The journey from today's 2000 papers to tomorrow's deployed systems will be challenging, requiring technical innovation, ethical reflection, and sustained collaboration. But the foundation laid by the researchers in our analysis – their creativity, dedication, and vision – gives me confidence that edge computing will fulfil its promise of distributed intelligence that enhances rather than replaces human capability.

The edge beckons not as an endpoint but as a beginning – of new computational paradigms, new forms of intelligence, and new relationships between humans and machines. Our analysis captures a moment in this journey, but the destination remains unwritten. That story will be authored not by any individual but by the community revealed in these pages – 8683 researchers and counting, each contributing to computing's distributed future.

May we write that future wisely.

Acknowledgments: I am deeply grateful to the Editor-in-Chief of the Journal of Edge Computing, Professor Tetiana A. Vakaliuk, for her gracious invitation to contribute this editorial, her invaluable support throughout the survey process, and the fruitful discussions that enriched this work. Additionally, I extend my sincere appreciation to the journal's authors, editorial board members, reviewers, and all contributors whose dedication and expertise continue to advance the field of edge computing.

Declaration on generative AI: The author acknowledges the use of generative AI technologies in the development of this work. Claude Code was utilised as a coding assistant to accelerate the development of the Python-based scraping and bibliometric analysis pipeline, specifically for optimising API rate-limiting and network graph construction. Claude Opus 4.1 was employed to assist in synthesising initial drafts of technical descriptions and structuring the thematic analysis sections based on the statistical outputs. Grammarly was used for final copy-editing, grammar verification, and stylistic refinement. The author emphasises that AI tools were used solely for augmentation and data processing. All conceptualisation, study design, interpretation of bibliometric trends, validation of code logic, and conclusions are the original work of the human author. All AI-generated text and code underwent rigorous manual review to ensure accuracy, and the author assumes full responsibility for the integrity of the published work.

Data availability statement: In accordance with FAIR principles, the datasets and software generated during the current study are available in the Edge of arXiv repository on GitHub at https://github.com/ssemerikov/arXivedge. This repository includes:

1. The raw and processed metadata of all 2025 arXiv papers analysed.
2. The complete Python source code for the scraping, analysis, and paper generation pipeline.
3. The generated bibliometric networks and LaTeX tables.

The raw data source is publicly available via the arXiv API (https://arXiv.org/help/api).

References

[1] Boboris, K., 2025. Attention Authors: Updated Practice for Review Articles and Position Papers in arXiv CS Category. Available from: https://blog.arxiv.org/2025/10/31.

[2] Burbano, J.S., Abdullah, A., Zhantileuov, E., Liyanage, M. and Schuster, R., 2025. Dynamic Edge Server Selection in Time-Varying Environments: A Reliability-Aware Predictive Approach. Available from: https://doi.org/10.48550/arXiv.2511.10146.

[3] Chaban, O., Manziuk, E. and Radiuk, P., 2025. Method of adaptive knowledge distillation from multi-teacher to student deep learning models. *Journal of Edge Computing*, 4(2), pp.159–178. Available from: https://doi.org/10.55056/jec.978.

[4] Chen, Y., Lu, J., Cao, S., Wang, W., Li, G. and Wen, G., 2025. FedCure: Mitigating Participation Bias in Semi-Asynchronous Federated Learning with Non-IID Data. Available from: https://doi.org/10.48550/arXiv.2511.10227.

[5] Cheong, C., Davis, G. and Choi, S., 2025. Weaver: Kronecker Product Approximations of Spatiotemporal Attention for Traffic Network Forecasting. Available from: https://doi.org/10.48550/arXiv.2511.08888.

[6] Guo, T., Zhao, S., Zhu, S. and Ma, C., 2025. SPEED-Q: Staged Processing with Enhanced Distillation towards Efficient Low-bit On-device VLM Quantization. Available from: https://doi.org/10.48550/arXiv.2511.08914.

[7] Ji, M., Souza, A.H. and Garg, V., 2025. On topological descriptors for graph products. Available from: https://doi.org/10.48550/arXiv.2511.08846.

[8] Jin, S., Jiang, Y., Liu, Y., Ma, T., Cao, D., Wei, L., Liu, X. and Zeng, X., 2025. DeepDR: an integrated deep-learning model web server for drug repositioning. Available from: https://doi.org/10.48550/arXiv.2511.08921.

[9] Konstantinidis, A.L., Papadopoulos, C. and Velissaris, G., 2025. Algorithms and Complexity of Hedge Cluster Deletion Problems. Available from: https://doi.org/10.48550/arXiv.2511.10202.

[10] Li, Z., Lin, C., Zheng, L., Wei, W.D., Liang, J. and Song, Q., 2025. GraphIF: Enhancing Multi-Turn Instruction Following for Large Language Models with Relation Graph Prompt. Available from: https://doi.org/10.48550/arXiv.2511.10051.

[11] Mehry, S. and Molaeinejad, M., 2025. The Homomorphism Submodule Graph. Available from: https://doi.org/10.48550/arXiv.2511.07837.

[12] Njoroge, T., Kibuku, R. and Mugoye, K., 2025. Comparative and edge-hybrid modeling of EfficientNetV2 and MobileNetV2 for multi-class crop disease classification with statistical validation. *Journal of Edge Computing*, 4(2), pp.234–262. Available from: https://doi.org/10.55056/jec.905.

[13] Onuchin, A., Sorokin, K., Beketov, M. and Tupikina, L., 2025. Iterative Ricci-Foster Curvature Flow with GMM-Based Edge Pruning: A Novel Approach to Community Detection. Available from: https://doi.org/10.48550/arXiv.2511.08919.

[14] Ren, Y., Liu, Y., Zhou, Y., Zheng, Z., Shang, L., Yang, F. and Wang, Z., 2025. Bridging the Initialization Gap: A Co-Optimization Framework for Mixed-Size Global Placement. Available from: https://doi.org/10.48550/arXiv.2511.10073.

[15] Riabko, A.V., Vakaliuk, T.A., Zaika, O.V. and Kukharchuk, R.P., 2025. Acoustic Doppler localisation and tracking in 3D space with retardation correction. *Journal of Edge Computing*, 4(2), pp.263–287. Available from: https://doi.org/10.55056/jec.1201.

[16] Semerikov, S.O., Vakaliuk, T.A., Kanevska, O.B., Ostroushko, O.A. and Kolhatin, A.O., 2025. Edge intelligence unleashed: a survey on deploying large language

models in resource-constrained environments. *Journal of Edge Computing*, 4(2), pp.179–233. Available from: https://doi.org/10.55056/jec.1000.

[17] Song, H. and Xu, H., 2025. Analytical estimations of edge states and extended states in large finite-size lattices. Available from: https://doi.org/10.48550/arXiv.2511.07875.

[18] Xu, H., Chen, Y. and Zhang, D., 2025. Beyond empirical models: Discovering new constitutive laws in solids with graph-based equation discovery. Available from: https://doi.org/10.48550/arXiv.2511.09906.

[19] Yao, Y., Yang, G., Xu, R., Tu, Y. and Mo, H., 2025. An Improved Dual-Attention Transformer-LSTM for Small-Sample Prediction of Modal Frequency and Actual Anchor Radius in Micro Hemispherical Resonator Design. Available from: https://doi.org/10.48550/arXiv.2511.08900.