

A Survey of Trust-Enhanced Blockchain-Assisted Federated Learning for Secure Collaboration in Vehicular Networks

Dinesh Kumar P¹, Varshini G²

¹Assistant Professor, Department of Computer Science and Engineering, Vivekanandha College of Engineering for Women, Tamilnadu, India

²PG Scholar, Department of Computer Science and Engineering, Vivekanandha College of Engineering for Women, Tamilnadu, India

Abstract—Connected and autonomous vehicles generate large volumes of sensor, trajectory and driving-behaviour data, and federated learning (FL) allows them to build shared models without uploading raw data to a central server. Open vehicular environments, however, combine high mobility, unstable links and participants of unknown reliability, which exposes FL to poisoned updates, free-riding, a vulnerable central aggregator and contributions that nobody can verify. Blockchain supplies tamper-evident records and decentralised coordination, while trust and reputation models help decide which vehicles deserve influence over the global model. This paper surveys research that combines these three ingredients. We summarise the background, classify existing schemes by trust evaluation method, ledger design and aggregation strategy, compare representative proposals, and map common attacks to the defences that trust-enhanced architectures offer. We then discuss open problems, namely consensus latency under mobility, on-chain storage, privacy leakage from transparent ledgers, Sybil and collusion attacks on reputation, and the absence of shared benchmarks, and we outline directions toward lightweight, adaptive and standardised designs for secure vehicular collaboration.

Index Terms—Federated Learning, Blockchain, Trust Management, Reputation, Vehicular Networks, Internet of Vehicles, Secure Collaboration, Poisoning Attacks.

I. INTRODUCTION

Modern vehicles are rolling sensor platforms. Cameras, radar, lidar, positioning units and in-vehicle diagnostics together produce data that can improve perception, traffic prediction, route planning and collision avoidance. Learning from this data at fleet

scale is attractive, but gathering it in one place raises obvious difficulties: the volume is large, wireless bandwidth is limited, and location traces and driving habits are personal information that owners and regulators are reluctant to release.

Federated learning addresses part of this problem. Each vehicle trains on its own data and sends only model parameters to an aggregator, which merges them into a global model [1]. The idea has been studied extensively for mobile and edge settings [2]–[4] and for vehicular systems in particular [5]. Yet the same openness that makes vehicular FL useful also makes it fragile. A participant can submit deliberately corrupted parameters, a vehicle with poor sensors can unknowingly degrade the model, and a compromised aggregator can alter or discard updates. Because vehicles move between roadside units and appear in and out of coverage, there is rarely a long interaction history on which to base confidence in any single participant.

Two complementary tools have been proposed to close this gap. Blockchain removes reliance on a single trusted party by keeping an append-only, jointly maintained record of who contributed what [6], [7]. When combined, the ledger provides evidence that is hard to alter and the trust model provides a rule for acting on it. The resulting family of schemes is what we call trust-enhanced blockchain-assisted federated learning.

Existing surveys cover blockchain and FL for vehicular networks in general [6], [7], but they treat trust only as one of several topics. The present paper concentrates on the trust dimension. Our contributions

are as follows: (i) a compact background on FL, blockchain and trust models in vehicular settings; (ii) a classification and review of representative schemes; (iii) a comparison table and an attack-to-defence mapping; (iv) a reference architecture that captures the common structure of the surveyed work; and (v) a discussion of unresolved challenges and future directions.

The structure of this research article is outlined in the following sections. Section II gives background and the threat model. Section III reviews the literature. Section IV compares the surveyed schemes. Section V presents the reference architecture. Section VI presents the major challenges and prospective research opportunities, while Section VII summarizes the principal findings and concludes the paper.

II. BACKGROUND AND THREAT MODEL

A. FEDERATED LEARNING IN VEHICULAR NETWORKS

In the standard FedAvg procedure, a server selects a set of clients in each round, sends them the current global model, receives their locally trained models and averages them, usually in proportion to local data size [1]. Vehicular deployments differ from the cloud-device setting in several ways. Vehicles are fast-moving and may leave a roadside unit's range mid-round; their data are strongly non-identically distributed because road type, climate and driver behaviour differ; and compute and energy budgets vary widely between vehicle models. Surveys of FL in mobile edge networks and in the vehicular Internet of Things highlight communication cost, statistical heterogeneity and participant selection as persistent obstacles [4], [5].

B. BLOCKCHAIN IN VEHICULAR NETWORKS

A blockchain is a replicated ledger whose entries are chained by cryptographic hashes, so that altering an old record invalidates all later ones [23]. Public chains using proof-of-work are generally too slow and too energy-hungry for vehicles, so the literature favours consortium or permissioned chains operated by roadside units, base stations or edge servers, with Byzantine-fault-tolerant consensus such as PBFT [22] or with lighter alternatives. Smart contracts on such chains can store reputation values, execute reward rules and register hashes of model updates, which

gives later verifiers an audit trail without placing large models on the ledger itself [8], [9].

C. TRUST AND REPUTATION MODELS

Trust models estimate how likely a vehicle is to behave correctly. A popular choice is subjective logic, which represents an opinion as belief, disbelief and uncertainty and offers operators for combining direct experience with recommendations from others [21]. Other designs rely on Bayesian rating of exchanged messages [14] or on similarity between a participant's behaviour and that of its neighbours [15]. Whatever the formula, the usual inputs are the quality of past contributions, how often and how recently interactions took place, and whether they were positive or negative [8], [9].

D. THREAT MODEL

The attacks addressed by the surveyed work fall into five groups. Poisoning attacks alter the training data or the submitted update, either to degrade accuracy or to implant a backdoor [16], [17], [25]. Free-riders send little or no genuine work yet expect a reward. Sybil attackers create many identities to gain influence over voting or aggregation [24]. Inference attacks try to recover private data from gradients or model updates [19].

III. LITERATURE REVIEW

A. BLOCKCHAIN-ASSISTED FL FOR VEHICLES AND EDGE NETWORKS

Studies in this group use the ledger mainly to replace or supervise the central aggregator. The surveyed studies show a steady move from public-chain prototypes toward permissioned designs that fit roadside infrastructure.

^[13] Kim and co-authors proposed on-device federated learning that exchanges local updates through a blockchain. Miners verify the updates and are rewarded for doing so, which removes the central server. The authors analyse the end-to-end latency and derive the block generation rate that balances waiting time against the risk of forks, showing that ledger parameters directly affect learning speed.

^[11] Pokhrel and Choi examined blockchain-based FL for autonomous vehicles. Vehicles train locally and upload updates to blockchain-enabled miners, and the authors model communication and ledger delays together. Their analysis shows that the delay

introduced by block creation and by the vehicles' own wireless links must be considered together when such systems are designed.

^[10] Lu and co-authors addressed secure data sharing in the Internet of Vehicles by combining asynchronous federated learning with a permissioned blockchain, so that vehicles share learned models rather than raw data. Asynchrony lets slower vehicles contribute without stalling the round, which suits intermittent connectivity.

^[12] Li and colleagues introduced a decentralised framework in which a small committee validates submitted updates and then records the accepted ones on the chain. The committee changes between rounds, which reduces consensus cost compared with all-node voting and limits how long a malicious member can influence results. The design is general and was not tailored to vehicles, so mobility-specific committee selection remains open.

B. TRUST AND REPUTATION-DRIVEN FL

This group adds an explicit trust value for each participant. The central observation is that a ledger alone records misbehaviour but does not by itself prevent it, so a scoring rule is needed.

^[8] Kang and co-authors studied data sharing in vehicular edge computing. Although the work targets data sharing rather than learning, it supplied the reputation-plus-ledger pattern that later FL schemes adopted.

^[9] In Reliable Federated Learning for Mobile Networks, Kang and co-authors used the same pattern for FL worker selection. A task publisher computes each worker's reputation from belief, disbelief and uncertainty, adjusted for interaction frequency and recency, and stores it on a consortium blockchain maintained by edge nodes. In a simulation with ten workers, of which two were malicious and four unreliable, the proposed multi-weight model lowered the reputation of misbehaving workers more sharply than the baseline schemes.

^[14] Yang and co-authors proposed decentralised trust management for vehicular networks in which roadside units rate the messages that vehicles report using Bayesian inference and record the resulting trust offsets on a blockchain. The scheme is not a learning system, but it shows how infrastructure nodes can compute trust while the ledger keeps it consistent and auditable.

^[15] Chen and co-authors presented a reputation-driven asynchronous FL framework for trajectory prediction. Reputation is derived from the similarity of vehicles' motion data and combined with subjective-logic opinions; vehicles synchronise through a local directed acyclic graph, a reinforcement-learning agent groups them by reputation, and a permissioned blockchain with a reputation-based consensus records the transactions. The reported results on ApolloScope improved over the GRIP baseline (average displacement error 1.11 versus 1.25 at a three-second horizon), although the authors acknowledge that the method does not beat every state-of-the-art predictor.

C. SURVEYS AND CHALLENGE ANALYSES

^[6] Javed and co-authors published a broad survey on integrating blockchain with FL in vehicular networks. It organises the field through a layered view of smart transport infrastructure, background on FL and blockchain, security and privacy applications, and open issues. The challenges it lists include scalability, data quality, interpretability, near-real-time decision making, storage consistency between on-chain and off-chain data, and the combined overhead of consensus and FL communication.

^[7] Cai and co-authors examined the benefits and limits of blockchain-empowered FL more generally. They credit blockchain with decentralisation, transparent reputation and incentive mechanisms and security, while stressing that it is not a cure-all: consensus adds delay, public data can leak information, lightweight consensus can be exposed to Sybil attacks, and ledgers cannot hold large models. Proposed remedies include proof-of-reputation and committee consensus, DAG or sharded structures, homomorphic encryption and differential privacy, and off-chain storage with on-chain hashes.

^[5] Du and co-authors reviewed FL for the vehicular Internet of Things, covering recent advances and open issues in communication efficiency, learning quality and security, which frames the vehicular constraints that the trust-based schemes above attempt to meet.

D. ROBUST AGGREGATION AND PRIVACY TECHNIQUES

Trust scoring is usually complemented by statistical defences. Distance-based rules such as Krum keep only the submissions that agree most closely with the rest and can therefore withstand a limited share of

arbitrarily faulty participants [16], but later work showed that carefully crafted local model poisoning can still defeat several such rules [17], and backdoor updates can be hidden inside seemingly benign contributions [25].

These results motivate combining robust aggregation with long-term reputation rather than relying on either alone. For confidentiality, secure aggregation hides individual updates from the server [18], and differential privacy bounds what any single record can reveal [20]; both address the risk that gradients leak training data [19].

IV. COMPARATIVE ANALYSIS

Table I compares the representative schemes on their trust mechanism, ledger design, learning mode, main contribution and main limitation. The comparison is qualitative because the studies use different datasets, tasks and metrics, and no common benchmark yet exists for this field. Table II maps the attacks described in Section II to the defences that trust-enhanced architectures provide.

TABLE I. COMPARISON OF REPRESENTATIVE SCHEMES

Scheme	Trust Reputation Mechanism /	Ledger Consensus and	Learning Mode	Main Contribution	Main Limitation
Kang et al. [8]	Multi-weight subjective logic	Consortium blockchain run by roadside units	Data sharing (no FL)	Reputation-based secure data sharing in vehicular edge computing	Not a learning scheme; relies on roadside units
Kang et al. [9]	Subjective logic with interaction frequency and recency	Consortium blockchain at edge nodes	Synchronous FL	Reputation-driven selection of reliable FL workers	Small simulation (10 workers)
Lu et al. [10]	Permissioned node admission	Permissioned blockchain	Asynchronous FL	Model-based data sharing in the Internet of Vehicles	No explicit reputation score
Pokhrel and Choi [11]	Not central to the design	Blockchain-enabled miners	FL for autonomous vehicles	Joint analysis of learning and ledger delay	Analysis-focused; limited trust handling
Kim et al. [13]	Miner verification with rewards	Blockchain with proof-of-work	On-device FL	Server-free FL with latency analysis and block-rate tuning	Fork and latency cost of proof-of-work
Li et al. [12]	Committee validation of updates	Chain with committee consensus	Decentralised FL	Lower consensus cost and limited malicious influence	Not vehicle-specific; committee selection open
Yang et al. [14]	Bayesian rating of messages	Blockchain storing trust offsets	Trust management (no FL)	Decentralised trust for vehicular messages	Not a learning scheme
Chen et al. [15]	Subjective logic with motion similarity	Permissioned chain, reputation-based consensus	Asynchronous FL with local DAG	Trajectory prediction with reputation-grouped aggregation	Not state-of-the-art on all metrics (as reported by the authors)

TABLE II. ATTACKS AND TRUST-ENHANCED DEFENCES

Attack	Effect on Vehicular FL	Trust-Based Countermeasure	Role of Blockchain
Data or model poisoning	Reduced accuracy or hidden backdoor in the global model	Low-reputation clients excluded or down-weighted; anomaly checks on updates	Immutable record of contributions for later audit
Unreliable contributors	Noisy updates from poor sensors or unstable links	Reputation reflects quality, frequency and recency of past work	Shared, consistent reputation across edge nodes
Free-riding	Rewards paid without real training effort	Contribution validation before reward	Smart-contract reward settlement
Sybil identities	Extra identities inflate influence over voting or aggregation	Reputation tied to registered, accountable identities	Permissioned membership and identity registry
Compromised aggregator or validator	Updates altered, delayed or dropped	Trust evaluation spread over several edge nodes	Decentralised consensus and tamper-evident logs

Three patterns emerge from the comparison. First, consortium or permissioned ledgers dominate, since vehicles themselves cannot afford open mining and roadside infrastructure can host validators. Second, subjective logic is the most common trust formalism, probably because it handles uncertainty about newly encountered vehicles. Third, the strongest limitation shared across studies is evaluation: most use small simulations, so claims about scalability under realistic mobility remain largely unverified.

V. REFERENCE ARCHITECTURE

Across the reviewed studies, the architectures share a three-layer structure, shown in Fig. 1. At the vehicle layer, each vehicle trains on local data and reports updates together with evidence about its own behaviour. At the edge layer, roadside units or edge servers evaluate trust by combining direct observations with opinions from other nodes, validate updates, perform trust-weighted aggregation, select clients for the next round and settle incentives. At the blockchain layer, a consortium ledger keeps the reputation records, hashes of submitted models and an audit trail, and a consensus protocol keeps the edge nodes' view consistent.

Typical round proceeds as follows. The edge node selects vehicles with sufficient reputation and sends them the global model. Vehicles train and return updates.

Trust-enhanced blockchain-assisted FL: reference architecture




Fig. 1. Reference architecture of trust-enhanced blockchain-assisted federated learning.

The edge node checks each update against the others, discounts those that look anomalous and aggregates the rest with weights that depend on reputation. It then writes revised reputation values and the hash of the new model to the ledger through the consensus protocol, and finally releases rewards. Keeping the model itself off-chain and writing only its hash, as suggested in [7], keeps ledger size manageable.

VI. CHALLENGES AND FUTURE DIRECTIONS

Consensus latency under mobility. A vehicle may remain inside one unit's range for only a few seconds, whereas block confirmation can take longer [11], [13].

Future work should co-design the learning round and the consensus protocol, for example through asynchronous aggregation [10], [15] or committee-based validation [12], and should test them with realistic mobility traces.

Scalable storage. Full replication of the ledger on every node does not scale to large fleets. Off-chain storage with on-chain hashes, pruning, sharding and DAG-based structures are promising, but their effect on the trust of stored records needs systematic study [7].

Privacy of transparent records. Reputation values, update hashes and interaction histories could be used to link records to individual vehicles. Anonymous credentials, pseudonym rotation, differential privacy and zero-knowledge proofs of honest behaviour are candidate tools, but they complicate the verification that trust models rely on [6], [20].

Resilience of the trust mechanism itself. Reputation systems can be attacked directly through collusion, bad-mouthing, on-off behaviour and Sybil identities [24]. Few surveyed schemes evaluate these attacks, and adaptive adversaries who build reputation before poisoning are rarely modelled. Combining reputation with robust aggregation [16], [17] and with hardware-backed identities is a priority.

Cold start and fairness. Newly arriving vehicles have no history, yet excluding them wastes useful data and creates an unfair barrier to participation. Uncertainty-aware models such as subjective logic [21] offer a natural treatment, and incentive designs should reward both honesty and informative data from under-represented regions.

Heterogeneity and personalisation. Non-identical data across regions means that an update that looks anomalous may simply come from a different driving environment. Trust models that distinguish adversarial from legitimately different behaviour, together with personalised or clustered FL, are needed to avoid penalising honest vehicles.

Benchmarks and standards. Studies use different datasets and attack settings, so results cannot be compared directly. A shared benchmark that includes mobility, communication delay, attack scenarios and ledger cost, together with alignment to vehicular communication standards, would speed progress.

VII. CONCLUSION

Federated learning keeps raw driving data on the vehicle, but it does not tell the system whom to believe. This survey examined how blockchain and trust management can fill that gap. Blockchain supplies a shared, tamper-evident record and removes dependence on a single aggregator, while trust and reputation models turn recorded behaviour into decisions about selection and weighting. The reviewed work indicates that permissioned ledgers run by edge infrastructure, subjective-logic or similarity-based reputation, and asynchronous or committee-based processing form the prevailing design pattern. At the same time, evidence for the claimed benefits remains thin: evaluations are mostly small simulations, the trust mechanism is rarely attacked directly, and privacy, storage and latency costs are insufficiently quantified. Progress will depend on lightweight consensus matched to vehicle mobility, privacy-preserving yet verifiable reputation, attack-aware trust models and common benchmarks. If those gaps are closed, trust-enhanced blockchain-assisted federated learning can become a practical basis for secure collaboration among connected vehicles.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2017, pp. 1273–1282.
- [2] P. Kairouz *et al.*, "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, nos. 1–2, pp. 1–210, 2021.
- [3] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1–19, 2019.
- [4] 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.
- [5] Z. Du, C. Wu, T. Yoshinaga, K.-L. A. Yau, Y. Ji, and J. Li, "Federated learning for vehicular Internet of Things: Recent advances and open

- issues,” *IEEE Open Journal of the Computer Society*, vol. 1, pp. 45–61, 2020.
- [6] A. R. Javed, M. A. Hassan, F. Shahzad, W. Ahmed, S. Singh, T. Baker, and T. R. Gadekallu, “Integration of blockchain technology and federated learning in vehicular (IoT) networks: A comprehensive survey,” *Sensors*, vol. 22, no. 12, Art. no. 4394, 2022, doi: 10.3390/s22124394.
- [7] Z. Cai, J. Chen, Y. Fan, Z. Zheng, and K. Li, “Blockchain-empowered federated learning: Benefits, challenges, and solutions,” *arXiv preprint arXiv:2403.00873*, 2024.
- [8] J. Kang, R. Yu, X. Huang, M. Wu, S. Maharjan, S. Xie, and Y. Zhang, “Blockchain for secure and efficient data sharing in vehicular edge computing and networks,” *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 4660–4670, Jun. 2019.
- [9] J. Kang, Z. Xiong, D. Niyato, Y. Zou, Y. Zhang, and M. Guizani, “Reliable federated learning for mobile networks,” *IEEE Wireless Communications*, vol. 27, no. 2, pp. 72–80, Apr. 2020, arXiv:1910.06837.
- [10] Y. Lu, X. Huang, K. Zhang, S. Maharjan, and Y. Zhang, “Blockchain empowered asynchronous federated learning for secure data sharing in Internet of Vehicles,” *IEEE Transactions on Vehicular Technology*, vol. 69, no. 4, pp. 4298–4311, Apr. 2020.
- [11] S. R. Pokhrel and J. Choi, “Federated learning with blockchain for autonomous vehicles: Analysis and design challenges,” *IEEE Transactions on Communications*, vol. 68, no. 8, pp. 4734–4746, Aug. 2020.
- [12] Y. Li, C. Chen, N. Liu, H. Huang, Z. Zheng, and Q. Yan, “A blockchain-based decentralized federated learning framework with committee consensus,” *IEEE Network*, vol. 35, no. 1, pp. 234–241, Jan./Feb. 2021.
- [13] H. Kim, J. Park, M. Bennis, and S.-L. Kim, “Blockchained on-device federated learning,” *IEEE Communications Letters*, vol. 24, no. 6, pp. 1279–1283, Jun. 2020.
- [14] Z. Yang, K. Yang, L. Lei, K. Zheng, and V. C. M. Leung, “Blockchain-based decentralized trust management in vehicular networks,” *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 1495–1505, Apr. 2019.
- [15] W. Chen, L. Jia, Y. Zhou, and Q. Ren, “Reputation-driven asynchronous federated learning for enhanced trajectory prediction with blockchain,” *arXiv preprint arXiv:2407.19428*, 2024.
- [16] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, “Machine learning with adversaries: Byzantine tolerant gradient descent,” in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2017, pp. 119–129.
- [17] M. Fang, X. Cao, J. Jia, and N. Gong, “Local model poisoning attacks to Byzantine-robust federated learning,” in *Proc. 29th USENIX Security Symp.*, 2020, pp. 1605–1622.
- [18] K. Bonawitz et al., “Practical secure aggregation for privacy-preserving machine learning,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, 2017, pp. 1175–1191.
- [19] L. Zhu, Z. Liu, and S. Han, “Deep leakage from gradients,” in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2019, pp. 14774–14784.
- [20] C. Dwork and A. Roth, “The algorithmic foundations of differential privacy,” *Foundations and Trends in Theoretical Computer Science*, vol. 9, nos. 3–4, pp. 211–407, 2014.
- [21] A. Jøsang, *Subjective Logic: A Formalism for Reasoning Under Uncertainty*. Cham, Switzerland: Springer, 2016.
- [22] M. Castro and B. Liskov, “Practical Byzantine fault tolerance,” in *Proc. 3rd Symp. Operating Syst. Design Implement. (OSDI)*, 1999, pp. 173–186.
- [23] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, “Blockchain challenges and opportunities: A survey,” *International Journal of Web and Grid Services*, vol. 14, no. 4, pp. 352–375, 2018.
- [24] J. R. Douceur, “The Sybil attack,” in *Proc. 1st Int. Workshop Peer-to-Peer Syst. (IPTPS)*, 2002, pp. 251–260.
- [25] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, “How to backdoor federated learning,” in *Proc. 23rd Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2020, pp. 2938–2948.