

Mapping the Evolution of TinyML: A Bibliometric and Science-Mapping Analysis of Global Research

Arpit Vyas

Department of Library, Bennett University, Greater Noida, Uttar Pradesh, India

Email: iamarpit44@gmail.com | ORCID: <https://orcid.org/0009-0006-8700-9068>

Received: 30 August 2026 | Revised: 14 September 2026 | Accepted: 26 September 2026 | Published online: 30 September 2026

How to cite: Vyas, A. (2026). Mapping the evolution of TinyML: A bibliometric and science-mapping analysis of global research. *Journal of Scholarly Research and Metrics*, 1(1), 45–65.

Abstract

Context: Tiny Machine Learning (TinyML) has emerged as an important edge-intelligence paradigm for deploying machine-learning models on resource-constrained and low-power devices. Its rapid expansion across diverse application domains has created a need for systematic mapping of the research landscape.

Aim: This study maps the development, scholarly impact, collaboration patterns, and thematic evolution of global TinyML research.

Approach: Bibliographic records were retrieved from Scopus using “TinyML” and “Tiny Machine Learning.” Following systematic screening and data cleaning, 1,291 publications covering 2020–2026 were analyzed using bibliometric and science-mapping techniques.

Key Findings: TinyML research exhibited rapid growth, with an annual growth rate of 69.84% and a peak complete-year output of 422 publications in 2025. India, China, and Italy were the leading corresponding-author countries, while the University of Bologna was the most productive institution. Keyword analysis revealed a thematic transition toward learning systems, edge computing, embedded systems, and real-time applications.

Interpretation: The findings demonstrate that TinyML has developed into a rapidly expanding, internationally collaborative, and increasingly application-oriented research domain.

Contribution: This study provides an updated quantitative mapping of TinyML scholarship and identifies influential contributors, collaboration patterns, established themes, and emerging research directions.

Keywords: TinyML; tiny machine learning; bibliometric analysis; scientometrics; edge artificial intelligence; research trends

1. Introduction

The rapid expansion of the Internet of Things (IoT) and edge computing has increased the need for intelligent processing directly on low-power edge devices. Tiny Machine Learning (TinyML) refers to the deployment of machine-learning technologies, including models, algorithms, software, and supporting hardware, for on-device data processing on highly resource- and power-constrained devices, particularly microcontrollers and embedded systems (Warden & Situnayake, 2019). TinyML enables local inference with reduced dependence on cloud infrastructure, thereby supporting low-latency, energy-efficient, and privacy-aware applications. These properties are particularly relevant to real-time and connectivity-limited applications (Heydari & Mahmoud, 2025).

With the rapid evolution of hardware technology, lightweight neural architectures, model compression, and embedded-AI technology, TinyML has significantly expanded its capabilities. Currently, applications are being developed in various sectors, such as healthcare, agriculture, environmental monitoring, industrial automation, security, and smart-city systems (Pazmiño Ortiz et al., 2025). For instance, TinyML for smart-city research encompasses urban mobility, public safety, environmental sensing, waste management, and infrastructure monitoring (Trigkas et al., 2025). However, deployment is still limited by memory, computing power, energy consumption, reliability, and the tradeoff between model accuracy and resource efficiency (Immonen & Hämäläinen, 2022; Pazmiño Ortiz et al., 2025).

Han et al. (2026) provide an important bibliometric foundation for understanding the development of TinyML. However, further analysis is needed to characterize the field across multiple complementary dimensions, particularly the relationships among publication productivity, source impact, institutional and geographic participation, international collaboration, citation influence, and thematic evolution. Examining these dimensions together can provide a more integrated understanding of the scholarly, collaborative, and thematic structure of the rapidly developing TinyML knowledge domain.

In this context, the present study aims to map the worldwide research landscape around TinyML, based on 1,291 publications found in Scopus (2020–2026). It covers changes in publication volume, key sources, institutions and countries, citation impact, international collaboration and thematic evolution. The study uses bibliometric and science-mapping methods to offer a current view of the intellectual, collaborative, and thematic organization of the TinyML research field, as well as to uncover new avenues for future research.

1.1 Aim and Research Questions

The main objective of this study is to systematically map the global research landscape of Tiny Machine Learning (TinyML) using bibliometric methods. Specifically, the study aims to analyze the development of TinyML scholarship; identify key publication sources, institutions, and countries; evaluate publication impact and international research collaboration; and trace the evolution of the field's intellectual structure and themes. It analyzes 1,291 Scopus-indexed publications from 2020 to 2026, combining productivity, impact, collaboration, and science-mapping metrics to provide a comprehensive quantitative evaluation of TinyML research.

Accordingly, the study addresses the following research questions:

RQ1: How has the scientific production of TinyML research evolved over the study period?

RQ2: Which publication sources have contributed most prominently to TinyML research, and what is their scholarly impact?

RQ3: Which institutions and countries are the leading contributors to TinyML research?

RQ4: What patterns of citation impact and international collaboration characterize the global TinyML research landscape?

RQ5: What are the dominant, evolving, and emerging research themes within the TinyML knowledge domain?

2. Literature Review

2.1 Evolution and Development of TinyML

Tiny Machine Learning (TinyML) has evolved from the fusion of embedded systems, edge computing, the Internet of Things (IoT), and lightweight machine learning technologies. Earlier IoT architectures relied mainly on central cloud platforms and continuously sent data to distant

servers, resulting in latency, bandwidth use, connectivity dependence, energy consumption, and privacy concerns. Edge computing brought computation closer to the data sources while TinyML took it a step further by enabling inference on low-power microcontrollers or sensor-class devices (Abadade et al., 2023; Wainbuch & Samuel, 2024).

Major advances in microcontroller performance, low-power neural networks, and model quantization have driven the technology forward. Modern MCUs integrate computational instructions, dedicated on-chip SRAM, embedded flash, and in some cases dedicated accelerators, enabling computation that would previously not have been possible on such constrained hardware (Abadade et al., 2023). In the software ecosystem, the development and optimization of edge-deployment workflows have been facilitated by frameworks such as TensorFlow Lite for Microcontrollers, CMSIS-NN, Edge Impulse, and Apache TVM (Wainbuch & Samuel, 2024). To further reduce memory and computation, quantization, pruning, knowledge distillation, and efficient neural architectures are adopted. Elhanashi et al. (2024), for instance, describe TensorFlow Lite for Microcontrollers, which can be deployed on MCUs with very low memory requirements, and highlight the importance of quantization and pruning for improving inference efficiency.

TinyML has also moved beyond classic static inference, which involves training a model on more powerful platforms before deploying it with a fixed set of parameters. Rajapakse et al. (2023) found that TinyML was gradually moving toward adaptive and continuously learning systems, with notable growth in reformable TinyML, in which deployed models can be updated or improved after deployment. This development also reflects a growing move toward distributed, cooperative edge intelligence. The transition has been illustrated by Sanchez-Iborra et al. (2023) for smart agriculture using a hierarchical TinyML–edge computing architecture that resulted in lower wireless transmissions, energy usage, and response time, as well as improved privacy and security.

More recently, TinyML development has turned toward advanced neural-network optimization and integration with next-generation communication systems. Lê et al. (2026) highlight critical directions for deploying deeper models on ultra-low-power MCUs, including hardware acceleration, algorithm-architecture co-design, pruning, quantization, and low-rank factorization. In addition, Vu et al. (2026) situate TinyML in the context of emerging 6G ecosystems, in which resource-limited intelligence at the edge can work in tandem with large models to enable scalable, distributed intelligent services. Thus, TinyML has gradually evolved from lightweight embedded inference to a more comprehensive paradigm of adaptive, distributed edge intelligence using energy-efficient resources.

2.2 TinyML Technologies and Application Domains

Together, TinyML's resource-efficient hardware, lightweight software frameworks, and model-compression techniques allow the execution of machine-learning inference directly on edge devices with constrained resources. Its technological ecosystem is mostly centered on microcontroller units (MCUs), embedded platforms with sensors and optimized neural networks. Common software environments include TensorFlow Lite for Microcontrollers and Edge Impulse, as well as other lightweight deployment toolchains. Quantization and pruning decrease the model size, computational requirements and power needs. This technological base enables TinyML applications that emphasize low latency, energy efficiency, localized decision-making, and reduced reliance on the cloud.

Healthcare is one of the prominent application domains. In a systematic review of 136 peer-reviewed studies published from 2020 to 2025, Yahyati et al. (2025) found TinyML applications in patient monitoring, medical diagnostics, wearable sensing, early disease detection, education and transportation. Abadade et al. (2024) created quantized TinyML models for real-time respiratory-condition detection via lung sounds. Their custom CNN achieved 96% accuracy and 97% precision, showing that low-power, digital-stethoscope-type devices can potentially perform diagnostic inference. TinyML can be applied to biometric recognition, such as fingerprints, faces, irises, gait, voice, and keystrokes, among others, and has been used in wearables, healthcare monitoring, access control, surveillance and in IoT security (Essahraoui & Lamaakal, 2026).

Another significant application is agriculture, where TinyML can be used for crop monitoring, disease detection, irrigation, and resource management. Gookyi et al. (2024) proposed a maize disease detection model using CNN with an accuracy of 94.60%, a classification latency of 7.60 ms, RAM usage of 344.70 KB, and flash usage of 726.60 KB, showing the practical deployability of the model at the MCU scale. Their comparison also showed that Edge Impulse offered advantages in ease of use and memory usage, while TensorFlow offered greater flexibility and customization. Likewise, Hinostroza et al. (2026) proposed on-device adaptive sampling for agricultural wireless sensor nodes achieving a mean absolute error (MAE) of 0.537 °C in temperature reconstruction and a 77.8% reduction in energy consumption compared with fixed-interval sampling.

The application base has expanded to include smart cities, transportation and industrial IoT. Alharthi et al. (2026) reviewed 35 Industrial IoT (IIoT) studies published between 2018 and 2026 and found that the main industrial applications are predictive maintenance, equipment monitoring, anomaly detection, and energy management, while the most frequent sensing modality is vibration. For smart-city research, Trigkas et al. (2025) analyzed 66 studies, including 29 on mobility and transportation, 17 on public safety, 10 on environmental sensing, six on waste management, and four on infrastructure monitoring. Together, these applications illustrate that TinyML has evolved beyond an embedded-AI application into a cross-domain technological platform that enables autonomous, energy-efficient, privacy-aware, and real-time intelligence at the edge.

2.3 Challenges and Limitations of TinyML

Although TinyML brings machine learning to edge devices with limited hardware resources, significant hardware constraints still limit its applications. Microcontrollers usually have limited memory, storage space, processing power and energy resources, thus limiting the size and complexity of deployable models. For these reasons, quantization, pruning, and model compression are often necessary, and these optimizations can involve trade-offs between resource usage and prediction accuracy. Real-time applications add to these challenges, as inference needs to be conducted within tight latency bounds with satisfactory accuracy and energy use (Heydari & Mahmoud, 2025). The same limiting factors have been reported in practical embedded applications, where memory, computation, energy usage and deployment complexity are key issues.

Reliability and robustness are other limitations. TinyML models deployed in real-world settings can be subject to sensor noise, environmental variability, domain shifts, imbalanced data, hardware nonidealities and adversarial perturbations. Pereira et al. (2026) highlight memory

efficiency as a critical hardware limitation and note that standardized robustness benchmarking is still underdeveloped. Furthermore, existing evaluation approaches do not always consider robustness alongside resource usage, including memory and energy consumption, making comprehensive comparisons between TinyML systems difficult.

Security and privacy are growing in importance as TinyML enters healthcare, industrial, biometric, and other safety-critical fields. While local inference can limit the amount of sensitive raw information sent from deployed devices and systems, attackers can still tamper with, steal, attack, or modify models without authorization. Standard security methods can also have computational and memory costs that can be difficult to accommodate on very resource-constrained devices. Huckelberry et al. (2026) note that TinyML security is comparatively understudied and that many traditional security measures are hard to implement under tight resource constraints. In distributed TinyML-enabled edge systems, Baciú et al. (2025) also point out the problems of model tampering and access control.

Finally, the TinyML ecosystem is heterogeneous, with challenges concerning portability, scalability, interoperability, development workflows and benchmarking. Architectural differences among the microcontroller, software framework, sensor and deployment environments can make optimizing for each platform and achieving reproducibility difficult. The deployment complexity and scalability of these IoT platforms, as well as the lack of standardized benchmarking, have been identified as current limitations in Industrial TinyML studies (Alharthi et al., 2026). These challenges highlight the need for standardized evaluation protocols, secure systems, hardware–software interoperability and the ability to develop systems that meet the needs for accuracy, latency, memory, energy efficiency and reliability in a wide range of deployments.

2.4 Research Gaps and Need for the Present Study

Technological foundation, optimization, application, and deployment aspects of TinyML have seen significant progress, as detailed in the existing literature. To the author's knowledge, existing reviews have focused mainly on the technical or application aspects of TinyML. Abadade et al. (2023) gave a general overview of TinyML architectures and their hardware/software and optimization methods, and Rajapakse et al. (2023) focused on reformable TinyML and the adaptability of models after deployment. More recent reviews have focused on efficient neural-network design (Lê et al., 2026), robustness (Pereira et al., 2026), and security threats and defenses (Huckelberry et al., 2026). All of these studies give insights into individual technological aspects of TinyML, but mostly follow a qualitative or systematic review approach.

The body of literature has also become increasingly application-specific. Yahyati et al. (2025) provided a survey of TinyML applications for healthcare, education and transportation, while Trigkas et al. (2025) focused on smart-city applications. In a similar vein, Alharthi et al. (2026) comprehensively analyzed TinyML in the context of Industrial IoT regarding its applications, hardware and software components, optimization strategies, performance metrics, and deployment issues. The available evidence, therefore, offers a detailed understanding of specific domains but the increased specialization of such reviews renders it challenging to gain an integrated quantitative understanding of the evolution of the overall TinyML knowledge domain.

Another limitation is the lack of systematic documentation of the scholarly organization of the field. While existing reviews recognize the key developments and problems, such as the challenges of resource limitations, model optimization, energy efficiency, security, robustness, interoperability, and benchmarking, they do not capture the evolution of research productivity and

scholarly influence, nor the role of institutional participation, geographic contribution, international cooperation, publication outlets, or thematic relationships. This limitation is significant because TinyML research is growing rapidly and is highly cross-disciplinary.

Therefore, the work presented in this paper employs a bibliometric method to provide a more extensive quantitative, descriptive mapping of the TinyML research landscape. The study's findings are based on 1,291 publications indexed in the Scopus database, spanning the period 2020–2026, and it identifies trends in publications in the subject area, the most influential sources, publications by institutions and countries, citation impact, international collaboration of authors, and the evolution and interconnection of research themes. Additionally, the use of trend-topic and keyword co-occurrence analysis helps determine both well-established and emerging thematic trends. The study aims to present an overview of the intellectual and developmental landscape of TinyML research, including productivity, impact, collaboration metrics, and thematic indicators, all within a single investigation, to create an evidence base for future research opportunities.

Thus, the value of the present study lies not merely in updating publication counts or expanding the Scopus dataset. It offers a holistic perspective on the field, combining analyses of source concentration and impact, institutional and country-level involvement, international collaboration, and the evolution and interaction of research themes over time. The analysis of SCP/MCP patterns, country impact, trend topics, and keyword co-occurrence, in particular, facilitates the identification of differences between research productivity, international connectivity, scholarly impact, and thematic evolution. These dimensions complement each other in the structure of the global TinyML research landscape and in the evolution of its research priorities. The study thus adds a multidimensional science-mapping perspective that complements, rather than replaces, the earlier bibliometric evaluation of TinyML by Han et al. (2026).

3. Methodology

3.1 Research Design

This study employed bibliometric and science-mapping methods to examine the development, scholarly impact, collaboration patterns, and thematic evolution of Tiny Machine Learning (TinyML) research. Bibliometric analysis was used to assess publication productivity, sources, institutions, countries, citations, and international collaboration, while science-mapping techniques were used to examine the thematic structure and evolution of the field.

3.2 Database Selection

Scopus was chosen as the sole bibliographic database because it provides structured, standardized metadata suitable for bibliometric analysis, including information on authors, affiliations, countries, publication sources, citations, references, and keywords. A single database ensured consistency of bibliographic metadata fields and citation-link structures during data extraction and analysis. However, reliance on Scopus alone may exclude relevant publications indexed only in other databases.

3.3 Search Strategy

The Scopus search was performed on 5 July 2026 using the phrases “TinyML” and “Tiny Machine Learning” in article titles, abstracts, and keywords. The initial search string was: TITLE-ABS-KEY ("TinyML" OR "Tiny Machine Learning")

The preliminary search retrieved 2,275 records, which were then narrowed by subject area, document type, language, and source type. The complete Scopus search string with all filters was:

Final search string: TITLE-ABS-KEY ("TinyML" OR "Tiny Machine Learning") AND (EXCLUDE (SUBJAREA , "NURS") OR EXCLUDE (SUBJAREA , "VETE") OR EXCLUDE (SUBJAREA , "IMMU") OR EXCLUDE (SUBJAREA , "PHAR") OR EXCLUDE (SUBJAREA , "PSYC") OR EXCLUDE (SUBJAREA , "ARTS") OR EXCLUDE (SUBJAREA , "HEAL") OR EXCLUDE (SUBJAREA , "ECON") OR EXCLUDE (SUBJAREA , "NEUR") OR EXCLUDE (SUBJAREA , "EART") OR EXCLUDE (SUBJAREA , "ENVI") OR EXCLUDE (SUBJAREA , "AGRI") OR EXCLUDE (SUBJAREA , "MULT") OR EXCLUDE (SUBJAREA , "CENG") OR EXCLUDE (SUBJAREA , "BUSI") OR EXCLUDE (SUBJAREA , "BIOC") OR EXCLUDE (SUBJAREA , "CHEM") OR EXCLUDE (SUBJAREA , "SOCI") OR EXCLUDE (SUBJAREA , "MEDI") OR EXCLUDE (SUBJAREA , "MATE") OR EXCLUDE (SUBJAREA , "ENER")) AND (EXCLUDE (DOCTYPE , "le") OR EXCLUDE (DOCTYPE , "tb") OR EXCLUDE (DOCTYPE , "no") OR EXCLUDE (DOCTYPE , "sh") OR EXCLUDE (DOCTYPE , "ed") OR EXCLUDE (DOCTYPE , "bk") OR EXCLUDE (DOCTYPE , "ch") OR EXCLUDE (DOCTYPE , "cr")) AND (LIMIT-TO (LANGUAGE , "English")) AND (EXCLUDE (SRCTYPE , "k") OR EXCLUDE (SRCTYPE , "d"))

Because the search was conducted on 5 July 2026, publications for 2026 were interpreted as a partial year.

3.4 Inclusion and Exclusion Criteria

Records were screened sequentially according to predefined criteria. Publications within Computer Science, Engineering, Mathematics, Physics and Astronomy, and Decision Sciences were retained, reducing the dataset to 1,587 records. Screening by document type retained conference papers, articles, and reviews, resulting in 1,451 records. Restricting the dataset to English-language publications resulted in 1,446 records. Screening by source type retained conference proceedings and journals, leaving 1,304 records. Document type and source type were treated separately: document type refers to the form of publication (conference paper, article, or review), whereas source type refers to the publication venue (conference proceeding or journal). After removing 11 incomplete or irrelevant records and two duplicates, the final dataset comprised 1,291 publications.

3.5 PRISMA-Based Study Selection

The complete selection procedure is shown in Figure 1.

Figure 1. PRISMA-based study selection process for TinyML publications.



Source: Author's own compilation based on the Scopus search and study-selection process (search conducted on 5 July 2026).

3.6 Data Extraction and Cleaning

Bibliographic records were downloaded from Scopus and checked for completeness, relevance, and duplication. The retrieved metadata included publication year, authors, author affiliation, document type, source type, countries, citations, references and keywords. Duplicate, incomplete, and irrelevant records were excluded. The final dataset for bibliometric and science-mapping analysis comprised 1,291 records.

3.7 Bibliometric and Science-Mapping Analysis

The bibliometric analysis and thematic exploration were performed using Biblioshiny (version 5.5.0), and keyword co-occurrence mapping was performed using VOSviewer (version 1.6.20). The analyses included scientific production by year, characteristics of scientific production, source productivity and impact, institutional and country productivity, citation impact, international scientific collaboration, and thematic development.

Bradford's Law was used to identify core publication sources while the h-index, g-index, m-index, total citations (TC) and number of publications (NP) were used to measure the impact of the sources. Accordingly, Table 2 reports source productivity, while Table 3 reports source impact based on citations.

International collaboration was analyzed using single-country publications (SCP), multiple-country publications (MCP), and the MCP percentage, that is, the share of publications involving international co-authorship.

Trend-topic and keyword co-occurrence analysis were used to examine the thematic development. In VOSviewer, all keywords were analyzed using full counting and a minimum of 10 occurrences, yielding 297 keywords in five clusters, with 16,115 links and a total link strength of 54,903. The temporal evolution of and relationships between TinyML research themes were explored through an overlay visualization.

A thematic strategy map was generated in Biblioshiny version 5.5.0 using All Keywords and the Louvain clustering algorithm, with themes positioned according to centrality (relevance degree) and density (development degree) across the Motor, Basic, Niche, and Emerging or Declining quadrants.

4. Results

Table 1. Descriptive Characteristics of the TinyML Bibliometric Dataset


Category	Indicator	Result
Main Information	Timespan	2020–2026
	Sources (journals, books, etc.)	660
	Documents	1,291
	Annual growth rate (%)	69.84
	Document average age (years)	1.79
	Average Citations Per Document	9.205
	References	76,158
Document Contents	Keywords Plus	6,561
	Author keywords	3,063
Authors	Authors	3,847
	Authors of Single-Authored Documents	58
Author Collaboration	Single-Authored Documents	60
	Co-authors Per Document	4.33
	International Co-Authorships (%)	25.41
Document Types	Articles	473
	Conference Papers	790
	Reviews	28

Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Table 1 shows the main features of the TinyML bibliometric data for 2020–2026. The dataset comprises 1,291 documents from 660 sources, with an annual growth rate of 69.84%. The publications received an average of 9.205 citations each and cited 76,158 references in total. The 6,561 Keywords Plus and 3,063 author keywords indicate considerable topical diversity. Regarding authorship patterns, 3,847 authors contributed to the documents, the average number of co-authors per document was 4.33, indicating strong collaboration, and the international co-authorship percentage was 25.41%. Most documents are conference papers (790), followed by

articles (473) and reviews (28), reflecting the centrality of dissemination through conferences in TinyML research.

Figure 2. Annual Scientific Production of TinyML Research



Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Figure 2 shows that the number of TinyML publications has risen steadily since 2020, indicating a growing research field. Annual output grew from 9 documents in 2020 to 43 in 2021 and 111 in 2022, reflecting early growth in scholarly interest. Output continued to increase, reaching 200 publications in 2023 and 290 in 2024, and peaked at 422 publications in 2025. The apparent decrease in publications in 2026 (216) should not be seen as a decrease in research activity since this is only a partial year. This publication pattern suggests that TinyML has evolved from an emerging research topic into a fast-growing research area.

Table 2. Core Publication Sources in TinyML Research Based on Bradford's Law

Rank	Source	Publications (n)	Cumulative Publications	Bradford Zone
1	IEEE Internet of Things Journal	51	51	Zone 1
2	ACM International Conference Proceeding Series	29	80	Zone 1
3	Electronics (Switzerland)	23	103	Zone 1
4	IEEE Sensors Journal	18	121	Zone 1
5	Proceedings of the International Joint Conference on Neural	16	137	Zone 1

	Networks			
6	Future Internet	15	152	Zone 1
7	IEEE Transactions on Consumer Electronics	15	167	Zone 1
8	Internet Technology Letters	15	182	Zone 1
9	ACM Transactions on Embedded Computing Systems	14	196	Zone 1
10	Proceedings – IEEE International Symposium on Circuits and Systems	14	210	Zone 1

Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Table 2 shows the main publication sources in TinyML research identified using Bradford's Law, which ranks sources by productivity and publication concentration. IEEE Internet of Things Journal is the most prolific with 51 publications, followed by ACM International Conference Proceeding Series (29) and Electronics (Switzerland) (23). IEEE Sensors Journal contributed 18 publications and the Proceedings of the International Joint Conference on Neural Networks 16, while the other core sources accounted for 14–15 publications each. In total, these ten sources have published 210 papers and form a central group of highly concentrated sources for TinyML research, represented in Bradford Zone 1. The interdisciplinary nature of TinyML research is also reflected in the presence of journals and conference proceedings on topics such as IoT, sensors, embedded computing, neural networks, circuits, and consumer electronics. Importantly, this ranking is based on a measure of publication productivity and source concentration, not on citation impact (which is measured separately in Table 3, based on citation data).

Table 3. Local Impact of Leading Publication Sources in TinyML Research

Rank	Source	h-index	g-index	m-index	Total Citations (TC)	Publications (NP)	Starting Year
1	IEEE Internet of Things Journal	16	28	2.286	856	51	2020
2	Electronics (Switzerland)	10	16	2.5	269	23	2023
3	IEEE Transactions on Instrumentation and Measurement	10	11	2	296	11	2022
4	ACM International Conference Proceeding Series	9	15	1.5	279	29	2021
5	IEEE Sensors Journal	9	16	2.25	273	18	2023
6	ACM Transactions on	8	11	2	141	14	2023

	Embedded Computing Systems						
7	IEEE Transactions on Consumer Electronics	8	14	2.667	200	15	2024
8	Future Internet	7	15	1.4	355	15	2022
9	IEEE Sensors Letters	7	8	1.4	104	8	2022
10	2024 IEEE Sensors Applications Symposium (SAS 2024) – Proceedings	6	8	2	89	8	2024

Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Table 3 shows the local scholarly impact of the main sources of publications on TinyML based on the h-index, g-index, m-index, total citations (TC) and number of publications (NP). The IEEE Internet of Things Journal has the highest h-index (16), g-index (28), total citations (856), and publications (51) among all journals. Both Electronics (Switzerland) and IEEE Transactions on Instrumentation and Measurement have an h-index of 10, but their profile in terms of citations and publications varies. IEEE Transactions on Consumer Electronics has the highest m-index (2.667), meaning that its h-index has grown fastest relative to the year of its first TinyML publication, whereas Future Internet has 15 TinyML papers, with 355 citations. In general, these findings show that productivity and citation-based impact are distinct dimensions of source performance, and that high publication output does not necessarily entail the highest impact across all indicators. Table 3 therefore identifies core sources based on citation metrics, complementing the productivity-based core sources identified in Table 2 using Bradford's Law.

Figure 3. Most Productive Institutional Affiliations in TinyML Research



Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Figure 3 shows the 10 most productive institutional affiliations in TinyML research. The University of Bologna had the highest number of publications, with 33, far exceeding the number published by other institutions, followed by Nirma University with 19 publications and Xidian University with 16 publications. The Federal University of Rio Grande do Sul is in fourth place, with 14 publications. Mid Sweden University, Mohammed Premier University, Pukyong National University, and STMicroelectronics each contributed 13 publications, and the top ten is completed by the National Institute of Technology and the University of Porto (11 publications each). The spread of affiliations across Europe, Asia, South America, and Africa illustrates the geographically distributed institutional involvement in TinyML research, while the presence of STMicroelectronics indicates notable industrial participation alongside academic research.

Table 4. Top 10 Corresponding Author Countries and International Collaboration in TinyML Research

Rank	Country	Publications (n)	Publications (%)	SCP	MCP	MCP (%)
1	India	140	10.8	112	28	20
2	China	115	8.9	87	28	24.3
3	Italy	79	6.1	58	21	26.6
4	USA	62	4.8	53	9	14.5
5	Germany	30	2.3	22	8	26.7

6	Korea	30	2.3	19	11	36.7
7	Morocco	29	2.2	25	4	13.8
8	Brazil	27	2.1	15	12	44.4
9	Switzerland	27	2.1	10	17	63
10	United Kingdom	21	1.6	12	9	42.9

Note: SCP = single-country publications; MCP = multiple-country publications; MCP (%) = proportion of publications involving international collaboration.

Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Table 4 shows the geographical distribution and international collaboration pattern of the corresponding authors in TinyML research. India has the highest number of publications (140; 10.8%), followed by China (115; 8.9%) and Italy (79; 6.1%). The USA ranks fourth with 62 publications, most of which are single-country publications (53). Levels of international collaboration vary across countries: Switzerland has the highest share of multiple-country publications (63.0%) despite only 27 documents, reflecting strong international research collaboration. Other countries with relatively high international collaboration rates include Brazil (44.4%) and the United Kingdom (42.9%). By contrast, Morocco (13.8%) and the USA (14.5%) show the lowest MCP shares. These results show that higher research productivity does not necessarily coincide with higher international collaboration in TinyML research.

Table 5. Top 10 Most Cited Countries in TinyML Research


Rank	Country	Total Citations (TC)	Average Citations per Article
1	Italy	1,078	13.6
2	China	1,074	9.3
3	India	807	5.8
4	Switzerland	506	18.7
5	Greece	420	32.3
6	USA	344	5.5
7	Saudi Arabia	295	17.4
8	United Kingdom	259	12.3
9	Sweden	194	10.2
10	Korea	162	5.4

Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Table 5 shows the 10 most cited countries in TinyML research. Italy has the highest total citations (1,078), followed closely by China (1,074) and India (807). Although Greece has fewer total citations (420), its average of 32.3 citations per article is the highest among these countries.

Switzerland also shows strong impact, with 506 citations and 18.7 citations per article, followed by Saudi Arabia (17.4 citations per article). In contrast, although India, the USA, and Korea were highly productive, they had relatively low average citation rates. These results reveal that publication volume and citation influence vary across countries, with some smaller contributors achieving relatively high per-article scholarly impact.


Figure 4. Evolution of Trend Topics in TinyML Research



Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Figure 4 shows the time-series evolution and frequency of key topics in TinyML research. Earlier studies focused on learning algorithms and controllers, with median prominence around 2023. In 2024, the focus expanded to machine learning, Internet of Things, tiny machine learning, and deep learning. Of these, TinyML has the highest frequency (771), with its median prominence in 2025, indicating that the topic is becoming more central in the literature. Learning systems, edge computing, embedded systems and real-time applications are more recent themes, with median activity around 2025. The overall evolution across time reflects a transition from basic hardware and algorithmic areas to integrated TinyML applications with a focus on edge intelligence, embedded deployment and real-time machine learning applications.


Figure 5. Overlay Visualization of Keyword Co-occurrence in TinyML Research



Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Figure 5 depicts the temporal evolution of co-occurrence in TinyML research using VOSviewer. Using all keywords, full counting, and a minimum of 10 occurrences, the analysis produced 297 keywords grouped into five clusters. The network contains 16,115 links with a total link strength of 54,903, indicating strong thematic interconnection. Several key terms are featured prominently, including “TinyML,” “machine learning,” “tiny machine learning,” “microcontrollers,” “Internet of Things,” “edge computing” and “deep learning.” The overlay visualization shows changes over time, from earlier to more recent topics. Significant emerging research interest can be seen in the areas of embedded systems, learning systems, edge intelligence, energy utilization, quantization, and application-oriented topics. The visualization illustrates the evolution of TinyML toward more resource-efficient, intelligent, and edge-focused research.

Figure 6. Thematic Strategy Matrix of TinyML Research



Source: Author's own analysis based on Scopus bibliographic data (2020–2026).

Figure 6 shows the thematic structure of TinyML research based on the centrality and density of each theme. The research areas that are well-developed and influential can be seen as Motor Themes, such as neural networks and energy efficiency. TinyML, machine learning, and learning systems are Basic Themes, reflecting high centrality but comparatively lower development. Specialized topics with lower centrality, such as tiny machine learning, object detection, and antennas, are Niche Themes. Emerging or Declining Themes include embedded systems and real-time systems, which have relatively low development and centrality. Deep learning, microcontrollers, and convolutional neural networks occupy a more central transitional position, suggesting connections to several thematic areas.

5. Discussion

The findings indicate that TinyML has progressed from an emerging research topic into a rapidly expanding and increasingly structured research domain. Regarding RQ1, the strong growth in scientific production demonstrates sustained scholarly interest in deploying machine learning on resource-constrained edge devices. Rather than interpreting the partial 2026 output as a decline, the publication trajectory through the last complete year, 2025, suggests continued expansion of the field. This growth is consistent with the broader technological development of TinyML from lightweight embedded inference toward adaptive, distributed, and resource-efficient edge intelligence discussed in the existing literature (Abadade et al., 2023; Lê et al., 2026; Rajapakse et al., 2023). Thus, publication growth appears to reflect not only increased academic attention but also the widening technological and application scope of TinyML.

For RQ2, the source analysis reveals an important distinction between publication productivity and scholarly impact. The concentration of publications in outlets associated with IoT, sensors, embedded computing, neural networks, and electronic systems demonstrates the interdisciplinary character of TinyML. However, the most productive sources are not uniformly the strongest across all citation-based indicators. This distinction suggests that publication volume alone provides an incomplete representation of scholarly influence. The prominence of the IEEE

Internet of Things Journal across both productivity and impact measures nevertheless indicates that the intersection of TinyML with IoT and edge intelligence has become an important component of the field's scholarly structure. More generally, the source patterns show that TinyML research is distributed across both journals and conference proceedings, reflecting its combination of rapidly developing technological research and more established scholarly dissemination.

The institutional and geographic findings address RQ3 and demonstrate that TinyML research is globally distributed rather than concentrated within a single national or regional research system. Leading institutions span Europe, Asia, South America, and Africa, while the presence of an industrial organization among the productive affiliations indicates participation beyond academia. At the country level, India, China, and Italy occupy leading positions in publication output. These patterns suggest that TinyML development is being shaped by geographically diverse research communities, which is particularly relevant for a field involving heterogeneous hardware platforms, IoT environments, sensing systems, and deployment contexts. Such geographic diversity can broaden the range of applications and technical problems addressed by the field.

Regarding RQ4, the collaboration and citation results demonstrate that research productivity, international connectivity, and citation influence are distinct dimensions of national research performance. Highly productive countries do not necessarily exhibit the highest proportion of internationally co-authored publications, while countries with smaller publication portfolios can demonstrate stronger international collaboration or greater average citation influence. For example, Switzerland combines a comparatively smaller publication output with a high MCP share, whereas India and the USA show substantially lower MCP proportions despite their visible contribution to the literature. Similarly, Greece and Switzerland achieve relatively high average citation rates despite producing fewer publications than the largest contributors. These differences indicate that the global influence of TinyML research cannot be assessed solely through publication counts; collaboration intensity and citation impact provide complementary perspectives on national research positioning.

The thematic findings provide the answer to RQ5 and indicate a transition from foundational learning and hardware-oriented topics toward more integrated, deployment-oriented TinyML research. The increasing prominence of learning systems, edge computing, embedded systems, real-time applications, energy utilization, and quantization suggests that the field is moving toward practical implementation under strict resource constraints. This interpretation is reinforced by the thematic strategy matrix. Neural networks and energy efficiency appear as motor themes, indicating that efficient learning architectures remain both well developed and structurally important. TinyML, machine learning, and learning systems occupy the basic-theme quadrant, suggesting that they provide broad foundations connecting multiple areas of the knowledge domain. At the same time, the positioning of embedded and real-time systems in the emerging-or-declining quadrant should be interpreted cautiously: in combination with their recent appearance in the temporal analysis, their position is more consistent with emerging thematic development than with established decline.

Taken together, the five research questions reveal a field characterized by rapid growth, interdisciplinary publication structures, geographically distributed participation, uneven international connectivity and citation influence, and a shift toward resource-efficient real-world deployment. These dimensions are interrelated but should not be treated as interchangeable: high publication productivity does not automatically imply high citation influence or international collaboration, just as an emerging theme may be increasingly important despite currently having lower thematic density. This multidimensional interpretation strengthens the contribution of the present bibliometric analysis by showing not merely how much TinyML research has been produced, but how its scholarly, collaborative, and thematic structures are developing. The resulting picture complements previous technically oriented reviews by providing a broader

science-mapping perspective on the organization and evolution of the TinyML knowledge domain.

5.1 Implications for Researchers and Practitioners

The results have significant research and development implications for TinyML. The literature continues to grow each year and is emerging as an interdisciplinary research field that combines machine learning, IoT, embedded systems, sensors, and edge computing. Identifying influential publication sources, institutions and countries can help researchers choose the right publication venue and collaborators. The varied patterns of international collaboration also suggest opportunities for stronger cross-country research networks. The increasing focus on edge computing and embedded systems, energy usage, quantization and real-time applications suggests a shift toward deployed and resource-efficient solutions for practitioners. These developments are especially pertinent to organizations that require low-latency, low-energy intelligence in smart cities, environmental monitoring, agriculture, industrial IoT, and healthcare.

5.2 Future Research Directions

The bibliometric results suggest several avenues for future TinyML research. Given recent trends in learning systems, edge computing, embedded systems, real-time applications, energy utilization, and quantization, future research should focus on increasingly adaptive and resource-efficient on-device intelligence. Further work is needed on hardware–software co-design, model compression, energy-aware inference, and efficient neural architectures that can be configured to achieve the desired accuracy while reducing memory, latency, and power consumption. In addition, robustness and security should be key research areas as resource-constrained TinyML systems are vulnerable to sensor noise, domain shifts, adversarial attacks, and physical tampering (Huckelberry et al., 2026; Pereira et al., 2026). Further studies are also needed to develop standardized benchmarks that improve reproducibility and cross-platform comparability. Finally, greater international cooperation would enable access to shared datasets, tools that are interoperable, and validation in diverse deployment contexts, potentially supporting the transition from experimental prototypes to scalable real-world TinyML systems.

5.3 Limitations

The following limitations should be considered when interpreting the results of this study. First, the analysis is based only on the Scopus database, so relevant publications indexed in other databases, such as Web of Science, IEEE Xplore, and Google Scholar, may not have been captured. Second, the search included only the phrases “TinyML” and “Tiny Machine Learning” in titles, abstracts, and keywords; as a result, conceptually similar studies using other terms may not have been included. Third, citation-based indicators may be affected by publication time, disciplinary citation practices, database coverage and indexing time, potentially disadvantaging more recent studies. Fourth, variations in author names, institutional affiliations, and keywords may affect bibliometric network structures even after data cleaning. Finally, the data cover 2020–2026, with 2026 being a partial publication year. The findings thus reflect the state of the field of research at the time of data collection, and do not provide a definitive or fixed account of this fast-changing research area.

6. Conclusion

This study mapped the development and intellectual structure of TinyML research based on 1,291 publications indexed in Scopus between 2020 and 2026. The results indicate rapid growth in scholarly activity, with an annual growth rate of 69.84% and 422 publications in 2025. The IEEE Internet of Things Journal was the most productive source, and the University of Bologna was the most productive institution. The leading corresponding-author countries were India, China, and Italy, although international collaboration and citation influence varied considerably across

countries. Thematic analysis also shows a shift from basic learning algorithms and controllers to edge computing, embedded systems, learning systems, energy efficiency and real-time applications. The findings in this study offer a new quantitative overview of TinyML scholarship and serve as a basis for researchers and practitioners to identify influential sources, institutions, research themes, collaboration opportunities, and future research directions.

Declarations

Funding

The author received no specific funding for this research.

Conflict of Interest

The author declares no conflict of interest.

Author Contributions (CRediT)

Arpit Vyas: Conceptualization, Methodology, Data Curation, Formal Analysis, Investigation, Visualization, Writing – Original Draft, Writing – Review & Editing.

Ethics Approval and Informed Consent

Not applicable. This study is based exclusively on bibliographic data obtained from Scopus and did not involve human participants or animals.

Data Availability

The bibliographic data analyzed in this study were retrieved from the Scopus database on 5 July 2026. The analyzed data may be made available by the author upon reasonable request, subject to Scopus licensing and database-access conditions.

Acknowledgements

Not applicable.

References

- Abadade, Y., Benamar, N., Bagaa, M., & Chaoui, H. (2024). Empowering healthcare: TinyML for precise lung disease classification. *Future Internet*, 16(11), 391. <https://doi.org/10.3390/fi16110391>
- Abadade, Y., Temouden, A., Bamoumen, H., Benamar, N., Chtouki, Y., & Hafid, A. S. (2023). A comprehensive survey on TinyML. *IEEE Access*, 11, 96892–96922. <https://doi.org/10.1109/ACCESS.2023.3294111>
- Alharthi, S., Rashid, M., & Aljabri, M. (2026). TinyML in industrial IoT: A systematic review of applications, system components, and methodologies. *Sensors*, 26(8), 2550. <https://doi.org/10.3390/s26082550>
- Baciu, V.-E., Braeken, A., Segers, L., & Silva, B. D. (2025). Secure tiny machine learning on edge devices: A lightweight dual attestation mechanism for machine learning. *Future Internet*, 17(2), 85. <https://doi.org/10.3390/fi17020085>
- Elhanashi, A., Dini, P., Saponara, S., & Zheng, Q. (2024). Advancements in TinyML: Applications, limitations, and impact on IoT devices. *Electronics*, 13(17), 3562. <https://doi.org/10.3390/electronics13173562>
- Essahraoui, S., & Lamaakal, I. (2026). A comprehensive survey of TinyML-based biometric recognition for IoT edge devices. *IEEE Internet of Things Journal*, 13(6), 10564–10588. <https://doi.org/10.1109/JIOT.2026.3656932>
- Gookyi, D. A. N., Wulnye, F. A., Arthur, E. A. E., Ahiadormey, R. K., Agyemang, J. O., Agyekum, K. O.-B. O., & Gyaang, R. (2024). TinyML for smart agriculture:

- Comparative analysis of TinyML platforms and practical deployment for maize leaf disease identification. *Smart Agricultural Technology*, 8, 100490. <https://doi.org/10.1016/j.atech.2024.100490>
- Han, H., Trimi, S., & Lee, S. M. (2026). Tiny Machine Learning (TinyML): Research trends and future application opportunities. *Array*, 29, 100674. <https://doi.org/10.1016/j.array.2025.100674>
- Heydari, S., & Mahmoud, Q. H. (2025). Tiny machine learning and on-device inference: A survey of applications, challenges, and future directions. *Sensors*, 25(10), 3191. <https://doi.org/10.3390/s25103191>
- Hinostroza, A., Tarrillo, J., & Nuñez, M. (2026). Smart energy management in agricultural wireless sensor nodes using TinyML-based adaptive sampling. *Sensors*, 26(7), 2014. <https://doi.org/10.3390/s26072014>
- Huckelberry, J., Zhang, Y., Sansone, A., Mickens, J., Bereel, P., & Janapa Reddi, V. (2026). TinyML security: Attacks, defenses, and open challenges in resource-constrained machine learning systems. *ACM Computing Surveys*. <https://doi.org/10.1145/3793197>
- Immonen, R., & Hämäläinen, T. (2022). Tiny machine learning for resource-constrained microcontrollers. *Journal of Sensors*, 2022, 1–11. <https://doi.org/10.1155/2022/7437023>
- Lê, M. T., Wolinski, P., & Arbel, J. (2026). Efficient neural networks for tiny machine learning: A comprehensive review. *ACM Transactions on Intelligent Systems and Technology*, 17(4), 1–41. <https://doi.org/10.1145/3798276>
- Pazmiño Ortiz, L. A., Maldonado Soliz, I. F., & Guevara Balarezo, V. K. (2025). Advancing TinyML in IoT: A holistic system-level perspective for resource-constrained AI. *Future Internet*, 17(6), 257. <https://doi.org/10.3390/fi17060257>
- Pereira, E., Barboza, E., Araújo, Í., dos Santos, S., de Andrade, G., da Silva, I., & Martins, A. (2026). Robustness in TinyML: A systematic literature review. *ACM Transactions on Embedded Computing Systems*, 25(5), Article 92. <https://doi.org/10.1145/3820656>
- Rajapakse, V., Karunanayake, I., & Ahmed, N. (2023). Intelligence at the extreme edge: A survey on reformable TinyML. *ACM Computing Surveys*, 55(13s), 1–30. <https://doi.org/10.1145/3583683>
- Sanchez-Iborra, R., Zoubir, A., Hamdouchi, A., Idri, A., & Skarmeta, A. (2023). Intelligent and efficient IoT through the cooperation of TinyML and edge computing. *Informatica*, 34(1), 147–168. <https://doi.org/10.15388/22-INFOR505>
- Trigkas, A., Piromalis, D., & Papageorgas, P. (2025). Edge intelligence in urban landscapes: Reviewing TinyML applications for connected and sustainable smart cities. *Electronics*, 14(14), 2890. <https://doi.org/10.3390/electronics14142890>
- Vu, T.-H., Tu, N. H., Huynh-The, T., Voznak, M., Lee, K., Kim, S., & Pham, Q.-V. (2026). Integration of TinyML and LargeML: A survey of 6G and beyond. *IEEE Internet of Things Journal*, 13(10), 20558–20582. <https://doi.org/10.1109/JIOT.2026.3674358>
- Wainbuch, R., & Samuel, A. J. (2024). TinyML: Deploying machine learning on microcontrollers for IoT applications. *Journal of Science, Technology and Engineering Research*, 2(2), 44–57. <https://doi.org/10.64206/pfpcet13>
- Warden, P., & Situnayake, D. (2019). *TinyML: Machine learning with TensorFlow Lite on Arduino and ultra-low-power microcontrollers*. O'Reilly Media.
- Yahyati, C., Lamaakal, I., Maleh, Y., El Makkaoui, K., Ouahbi, I., Almousa, M., & Abd El-Latif, A. A. (2025). A systematic review of state-of-the-art TinyML applications in healthcare, education, and transportation. *IEEE Access*, 13, 204513–204562. <https://doi.org/10.1109/ACCESS.2025.3633575>