Federated Learning for Robotic and Autonomous Systems: A Survey on Architectures, Synergies with Distributed Ledger Technologies, and Future Directions

التعلم الفيدرالي للأنظمة الروبوتية والمستقلة: دراسة استقصائية حول البنيات، والتآزر مع تقنيات دفتر الأستاذ الموزع، والتوجهات المستقبلية

Abdelrazak A. Yousef Elbunan
Computer Science Department Higher Institute of Science and Technology Mesallata, Libya

عبدالرزاق يوسف البونان
قسم علوم الحاسب
المعهد العالي للعلوم والتقنية / مسلاته

Abdelrazak.elbunan1982@gmail.com

Nuradeen K. Emhemed Fethalla
Electrical and Computer Engineering Department Elmergib University, Faculty of Engineering Alkhoms, Libya

نورالدين امحمد فتح الله
قسم الهندسة الكهربائية والحاسوب
جامعة المرقب / كلية الهندسة

nkfethalla@elemergib.edu.ly

Badriya Abdullah Altarhuni
Department of Information Technology, Higher Institute of Science and Technology Tajoura Libya

بدرية عبدالله الترهوني
قسم تقنيات الحاسوب
المعهد العالي للعلوم والتقنية / تاجوراء

badriyaabdullh@gmail.com

Abstract

The rapid proliferation of autonomous robotic systems, ranging from nano-drones to industrial collaborative robots (cobots), is generating massive, distributed datasets. While deep learning (DL) serves as the cornerstone of modern robotic intelligence, the conventional approach of centralizing this data for training poses insurmountable challenges related to privacy, security, bandwidth, and latency. Federated Learning (FL) has emerged as a disruptive paradigm that enables collaborative

model training across distributed devices without the need for raw data exchange. However, the integration of FL into real-world robotic swarms—characterized by extreme heterogeneity, dynamic connectivity, and stringent resource constraints—introduces a unique set of complexities that extend far beyond those of conventional edge devices. This survey provides a comprehensive and critical examination of the burgeoning field of FL within robotic and autonomous systems. We move beyond a mere overview to present a novel taxonomy that classifies FL architectures for robotics based on communication topology, learning paradigm, and application criticality. A significant portion of our analysis is dedicated to the potent synergy between FL and Distributed Ledger Technologies (DLTs), particularly blockchain, for achieving decentralized trust, auditability, and robust aggregation in the presence of potentially malicious agents. We extensively review applications across perception, control, and collaborative tasks, highlighting pioneering works in multi-robot SLAM, federated reinforcement learning, and human-robot interaction. Furthermore, we identify and discuss pressing open challenges, including communication efficiency in mobile swarms, energy-aware client selection, personalized learning for non-IID data, and defense mechanisms against sophisticated adversarial attacks. This paper serves as a foundational reference for researchers and practitioners aiming to develop the next generation of private, secure, and collectively intelligent robotic systems.

Keywords: Autonomous Robotic Systems, Federated Learning, Distributed Ledger Technologies, Deep Learning**,** Collaborative Robots**,** Human-robot interaction

الملـخـص

يُولّد الانتشار السريع للأنظمة الروبوتية ذاتية التشغيل بدءًا من الطائرات النانوية بدون طيار وصولًا إلى الروبوتات التعاونية الصناعية (cobots) مجموعات بيانات هائلة وموزعة. وبينما يُشكّل التعلم العميق (DL) حجر الزاوية في الذكاء الاصطناعي الروبوتي الحديث فإن النهج التقليدي لتجميع هذه البيانات لأغراض التدريب يُشكّل تحدياتٍ جسيمة تتعلق بالخصوصية والأمان وعرض النطاق الترددي وزمن الوصول. وقد برز التعلم الفيدرالي

(FL) كنموذجٍ مُبتكر يُتيح تدريب النماذج التعاونية عبر الأجهزة الموزعة دون الحاجة إلى تبادل البيانات الخام. ومع ذلك، فإن دمج التعلم الفيدرالي في أسراب الروبوتات الواقعية التي تتميز بتباينٍ شديد، وترابطٍ ديناميكي، وقيودٍ صارمة على الموارد - يُقدّم مجموعةً فريدةً من التعقيدات التي تتجاوز بكثير تلك الموجودة في أجهزة الحافة التقليدية. تُقدّم هذه الدراسة دراسةً شاملةً ونقديةً لمجال التعلم الفيدرالي المزدهر ضمن الأنظمة الروبوتية والأنظمة ذاتية التشغيل. نتجاوز مجرد النظرة العامة لنقدم تصنيفًا جديدًا يُصنّف بنيات FL للروبوتات بناءً على طوبولوجيا الاتصالات، ونموذج التعلم، وأهمية التطبيق. يُكرّس جزء كبير من تحليلنا للتآزر الفعّال بين FL وتقنيات دفتر الأستاذ الموزع (DLTs)، وخاصةً تقنية البلوك تشين، لتحقيق ثقة لامركزية، وقابلية تدقيق، وتجميع قوي في ظل وجود عوامل ضارة محتملة. نستعرض بشكل موسع التطبيقات في مهام الإدراك والتحكم والتعاون، مع تسليط الضوء على الأعمال الرائدة في SLAM متعدد الروبوتات، والتعلم التعزيزي الفيدرالي، والتفاعل بين الإنسان والروبوت. علاوة على ذلك، نحدد ونناقش التحديات المفتوحة المُلحة، بما في ذلك كفاءة الاتصالات في أسراب الروبوتات المتنقلة، واختيار العملاء الواعي للطاقة، والتعلم المُخصّص للبيانات غير المُعرّفة بالبيانات (IID)، وآليات الدفاع ضد الهجمات العدائية المُعقدة. تُعدّ هذه الورقة مرجعًا أساسيًا للباحثين والممارسين الذين يسعون إلى تطوير الجيل القادم من أنظمة الروبوتات الخاصة والآمنة والذكية جماعيًا. الكلمات المفتاحية: أنظمة الروبوتات المستقلة، التعلم الفيدرالي، تقنيات دفتر الأستاذ الموزع، التعلم العميق، الروبوتات التعاونية، التفاعل بين الإنسان والروبوت.

1. Introduction

The convergence of robotics, artificial intelligence, and pervasive connectivity is ushering in an era of ubiquitous autonomous systems. From smart factories and precision agriculture to autonomous vehicles and search-and-rescue missions, networks of robots are expected to perform increasingly complex tasks in unstructured and often cooperative environments [1]. This shift is underpinned by advances in Deep Learning (DL), which has become the de facto standard for enabling high-level cognitive functions such as visual perception, natural language understanding, and adaptive control [2, 3]. Traditionally, the development of these DL models relies on the collection of vast amounts of data into centralized data centers for training a paradigm that is rapidly becoming untenable. The challenges are multifaceted:

(1) **Data Privacy and Security:** Robotic systems often operate in sensitive environments (e.g., homes, hospitals, industrial facilities), generating data that is proprietary, personal, or mission-critical. Transmitting and centralizing this data creates significant attack surfaces and risks of leakage [4].

(2) **Network Bandwidth:** The volume of sensor data generated by a single robot (e.g., high-resolution point clouds, video streams) can be enormous, making constant transmission to a cloud infrastructure prohibitively expensive and inefficient [5].

(3) **Latency:** For time-sensitive applications like autonomous navigation or real-time manipulation, the round-trip delay to a remote cloud server can be catastrophic [6].

(4) **Resource Constraints:** Many robots, especially smaller drones or mobile robots, have limited onboard computational, storage, and energy resources, making local processing of large models difficult [7].

Federated Learning (FL) presents a compelling alternative to this centralized approach. As first conceptualized by Google [8], FL is a machine learning setting where multiple clients (e.g., mobile phones, robots) collaboratively train a model under the orchestration of a central server, while keeping the training data decentralized. Each client performs local computation on its private data and shares only model updates (e.g., gradients, weights) with the server, which aggregates them to form a global model. This process inherently enhances privacy, reduces bandwidth consumption, and can leverage distributed compute resources [9].

While significant research has explored FL in the context of mobile and IoT devices [10, 11], its application within robotic and autonomous systems introduces distinct dimensions of complexity. Robots are not merely data sources; they are active agents whose actions influence the data they collect (leading to non-IID data distribution), they operate in dynamically changing environments, and they often have critical real-time performance requirements. Furthermore, the collaborative nature of multi-robot systems (MRS) demands learning frameworks that are not only privacy-preserving but also robust, scalable, and capable of

handling systemic heterogeneity. A conceptual illustration showing potential application areas is shown in Figure 1.

This survey aims to provide a thorough and nuanced analysis of the integration of FL into robotic and autonomous systems. We move beyond existing surveys focused solely on general FL security [12] or communication efficiency [13] by focusing specifically on the robotic domain's unique challenges and opportunities. Our contributions are fourfold:

1. We present a detailed background covering the essential pillars of modern robotic learning: Cloud/Fog Robotics, Distributed DL, and the critical security/privacy landscape.
2. We propose a novel taxonomy for FL in robotics, categorizing systems by their architecture, learning task, and operational constraints.
3. We provide an in-depth exploration of the integration of DLTs with FL, arguing that this combination is particularly well-suited for building trust and decentralization in adversarial multi-robot environments.
4. We conduct a comprehensive review of state-of-the-art applications, from perception and SLAM to control and human-robot collaboration, and delineate a detailed roadmap of future research directions.

The remainder of this paper is organized as follows. Section 2 establishes the necessary technical background. Section 3 delves into the deployment of FL at the edge, analyzing key sub-problems. Section 4 explores the synergies with DLTs. Section 5 surveys concrete applications. Finally, Section 6 concludes and outlines future work.

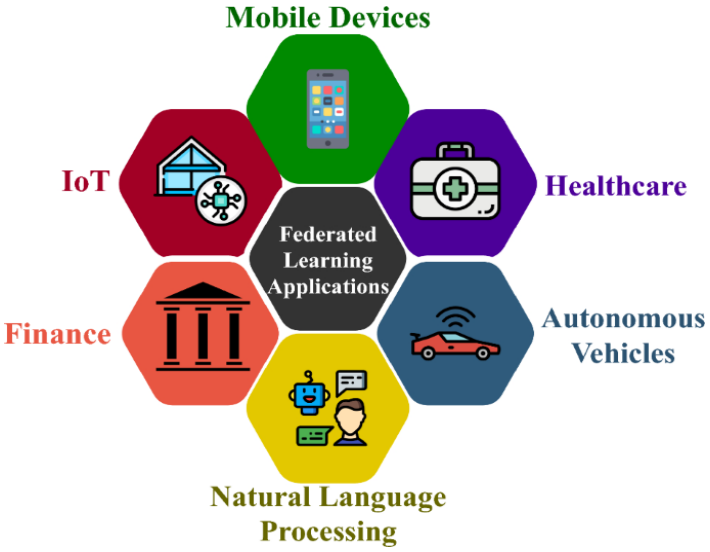

Figure 1 A conceptual illustration showing potential application areas

2. Background:

The successful adoption of FL in robotics is built upon the convergence of several technological trends. This section provides a foundational understanding of these areas, highlighting their relevance to the development of federated robotic systems.

2.1. Cloud, Fog, and Edge Robotics:

The paradigm of cloud robotics, first prominently discussed by Kuffner [14], leverages remote data centers to provide robots with vast computational resources and shared knowledge. This enables capabilities like offloading intensive DL inference, accessing massive datasets for learning, and enabling collective robot learning [15]. However, the latency and bandwidth limitations of cloud-only approaches led to the emergence of fog robotics [16] and edge robotics [17]. These paradigms distribute computation across the cloud-to-thing continuum, pushing processing closer to the data source. A fog node (e.g., a local server, a powerful robot) can aggregate data from multiple nearby robots, perform intermediate processing, and reduce the load on the cloud. FL aligns perfectly with this vision, acting as a structured framework for distributed learning across this hierarchical architecture, from the deep cloud to the extreme edge on the robot itself.

2.2. Distributed Deep Learning:

Training large DL models is computationally expensive. Distributed Deep Learning (DDL) seeks to accelerate this process by parallelizing computation across multiple devices [18]. Traditional DDL often involves data parallelism, where the global dataset is partitioned across workers who synchronize their gradients frequently. While FL shares the goal of distributed training, its fundamental constraint is that the data partitions are inherently private and non-IID, residing on separate edge devices. This makes direct application of many DDL synchronization protocols inefficient or infeasible. Research in FL has thus focused on developing specialized aggregation algorithms (like Federated Averaging - FedAvg [8]) and communication compression techniques [19] that are robust to these conditions, forming a distinct subfield of DDL tailored for privacy and decentralization.

2.3. Security and Privacy in Machine Learning:

The widespread use of DL has exposed critical vulnerabilities. Adversarial attacks can be broadly categorized into:

(1) **Evasion attacks**: (at inference time), where carefully crafted perturbations cause misclassification [20].

(2) **Poisoning attacks** (at training time), where malicious actors inject corrupted data or model updates to compromise the global model [21]. FL is vulnerable to both, though poisoning poses a more direct threat. Furthermore, even without malicious intent, privacy risks exist. Model inversion attacks can reconstruct parts of the training data from the shared model updates [22], and membership inference attacks can determine if a specific data sample was part of the training set [23]. Mitigation techniques are multi-layered, including **differential privacy (DP)** [24], which adds calibrated noise to updates to obscure any single data point's contribution; **homomorphic encryption (HE)** [25], which allows computation on encrypted data; and **secure multi-party computation (SMPC)** [26], which enables multiple parties to jointly compute a function without revealing their private inputs. A robust FL system for robotics must integrate these techniques to provide end-to-end security and privacy guarantees.

2.4. Federated and Distributed Reinforcement Learning:

Reinforcement Learning (RL), where an agent learns optimal behavior through trial-and-error interactions with an environment, is highly relevant to robotics. Multi-Agent RL (MARL) extends this to settings with multiple cooperating or competing agents [27]. Federated Reinforcement Learning (FRL) is a natural fusion of FL and RL/MARL, where multiple agents learn policies from their local experiences and periodically aggregate their knowledge (e.g., value functions, policy parameters) to form a global policy without sharing their raw experience trajectories [28]. This is particularly valuable for robots operating in similar but distinct environments, as it allows for collective learning while preserving the privacy of their specific operational data. FRL has been applied to problems ranging from resource allocation in networks [29] to learning robust control policies for autonomous systems [30].

3. Federated Learning at the Edge: Robotic Considerations:

Deploying FL in robotic networks amplifies challenges observed in general edge networks and introduces new ones. This section dissects these challenges and reviews proposed solutions.

3.1. System Heterogeneity and Task Allocation:

Robotic fleets are highly heterogeneous. Members may differ in computational capability (CPU/GPU), battery life, sensor suite, and mobility. This heterogeneity makes efficient task allocation—deciding which robots participate in which training rounds—a critical problem. Simple random selection can lead to stragglers that slow down the entire federated round. Matching-theoretic approaches [31] and reinforcement learning-based schedulers [32] have been proposed to dynamically match learning tasks with robots that have sufficient resources, thereby improving overall training efficiency and system throughput.

3.2. Communication Efficiency in Mobile Swarms:

Communication is the primary bottleneck in FL. For mobile robotic swarms, this is exacerbated by dynamic network topologies, packet loss, and limited bandwidth. Research focuses on:

- **Communication Reduction:** Techniques like gradient sparsification [33], quantization [34], and knowledge distillation [35] reduce the size of model updates transmitted.
- **Asynchronous FL:** Relaxing the synchronous aggregation requirement of FedAvg allows robots to contribute updates as they become available, mitigating the straggler problem and better suited for real-world deployments with intermittent connectivity [36].
- **Over-the-Air Computation (AirComp):** This innovative technique leverages the superposition property of wireless channels to allow multiple devices to transmit their model updates simultaneously, naturally achieving a form of analog aggregation and significantly reducing communication latency [37].

3.3. Energy-Aware Learning:

Battery life is a paramount concern for mobile robots. The energy cost of FL comes from both local computation (training) and communication (transmitting updates). Strategies to improve energy efficiency include:

- **Adaptive Local Training:** Dynamically adjusting the number of local training epochs based on battery level and channel conditions [38].
- **Joint Communication-Computation Optimization:** Formulating the training process as an optimization problem that minimizes total energy consumption under accuracy and latency constraints [39].
- **Hardware-Aware Design:** Utilizing specialized low-power AI accelerators on robots to improve the energy efficiency of the local training process itself [40].

3.4. Client Selection and Resource Management:

Beyond task allocation, the problem of client selection is crucial for robustness. Selecting clients based solely on resource availability might lead to a biased global model if the selected subset is not representative of the entire population. FedCS [41] is a framework that manages client selection in heterogeneous environments. More advanced approaches

use reinforcement learning to learn an optimal selection policy that balances resource constraints with statistical utility [42].

3.5. Enhanced Privacy and Security Mechanisms:

The standard FL protocol is vulnerable. Enhanced mechanisms for robotics include:

- **Byzantine-Robust Aggregation:** Aggregation rules like Krum [43] and Multi-Krum are designed to be robust against a limited number of malicious clients sending poisoned model updates.
- **Hybrid Privacy Techniques:** Combining DP with cryptography, such as using DP to add noise before applying HE, provides a layered defense [44].
- **Anomaly Detection:** Employing auxiliary models or statistical tests to detect and filter out anomalous model updates before aggregation [45].

4. Synergies with Distributed Ledger Technologies:

The centralized aggregation server in traditional FL represents a single point of failure and a trust bottleneck. Distributed Ledger Technologies (DLTs), with blockchain as the most prominent example, offer a paradigm shift towards fully decentralized and trustless systems, making them a perfect match for ad-hoc robotic swarms.

4.1. Blockchain as a Decentralized Aggregator:

A blockchain can replace the central server entirely. In this model, local model updates are broadcast to the peer-to-peer network. Smart contracts—self-executing code deployed on the blockchain—can be programmed to perform aggregation algorithms (like FedAvg) once a sufficient number of updates have been submitted and verified [46]. This eliminates the need for a trusted central authority, as the execution of the smart contract is immutable and transparent to all participants. The global model can then be stored on the blockchain or in a distributed hash table (DHT), accessible to all robots.

4.2. Incentivization and Audibility:

Blockchains naturally enable token-based incentivization mechanisms. Robots can be rewarded with cryptographic tokens for contributing high-quality model updates, encouraging participation and discouraging lazy or malicious behavior [47]. Furthermore, every transaction (model update submission, aggregation event) is timestamped and recorded on an immutable ledger, providing complete auditability and traceability. This is crucial for debugging, compliance, and forensic analysis in case of a failure or attack.

4.3. Enhanced Security and Identity Management:

The cryptographic foundations of DLTs provide robust identity management. Each robot can have a unique cryptographic identity, preventing sybil attacks. Moreover, the consensus mechanism (e.g., Proof of Work, Proof of Authority, Proof of Stake) secures the network against malicious actors trying to tamper with the aggregation process or the stored global model.

4.4. Applications in Robotic Swarms:

The combination of FL and blockchain is particularly suited for:

- **Vehicular Networks (V2X):** Building trusted collaborative perception models between autonomous vehicles without a central coordinator [48].
- **Drone Swarms:** Secure and auditable sharing of learned navigation maps or object detection models in a fleet of drones for disaster response [49].
- **Industrial IoT:** Collaborative predictive maintenance between robots from different manufacturers on a factory floor, with data privacy guaranteed by FL and business logic enforced by smart contracts [50].

While promising, this integration faces challenges, primarily the significant computational overhead and latency introduced by consensus mechanisms, which may be unsuitable for real-time control loops in robotics.

5. Applications in Robotic and Autonomous Systems:

FL is being actively researched across a spectrum of robotic applications. We categorize and highlight key works below.

5.1. Perception and Scene Understanding:

- **Visual Place Recognition:** Robots can collaboratively learn a shared visual place recognition model without sharing images from their specific operating environments, improving generalization [51].
- **Multi-Robot SLAM:** Fe-SLAM [52] is a framework where robots collaboratively build and update a global map by sharing learned feature descriptors or sub-maps via FL, rather than raw sensor data, preserving the privacy of their precise locations and environments.
- **Object Detection and Recognition:** Teams of robots can improve object detection accuracy for rare objects by learning from each other's federated experiences [53].

5.2. Control and Navigation

- **Federated Reinforcement Learning (FRL):** As discussed in 2.4, FRL allows robots to learn robust control policies. For example, a fleet of warehouse robots can learn optimal navigation policies adapted to different floor layouts and dynamic obstacles [54].
- **Imitation Learning:** FL enables learning from demonstration across multiple human teachers and robots, aggregating different driving styles or manipulation strategies into a robust global policy while keeping user-specific data private [55].

5.3. Human-Robot Interaction (HRI):

- **Activity Recognition:** Robots in smart homes can collaboratively learn models of human activity from on-board sensors without centrally pooling private video or audio data [56].
- **Natural Language Processing:** Personal assistant robots can improve their language understanding models by learning from

interactions with multiple users in a privacy-preserving manner via FL [57].

5.4. Collaborative Task Execution:

- **Distributed Formation Control:** Swarms of drones can use FL to adapt their formation control algorithms based on shared wind conditions or obstacle avoidance experiences.
- **Collaborative Manipulation:** Multiple robotic arms can learn to coordinate on a complex assembly task by sharing policy updates through an FL framework.

6. Conclusion and Future Directions:

Federated Learning represents a fundamental shift in how intelligent systems, particularly robotic swarms, can learn from distributed data. It directly addresses the critical constraints of privacy, bandwidth, and latency that hinder centralized cloud-based approaches. This survey has outlined the core concepts, reviewed the current state-of-the-art, and highlighted the powerful synergy with DLTs for building decentralized, trustable systems.

However, the field is still in its relative infancy. Several compelling research directions demand attention:

1. **Personalized FL for Robotics:** Developing algorithms that can produce a single global model that performs well on average is insufficient. Future work must focus on personalization—creating mechanisms for each robot to fine-tune the global model to its specific environment, dynamics, and tasks without compromising the collaborative benefit.
2. **Lifelong and Meta-Learning:** Robotic environments are non-stationary. Integrating FL with lifelong learning and meta-learning frameworks will be essential for systems that can continuously adapt to new tasks and environments without catastrophic forgetting.
3. **Standardized Benchmarks and Datasets:** The community lacks standardized benchmarks and realistic non-IID datasets for

evaluating FL algorithms in robotic contexts. Creating such resources will be vital for fair comparison and rapid progress.

4. **Cross-Modal and Multi-Task FL:** Robots possess multi-modal sensors. Research into FL frameworks that can effectively fuse and learn from heterogeneous data types (vision, lidar, proprioception) across different robots, and even for different but related tasks, is a promising avenue.
5. **Tackling System Dynamics:** The impact of robot mobility, network dynamics, and resource fluctuation on FL convergence and performance needs more rigorous theoretical and experimental analysis. Developing FL protocols that are inherently adaptive to these dynamics is crucial.

In conclusion, while significant challenges remain, the potential of FL to enable secure, scalable, and collective intelligence in heterogeneous robotic swarms is immense. Its continued development, especially in concert with technologies like DLT and advanced wireless communication (5G/6G), will be a key enabler for the truly autonomous and collaborative systems of the future.

References

[1] B. P. Gerkey and M. J. Matarić, "A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems," The International Journal of Robotics Research, vol. 23, no. 9, pp. 939–954, 2004.

[2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[3] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," The Journal of Machine Learning Research, vol. 17, no. 1, pp. 1334–1373, 2016.

[4] M. Simsek, A. Aijaz, M. Dohler, J. Sachs, and G. Fettweis, "5G-enabled tactile internet," IEEE Journal on Selected Areas in Communications, vol. 34, no. 3, pp. 460–473, 2016.

[5] R. Szeliski, Computer Vision: Algorithms and Applications. Springer Science & Business Media, 2010.

[6] M. Quigley et al., "ROS: an open-source Robot Operating System," in ICRA workshop on open source software, vol. 3, no. 3.2, 2009, p. 5.

[7] S. Liu, L. Li, J. Tang, S. Wu, and J.-L. Gaudiot, "Creating autonomous vehicle systems," Synthesis Lectures on Computer Science, vol. 6, no. 1, pp. 1–186, 2017.

[8] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273–1282.

[9] P. Kairouz et al., "Advances and Open Problems in Federated Learning," Foundations and Trends® in Machine Learning, vol. 14, no. 1–2, pp. 1–210, 2021.

[10] W. Y. B. Lim et al., "Federated learning in mobile edge networks: A comprehensive survey," IEEE Communications Surveys & Tutorials, vol. 22, no. 3, pp. 2031–2063, 2020.

[11] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Processing Magazine, vol. 37, no. 3, pp. 50–60, 2020.

[12] M. Fang, X. Cao, J. Jia, and N. Gong, "Local Model Poisoning Attacks to Byzantine-Robust Federated Learning," in 29th USENIX Security Symposium (USENIX Security 20), 2020, pp. 1605–1622.

[13] S. Niknam, H. S. Dhillon, and J. H. Reed, "Federated Learning for Wireless Communications: Motivation, Opportunities, and Challenges," IEEE Communications Magazine, vol. 58, no. 6, pp. 46–51, 2020.

[14] J. J. Kuffner, "Cloud-Enabled Robots," in IEEE-RAS International Conference on Humanoid Robots, 2010.

[15] R. Arumugam, V. R. Enti, L. Bingbing, W. Xiaojun, K. Baskaran, and F. F. Kong, "DAvinCi: A cloud computing framework for service robots," in 2010 IEEE International Conference on Robotics and Automation, 2010, pp. 3084–3089.

[16] L. G. Jaimes, J. E. Siegel, and A. M. Sarma, "Fog Computing: The Lightweight Companion of Cloud Computing in the Internet of Things and Robotics," in Cloud Computing: Principles and Paradigms, R. Buyya, J. Broberg, and A. M. Goscinski, Eds. John Wiley & Sons, Inc., 2011, pp. 457–482.

[17] S. K. Datta, C. Bonnet, and J. Haerri, "Fog Computing architecture to enable consumer centric Internet of Things services," in 2015 International Conference on Consumer Electronics (ICCE), 2015, pp. 564–567.

[18] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, and V. Josifovski, "Scaling Distributed Machine Learning with the Parameter Server," in 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14), 2014, pp. 583–598.

[19] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and Communication-Efficient Federated Learning from Non-IID Data," IEEE Transactions on Neural Networks and Learning Systems, vol. 31, no. 9, pp. 3400–3413, 2020.

[20] C. Szegedy et al., "Intriguing properties of neural networks," arXiv preprint arXiv:1312.6199, 2013.

[21] B. Biggio, B. Nelson, and P. Laskov, "Poisoning attacks against support vector machines," in Proceedings of the 29th International Coference on International Conference on Machine Learning, 2012, pp. 1467–1474.

[22] M. Fredrikson, S. Jha, and T. Ristenpart, "Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures," in Proceedings of the 22nd ACM SIGSAC

Conference on Computer and Communications Security, 2015, pp. 1322–1333.

[23] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," in 2017 IEEE Symposium on Security and Privacy (SP), 2017, pp. 3–18.

[24] C. Dwork, A. Roth et al., "The algorithmic foundations of differential privacy," Foundations and Trends® in Theoretical Computer Science, vol. 9, no. 3–4, pp. 211–407, 2014.

[25] C. Gentry, "A fully homomorphic encryption scheme," Stanford university, 2009.

[26] A. C.-C. Yao, "Protocols for secure computations," in 23rd annual symposium on foundations of computer science (sfcs 1982), 1982, pp. 160–164.

[27] L. Busoniu, R. Babuska, and B. De Schutter, "A comprehensive survey of multiagent reinforcement learning," IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), vol. 38, no. 2, pp. 156–172, 2008.

[28] Y. Zhu, Y. Liu, J. J. Q. Yu, and X. Yuan, "Federated Reinforcement Learning: Techniques, Applications, and Open Challenges," arXiv preprint arXiv:2108.11887, 2021.

[29] Z. Yang, M. Chen, K.-K. Wong, H. V. Poor, and S. Cui, "Federated Learning for 6G: Applications, Challenges, and Opportunities," Engineering, 2021.

[30] B. Liu, L. Wang, and M. Liu, "Lifelong Federated Reinforcement Learning: A Learning Architecture for Navigation in Cloud Robotic Systems," IEEE Robotics and Automation Letters, vol. 4, no. 4, pp. 4555–4562, 2019.

[31] Y. Zhan, P. Li, Z. Qu, D. Zeng, and S. Guo, "A Learning-Based Incentive Mechanism for Federated Learning," IEEE Internet of Things Journal, vol. 7, no. 7, pp. 6360–6368, 2020.

[32] M. R. Sprague, A. Jalalirad, M. Scavuzzo, C. Capota, M. Neun, and L. Do, "Asynchronous Federated Learning for Geospatial Applications," in ECML PKDD 2018 Workshops, 2018.

[33] A. F. Aji and K. Heafield, "Sparse Communication for Distributed Gradient Descent," in Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, 2017, pp. 440–445.

[34] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "QSGD: Communication-Efficient SGD via Gradient Quantization and Encoding," in Advances in Neural Information Processing Systems, 2017, vol. 30.

[35] J. H. Cho, J. Wang, and G. Joshi, "Client Selection in Federated Learning: Convergence Analysis and Power-of-Choice Selection Strategies," arXiv preprint arXiv:2010.01243, 2020. [36] C. Xie, S. Koyejo, and I. Gupta, "Asynchronous Federated Learning," arXiv preprint arXiv:1903.03934, 2019.

[37] G. Zhu, Y. Wang, and K. Huang, "Broadband Analog Aggregation for Low-Latency Federated Edge Learning," IEEE Transactions on Wireless Communications, vol. 19, no. 1, pp. 491–506, 2020.

[38] S. Wang et al., "Adaptive Federated Learning in Resource Constrained Edge Computing Systems," IEEE Journal on Selected Areas in Communications, vol. 37, no. 6, pp. 1205–1221, 2019.

[39] N. H. Tran, W. Bao, A. Zomaya, M. N. H. Nguyen, and C. S. Hong, "Federated Learning over Wireless Networks: Optimization Model Design and Analysis," in IEEE INFOCOM 2019 - IEEE Conference on Computer Communications, 2019, pp. 1387–1395.

[40] B. Moons, K. Goetschalckx, N. Van Berckelaer, and M. Verhelst, "Minimum Energy Quantized Neural Networks," in 2017 51st Asilomar Conference on Signals, Systems, and Computers, 2017, pp. 1921–1925.

[41] T. Nishio and R. Yonetani, "Client Selection for Federated Learning with Heterogeneous Resources in Mobile Edge," in ICC 2019

- 2019 IEEE International Conference on Communications (ICC), 2019, pp. 1–7.

[42] L. U. Khan, W. Saad, Z. Han, E. Hossain, and C. S. Hong, "Federated Learning for Internet of Things: Recent Advances, Taxonomy, and Open Challenges," IEEE Communications Surveys & Tutorials, vol. 23, no. 3, pp. 1759–1799, 2021.

[43] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent," in Advances in Neural Information Processing Systems, 2017, vol. 30.

[44] R. C. Geyer, T. Klein, and M. Nabi, "Differentially Private Federated Learning: A Client Level Perspective," arXiv preprint arXiv:1712.07557, 2017.

[45] C. Fung, C. J. M. Yoon, and I. Beschastnikh, "The Limitations of Federated Learning in Sybil Settings," in 23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020), 2020, pp. 301–316.

[46] Y. Qu, S. R. Pokhrel, S. Garg, L. Gao, and Y. Xiang, "A Blockchained Federated Learning Framework for Cognitive Computing in Industry 4.0 Networks," IEEE Transactions on Industrial Informatics, vol. 17, no. 4, pp. 2964–2973, 2021.

[47] J. Kang, Z. Xiong, D. Niyato, S. Xie, and J. Zhang, "Incentive Mechanism for Reliable Federated Learning: A Joint Optimization Approach to Combining Reputation and Contract Theory," IEEE Internet of Things Journal, vol. 6, no. 6, pp. 10700–10714, 2019.

[48] Y. Lu, X. Huang, K. Zhang, S. Maharjan, and Y. Zhang, "Blockchain and Federated Learning for 5G Beyond: Vision, Emerging Paradigms, and Challenges," IEEE Wireless Communications, vol. 28, no. 4, pp. 140–146, 2021.

[49] A. Ferdowsi and W. Saad, "Deep Learning for Signal Authentication and Security in Massive Internet of Things Systems," IEEE Transactions on Communications, vol. 67, no. 2, pp. 1371–1387, 2019.

[50] M. S. Jamil, F. Ahmad, and J. Qadir, "A Blockchain-Based Federated Learning Framework for Industrial IoT," in 2020 IEEE 17th International Conference on Smart Communities: Improving Quality of Life Using ICT, IoT and AI (HONET), 2020, pp. 147–151.

[51] T. Cieslewski, S. Choudhary, and D. Scaramuzza, "Data-Efficient Decentralized Visual SLAM," in 2018 IEEE International Conference on Robotics and Automation (ICRA), 2018, pp. 2466–2473.

[52] Y. Liu, J. Yu, and J. J. Q. Yu, "Fe-SLAM: A Lightweight and Efficient Federated Visual SLAM Framework," in 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021.

[53] L. Zhang, L. Shen, L. Ding, D. Tao, and L.-Y. Duan, "Fine-Tuning Global Model via Data-Free Knowledge Distillation for Non-IID Federated Learning," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 10174–10183.

[54] Z. Wang, M. Li, and Y. Wang, "Federated Reinforcement Learning for Autonomous Driving," in 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021.

[55] A. Mandlekar et al., "Scaling Robot Learning with Semantically Imagined Experience," arXiv preprint arXiv:2202.11565, 2022.

[56] Y. Chen, X. Qin, J. Wang, C. Yu, and W. Gao, "FedHealth: A Federated Transfer Learning Framework for Wearable Healthcare," IEEE Intelligent Systems, vol. 35, no. 4, pp. 83–93, 2020.

[57] A. Hard et al., "Federated Learning for Mobile Keyboard Prediction," arXiv preprint arXiv:1811.03604, 2018.