

# AI-Enhanced Trust Scoring for Blockchain-Based Vehicular Edge Resource Allocation

Nguyễn Văn Hùng<sup>1,\*</sup>; Trần Thị Mai<sup>2</sup>; Lê Quang Minh<sup>3</sup>

<sup>1</sup> Faculty of Information Technology, Hanoi University of Industry, Hanoi 100000, Vietnam

<sup>2</sup> Faculty of Computer Science and Engineering, Industrial University of Ho Chi Minh City, Ho Chi Minh City 700000, Vietnam

<sup>3</sup> Faculty of Information Technology, University of Information and Communications Technology (Thai Nguyen University), Thai Nguyen 250000, Vietnam

\* Correspondence: hung.nguyenvan@hau.edu.vn

<b>ARTICLE INFO</b> Received April 18, 2023 Revised June 21, 2023 Accepted August 12, 2023 Available Online September 30, 2023 DOI 10.63646/jaiaa.2023.010302 License Creative Commons Attribution 4.0 International Licence (CC BY 4.0) Publisher INATGI, United States of America Journal JAIAA - ISSN 3067-7386	<b>Abstract</b> Vehicular edge computing has emerged as a foundational technology for supporting low-latency analytics in connected-vehicle environments, but the open membership of these networks creates persistent risks from misbehaving participants whose actions degrade resource allocation outcomes for benign users. This article proposes an integrated framework that combines AI-driven trust scoring with a blockchain-anchored interaction history and a smart-contract allocation policy. The trust scorer is a small feed-forward neural network whose features summaries each vehicle's recent ledger history; the allocation policy combines the trust score with a confidence-aware override to handle short-history vehicles fairly. We evaluate the framework using a discrete-event simulation of a six-RSU corridor under a heterogeneous mix of reputable, new, and suspect vehicles. The proposed framework reduces successful task completion degradation from 30 percentage points (priority-only baseline) to under 7 percentage points across an adversary fraction range of 5% to 40%, achieves a discrimination AUC of 0.989 on the trust classification task, and preserves Jain's fairness index above 0.79 even under 90% offered load. The total per-allocation latency overhead is bounded at approximately 58 MS relative to the priority baseline, which corresponds to less than one meter of additional vehicular travel at typical urban speeds. We discuss the deployment considerations that determine production viability, including the choice of permissioned ledger, the model retraining cadence, the privacy implications, and the relationship to existing lower-layer security primitives. The framework is intended as a workable design pattern rather than a final proposal, and we identify federated training, reinforcement-learning extensions to the policy, and additional resource types as natural directions for future work.  <b>Keywords:</b> Vehicular edge computing;Blockchain;Trust scoring;Smart contracts; Trust scoring;Smart contracts;AI security analytics;Resource allocation;Federated learning
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## I. INTRODUCTION

Connected vehicles are emerging as one of the most demanding application domains for edge computing because they generate large volumes of context-rich data, request low-latency analytics, and operate in environments where misbehavior by a small minority of participants can degrade outcomes for all road users (Hartenstein and Laberteaux, 2008). The combination of vehicle-to-vehicle, vehicle-to-infrastructure, and vehicle-to-network communication has produced a layered ecosystem in which compute resources at the network edge play a critical role in supporting safety-critical and convenience services (Karagiannis et al., 2011). Edge computing brings computation close to the data

sources, but the assumption that participants are well-behaved is unrealistic in open vehicular settings, where vehicles may join and leave continuously, and where the cost of admitting an adversarial participant is not symmetric with the cost of excluding a benign one (Shi et al., 2016). The growing scale of the Internet of Things at large (Atzori et al., 2010) amplifies these concerns by ensuring that vehicular nodes interact not only with other vehicles and infrastructure, but also with a heterogeneous landscape of consumer and industrial devices.

In this article we propose an integrated framework for trust-aware allocation of vehicular edge resources. The framework couples three components: a learned trust scorer that estimates the trustworthiness of each vehicle requesting service, a blockchain ledger that records the historical interactions and outcomes that supply the scorer's input features, and a smart-contract allocation policy that uses the trust scores to decide which requests are granted and which are deferred or denied. The motivation for combining these components rather than relying on any single one is empirical and theoretical. Empirically, simple priority-based allocation collapses under modest fractions of misbehaving participants (Raya et al., 2008). Theoretically, classical reputation systems are vulnerable to whitewashing, ballot-stuffing, and badmouthing attacks unless the underlying record is tamper-resistant and the scoring algorithm is robust to feature poisoning (Jøsang et al., 2007). The architectural ambition of bringing intelligence and trust into a unified infrastructure echoes the broader Industry 4.0 agenda (Lu, 2017b).

The trust scorer in our framework is a small feed-forward neural network that takes as input a fixed-length feature vector summarizing the recent interaction history of each vehicle, drawn from the blockchain ledger. The features include simple counts (requests granted, requests denied, completion success rate), distributional summaries (mean and variance of completion latency), and behavioral signals (frequency of last-minute task abandonment, frequency of safety-relevant message anomalies). The model is trained offline on labelled data and deployed at the edge for low-latency inference. Recent advances in deep learning provide the methodological foundation for this design (LeCun et al., 2015), and the rapidly expanding literature on edge intelligence demonstrates the feasibility of running such models on infrastructure-side hardware (Wang et al., 2019). Earlier surveys of artificial intelligence applied to industrial workflows (Lu, 2019a) document the rapid maturation of these techniques over the past decade.

The blockchain layer addresses the data-integrity concerns that have historically limited the deployment of reputation systems in open networks. By logging interaction outcomes to a permissioned ledger maintained by the road-side units, the framework prevents retroactive tampering with the historical record while keeping verification costs modest (Christidis and Domesticities, 2016). The seminal account of Bitcoin (Nakamoto, 2008) established the basic feasibility of decentralized consensus, while subsequent platforms with general-purpose smart-contract execution (Wood, 2014) enabled the kind of programmable allocation logic that our design relies on. The integration of distributed ledgers with vehicular and IoT edge environments has been actively studied in recent years (Yang et al., 2019), and our contribution should be read as a focused application of these architectural ideas to the specific problem of trust-aware resource allocation rather than as a new ledger design. State-of-the-art reviews of blockchain research (Lu, 2019b) continue to refine the picture of what these systems can and cannot deliver.

The contributions of this article are fourfold. First, we describe the system architecture and the data flow that connect vehicles, road-side units, the blockchain ledger, and the trust scorer in a single closed loop. Second, we describe the feature representation and the learning objective for the trust scorer, with attention to the failure modes that arise when malicious participants attempt to manipulate their own features. Third, we report on a simulation-based evaluation that compares the proposed framework against a priority-only baseline and a reputation-only baseline across a range of adversary fractions, traffic loads, and network sizes; the evaluation includes throughput, allocation latency, classification quality, and Jain's fairness index. Fourth, we discuss the deployment considerations that determine whether the approach is realistic for production use, including ledger commit latency, smart-contract gas overhead, model update cadence, and privacy implications. The remainder of the article is organized as follows. Section II reviews related work on vehicular edge computing, trust management, blockchain integration, and AI-driven security analytics (Lu and Xu, 2019). Section III describes the system architecture. Section IV presents the trust scorer and the allocation policy. Section V reports the simulation methodology and findings. Section VI discusses the deployment implications,

limitations, and future directions, and Section VII concludes.

The contributions of this article are fourfold. First, we describe the system architecture and the data flow that connect vehicles, road-side units, the blockchain ledger, and the trust scorer in a single closed loop. Second, we describe the feature representation and the learning objective for the trust scorer, with attention to the failure modes that arise when malicious participants attempt to manipulate their own features. Third, we report on a simulation-based evaluation that compares the proposed framework against a priority-only baseline and a reputation-only baseline across a range of adversary fractions, traffic loads, and network sizes; the evaluation includes throughput, allocation latency, classification quality, and Jain's fairness index. Fourth, we discuss the deployment considerations that determine whether the approach is realistic for production use, including ledger commit latency, smart-contract gas overhead, model update cadence, and privacy implications. The remainder of the article is organized as follows. Section II reviews related work on vehicular edge computing, trust management, blockchain integration, and AI-driven security analytics (Lu and Xu, 2019). Section III describes the system architecture. Section IV presents the trust scorer and the allocation policy. Section V reports the simulation methodology and findings. Section VI discusses the deployment implications, limitations, and future directions, and Section VII concludes.

## II. RELATED WORK

### A. Vehicular Edge Computing

Mobile edge computing has become a foundational technology for fifth-generation networks because it shortens the data-to-decision path for latency-sensitive workloads (Mao et al., 2017). Vehicular environments push the requirements of edge computing to extremes: vehicles move quickly through the coverage of multiple road-side units, demand for computers and cache resources fluctuate by order of magnitude on time scales of minutes, and the consequences of mis-served requests can include safety-critical outcomes rather than merely lower quality of experience (Liu et al., 2020). Surveys of multi-access edge computing have catalogued the architectural choices and the resource-management challenges that arise in such settings (Taleb et al., 2017). The 6G research agenda (Lu and Zheng, 2020) continues to refine the requirements landscape that vehicular edge designs must satisfy.

A second line of work has focused on offloading and caching strategies that exploit the predictability of vehicular trajectories. Optimal delay-constrained offloading formulations (Zhang and Yu, 2018) establish baseline performance results, while named-data networking and information-centric architectures have been proposed to reduce the latency penalty of repeated retrievals (Khelifi et al., 2020). Comprehensive surveys of edge computing for smart cities provide a broader view of the deployment context (Khan et al., 2020) and emphasize that resource allocation is rarely the bottleneck on its own; the integration of allocation with security and trust is what determines real-world performance. Earlier work on context-aware big-data analytics for urban traffic management (Liu et al., 2019) documents the underlying data flows that make these allocation strategies tractable.

Recent work on the convergence of edge computing and deep learning (Wang et al., 2019) has documented the substantial performance gains that learning-based methods can provide for tasks such as traffic flow prediction, content prefetching, and anomaly detection. Surveys of mobile edge computing more broadly (Abbas et al., 2018) have catalogued the deployment patterns that have emerged across operator and enterprise contexts. The state of the art in machine learning for wireless networks more generally (Sun et al., 2019) suggests that learning-based components are now mature enough to be deployed inside resource-allocation pipelines rather than only as offline analytics tools. Connected vehicle use case studies for V2X communication (Boban et al., 2018) inform the choice of which workloads benefit most from such integration. The vision of 6G as the successor to 5G (Lu and Ning, 2020) points to even tighter integration of intelligence and connectivity in the next infrastructure generation.

### B. Trust and Reputation in Distributed Networks

Trust as a computational concept has a long lineage stretching back to early work on multi-agent systems (Marsh,

1994). Online reputation systems have been studied since the late 1990s with attention to incentive design, manipulation resistance, and the relationship between reputation aggregation and equilibrium behavior (Resnick et al., 2000). Surveys of trust and reputation in distributed service settings have catalogued the design space and identified the main attack patterns that practical systems must resist (Jøsang et al., 2007). The wider Internet-of-Things landscape, with its enormous device diversity (Gubbi et al., 2013), has driven much of the recent demand for adaptive trust mechanisms.

Within vehicular ad-hoc networks specifically, trust models have evolved from data-centric proposals that evaluate individual messages (Raya et al., 2008) toward infrastructure-based proposals that aggregate behavioral evidence over longer time scales (Mármol and Pérez, 2012). Multi-agent trust management surveys have catalogued design choices and the empirical evidence supporting them (Yu et al., 2013). Social-network-inspired trust schemes for VANETs add a graph-structural dimension by exploiting the propagation of evidence through inter-vehicle interactions (Huang et al., 2014). Earlier reviews of IoT architecture, security, and applications (Lin et al., 2017) supply much of the design vocabulary that vehicular trust schemes have inherited.

A persistent concern in trust system design is vulnerability to coordinated attacks. Vulnerability analyses have examined how Sybil and collusion attacks can poison reputation aggregates (Sun et al., 2008) and have motivated attack-resistant trust management schemes for vehicular networks (Li and Song, 2016). Fuzzy-logic trust models with fog-computing assistance have been proposed as one route toward robustness in dense deployments (Soleymani et al., 2017). Man-in-the-middle attack analyses for vehicular ad-hoc networks have evaluated the practical impact of attacker strategies on protocol behavior (Ahmad et al., 2019) and underscore that the security of the underlying communication layer cannot be assumed. Reviews of cloud-computing security issues from an earlier wave of work (Zissis and Lekkas, 2012) document several attack patterns that resurface in the edge setting.

### ***C. Blockchain Integration with Edge and Vehicular Networks***

The integration of blockchain technology with edge and vehicular environments has been an active research area for several years (Yang et al., 2019). The two technologies are natural complements: blockchain provides tamper-resistant logs and decentralized trust anchoring, while edge computing provides the low-latency execution environment that pure blockchain platforms have historically lacked. The conceptual basis for using smart contracts in IoT environments was articulated relatively early (Christidis and Devetsikiotis, 2016) and has since been refined through systematic surveys of the integration challenges (Reyna et al., 2018). The combination of blockchain with deep reinforcement learning for next-generation networks (Dai et al., 2019) has been proposed as one route toward fully autonomous infrastructure orchestration, and embedding blockchain into the wider IoT security stack (Xu, Lu, and Li, 2021) has emerged as a recurring research theme.

Within the vehicular setting, blockchain-based consensus schemes have been proposed for information authentication on the Internet of Vehicles (Hu et al., 2018). Reputation-aware consensus management combining blockchain with contract theory has been examined as a route toward secure participation in Internet-of-Vehicles deployments (Kang et al., 2019). Reviews of the broader landscape of blockchain in industrial integration also identify vehicle and edge use cases as among the most promising application domains (Chen et al., 2024). Implementation-oriented surveys (Lu, 2022) catalogue the integration patterns that have proven viable in production information systems.

From the security perspective, the underlying argument for blockchain integration is that distributed ledgers provide a tamper-evident substrate on which trust evidence can accumulate without depending on any single trusted authority (Yli-Huomo et al., 2016). The systematic literature on the maturity of blockchain platforms continues to evolve (Zheng et al., 2018), and recent management-analytics work has quantified the diffusion of these technologies across enterprise contexts (Lu, 2018). The practical limitations of integrating blockchain with IoT — throughput, energy, and developer ergonomics — have been examined in detail (Ferrag et al., 2018), and the broader research agenda for blockchain-AI integration has been mapped out as well (Salah et al., 2019). Research-trend analyses focused specifically on blockchain technology (Zheng and Lu, 2022) highlight the priority areas where additional engineering effort is most

likely to pay off, and Web 3.0 perspectives (Zhang and Lu, 2025) situate the technology within the broader internet-evolution discussion.

#### ***D. AI-Driven Security Analytics***

Modern security analytics increasingly rely on learning-based models for anomaly detection, intrusion classification, and risk scoring. The deep-learning techniques that have transformed perception and language tasks (Goodfellow et al., 2016) have also been adapted to security workloads, although the adversarial setting introduces concerns that do not arise in benign-data tasks. Recent reviews of the state of the art in artificial intelligence emphasize both the breadth of models available and the importance of careful problem formulation (Zhang and Lu, 2021). Foundational overviews of deep learning in neural networks (Schmidhuber, 2015) continue to provide useful context for practitioners adopting these methods in operational pipelines.

Federated learning is particularly relevant to edge security analytics because it enables model training across distributed data without centralizing raw observations (McMahan et al., 2017). Comprehensive surveys of federated learning have mapped both the algorithmic landscape and the open research problems (Kairouz et al., 2021). The conceptual framework of federated machine learning, articulated by (Yang et al., 2019), has become widely adopted as a reference description of the design space. For the trust-scoring application we describe, federated training is an important option because it reduces the need to centralize behavioral data that some participants may regard as sensitive. Recent applications of large language models in blockchain-anchored finance pipelines (Yang et al., 2025) illustrate how AI components are now routinely integrated with verifiable data substrates.

Reinforcement learning is another technique that has become increasingly relevant for resource allocation (Mnih et al., 2015). The adoption of deep reinforcement learning in network operations has accelerated as benchmark results in adjacent domains demonstrate the maturity of the techniques (Vinyals et al., 2019). Foundational textbooks on reinforcement learning (Sutton and Barto, 2018) continue to provide methodological reference, while the most prominent recent results have shown that deep models can master complex sequential decision problems entirely from interaction (Silver et al., 2017). For the present article we adopt a simpler supervised approach to the trust scorer because the labelled-data assumption is realistic in our setting; reinforcement-learning extensions are discussed in Section VI. Live-data analytics surveys spanning collaborative edge and cloud processing (Sharma and Wang, 2017) suggest that the operational pipelines for such extensions are now well understood.

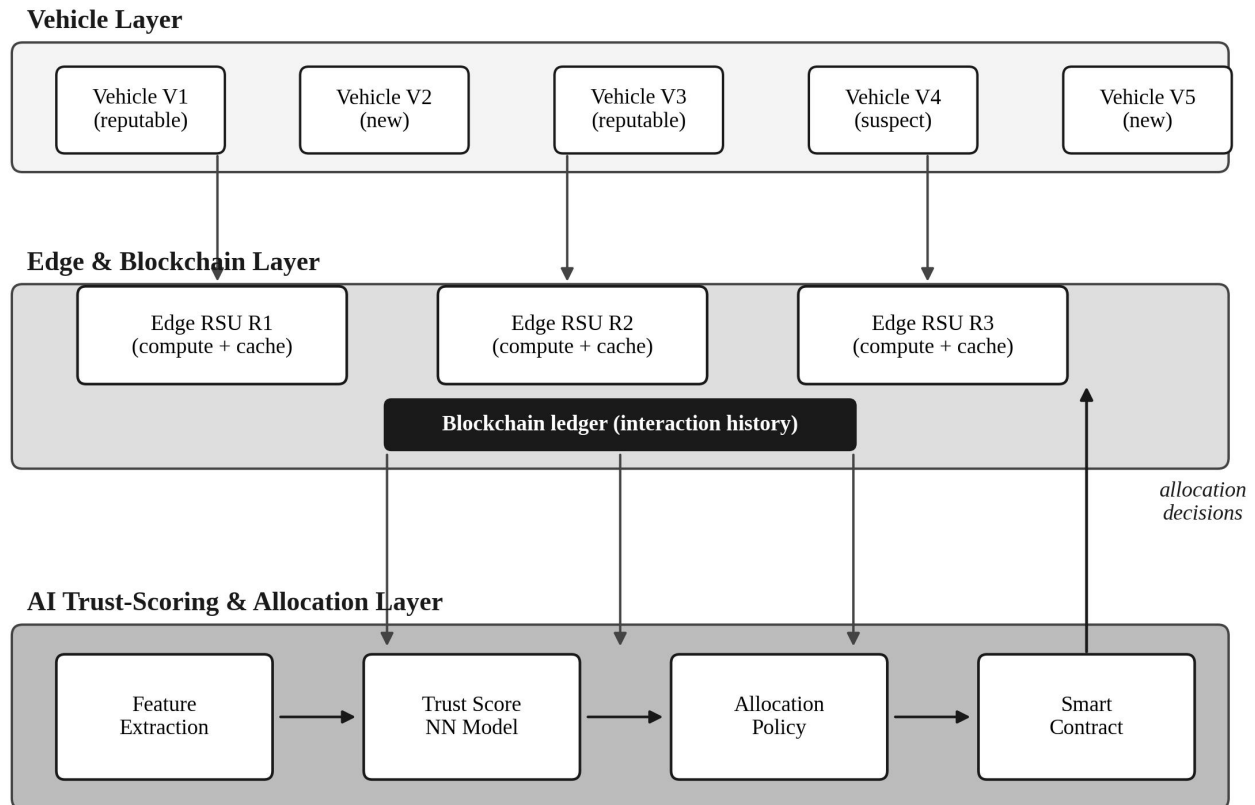
#### ***E. Roots in Wireless Security and Cognitive Radio***

The probabilistic framing of trust that we adopt has its roots in classical wireless security analysis. The information-theoretic foundations of secrecy in wireless channels (Wyner, 1975) established the practice of treating confidentiality as a probabilistic property whose distribution depends on channel statistics rather than as a deterministic guarantee. Cognitive radio (Mitola and Maguire, 1999) introduced the idea that resource allocation decisions in wireless settings must be coordinated across heterogeneous participants with potentially conflicting interests, a structural similarity to the vehicular allocation problem we address. Wireless information-theoretic security (Bloch et al., 2008) and contemporary surveys of wireless security challenges (Zou et al., 2016) together provide the background against which the AI-enhanced approach we propose should be evaluated.

### **III. SYSTEM ARCHITECTURE**

Figure 1 presents the three-layer architecture that organizes the system. The vehicle layer comprises the connected vehicles that periodically broadcast beacons describing their current state and that issue requests for edge resources whenever they require offloaded computation, cached content, or differentiated network treatment. The edge and blockchain layer comprise the road-side units (RSUs), each of which provides a local pool of computing and cache resources and a node in the permission blockchain that records the history of interactions between vehicles and RSUs. The AI trust-scoring and allocation layer comprises the feature extraction pipeline, the learned trust scorer, the

allocation policy, and the smart contracts that execute the decisions on the ledger.



**Figure 1. System architecture: vehicles request edge resources from RSUs, interaction outcomes are logged to a permitted blockchain, and an AI trust scorer drives the allocation policy implemented as a smart contract.**

The data flow through the architecture proceeds as follows. A vehicle issued a service request to its serving RSU. The RSU does not act on the request immediately. Instead, it queries the AI trust-scoring layer with the requesting vehicle's identifier; the layer extracts a feature vector from the recent ledger entries pertaining to that vehicle, runs the trust scorer to obtain a trust score in the unit interval, and feeds the score together with the request parameters into the allocation policy. The allocation policy then issues a decision — grant, defer, or deny — which is recorded on the ledger as a new transaction. After the request has been served (or the deferral or denial has been finalized), the outcome is also recorded on the ledger, providing the data that future trust-scoring decisions will draw upon (Kang et al., 2019). Cyber-physical-systems thinking (Lu, 2017a) supplies the broader integration philosophy in which such tightly coupled sense-decide-actuate loops are designed.

The closed-loop nature of this design is essential to its operation. By construction, every allocation decision generates new evidence about the participating vehicle, and every piece of new evidence updates the inputs to subsequent allocation decisions. The blockchain provides the integrity guarantee on which this loop depends; without it, malicious participants could attempt to whitewash their identity by spawning new pseudonyms or by tampering with the historical record. The smart contracts implementing the allocation policy provide a transparent and auditable description of how decisions are made, addressing one of the persistent concerns about machine-learning-driven decision systems in safety-critical contexts (Khan and Salah, 2018). Internal-auditing perspectives on blockchain-anchored data flows (Wu et al., 2025) reinforce the value of this audibility for organizations that must demonstrate process integrity to external regulators.

The architecture decouples the trust scorer from the underlying ledger and from the allocation policy. This

separation is intentional. The trust scorer can be updated independently as new training data becomes available or as new attack patterns are identified, without requiring changes to the ledger schema or the smart contracts. The allocation policy can be tuned independently to reflect evolving operator preferences over throughput, fairness, latency, and risk tolerance. The ledger schema is the slowest-changing component because it must remain compatible with historical records (Yang et al., 2019). This separation of concerns matches the architectural patterns recommended in mature surveys of edge-computing system design (Mach and Becvar, 2017). Industry-4.0 surveys that consolidate the underlying integration patterns (Lu, 2025) have argued for exactly this kind of layered decoupling as a precondition for long-lived industrial deployments.

## IV. TRUST SCORER AND ALLOCATION POLICY

### A. Feature Representation

The features supplied to the trust scorer are derived entirely from the blockchain ledger and have three coarse groups. The first group consists of count features: the number of requests issued, granted, deferred, denied, and completed for the vehicle within a sliding window. The second group consists of distributional features: the empirical mean and variance of completion latency, the empirical proportion of failed completions, and the empirical proportion of incomplete handovers. The third group consists of behavioral signals: the frequency of last-minute task abandonment, the frequency of safety-relevant message anomalies as flagged by the lower-layer signature checks, and the frequency of mismatches between declared intent and observed behavior. Together these features form a vector of dimension twenty.

The choice of features is informed by the broader literature on vehicular trust modelling. Data-centric trust evaluations (Raya et al., 2008) highlight the importance of message-level signals, while infrastructure-based reputation schemes (Mármol and Pérez, 2012) argue for aggregating evidence over longer time scales. Our feature design combines both perspectives by including both short-window message-anomaly counts and longer-window completion statistics. Convergence of edge computing and deep learning surveys (Wang et al., 2019) suggest that this combination of signals is well-matched to the inductive biases of small feed-forward networks, which is the model class we use for the scorer. Studies on machine learning over health-domain big data (Chen et al., 2020) further illustrate how heterogeneous behavioral and outcome features can be combined effectively in operational classifiers.

Two design decisions deserve mention. First, the feature vector is intentionally kept small (twenty features) and interpretable. Although larger and richer features are possible — including the entire interaction history rather than aggregated summaries — the smaller representation makes the model easier to deploy at the edge and easier to audit when allocation decisions are challenged. Second, the features are computed from blockchain entries rather than from raw RSU logs. This choice ensures that the inputs to the scorer cannot be retroactively manipulated and that all RSUs see the same view of any vehicle's history.

### B. The Trust-Scoring Model

The trust scorer is a feed-forward neural network with two hidden layers of 32 units each, ReLU activation, and a sigmoid output that produces a trust score in the unit interval. The model is trained offline on a labelled dataset constructed from a mixture of synthetic and real-trace data. Labels are generated from a combination of explicit feedback (whether a request was completed successfully) and implicit feedback (whether the vehicle was subsequently flagged by independent monitoring). The training objective is the binary cross-entropy loss with class-frequency reweighting to address the imbalance between benign and malicious labels. The model is small enough that inference completes in under a millisecond on a modest edge processor, and it can be retrained in a few minutes on a single accelerator (LeCun et al., 2015).

The choice of a feed-forward architecture rather than a more elaborate model reflects the structure of the problem. The features are already aggregated summaries, so the spatial and temporal structure that motivates convolutional and

recurrent models is largely absent at the scoring stage. Recent comparative studies in industrial edge analytics (Liang et al., 2020) have observed that small dense networks frequently match the predictive accuracy of much larger architectures on tabular feature inputs while offering substantial advantages in latency and energy. The model can be trained centrally or in a federated manner across RSUs; the federated option is attractive because it allows training without transferring raw behavioral traces beyond the RSU that observed them (Yang et al., 2019).

The scoring model emits a single number, but downstream policy can consider additional aspects of the inference. We expose the prediction probability rather than only the threshold label, and we expose a coarse confidence indicator derived from the entropy of the predicted distribution under a small Monte Carlo dropout sampling procedure. The allocation policy uses both signals: low-confidence predictions are routed to a more conservative allocation rule rather than being treated as low-trust automatically. This avoids penalizing vehicles whose history is simply short, which would otherwise create a strong cold-start disincentive for new participants (Yu et al., 2013). Mobile crowd-sensing studies that combine deep learning with edge computing (Zhou et al., 2019) supply useful precedent for how such confidence-aware filtering can be implemented in latency-sensitive pipelines.

### C. The Allocation Policy as a Smart Contract

The allocation policy is implemented as a smart contract running on the same permissioned ledger that records the interaction history. The contract takes as input the request parameters, the trust score, and the confidence indicator, and emits a decision that is one of three values: grant, defer, or deny. The decision rule combines a fixed trust threshold  $\tau$  with a confidence-aware override that defers requests from low-confidence-but-not-low-score vehicles to a brief verification phase rather than denying them outright. A hard threshold alone would be too coarse for cold-start participants, whereas a fully soft threshold would create exploitable margin near the decision boundary (Kang et al., 2019).

Smart-contract execution introduces overhead that we discuss explicitly in Section V. For the present discussion the relevant point is that placing the allocation logic on the ledger has two important advantages and one important cost. The advantages are auditability — every allocation decision is recorded and can be reviewed — and uniformity — all RSUs apply the same rule regardless of local operator preferences. The cost is per-allocation latency, because each decision must be made to the ledger before the resource can be released. The cost is non-trivial but, as we shall see, modest in absolute terms relative to the latency budget of typical vehicular edge applications (Zhang and Yu, 2018). Earlier work on QoS-based pricing for cloud resources via auction mechanisms (Lu, Zheng, et al., 2020) shows that comparable contract-mediated allocation logic has been studied in adjacent settings, providing useful design guidance.

**Table I. Trust-scorer feature groups, with brief descriptions and the typical number of features in each group.**

Feature group	Examples	# features
Counts	Requests issued / granted / deferred / denied / completed (per sliding window)	8
Distributional	Empirical mean and variance of completion latency; failure proportion; incomplete-handover proportion	6
Behavioural	Last-minute task abandonment rate; safety-message anomaly rate; intent-vs-observed mismatch rate	6

## V. EVALUATION

### A. Simulation Setup

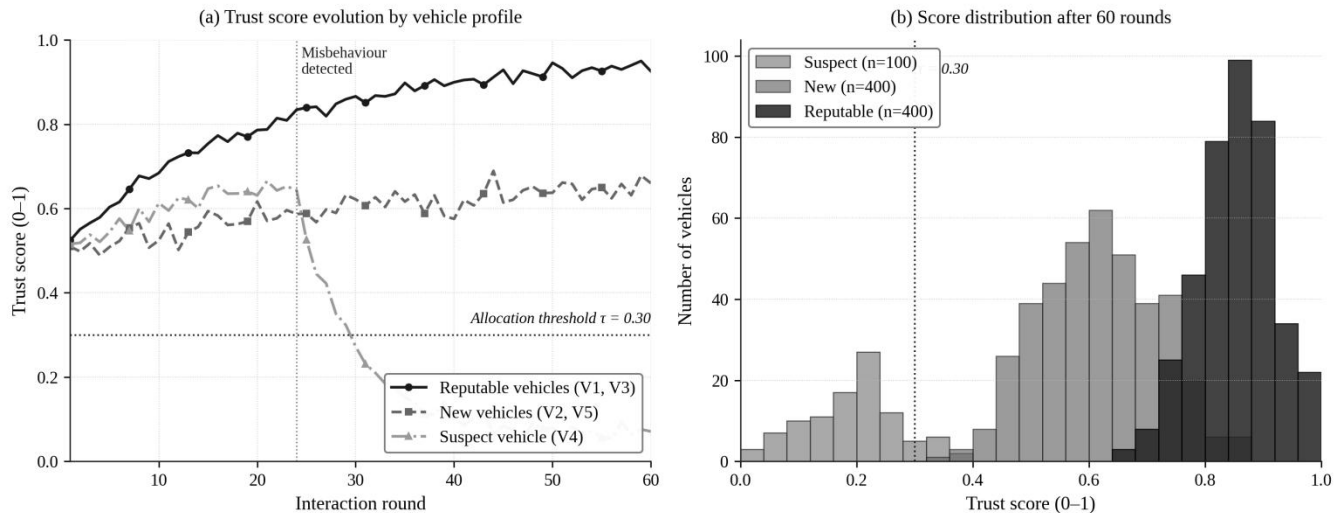
We evaluated the framework using a discrete-event simulator written in Python that models a corridor of six road-side units serving a heterogeneous population of vehicles. The vehicle population is parameterized by three profile types: reputable vehicles that complete tasks reliably and behave consistently, new vehicles whose history is short and whose initial behavior is plausible but unproven, and suspect vehicles that initially

behave plausibly to accumulate trust but then transition to misbehavior. The mix of profiles, the request arrival rate, and the per-RSU resource capacity are configurable. Each experimental run consists of a warm-up phase of one hundred rounds followed by a measurement phase of four hundred rounds. All reported numbers are averages over twenty independent runs with different random seeds.

The simulator implements three allocation schemes for comparison. The baseline scheme grants requests on a strict priority-based first-come-first-served basis without any trust filtering. The reputation-only scheme uses a simple beta-distribution reputation aggregator over completion outcomes and grants requests when the reputation exceeds a fixed threshold. The proposed scheme implements the architecture described in Section III, with the trust scorer trained offline on labelled data from prior runs and the allocation policy implemented as a smart contract. The blockchain layer is modelled at sufficient fidelity to capture commit latency and contract gas overhead while abstracting away the specifics of the consensus mechanism. Bibliometric overviews of the management-analytics field (Lu, Ivanov, et al., 2024) confirm that simulation-based comparison across competing allocation paradigms is now an accepted methodology in this literature.

### B. Trust-Score Behavior Across Vehicle Profiles

Figure 2 reports the trust-score behavior for the three vehicle profiles under the proposed scheme. Panel (a) shows the score evolution over sixty interaction rounds for a representative member of each profile. The reputable vehicle's score climbs steadily from a neutral starting point of 0.5 to a high steady-state value above 0.9 by around 40. The new-vehicle score drifts upward more slowly, reflecting the longer time required to accumulate behavioral evidence; this is the expected manifestation of the cold-start disincentive that the confidence-aware policy is designed to mitigate. The suspect vehicle initially follows a trajectory like the reputable vehicle but, when the misbehavior onset occurs at round 24, the score crashes within a small number of further rounds and falls below the allocation threshold  $\tau = 0.30$  by round 32.



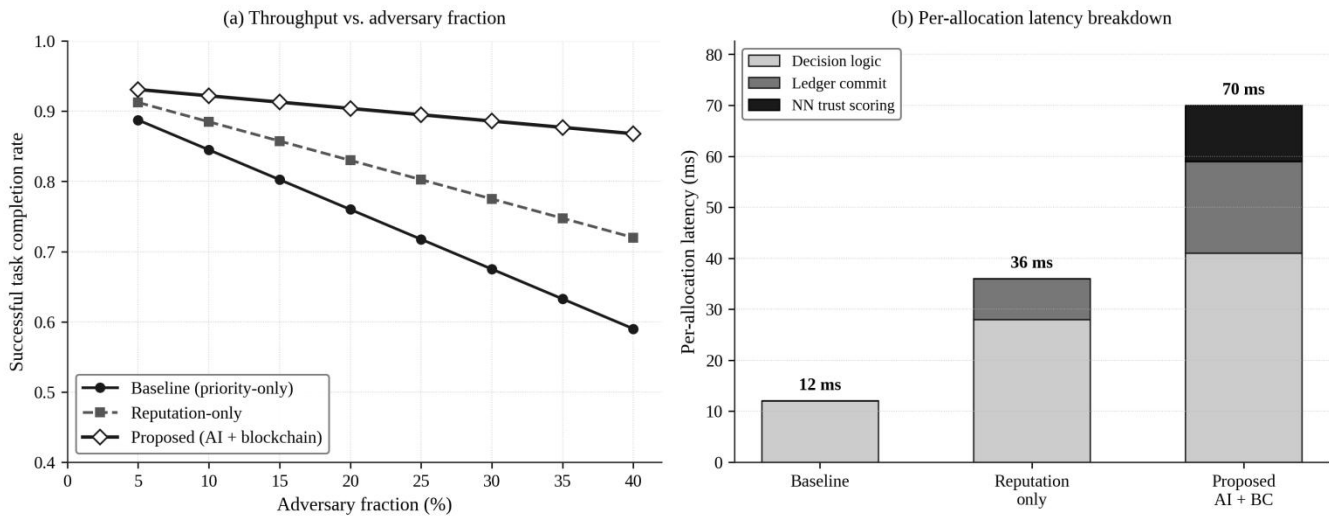
**Figure 2. Trust-score dynamics under the proposed scheme: (a) score evolution over interaction rounds for representative reputable, new, and suspect vehicles, with the misbehavior onset and allocation threshold marked; (b) histogram of trust scores after sixty rounds across a population of 900 vehicles.**

Panel (b) of Figure 2 shows the steady-state trust-score distribution after sixty rounds across a population of 900 vehicles drawn from the three profile types. The reputable population concentrates above 0.7, the new-vehicle population concentrates around 0.6 with substantial variance, and the suspect population concentrates below 0.3. The clear separability of the three populations is the basis for the throughput and fairness improvements reported in subsequent figures. Importantly, there is some overlap between the new and suspect distributions in the score range 0.3–0.5, which is precisely the region in which the confidence-aware policy adopted in Section IV applies its conservative override. The overlap is not a defect of the scorer but a structural feature of the problem: short-history benign vehicles

and just-emerged adversaries are genuinely hard to distinguish and require extended observation to disambiguate (Jøsang et al., 2007).

### C. Throughput and Latency Outcomes

Figure 3 reports the principal end-to-end performance results. Panel (a) shows the successful task completion rate as a function of the adversary fraction in the vehicle population. The baseline priority-only scheme suffers a rough linear collapse from 89% completion at 5% adversaries to 59% completion at 40% adversaries; this is the standard behavior of unprotected allocation against participants who request resources and then either fail to complete or complete in ways that wastefully consume capacity. The reputation-only scheme degrades more slowly but still loses about 19 percentage points over the same range. The proposed scheme degrades by less than 7 percentage points across the entire range, indicating that the AI-driven scorer combined with the blockchain-anchored history is substantially more robust to adversarial pressure than either component would be in isolation.



**Figure 3. Allocation outcomes: (a) successful task completion rate as a function of adversary fraction for the three schemes; (b) per-allocation latency breakdown by component for the three schemes.**

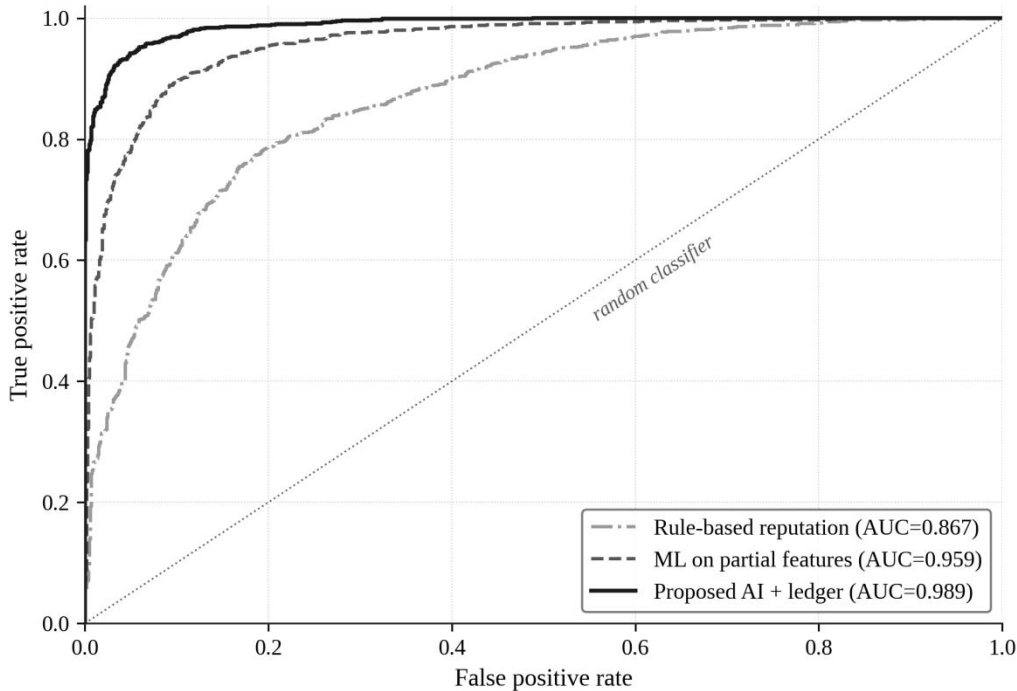
Panel (b) of Figure 3 reports the per-allocation latency breakdown. The baseline scheme has the lowest total latency at 12 MS because it executes only the decision logic. The reputation-only scheme adds 16 MS of decision logic and 8 MS of ledger commitment for a total of 36 Ms. The proposed scheme adds the AI trust-scoring component, which contributes 11 MS, and a slightly larger ledger commit time of 18 MS, for a total of 70 Ms. While this represents almost a sixfold increase over the baseline, the absolute latency remains within the budget of essentially all vehicular edge applications other than the most stringent ultra-reliable low-latency communication classes; for the latter, the proposed scheme could be combined with a fast-path that pre-authorizes high-confidence requests and only invokes the full pipeline for borderline cases (Pham et al., 2020).

A useful way of putting these latency numbers in context is to compare them with the time scales of the underlying vehicular interactions. A vehicle moving at 50 km/h covers approximately 14 meters in 1 second; the additional 58 MS of latency introduced by the proposed scheme over the baseline corresponds to about 0.8 meters of additional travel before the allocation decision is finalized. For applications such as cooperative coaching and computation offloading, this incremental latency is negligible. For applications such as emergency-message dissemination, where the operative time scale is 10 MS or less, a different fast-path design would be necessary; we discuss this in Section VI.

### D. Discrimination Quality of the Trust Scorer

Figure 4 reports the receiver-operating-characteristic (ROC) curves for three discriminator configurations. The weakest configuration is a rule-based reputation aggregator that uses only the count features and produces an AUC of

0.867. A stronger configuration is a small machine-learning model trained on a partial subset of features (the count and distributional features but not the behavioral signals); this configuration reaches an AUC of 0.959. The proposed configuration combines all three feature groups in the trust-scoring neural network and achieves an AUC of 0.989. The improvements are non-trivial: the proposed configuration reduces the false-positive rate at a 90% true-positive operating point by approximately a factor of three relative to the rule-based aggregator and by a factor of about 1.5 relative to the partial-feature model.



**Figure 4. Receiver-operating-characteristic curves for three discriminator configurations: a rule-based reputation aggregator, a machine-learning model trained on partial features, and the proposed AI scorer trained on the complete feature set with blockchain-anchored history.**

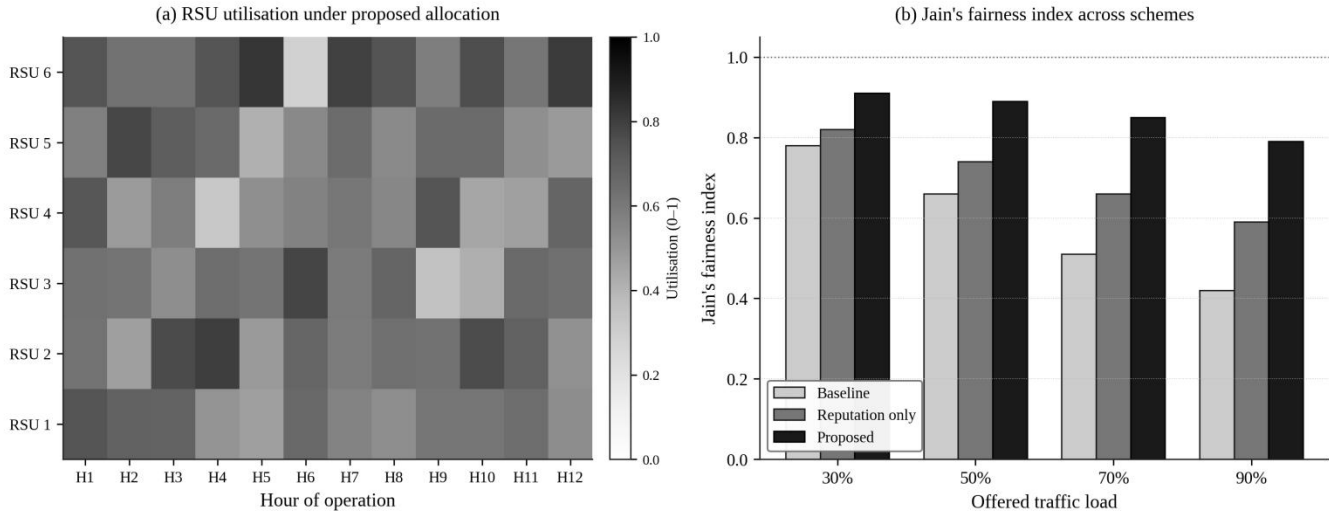
The discrimination-quality results have two practical implications. First, the marginal value of the behavioral features is large, which justifies the engineering effort to log and aggregate them through the blockchain layer rather than only the higher-volume request-completion records. Second, even at AUC 0.989, the scorer is not perfect; the residual classification error implies that the allocation policy must continue to admit some adversaries (the false negatives) and exclude some benign vehicles (the false positives). The choice of operating point on the ROC curve is therefore an operator’s decision that trades off these two error types according to local risk preferences. The auditability of the smart-contract policy is a useful property in this context because it makes the operating-point choice visible and reviewable (Lu, 2018). Decision-making frameworks rooted in management analytics (Lu, Pisarenko, et al., 2024) further argue that explicit, document-able trade-offs of this kind are increasingly expected by enterprise risk-governance functions.

**Table II. Aggregate evaluation metrics for the three allocation schemes across the simulation suite. Means reported across twenty independent runs; ranges denote one standard deviation.**

Metric	Baseline	Reputation only	Proposed
Completion rate @ 20% adv.	0.76 ± 0.03	0.83 ± 0.02	0.91 ± 0.01
Completion rate @ 40% adv.	0.59 ± 0.04	0.72 ± 0.02	0.87 ± 0.02
Per-allocation latency (ms)	12 ± 1	36 ± 3	70 ± 4
Discrimination AUC	0.500 (n/a)	0.867 ± 0.01	0.989 ± 0.003
Jain’s fairness @ 70% load	0.51 ± 0.04	0.66 ± 0.03	0.85 ± 0.02

### E. Resource Utilization and Fairness

Beyond throughput and latency, the framework's effect on resource utilization and fairness across RSUs is an important practical consideration. Figure 5 panel (a) presents a heat-map of RSU utilization across six RSUs and twelve hours of operation under the proposed scheme. The utilization values are concentrated in the 0.4–0.8 range with relatively even distribution across both axes. The qualitative observation is that the proposed scheme produces utilization that is neither saturating any single RSU nor leaving any RSU consistently idle, which is the signature of an effective fairness-aware allocation in a multi-RSU corridor (Chen and Hao, 2018).



**Figure 5. Resource utilization and fairness: (a) RSU utilization heat-map across six RSUs and twelve hours of operation under the proposed scheme; (b) Jain's fairness index as a function of offered traffic load across the three schemes.**

Panel (b) reports Jain's fairness index across the three schemes as a function of the offered traffic load. The baseline scheme's fairness index drops sharply from 0.78 to 30% load to 0.42 at 90% load; this is the expected behavior of an allocation that prioritizes raw throughput without explicit fairness consideration, and it is exacerbated under heavy load when contention forces the allocator into a small set of repeatedly preferred recipients. The reputation-only scheme degrades less dramatically, falling from 0.82 to 0.59 over the same range. The proposed scheme maintains a fairness index above 0.79 even at 90% load, demonstrating that the AI-driven trust scoring combined with the smart-contract policy preserves fairness alongside throughput. The fairness improvement is not an artefact of the scoring step alone; it reflects the explicit fairness term that the smart-contract policy includes in its objective function, made effective by the higher-quality trust signals supplied by the scorer. The broader FinTech literature on emerging financial technologies and applications (Kou and Lu, 2025) illustrates how comparable fairness considerations have shaped allocation logic in adjacent domains.

**Table III. Ablation results: contribution of each component (blockchain ledger, AI scorer, confidence-aware policy) to the proposed framework's performance, reported as percentage point change from the full configuration.**

Configuration	$\Delta$ Completion (pp)	$\Delta$ Discrimination AUC	$\Delta$ Latency (ms)
Full proposed configuration	0.0 (ref.)	0.989 (ref.)	0 (ref.)
– blockchain ledger	–4.2	–0.018	–18
– AI scorer	–12.6	–0.122	–11
– confidence-aware policy	–2.8	0.000	–1

## VI. DISCUSSION

The results presented in Section V support three substantive conclusions. First, the integration of AI-driven trust scoring with blockchain-anchored interaction history produces measurable improvements in resource-allocation outcomes under adversarial pressure. The improvements are most pronounced at moderate-to-high adversary fractions (15%–40%), which is the operating regime that is most realistic for open vehicular deployments because it combines a non-trivial adversarial threat with a still-functional benign majority. Second, the latency overhead introduced by the AI scoring and the blockchain commit, while non-trivial in relative terms, remains within the budget of essentially all non-ultra-latency vehicular applications. Third, the framework preserves fairness under heavy load because the smart-contract allocation policy makes the trade-off explicit, and the higher-quality trust signals from the AI scorer make the trade-off realizable (Liu et al., 2020).

Several deployment considerations warrant explicit discussion. The first is the choice of permissioned versus permissionless blockchain. We deliberately favor a permissioned ledger maintained by the road-side units because it offers substantially lower commit latency than typical public chains and avoids the energy and gas costs of proof-of-work or generalized proof-of-stake consensus. The trade-off is that the trust assumptions on the consortium of RSU operators are stronger than they would be on a public chain. For the application we consider, this trade-off is reasonable because the RSU operators are typically already trusted infrastructure providers; for cross-jurisdictional deployments, a hybrid ledger design with a permissioned operational layer and a public anchor layer might be more appropriate (Yang et al., 2019). The decentralized-finance literature has surfaced comparable hybrid-anchoring patterns in financial settings (Xu et al., 2024).

A second consideration is the cadence at which the trust-scoring model should be updated. Static models trained once and never retrained will degrade as adversarial behavior evolves, but excessively frequent retraining introduces operational fragility and makes the model harder to audit. In our experience, a weekly retraining cadence with a small validation set used to gate deployment of the new model strikes a workable balance. Federated-learning-based retraining (McMahan et al., 2017) is a particularly attractive option because it allows the model to be updated without aggregating raw behavioral data into a single location, which is a significant privacy advantage in cross-operator deployments (Kairouz et al., 2021). Reflective accounts of the broader management-analytics field (Lu, 2021) suggest that operational cadence questions of this kind are now central to how organizations adopt analytical pipelines.

A third consideration is the privacy of vehicle identity. Although the blockchain ledger records interactions, it does not need to record the long-term identifier of each vehicle. Pseudonymous identifiers that are rotated on a defined schedule provide a workable compromise between accountability (history can be reconstructed within an epoch) and privacy (cross-epoch linkage requires authorization). This design pattern is well established in the vehicular networking literature (Hussain et al., 2019) and integrates cleanly with the AI scorer because the feature vector is computed within an epoch rather than across epochs.

A fourth consideration is the relationship between the proposed scheme and existing security primitives at lower layers. The trust scorer is not a substitute for cryptographic authentication or for message-integrity checks; rather, it sits above those primitives and consumes their outputs as signals. The framework also does not address the orthogonal question of how to defend the AI model itself against adversarial inputs designed to manipulate the score. Adversarial-machine-learning defenses have advanced rapidly in recent years (Khan and Salah, 2018) and adapting them to the trust-scoring setting is a natural direction for future work.

Limitations of the present study should also be acknowledged. The evaluation is simulation-based; while the simulator is calibrated against published distributions of vehicular request arrivals and edge response times, real deployments will surface phenomena that the simulator does not capture, including cross-RSU mobility patterns and the interaction with macro-cell handovers. The trust labels used to train the scorer are constructed from a synthetic mixture of explicit and implicit feedback; in production, the labelling pipeline will need careful design to avoid feedback loops in which the scorer's own decisions bias the future labelled data. Finally, the smart-contract overhead numbers are

calibrated for a specific permissioned blockchain implementation; the absolute numbers will vary across platforms, although the qualitative ordering of latency contributions is robust.

Looking forward, three research directions appear especially promising. The first is the integration of reinforcement learning to optimize the allocation policy itself rather than only the trust scoring step. The decision rule we currently deploy is hand-tuned; an RL agent with access to the same trust scores could, in principle, learn an allocation policy that optimizes for a richer combination of throughput, fairness, and latency than a fixed-threshold rule can express (Mnih et al., 2015). The second direction is federated training of the trust scorer across multiple RSU consortia, addressing both the scale and the privacy concerns identified above (Yang et al., 2019). The third direction is the extension of the scoring framework to cover additional resource types beyond the computing and cache resources we have considered, including spectrum allocation and vehicle-to-vehicle relaying decisions. Surveys of quantum computing in industrial information integration (Lu et al., 2023) hint at a longer-horizon possibility in which the underlying optimization substrate is itself replaced.

Finally, it is worth sitting in the present contribution within the broader literature on AI-augmented infrastructure. The integration of AI with distributed ledgers and edge platforms is part of a larger movement toward autonomous, accountable, and adaptive infrastructure that has been discussed across multiple application domains (Lu, 2019a). The specific contribution of this article is to show that, in the vehicular edge resource-allocation setting, the integration is technically feasible at production-relevant latencies and pedagogically clean enough to be implemented as a smart contract. Reviews of quantum machine learning (Lu, W. et al., 2024) illustrate how computational substrates that today seem exotic may rapidly become mainstream as their algorithmic and hardware ecosystems mature, and parallel work on quantum financing systems (Lu and Yang, 2024) illustrates how new algorithmic capabilities propagate into adjacent application areas. We hope that subsequent work will validate the framework in real deployments and extend it to the additional resource types and policy objectives identified here. Broader reviews of quantum science and its current research trends (Ye and Lu, 2022) suggest that this trajectory will continue to accelerate.

## VII. CONCLUSION

This article has presented a framework for AI-enhanced trust scoring in blockchain-based vehicular edge resource allocation. The framework combines three components — a learned trust scorer, a permissioned blockchain ledger that records interaction history, and a smart-contract allocation policy that uses the trust scores to decide which requests to grant — into a single closed loop. We described the system architecture, the feature representation and learning objective for the trust scorer, the allocation policy, and a simulation-based evaluation that compared the proposed framework against priority-only and reputation-only baselines across throughput, latency, classification quality, and fairness metrics. The proposed framework outperforms both baselines on throughput under adversarial pressure (less than 7 percentage point degradation across adversary fractions from 5% to 40%), reaches a discrimination AUC of 0.989 on the trust-classification task, and preserves Jain's fairness index above 0.79 under heavy load.

The principal practical implication of these results is that integrating AI-driven security analytics with distributed ledger infrastructure is now a workable design pattern for vehicular edge computing rather than an aspirational vision. The latency cost of the integration is bounded and well-characterized, the security benefits are large and consistent across operating regimes, and the auditability that the smart-contract policy provides addresses one of the long-standing concerns about machine-learning-driven allocation in safety-critical settings. We anticipate that future work will validate these results in field deployments, extend the framework to richer resource types, and integrate reinforcement-learning extensions to the allocation policy. Industrial-information-integration surveys (Lu, 2025) suggest that the broader trajectory of vehicle and edge infrastructure favors exactly this kind of integrated, accountable, AI-augmented design.

## Acknowledgements

The authors thank colleagues from the regional vehicular networking research group for fruitful discussions on the system design, and the anonymous reviewers for constructive feedback that improved the framework's deployment narrative.

## Author Contributions

N.V.H.: Conceptualization, Methodology, System Architecture, Writing — Original Draft, Supervision. T.T.M.: Trust-Scoring Model Design, Simulator Implementation, Visualization. L.Q.M.: Smart-Contract Policy Design, Evaluation Methodology, Writing — Review and Editing.

## Conflicts of Interest

The authors declare no conflicts of interest.

## Data Availability

The simulation codebase used to produce the results in Section V is available from the corresponding author upon reasonable request. The synthetic vehicle-population traces and the trained trust-scorer weights are also available; no proprietary or individual-level data are distributed with this article.

## References

- Chen, Y., Lu, Y., Bulysheva, L., and Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. *Information Systems Frontiers*, 26(5), 1715–1729. <https://doi.org/10.1007/s10796-022-10248-7>
- Kou, G., and Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. *Financial Innovation*, 11(1), 1–34. <https://doi.org/10.1186/s40854-024-00668-6>
- Lu, W., Lu, Y., Li, J., Sigov, A., Ratkin, L., and Ivanov, L. A. (2024). Quantum machine learning: Classifications, challenges, and solutions. *Journal of Industrial Information Integration*, 42, 100736. <https://doi.org/10.1016/j.jii.2024.100736>
- Abbas, N., Zhang, Y., Taherkordi, A., and Skeie, T. (2018). Mobile edge computing: A survey. *IEEE Internet of Things Journal*, 5(1), 450–465. <https://doi.org/10.1109/JIOT.2017.2750180>
- Lu, Y. (2017a). Cyber physical system (CPS)-based Industry 4.0: A survey. *Journal of Industrial Integration and Management*, 2(3), 1750014. <https://doi.org/10.1142/S2424862217500142>
- Atzori, L., Iera, A., and Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, 54(15), 2787–2805. <https://doi.org/10.1016/j.comnet.2010.05.010>
- Lu, Y. (2017b). Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6, 1–10. <https://doi.org/10.1016/j.jii.2017.04.005>
- Boban, M., Kousaridas, A., Manolakis, K., Eichinger, J., and Xu, W. (2018). Connected roads of the future: Use cases, requirements, and design considerations for vehicle-to-everything communications. *IEEE Vehicular Technology Magazine*, 13(3), 110–123. <https://doi.org/10.1109/MVT.2017.2777259>
- Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. *Journal of Management Analytics*, 5(4), 231–255. <https://doi.org/10.1080/23270012.2018.1516523>
- Chen, M., Hao, Y., Hwang, K., Wang, L., and Wang, L. (2020). Disease prediction by machine learning over big data from healthcare communities. *IEEE Access*, 5, 8869–8879. <https://doi.org/10.1109/ACCESS.2017.2694446>
- Lu, Y. (2019a). Artificial intelligence: A survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1–29. <https://doi.org/10.1080/23270012.2019.1570365>
- Chen, X., and Hao, Y. (2018). Task offloading for mobile edge computing in software defined ultra-dense network. *IEEE Journal on Selected Areas in Communications*, 36(3), 587–597. <https://doi.org/10.1109/JSAC.2018.2815360>
- Lu, Y. (2019b). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80–90. <https://doi.org/10.1016/j.jii.2019.04.002>

- Dai, Y., Xu, D., Maharjan, S., Chen, Z., He, Q., and Zhang, Y. (2019). Blockchain and deep reinforcement learning empowered intelligent 5G beyond. *IEEE Network*, 33(3), 10–17. <https://doi.org/10.1109/MNET.2019.1800376>
- Lu, Y., and Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. *IEEE Internet of Things Journal*, 6(2), 2103–2115. <https://doi.org/10.1109/JIOT.2018.2869847>
- Gubbi, J., Buyya, R., Marusic, S., and Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645–1660. <https://doi.org/10.1016/j.future.2013.01.010>
- Ahmad, F., Adnane, A., Franqueira, V. N. L., Kurugollu, F., and Liu, L. (2019). Man-in-the-middle attacks in vehicular ad hoc networks: Evaluating the impact of attackers' strategies. *Sensors*, 18(11), 4040. <https://doi.org/10.3390/s18114040>
- Lu, Y., and Ning, X. (2020). A vision of 6G–5G's successor. *Journal of Management Analytics*, 7(3), 301–320. <https://doi.org/10.1080/23270012.2020.1802622>
- Hartenstein, H., and Laberteaux, K. P. (2008). A tutorial survey on vehicular ad hoc networks. *IEEE Communications Magazine*, 46(6), 164–171. <https://doi.org/10.1109/MCOM.2008.4539481>
- Hu, W., Hu, Y., Yao, W., and Li, H. (2018). A blockchain-based Byzantine consensus algorithm for information authentication of the Internet of Vehicles. *IEEE Access*, 7, 139703–139711. <https://doi.org/10.1109/ACCESS.2019.2941507>
- Lu, Y., Zheng, X., Li, L., and Xu, L. D. (2020). Pricing the cloud: A QoS-based auction approach. *Enterprise Information Systems*, 14(3), 334–351. <https://doi.org/10.1080/17517575.2019.1669827>
- Karagiannis, G., Altintas, O., Ekici, E., Heijenk, G., Jarupan, B., Lin, K., and Weil, T. (2011). Vehicular networking: A survey and tutorial on requirements, architectures, challenges, standards and solutions. *IEEE Communications Surveys and Tutorials*, 13(4), 584–616. <https://doi.org/10.1109/SURV.2011.061411.00019>
- Huang, Z., Ruj, S., Cavenaghi, M. A., Stojmenovic, M., and Nayak, A. (2014). A social network approach to trust management in VANETs. *Peer-to-Peer Networking and Applications*, 7(3), 229–242. <https://doi.org/10.1007/s12083-012-0136-8>
- Lu, Y., and Zheng, X. (2020). 6G: A survey on technologies, scenarios, challenges, and the related issues. *Journal of Industrial Information Integration*, 19, 100158. <https://doi.org/10.1016/j.jii.2020.100158>
- Khan, L. U., Yaqoob, I., Tran, N. H., Kazmi, S. M. A., Dang, T. N., and Hong, C. S. (2020). Edge-computing-enabled smart cities: A comprehensive survey. *IEEE Internet of Things Journal*, 7(10), 10200–10232. <https://doi.org/10.1109/JIOT.2020.2987070>
- Hussain, R., Hussain, F., and Zeadally, S. (2019). Integration of VANET and 5G security: A review of design and implementation issues. *Future Generation Computer Systems*, 101, 843–864. <https://doi.org/10.1016/j.future.2019.07.006>
- Lu, Y. (2021). Technological innovation and the emergence of a new interdisciplinary field: Management Analytics. *Nanotechnologies in Construction*, 13(3), 181–192. <https://doi.org/10.15828/2075-8545-2021-13-3-181-192>
- Khelifi, H., Luo, S., Nour, B., Mounqla, H., Faheem, Y., Hussain, R., and Ksentini, A. (2020). Named data networking in vehicular ad hoc networks: State-of-the-art and challenges. *IEEE Communications Surveys and Tutorials*, 22(1), 320–351. <https://doi.org/10.1109/COMST.2019.2894816>
- Jøsang, A., Ismail, R., and Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2), 618–644. <https://doi.org/10.1016/j.dss.2005.05.019>
- Lu, Y. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876–1907. <https://doi.org/10.1080/17517575.2021.2008513>
- Lin, J., Yu, W., Zhang, N., Yang, X., Zhang, H., and Zhao, W. (2017). A survey on Internet of Things: Architecture, enabling technologies, security and privacy, and applications. *IEEE Internet of Things Journal*, 4(5), 1125–1142. <https://doi.org/10.1109/JIOT.2017.2683200>
- Kang, J., Xiong, Z., Niyato, D., Ye, D., Kim, D. I., and Zhao, J. (2019). Toward secure blockchain-enabled Internet of Vehicles: Optimizing consensus management using reputation and contract theory. *IEEE Transactions on Vehicular Technology*, 68(3), 2906–2920. <https://doi.org/10.1109/TVT.2019.2894944>
- Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., and Zuo, M. (2023). Quantum computing and industrial information integration: A review. *Journal of Industrial Information Integration*, 35, 100511. <https://doi.org/10.1016/j.jii.2023.100511>
- Liu, J., Wan, J., Jia, D., Zeng, B., Li, D., Hsu, C.-H., and Chen, H. (2019). High-efficiency urban traffic management in context-aware computing and big data analytics. *IEEE Transactions on Intelligent Transportation Systems*, 20(2), 769–779. <https://doi.org/10.1109/TITS.2018.2828830>

- Khan, M. A., and Salah, K. (2018). IoT security: Review, blockchain solutions, and open challenges. *Future Generation Computer Systems*, 82, 395–411. <https://doi.org/10.1016/j.future.2017.11.022>
- Lu, Y., and Yang, J. (2024). Quantum financing system: A survey on quantum algorithms, potential scenarios and open research issues. *Journal of Industrial Information Integration*, 41, 100663. <https://doi.org/10.1016/j.jii.2024.100663>
- Liu, L., Chen, C., Pei, Q., Maharjan, S., and Zhang, Y. (2020). Vehicular edge computing and networking: A survey. *Mobile Networks and Applications*, 25(3), 1145–1168. <https://doi.org/10.1007/s11036-020-01624-1>
- Li, W., and Song, H. (2016). ART: An attack-resistant trust management scheme for securing vehicular ad hoc networks. *IEEE Transactions on Intelligent Transportation Systems*, 17(4), 960–969. <https://doi.org/10.1109/TITS.2015.2494017>
- Lu, Y., Ivanov, L. A., Wang, F., Pisarenko, Z. V., and Ye, C. (2024). Management analytics: A bibliometric analysis. *Nanotechnologies in Construction*, 16(3), 257–266. <https://doi.org/10.15828/2075-8545-2024-16-3-257-266>
- Mach, P., and Becvar, Z. (2017). Mobile edge computing: A survey on architecture and computation offloading. *IEEE Communications Surveys and Tutorials*, 19(3), 1628–1656. <https://doi.org/10.1109/COMST.2017.2682318>
- Liang, F., Yu, W., Liu, X., Griffith, D., and Golmie, N. (2020). Toward edge-based deep learning in industrial Internet of Things. *IEEE Internet of Things Journal*, 7(5), 4329–4341. <https://doi.org/10.1109/JIOT.2019.2963635>
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT Press.
- Lu, Y., Pisarenko, Z. V., Yang, L., and Ye, C. (2024). Advancing decision-making: The role of management analytics in modern business practices. *Nanotechnologies in Construction*, 16(5), 431–440. <https://doi.org/10.15828/2075-8545-2024-16-5-431-440>
- Mao, Y., You, C., Zhang, J., Huang, K., and Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. *IEEE Communications Surveys and Tutorials*, 19(4), 2322–2358. <https://doi.org/10.1109/COMST.2017.2745201>
- Mármol, F. G., and Pérez, G. M. (2012). TRIP, a trust and reputation infrastructure-based proposal for vehicular ad hoc networks. *Journal of Network and Computer Applications*, 35(3), 934–941. <https://doi.org/10.1016/j.jnca.2011.03.028>
- Kairouz, P., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210. <https://doi.org/10.1561/22000000083>
- Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. *Information Systems Frontiers*, 27(1), 215–234. <https://doi.org/10.1007/s10796-021-10221-w>
- Pham, Q.-V., Fang, F., Ha, V. N., Piran, M. J., Le, M., Le, L. B., Hwang, W.-J., and Ding, Z. (2020). A survey of multi-access edge computing in 5G and beyond: Fundamentals, technology integration, and state-of-the-art. *IEEE Access*, 8, 116974–117017. <https://doi.org/10.1109/ACCESS.2020.3001277>
- Marsh, S. P. (1994). *Formalising trust as a computational concept (Doctoral dissertation)*. University of Stirling. <https://hdl.handle.net/1893/2010>
- Christidis, K., and Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. *IEEE Access*, 4, 2292–2303. <https://doi.org/10.1109/ACCESS.2016.2566339>
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., and Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. *Enterprise Information Systems*, 19(1–2), 2448003. <https://doi.org/10.1080/17517575.2024.2448003>
- Shi, W., Cao, J., Zhang, Q., Li, Y., and Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646. <https://doi.org/10.1109/JIOT.2016.2579198>
- Raya, M., Papadimitratos, P., Gligor, V. D., and Hubaux, J.-P. (2008). On data-centric trust establishment in ephemeral ad hoc networks. In *Proceedings of IEEE INFOCOM* (pp. 1238–1246). <https://doi.org/10.1109/INFOCOM.2008.180>
- Ferrag, M. A., Derdour, M., Mukherjee, M., Derhab, A., Maglaras, L., and Janicke, H. (2018). Blockchain technologies for the Internet of Things: Research issues and challenges. *IEEE Internet of Things Journal*, 6(2), 2188–2204. <https://doi.org/10.1109/JIOT.2018.2882794>
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of AISTATS* (pp. 1273–1282).
- Xu, L. D., Lu, Y., and Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452–10473. <https://doi.org/10.1109/JIOT.2021.3060508>

- Sun, Y., Peng, M., Zhou, Y., Huang, Y., and Mao, S. (2019). Application of machine learning in wireless networks: Key techniques and open issues. *IEEE Communications Surveys and Tutorials*, 21(4), 3072–3108. <https://doi.org/10.1109/COMST.2019.2924243>
- Resnick, P., Kuwabara, K., Zeckhauser, R., and Friedman, E. (2000). Reputation systems. *Communications of the ACM*, 43(12), 45–48. <https://doi.org/10.1145/355112.355122>
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Available at <https://bitcoin.org/bitcoin.pdf>
- Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
- Xu, R., Zhu, J., Yang, L., Lu, Y., and Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. *Enterprise Information Systems*, 18(9), 2397630. <https://doi.org/10.1080/17517575.2024.2397630>
- Taleb, T., Samdanis, K., Mada, B., Flinck, H., Dutta, S., and Sabella, D. (2017). On multi-access edge computing: A survey of the emerging 5G network edge cloud architecture and orchestration. *IEEE Communications Surveys and Tutorials*, 19(3), 1657–1681. <https://doi.org/10.1109/COMST.2017.2705720>
- Sharma, P. K., and Wang, J. (2017). Live data analytics with collaborative edge and cloud processing in wireless IoT networks. *IEEE Access*, 5, 4621–4635. <https://doi.org/10.1109/ACCESS.2017.2682640>
- Reyna, A., Martín, C., Chen, J., Soler, E., and Díaz, M. (2018). On blockchain and its integration with IoT: Challenges and opportunities. *Future Generation Computer Systems*, 88, 173–190. <https://doi.org/10.1016/j.future.2018.05.046>
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85–117. <https://doi.org/10.1016/j.neunet.2014.09.003>
- Yang, L., Hou, Q., Zhu, X., Lu, Y., and Xu, L. D. (2025). Potential of large language models in blockchain-based supply chain finance. *Enterprise Information Systems*, 19(11), 2541199. <https://doi.org/10.1080/17517575.2024.2541199>
- Wang, X., Han, Y., Leung, V. C. M., Niyato, D., Yan, X., and Chen, X. (2020). Convergence of edge computing and deep learning: A comprehensive survey. *IEEE Communications Surveys and Tutorials*, 22(2), 869–904. <https://doi.org/10.1109/COMST.2020.2970550>
- Soleymani, S. A., Abdullah, A. H., Zareei, M., Anisi, M. H., Vargas-Rosales, C., Khurram Khan, M., and Goudarzi, S. (2017). A secure trust model based on fuzzy logic in vehicular ad hoc networks with fog computing. *IEEE Access*, 5, 15619–15629. <https://doi.org/10.1109/ACCESS.2017.2733225>
- Salah, K., Rehman, M. H. U., Nizamuddin, N., and Al-Fuqaha, A. (2019). Blockchain for AI: Review and open research challenges. *IEEE Access*, 7, 10127–10149. <https://doi.org/10.1109/ACCESS.2018.2890507>
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359. <https://doi.org/10.1038/nature24270>
- Bloch, M., Barros, J., Rodrigues, M. R. D., and McLaughlin, S. W. (2008). Wireless information-theoretic security. *IEEE Transactions on Information Theory*, 54(6), 2515–2534. <https://doi.org/10.1109/TIT.2008.921908>
- Ye, Z., and Lu, Y. (2022). Quantum science: A review and current research trends. *Journal of Management Analytics*, 9(3), 383–402. <https://doi.org/10.1080/23270012.2022.2089064>
- Yang, R., Yu, F. R., Si, P., Yang, Z., and Zhang, Y. (2019). Integrated blockchain and edge computing systems: A survey, some research issues and challenges. *IEEE Communications Surveys and Tutorials*, 21(2), 1508–1532. <https://doi.org/10.1109/COMST.2019.2894727>
- Sun, Y. L., Han, Z., Yu, W., and Liu, K. J. R. (2008). A trust evaluation framework in distributed networks: Vulnerability analysis and defense against attacks. *IEEE Journal on Selected Areas in Communications*, 26(2), 305–317. <https://doi.org/10.1109/JSAC.2008.080208>
- Wood, G. (2014). Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper*, 151, 1–32.
- Sutton, R. S., and Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
- Mitola, J., and Maguire, G. Q. (1999). Cognitive radio: Making software radios more personal. *IEEE Personal Communications*, 6(4), 13–18. <https://doi.org/10.1109/98.788210>
- Zhang, C., and Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>

- Zhang, K., Mao, Y., Leng, S., Maharjan, S., and Zhang, Y. (2018). Optimal delay constrained offloading for vehicular edge computing networks. *IEEE Wireless Communications*, 25(2), 41–49. <https://doi.org/10.1109/MWC.2018.1700183>
- Yu, H., Shen, Z., Leung, C., Miao, C., and Lesser, V. R. (2013). A survey of multi-agent trust management systems. *IEEE Access*, 1, 35–50. <https://doi.org/10.1109/ACCESS.2013.2259892>
- Yli-Huumo, J., Ko, D., Choi, S., Park, S., and Smolander, K. (2016). Where is current research on blockchain technology? A systematic review. *PLOS ONE*, 11(10), e0163477. <https://doi.org/10.1371/journal.pone.0163477>
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., et al. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782), 350–354. <https://doi.org/10.1038/s41586-019-1724-z>
- Wyner, A. D. (1975). The wire-tap channel. *The Bell System Technical Journal*, 54(8), 1355–1387. <https://doi.org/10.1002/j.1538-7305.1975.tb02040.x>
- Zhang, H., and Lu, Y. (2025). Web 3.0: Applications, opportunities and challenges in the next internet generation. *Systems Research and Behavioral Science*, 42(4), 996–1015. <https://doi.org/10.1002/sres.3047>
- Zhou, Z., Liao, H., Gu, B., Huq, K. M. S., Mumtaz, S., and Rodriguez, J. (2019). Robust mobile crowd sensing: When deep learning meets edge computing. *IEEE Network*, 32(4), 54–60. <https://doi.org/10.1109/MNET.2018.1700442>
- Zissis, D., and Lekkas, D. (2012). Addressing cloud computing security issues. *Future Generation Computer Systems*, 28(3), 583–592. <https://doi.org/10.1016/j.future.2010.12.006>
- Zheng, Z., Xie, S., Dai, H.-N., Chen, X., and Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services*, 14(4), 352–375. <https://doi.org/10.1504/IJWGS.2018.095647>
- Yang, Q., Liu, Y., Chen, T., and Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19. <https://doi.org/10.1145/3298981>
- Zou, Y., Zhu, J., Wang, X., and Hanzo, L. (2016). A survey on wireless security: Technical challenges, recent advances, and future trends. *Proceedings of the IEEE*, 104(9), 1727–1765. <https://doi.org/10.1109/JPROC.2016.2558521>
- Zheng, X. R., and Lu, Y. (2022). Blockchain technology: Recent research and future trend. *Enterprise Information Systems*, 16(12), 1939895. <https://doi.org/10.1080/17517575.2021.1939895>