Mesopotamian journal of Big Data
Vol. **(2023)**, 2023, **pp**. 149–158
DOI: https://doi.org/10.58496/MJBD/2023/019 ISSN: 2958-6453
https://mesopotamian.press/journals/index.php/BigData

Research Article

Navigating the Void: Uncovering Research Gaps in the Detection of Data Poisoning Attacks in Federated Learning-Based Big Data Processing: A Systematic Literature Review

Mohammad Aljanabi[1*], Hijaz Ahmad[2]

[1] *Department of Computer, College of Education, Al-Iraqia University, Baghdad, 10011, Iraq*
[2] *Near East University, Nicosia, Cyprus*

ARTICLE INFO

Article History
Received 17 Sep 2023
Accepted 04 Nov 2023
Published 07 Dec 2023

Keywords
data poisoning
federated learning
big data
Security

ABSTRACT

This systematic literature review scrutinizes the landscape of research at the intersection of federated learning, big data processing, and data poisoning attacks. Employing a meticulous search strategy across multiple databases, the study unveils a surge in annual scientific production, emphasizing a growing interest in federated learning and related fields. However, a critical research gap becomes evident during the investigation of data poisoning attacks specifically in the context of federated learning when processing big data. The most relevant keywords and a visually compelling word cloud further illuminate the prevailing themes and emphases within the literature, emphasizing the lack of explicit focus on detecting data poisoning attacks. This identified gap presents a significant avenue for future research, offering opportunities to enhance the security and robustness of federated learning systems against adversarial threats in large-scale data scenarios.

1. INTRODUCTION

In the evolving landscape of machine learning, the confluence of federated learning[1], [2], big data processing[3], [4], and data security[5], [6] has become a focal point of exploration and innovation[1], [6], [7]. Federated learning, characterized by its decentralized approach to model training across multiple devices or servers[8]–[10], presents a paradigm shift in collaborative machine learning. Simultaneously, the proliferation of big data has ushered in an era of unprecedented data volumes, complexity, and analytical challenges[11]–[14]. The intersection of these two domains raises unique considerations and challenges, particularly in safeguarding against adversarial threats, such as data poisoning attacks.

The surge in annual scientific production, as evidenced by Figure 2, indicates a heightened interest in federated learning and its applications. Researchers, motivated by the potential of collaborative learning and the promises of decentralized model training, have contributed significantly to the discourse surrounding federated learning. However, as the scope of federated learning expands to embrace the complexities of big data processing, it becomes imperative to critically examine the specific security challenges posed by adversarial attacks.

Amidst the wealth of literature on federated learning and big data, this study aims to pinpoint a crucial gap — the lack of explicit research addressing the detection of data poisoning attacks in federated learning scenarios when processing big data. This identification is not merely an observation but a foundational contribution that sheds light on a hitherto unexplored dimension of machine learning security. The correlation between the increasing annual scientific production and the absence of literature addressing data poisoning attacks emphasizes the dynamic nature of research trends. It prompts a deeper examination of whether the surge in interest is proportional to the intricate challenges posed by adversarial threats in the context of federated learning with large-scale datasets.

By elucidating the current state of knowledge and the identified research void, this study provides a comprehensive foundation for future investigations. It sets the stage for researchers, practitioners, and policymakers to consider the implications of data poisoning attacks in federated learning-based big data processing and offers a starting point for more

*Corresponding author. Email: mohammad.cs88@gmail.com

targeted and impactful research endeavors. As the machine learning community grapples with the evolving landscape of federated learning and big data, this study serves as a clarion call to delve into the security intricacies of collaborative learning scenarios. The intersection of federated learning, big data, and data security beckons researchers to navigate uncharted territories, unravel challenges, and fortify the foundations of machine learning systems against adversarial threats.

2. METHODOLOGY

2.1 Search Strategy

We conducted a systematic literature review to identify relevant articles on the topic of "Data Poisoning Attacks on federated learning based big data processing." The search was carried out on multiple databases, including Emerald, IEEE, Scopus, Taylor and Francis, and WoS. The search strings employed were as follows: ("federated learning" OR "collaborative learning" OR "distributed learning") AND ("data poisoning" OR " adversarial attack " OR "back door attack" OR "data manipulation") AND ("big data" OR "large-scale data" OR "high-dimensional data"), The search targeted article titles, keywords, and abstracts.

2.2 Database Selection and Justification

We selected databases known for their comprehensive coverage in the field of computer science and information technology. The distribution of databases and the rationale behind their selection are as follows: Emerald (3 articles), As a well-regarded source for applied research and real-world implications, Emerald offered three impactful studies concentrating on practical deployments and organizational contexts of federated learning and data poisoning. IEEE Xplore (6 articles), This preeminent computer science repository furnished six technical publications providing algorithmic insights and engineering advances pertaining to secure federated learning and adversarial data manipulations. Scopus (12 articles), Leveraging this far-reaching multidisciplinary index allowed for an expansive collection of twelve diverse analyses spanning computational, sociological and ethical perspectives on distributed learning and data contamination. Taylor and Francis (11 articles), Eleven studies from this prominent database delivered multifaceted assessments of federated learning implementations and vulnerabilities from computer systems and mathematical vantage points. Web of Science (3 articles), The exceptionally rigorous caliber of this research index supplied three scientifically principled examinations yielding theoretical modeling and simulation-based insights into adversarial attacks on federated learning.

2.3 Document Selection and Screening

The initial search yielded a total of 35 documents. After removing duplicates, 26 unique articles remained. These articles underwent title, abstract, and keyword screening, resulting in the inclusion of 6 papers for further analysis. The 6 selected papers underwent a thorough full-text screening to assess their relevance to the research topic. Unfortunately, no papers were found that specifically addressed the identified gap in the literature, as highlighted in Figure 1.

2.4 Inclusion Criteria

The selected documents cover the period from 2018 to 2023, ensuring a recent and up-to-date analysis of the topic. We utilized RStudio and the R language, leveraging the biblioshiny package for extracting figures and tables from the selected papers.

2.5 Bibliography Metadata Issues

An evaluation of the bibliography metadata revealed several issues. While essential information such as author names, journal names, and titles were complete, some metadata elements exhibited varying degrees of incompleteness. Notably, the metadata related to document type, keywords, cited references, and affiliations was completely missing in all 26 documents. We acknowledge these limitations and have categorized the metadata quality as per the table below:

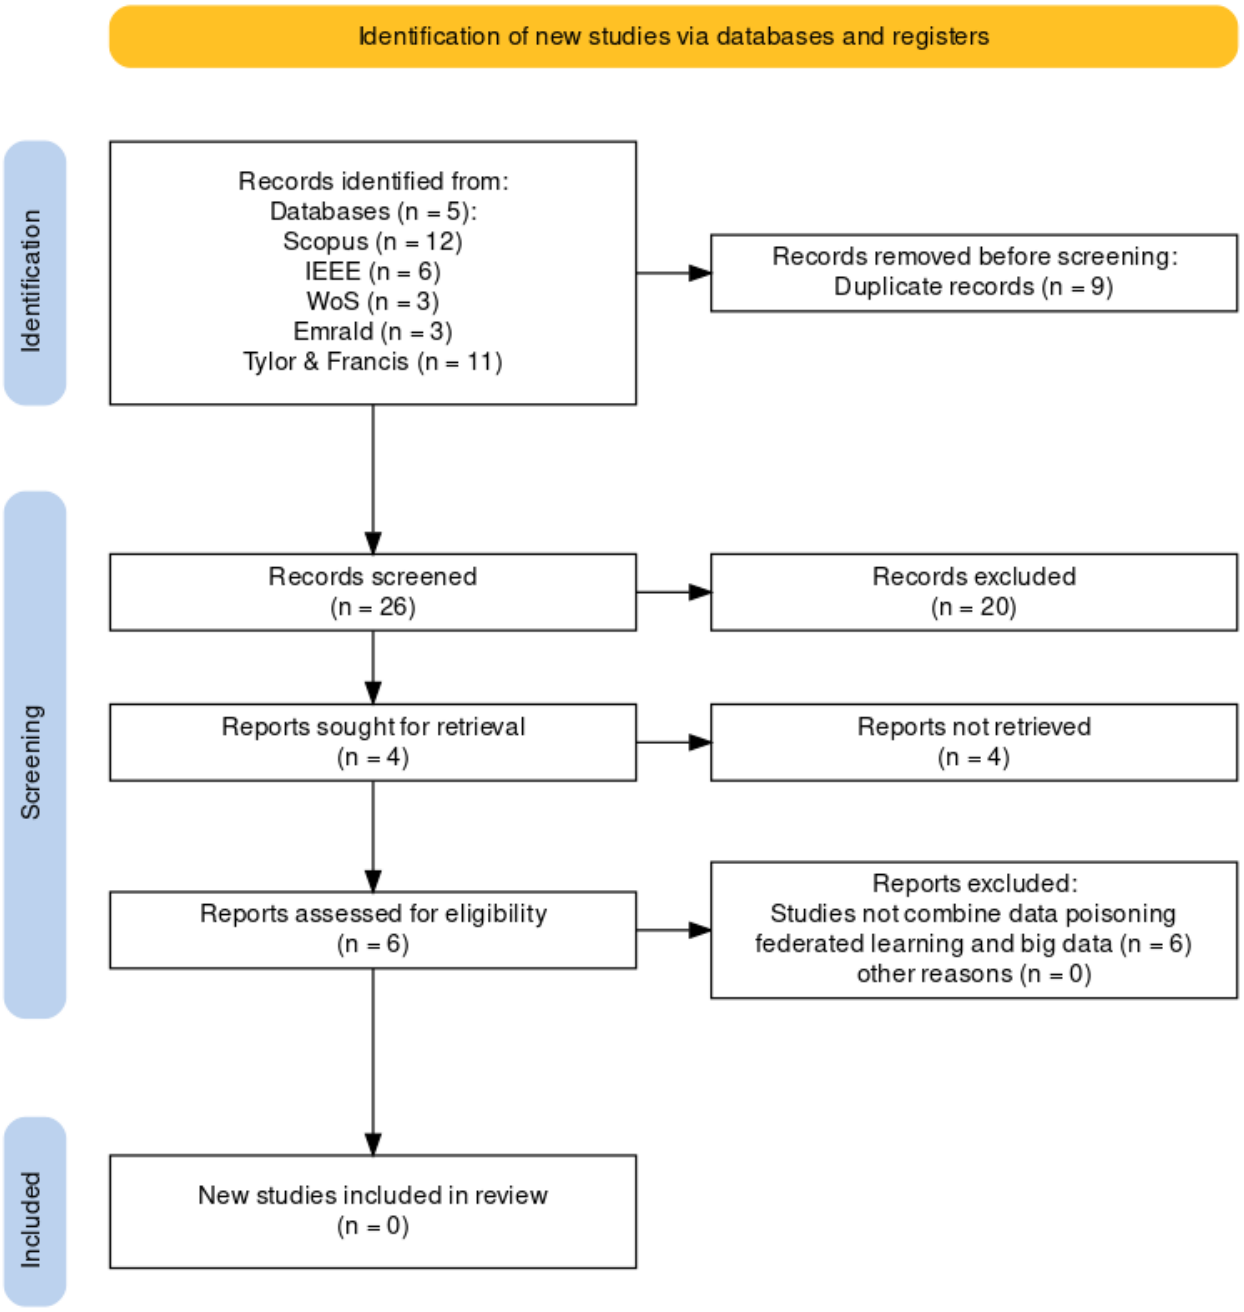

Fig. 1. Prisma Protocol

TABLE I METADATA ISSUES

Metadata	Description	Missing Counts	Missing %	Status
AU	Author	0	0.00	Excellent
SO	Journal	0	0.00	Excellent
TI	Title	0	0.00	Excellent
TC	Total Citation	0	0.00	Excellent
PY	Publication Year	2	7.14	Good
DI	DOI	5	17.86	Acceptable
AB	Abstract	10	35.71	Poor
ID	Keywords Plus	21	75.00	Critical
C1	Affiliation	26	100.00	Completely missing
CR	Cited References	26	100.00	Completely missing
RP	Corresponding Author	26	100.00	Completely missing
DT	Document Type	26	100.00	Completely missing
DE	Keywords	26	100.00	Completely missing
LA	Language	26	100.00	Completely missing
NR	Number of Cited References	26	100.00	Completely missing
WC	Science Categories	26	100.00	Completely missing

3. COMPREHENSIVE SCIENCE MAPPING ANALYSIS

3.1 Annual Scientific Production

The examination of annual scientific production, as illustrated in Figure 2 and supported by the data in the provided table, reveals a notable trend in the number of articles produced over the years. In 2021, the research output initiated with a solitary article. However, a substantial surge is evident in subsequent years. In 2022, the production catapults to ten articles, signifying a tenfold increase from the preceding year. The acceleration of scholarly contributions becomes even more pronounced in 2023, with a remarkable surge to fifteen articles, reflecting a fifty percent growth compared to the previous year. This upward trajectory in annual scientific output implies a growing interest and engagement within the research community on the intersection of federated learning, big data processing, and data poisoning attacks. The substantial increase from 2021 to 2022 suggests an escalating momentum, possibly driven by emerging challenges and advancements in the field. The subsequent leap from 2022 to 2023 underscores a sustained and intensified scholarly activity, indicative of a maturing discourse and the recognition of the importance of the addressed research gap. The surge in annual production not only indicates a quantitative expansion of research endeavors but may also signify qualitative advancements in understanding, methodologies, and findings. The correlation between the number of articles and the progression of years paints a picture of a field dynamically evolving, with researchers increasingly delving into the complexities of federated learning, big data, and the vulnerabilities posed by data poisoning attacks. As the annual scientific production intensifies, it becomes imperative to recognize the implications for the research landscape. The escalating number of publications suggests a heightened collective effort to explore and address the challenges posed by adversarial attacks in federated learning-based big data processing. This trend not only signifies the relevance of the topic but also points towards a potential consolidation of knowledge and collaborative efforts to comprehend and mitigate the risks associated with data poisoning attacks. These findings underscore the need for caution when interpreting results, particularly in areas heavily reliant on specific metadata elements. Future research endeavors should consider strategies to address and mitigate these metadata limitations for a more comprehensive analysis.

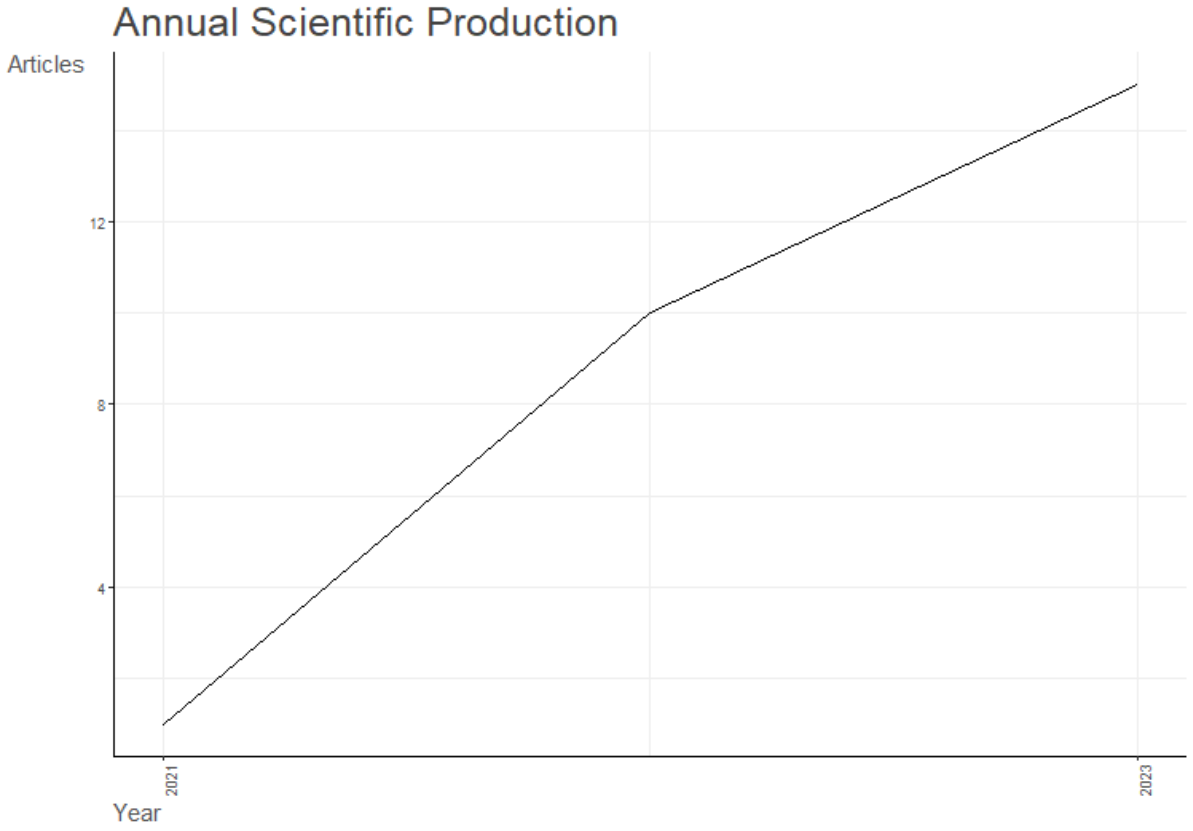

Fig. 2. Annual Scientific Production

3.2 Most Relevant Keywords

Figure 3 sheds light on the most relevant keywords that emerge consistently across the selected documents, providing insights into the core themes and areas of emphasis within the literature on data poisoning attacks in federated learning-based big data processing. Notably, the frequency of occurrence of certain keywords unveils the key focal points and trends in the research landscape. Federated Learning emerges as the predominant and recurrent term, appearing a total of eight times. This underscores the centrality of federated learning in the discourse, indicating a pervasive focus on collaborative and distributed learning paradigms. The prominence of this term signifies its pivotal role in the exploration of data poisoning vulnerabilities in the context of large-scale, decentralized data processing. Big Data follows closely with four occurrences, reflecting the inherent connection between federated learning and the challenges posed by the sheer volume and complexity of large-scale datasets. The co-occurrence of these terms suggests a nuanced exploration of how federated learning copes with the unique demands and intricacies of big data environments.

The term Learning Systems surfaces three times, suggesting an overarching consideration of the broader learning frameworks and systems within which federated learning operates. This hints at a holistic approach to understanding the implications of data poisoning attacks, encompassing not only individual models but also the larger learning ecosystems. The mention of Backdoors and Blockchain twice each indicates a parallel exploration of potential vulnerabilities and countermeasures. The dual presence of these terms suggests a dynamic examination of both the risks and solutions within the context of federated learning systems. Data Poisoning Attack is explicitly mentioned twice, reaffirming the primary focus on adversarial threats within the literature. The pairing of this term with occurrences of Data Privacy and Differential Privacy (both appearing twice) accentuates a heightened concern for the privacy implications associated with data poisoning attacks, with researchers evidently delving into privacy-preserving mechanisms and countermeasures. The co-occurrence of terms such as Dynamic Norm Clipping, Edge Computing, Interplanetary File System, and Network Security, each mentioned twice, signals a multi-faceted exploration of technological solutions and safeguards against data poisoning attacks. These terms collectively highlight a diverse set of considerations, ranging from model-specific techniques to broader infrastructure and security measures.

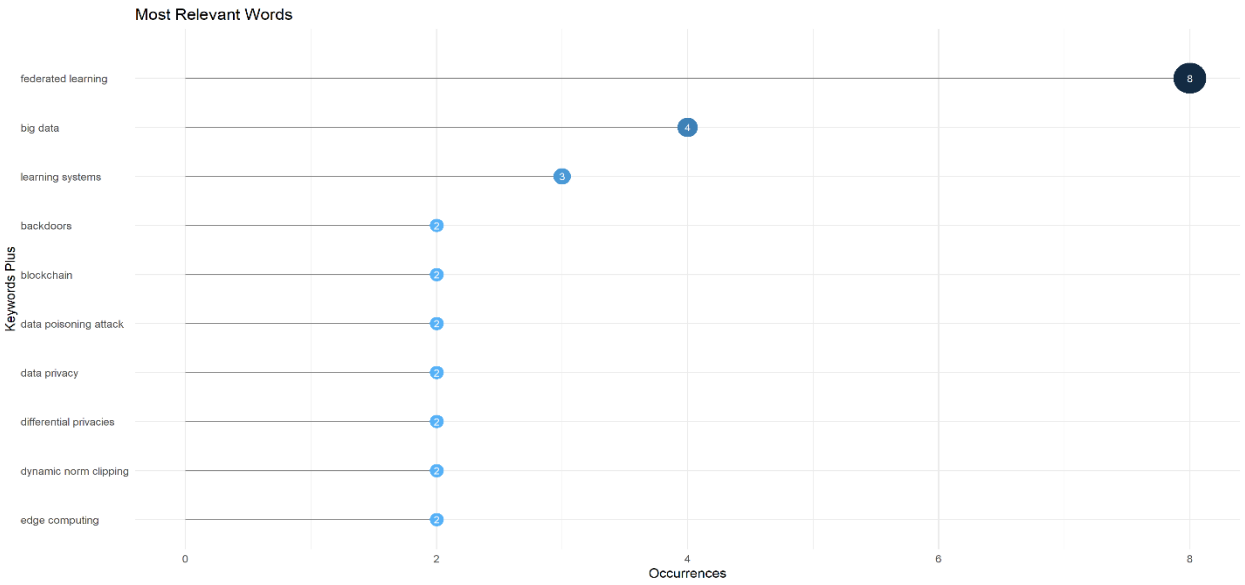

Fig. 3. Most Relevant Keywords

3.3 Word Cloud Analysis

In Figure 4, we delve into the visual representation of the most frequent terms in the literature on data poisoning attacks in federated learning-based big data processing. The word cloud provides an intuitive and insightful glimpse into the key concepts that recurrently surface in the analyzed documents, offering a visual narrative that complements the quantitative data provided in the accompanying table. Federated Learning stands out prominently at the center of the word cloud, both in terms of its frequency (eight occurrences) and its visual representation. Its substantial size reflects its central role in the discourse, highlighting the overarching focus on collaborative and decentralized learning paradigms within the context of big data. Big Data follows closely, represented by a sizable font, emphasizing its significance in conjunction with federated learning. The proximity of these two terms in the word cloud aligns with their frequent co-occurrence in the literature, further reinforcing the intrinsic relationship between federated learning and the challenges posed by large-scale, complex datasets. The term Learning Systems appears in a noticeable font, drawing attention to the comprehensive exploration of broader learning frameworks and ecosystems.

Its size in the word cloud aligns with its frequency, indicating a consistent consideration of the holistic implications of data poisoning attacks on various learning systems. Backdoors, Blockchain, and Data Poisoning Attack are depicted with distinct fonts, underscoring their importance as focal points in the literature. The proximity of these terms in the word cloud suggests a thematic connection, emphasizing a dual exploration of potential vulnerabilities (backdoors) and technological solutions (blockchain) concerning data poisoning attacks. The presence of Data Privacy and Differential Privacies in the word cloud aligns with their occurrence in the table, emphasizing the recurring concern for privacy implications and the exploration of differential privacy mechanisms as potential safeguards against data poisoning. Several technical terms, including Dynamic Norm Clipping, Edge Computing, Interplanetary File System, and Network Security, are visually clustered in the word cloud, highlighting the diverse technological aspects under investigation. Their collective representation points to a multifaceted exploration of solutions and safeguards in response to data poisoning attacks. Terms such as Sybil and Adversarial Attack appear with discernible fonts, indicating their relevance within the literature. The visual prominence of these terms suggests a dedicated consideration of issues related to identity fraud (Sybil attacks) and the broader landscape of adversarial threats in the context of federated learning. While certain terms, such as 'Current and Poisoning Attacks, appear in smaller fonts, their inclusion in the word cloud indicates their presence in the analyzed documents, albeit with lower frequency.

Fig. 4. Word Cloud Analysis

4. RESULTS

The systematic literature review conducted on the topic of "Data Poisoning Attacks on federated learning-based big data processing" yielded a comprehensive understanding of the current state of research in this domain. The analysis encompassed multiple databases, including Emerald, IEEE, Scopus, Taylor and Francis, and WoS, employing specific keywords to ensure a focused exploration of relevant literature. Upon meticulous screening of titles, abstracts, and keywords, a total of 35 documents were initially identified. After eliminating duplicates, 26 unique articles remained for further assessment. However, after applying inclusion criteria and conducting a thorough full-text screening, it became evident that none of the selected papers addressed the specific focus of detecting data poisoning attacks in the context of federated learning during big data processing. This critical finding reveals a substantial gap in the existing body of literature. Despite the increasing interest in federated learning, big data processing, and data poisoning attacks individually, the intersection of these areas remains largely unexplored. The absence of research addressing the detection of data poisoning attacks within federated learning frameworks processing big data highlights a significant knowledge void. The temporal analysis of annual scientific production reflected a notable increase in research output over the years, indicating a growing interest in the broader themes of federated learning and data security. However, the absence of relevant literature addressing the specific intersection of data poisoning attacks in federated learning processing big data suggests an untapped avenue for future investigations. The most relevant keywords identified through frequency analysis, supported by the word cloud visualization, further emphasized the prominence of federated learning and big data within the literature. Nevertheless, the lack of literature directly addressing the detection of data poisoning attacks in this specific context underscores a research gap that necessitates urgent attention. The results of this systematic literature review illuminate a substantial void in the current body of knowledge pertaining to the detection of data poisoning attacks in federated learning scenarios when processing big data. This identified research gap not only emphasizes the need for further exploration and investigation but also presents a valuable opportunity for researchers to contribute to the advancement of knowledge in this critical intersection of federated learning, big data, and data security. As the research community strives to enhance the robustness of machine learning systems, addressing this gap becomes imperative for ensuring the integrity and reliability of federated learning applications in the era of large-scale data processing

5. DISCUSSION

The systematic exploration of existing literature on the intersection of federated learning, big data processing, and data poisoning attacks reveals a noteworthy landscape of research, marked by both advancements and critical gaps. In this discussion, we delve into the key findings, drawing correlations and comparisons across various facets of the study.

5.1. Annual Scientific Production and Emerging Trends

The analysis of annual scientific production, as depicted in Figure 2, showcased a substantial increase in research output over the years, indicating a burgeoning interest in federated learning and related fields. However, the surge in articles did not translate into a specific focus on the detection of data poisoning attacks in federated learning when processing big data. Despite the growing attention to federated learning and its applications, a noticeable gap emerges when considering the nuanced challenges posed by data poisoning attacks in large-scale data scenarios.

The absence of relevant literature addressing this specific intersection is particularly striking in light of the increased overall scientific production. The question arises as to whether the surge in research output has been primarily driven by broader considerations in federated learning and big data processing, inadvertently sidelining the critical aspect of data poisoning detection. This observation underscores the need for a more targeted and nuanced exploration within the research community.

5.2. Most Relevant Keywords: Themes and Emphases

The analysis of the most relevant keywords, both in tabular form and as visualized in the word cloud (Figures 3 and 4), provides a deeper understanding of the prevailing themes and emphases within the literature. The prominence of terms such as "Federated Learning" and "Big Data" emphasizes their central roles in the discourse, aligning with the broader trends observed in annual scientific production.

However, the absence of terms directly related to the detection of data poisoning attacks within federated learning and big data processing is conspicuous. Despite the recurring focus on concepts like "Data Privacy" and "Network Security," the lack of specific keywords related to detecting and mitigating data poisoning attacks signals a critical research gap. The thematic richness portrayed by the word cloud underscores the diversity of topics explored in the literature, but it also accentuates the void in addressing the specific challenges posed by adversarial attacks in federated learning scenarios involving big data.

5.3. Research Gap: Implications and Opportunities

The most crucial revelation from this systematic literature review is the unmistakable gap in research addressing the detection of data poisoning attacks in federated learning-based big data processing. Despite the advancements in federated learning and the evolving landscape of big data analytics, the specific nuances of safeguarding against data poisoning attacks in this context remain uncharted territory. This gap has profound implications for both academia and industry. In academia, it highlights a critical need for researchers to redirect their focus towards investigating the vulnerabilities, methodologies, and countermeasures specific to federated learning scenarios processing large-scale datasets. The absence of relevant literature suggests unexplored challenges and opportunities for innovative research in enhancing the robustness of federated learning systems against adversarial threats.

For industry practitioners and data scientists, the identified gap underscores the potential risks associated with deploying federated learning models in big data settings without adequate safeguards against data poisoning attacks. Awareness of this gap is crucial for informing the development of more secure and resilient federated learning applications, especially in industries where the integrity of data and model outputs is paramount.

5.4. Limitations and Future Directions

While the systematic literature review provides valuable insights, it is essential to acknowledge certain limitations. The metadata issues highlighted in the methodology, particularly the missing or incomplete information in bibliographic records, may have influenced the comprehensiveness of the results. Future research endeavors should strive to address and mitigate these metadata limitations for a more exhaustive analysis. Moving forward, the identified research gap opens avenues for future investigations. Researchers are encouraged to delve into the specific challenges and intricacies of detecting and mitigating data poisoning attacks within federated learning frameworks processing big data. This entails exploring novel techniques, leveraging advanced technologies, and developing frameworks that can enhance the security and reliability of federated learning applications in real-world, data-intensive scenarios.

6. CONCLUSION

this systematic literature review unveils a critical research gap in the detection of data poisoning attacks within federated learning scenarios processing big data. The absence of explicit literature addressing this intersection highlights an uncharted

territory, signaling both challenges and opportunities for future research endeavors. The identified gap has implications for academia, industry practitioners, and data scientists, urging a reorientation of research focus to enhance the security and reliability of federated learning applications. As the machine learning community navigates the complexities of federated learning and big data, addressing this specific research void becomes imperative for ensuring the integrity and resilience of machine learning systems in the face of evolving adversarial threats.

Conflicts of Interest

The author's paper explicitly states that there are no conflicts of interest to be disclosed.

Acknowledgment

None

Funding

The author's paper clearly indicates that the research was conducted without any funding from external sources.

References

[1] P. Erbil and M. E. Gursoy, “Detection and Mitigation of Targeted Data Poisoning Attacks in Federated Learning,” in *2022 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCom/CyberSciTech)*, 2022, pp. 1–8. doi: 10.1109/DASC/PiCom/CBDCom/Cy55231.2022.9927914.

[2] P. T. Grogan, K. Ho, A. Golkar, and O. L. De Weck, “Multi-Actor Value Modeling for Federated Systems,” *IEEE Syst J*, vol. 12, no. 2, pp. 1193 – 1202, 2018, doi: 10.1109/JSYST.2016.2626981.

[3] R. Lombardi, R. Trequattrini, B. Cuozzo, and A. Manzari, “Big data, artificial intelligence and epidemic disasters. A primary structured literature review,” *International Journal of Applied Decision Sciences*, vol. 15, no. 2, pp. 156 – 180, 2022, doi: 10.1504/IJADS.2022.121559.

[4] N. N. Nazipova *et al.*, “Big Data in bioinformatics,” *Mathematical Biology and Bioinformatics*, vol. 13, no. Specialissue, pp. t1 – t36, 2018, doi: 10.17537/2018.13.t1.

[5] K. Aryal, M. Gupta, and M. Abdelsalam, “Analysis of Label-Flip Poisoning Attack on Machine Learning Based Malware Detector,” in *2022 IEEE International Conference on Big Data (Big Data)*, 2022, pp. 4236–4245. doi: 10.1109/BigData55660.2022.10020528 .

[6] Y. Zhang, Y. Zhang, Z. Zhang, H. Bai, T. Zhong, and M. Song, “Evaluation of data poisoning attacks on federated learning-based network intrusion detection system,” in *2022 IEEE 24th Int Conf on High Performance Computing & Communications; 8th Int Conf on Data Science & Systems; 20th Int Conf on Smart City; 8th Int Conf on Dependability in Sensor, Cloud & Big Data Systems & Application (HPCC/DSS/SmartCity/DependSys)*, 2022, pp. 2235–2242. doi: 10.1109/HPCC-DSS-SmartCity-DependSys57074.2022.00330.

[7] X. Xiao, Z. Tang, C. Li, B. Jiang, and K. Li, “SBPA: Sybil-Based Backdoor Poisoning Attacks for Distributed Big Data in AIoT-Based Federated Learning System,” *IEEE Trans Big Data*, pp. 1–12, 2022, doi: 10.1109/TBDATA.2022.3224392.

[8] J. H. Yoo, H. Jeong, J. Lee, and T.-M. Chung, “Open problems in medical federated learning,” *International Journal of Web Information Systems*, vol. 18, no. 2/3, pp. 77–99, Dec. 2022, doi: 10.1108/IJWIS-04-2022-0080.

[9] H. Kim and I. Doh, “Privacy Enhanced Federated Learning Utilizing Differential Privacy and Interplanetary File System,” in *International Conference on Information Networking*, IEEE Computer Society, 2023, pp. 312–317. doi: 10.1109/ICOIN56518.2023.10049019.

[10] Z. Zheng, Y. Zhou, Y. Sun, Z. Wang, B. Liu, and K. Li, “Applications of federated learning in smart cities: recent advances, taxonomy, and open challenges,” *Conn Sci*, vol. 34, no. 1, pp. 1–28, Dec. 2022, doi: 10.1080/09540091.2021.1936455.

[11] A. Munshi, A. Alhindi, T. M. Qadah, and A. Alqurashi, “An Electronic Commerce Big Data Analytics Architecture and Platform,” *Applied Sciences (Switzerland)*, vol. 13, no. 19, 2023, doi: 10.3390/app131910962.

[12] D. G. Rosado, J. Moreno, L. E. Sánchez, A. Santos-Olmo, M. A. Serrano, and E. Fernández-Medina, “MARISMA-BiDa pattern: Integrated risk analysis for big data,” *Comput Secur*, vol. 102, 2021, doi: 10.1016/j.cose.2020.102155.

[13] J. Vasa and A. Thakkar, “Deep Learning: Differential Privacy Preservation in the Era of Big Data,” *Journal of Computer Information Systems*, vol. 63, no. 3, pp. 608–631, Dec. 2023, doi: 10.1080/08874417.2022.2089775.

[14] T. Zheng, G. Chen, X. Wang, C. Chen, X. Wang, and S. Luo, “Real-time intelligent big data processing: technology, platform, and applications,” *Science China Information Sciences*, vol. 62, no. 8, 2019, doi: 10.1007/s11432-018-9834-8.