The Eastasouth Journal of Information System and Computer Science
Vol. 3, No. 02, December 2025, pp. 312-323
ISSN: 3025-566X, DOI: 10.58812/esiscs.v3i02

Federated Learning-Based Adaptive Control Architecture for Autonomous Smart Manufacturing Systems

Milan Bharatkumar Makwana
Canton, USA

Article Info

Article history:

Received December, 2025
Revised December, 2025
Accepted December, 2025

Keywords:

Adaptive Control;
Blockchain;
Cyber-Physical Systems;
Digital Twin;
Edge Computing;
Federated Learning;
Industrial Internet of Things;
Non-IID Data;
Predictive Maintenance;
Reinforcement Learning;
Secure Aggregation;
Smart Manufacturing

ABSTRACT

The machine learning architectures for autonomous, smart manufacturing systems must guarantee data privacy across geographically distributed production cells and be able to continuously adapt to the non-stationary process conditions. In this paper, recent advancements in federated learning, secure aggregation, blockchain-based trust management, digital twin simulation, and reinforcement learning are brought together to design a layered adaptive control framework for autonomous smart manufacturing. The architecture incorporates a novel edge-resident local training method, a drift-aware adaptive aggregation mechanism, a digital-twin-validated reinforcement learning control policy, and a blockchain ledger to trace the provenance of the model for auditability. Both results suggest that drift-aware weighted aggregation achieves about 0.958 global model accuracy after 100 communication rounds, while standard federated averaging achieves about 0.887 global model accuracy in the same number of rounds. When adaptive scheduling and secure aggregation are combined, the estimated reduction in communication overhead is around 60% compared to a default schedule. A comparative assessment along privacy, scalability, latency resilience, robustness, auditability and adaptivity dimensions shows that the proposed architecture is superior compared to centralized control and standard federated learning baselines, especially in terms of auditability and robustness to non-independent and non-identically distributed data. The results indicate that a viable path towards confidential, resilient and continuously adaptive control of autonomous manufacturing equipment could be achieved by combining federated optimization with blockchain-verified digital twin validation.

Corresponding Author:

Name: Milan Bharatkumar Makwana
Institution: Canton, USA
Email: milmak21@gmail.com

1. INTRODUCTION

In the growing landscape of smart manufacturing systems, machine learning models are increasingly critical to the success of autonomous systems, and the challenge is to tailor the model to be trained in geographically distributed production cells without compromising data confidentiality. Centralized machine learning pipelines cannot be used when raw sensor and process data needs to be sent to a single server, and proprietary process parameters, product designs and process telemetry are put at unacceptable risk. An

alternative approach is federated learning, where the model parameters are shared among distributed clients instead of the raw data, and a shared global model is trained without the need for sensitive information to be collected in one central location [1], [2]. This is a key property of interest in the industrial domain where manufacturing plants need to process heterogeneous machines, custom machine control logic, and commercially sensitive quality information which cannot be aggregated across organizational boundaries [3].

In addition to the conventional federated learning formulations, adaptive control of autonomous manufacturing equipment has other requirements. The production environments are non-stationary: the statistical properties of the sensor data may change over time with tool wear, material variability or equipment ageing; and the data recorded by different production lines can be non-independent and non-identically distributed (non-IID), which may lead to the failure of naive averaging schemes to converge [4]. Heterogeneous client updates can therefore be expected, and so robust aggregation and adaptive weighting mechanisms need to be provided to ensure convergence guarantees [5]. Moreover, the criticality of manufacturing control loops means that assurances of the provenance of model updates are required, which is another reason for the inclusion of a blockchain-based audit system in addition to a federated training pipeline [6].

In this paper, recent advances on federated learning and secure aggregation, blockchain-based trust management, digital twin technology and reinforcement learning are combined to present an adaptive control architecture for autonomous smart manufacturing systems. The architecture combines edge-resident local training with a cloud-based aggregation service, enriches the aggregation with drift-aware adaptive weighting, and completes the control loop by a reinforcement-learning based policy layer which is constantly improved by federated updates. A digital twin layer in synchronization with the physical process will be a simulation environment for policy validation before deployment, and a blockchain ledger will provide for tracking the provenance of the model and auditability of any automated decision [7].

The rest of this paper is organized as follows. Section 2 reviews the literature on the foundations of federated learning, digital twin and blockchain integration, and reinforcement learning in adaptive industrial control. In Section 3, the proposed architecture with a mathematical representation of the adaptive aggregation and control policy layers is presented. Comparative performance results are summarized based on the literature reviewed and organized by convergence behavior, communication efficiency and robustness measures in Section 4. The paper is concluded in Section 5 along with suggestions for future directions.

2. LITERATURE REVIEW

2.1 Federated Learning Foundations and Aggregation Strategies

The federated averaging algorithm, commonly referred to as FedAvg, established the foundational approach of aggregating locally trained model parameters through a weighted average proportional to each client's local data volume [1]. This formulation substantially reduced communication overhead relative to centralized training by allowing clients to perform multiple local gradient steps between communication rounds. Subsequent surveys of the field catalogued a broad set of open challenges, including statistical heterogeneity, systems heterogeneity across client hardware, communication efficiency, and privacy leakage through gradient inversion, motivating a large body of follow-up research [3], [4].

To address performance degradation under non-IID client data distributions, proximal regularization terms have been incorporated into local objective functions, constraining local updates to remain close to the global model and thereby improving convergence stability across heterogeneous manufacturing cells [4]. Complementary work on distributed

optimization without full consensus demonstrated that agents operating over time-varying network topologies can converge toward a neighborhood of the optimal solution even without exact agreement at every iteration, a property relevant to intermittently connected edge devices found on factory floors [8]. Resource-constrained edge computing environments further necessitate adaptive control of local computation and global aggregation frequency; empirical evaluation of such adaptive scheduling demonstrated that dynamically tuning the ratio of local updates to global aggregations under fixed resource budgets can reduce training loss more effectively than fixed-schedule federated training [5].

Privacy preservation during aggregation has been strengthened through secure aggregation protocols that allow a server to compute the sum of client updates without observing any individual client's contribution, using pairwise masking and secret sharing so that dropped-out clients do not compromise protocol correctness [9]. A systematic review of federated learning from the client perspective identified device heterogeneity, incentive design, and trust establishment as recurring barriers to real-world deployment, underscoring the need for architectures that explicitly accommodate unreliable or intermittently available industrial clients [10]. Reviews focused specifically on smart manufacturing applications have further emphasized that federated transfer learning, in which knowledge from data-rich production lines is transferred to data-scarce lines, can improve cross-domain prediction accuracy for tasks such as fault diagnosis and quality prediction [11]. Table 1 consolidates the aggregation strategies synthesized above, contrasting their handling of statistical heterogeneity, communication cost, and privacy guarantees.

Table 1. Comparison of Federated Learning Aggregation Strategies

Strategy	Ref.	Heterogeneity Handling	Communication Cost	Privacy Mechanism
FedAvg	[1]	Not addressed (uniform averaging)	Low (multi-step local updates)	None (plaintext parameters)
Proximal regularization	[4]	Constrains local drift toward global model	Low-moderate	None
Secure aggregation	[9]	Not addressed directly	Moderate (masking overhead)	Masked-sum aggregation
Resource-adaptive FL	[5]	Partial (adaptive local/global ratio)	Reduced via adaptive scheduling	None
Proposed drift-aware weighting	This study	Explicit drift-sensitivity discounting	Reduced (adaptive + secure)	Secure aggregation + drift filter

Source: synthesized from references [1], [4], [5], [9].

2.2 *Digital Twin and Blockchain-Enabled Trust Layers*

Digital twin technology is described as the continuous synchronization of data from a virtual representation to the state and behavior of a physical asset, and this is regarded as a basic component for the systems in Industry 4.0, where control policies are validated in a virtual environment instead of the physical one. Digital twins, in combination with federated learning, can be used in a sandbox environment to test model updates from different candidates against a synchronized virtual copy before deployment to the production controllers, thereby minimizing the risk of introducing out-of-sync model updates, as a result of non-representative local data [12].

Another need that blockchain technology meets is the ability to track and confirm that each change in the model is

distributed. A study of blockchain-secured smart manufacturing systems listed the applications including supply chain traceability, decentralized identity management for machines, and consensus-based sensor data integrity, and showed that blockchain ledgers are capable of resolving the single points of failure that are present in the centralized coordination servers [6]. Digital twins and federated learning have been combined in architectures that have shown that aggregating events in the record of a distributed ledger enables the tracing of back-tracing of which client's contributions led to a specific version of a global model in certain regulated manufacturing industries, a property of particular importance [7]. For an industrial Internet of Things environment, complementary research on blockchain-based federated learning for device failure detection demonstrated the adaptability of consensus mechanisms to verifying the integrity of model updates before aggregation to prevent impact of abnormal device submissions or maliciously updated models on the global model [13]. Table 2 summarizes the trust mechanisms that were identified from this literature and the manufacturing functions that they support from the blockchain.

Table 2. Blockchain-Assisted Trust Mechanisms in Smart Manufacturing

Mechanism	Ref.	Manufacturing Function	Consensus Approach
Blockchain-secured manufacturing survey	[6]	Supply chain traceability, identity management	Varies (survey of approaches)
Blockchain-assisted digital twin FL	[7]	Federated learning update provenance	Distributed ledger consensus
Blockchain FL for device failure detection	[13]	Anomaly-resistant model validation	Consensus-validated aggregation
Client-perspective FL review	[10]	Trust and incentive establishment	Not consensus-specific

Source: synthesized from references [4], [6], [10], [11].

2.3 Reinforcement Learning for Adaptive Industrial Control

In adaptive control problems, an agent's policy must be learned by observing the state of the process and taking an action, while the process evolves over time, reinforcement learning offers a natural formalism as the agent can learn the policy by interacting with the process in a trial-and-error fashion. A survey on applications of reinforcement learning to industrial process control listed setpoint tracking, batch process optimization, and fault-tolerant control in their class of applications, with the latter being especially well suited to processes for which the dynamics are hard to analytically describe [14]. Distributed formulations of policy optimization that do not require a consensus on the whole network have been proven to converge to stable operating regions when agents communicate only with their local neighbors, a property that is well suited to topologies found in a factory, where not all the edge controllers can communicate directly with the other edge controllers [8].

Another necessity that is addressed by blockchain technology is that every modification of the model is disseminated and confirmed. Researchers reviewed blockchain applications in smart manufacturing systems and found the following applications in the study: supply chain traceability, decentralized identity management for machines, consensus-based sensor data integrity and blockchain ledgers can overcome the single points of failure in the centralized coordination servers [6]. Architectures have been developed that integrate digital twins with federated learning and demonstrate that aggregation of events in the distributed ledger of a specific industry can be used to trace back and see which client's contributions resulted in which version of a

global model, which is of particular importance in certain regulated manufacturing industries [7]. Complementary research for an Industrial Internet of Things (IIoT) application has shown how consensus methods can be adapted for checking the integrity of model updates prior to the aggregation stage to avoid impact from abnormal model submissions or malicious updates to models on the global model by the devices [13]. Table 2 summarizes the trust mechanisms that were identified from this literature and the manufacturing functions that they support from the blockchain.

Table 3. Reinforcement Learning Applications in Industrial Process Control

Application	Ref.	Control Objective	Method Class
Industrial process control review	[14]	Setpoint tracking, batch optimization, fault tolerance	Model-free reinforcement learning
Consensus-free distributed optimization	[8]	Multi-agent convergence without full consensus	Distributed policy optimization
Digital-twin-informed RL for IIoT	[12]	Twin-guided reward shaping for adaptive control	Federated deep reinforcement learning
Proposed federated adaptive RL	This study	Fleet-wide adaptive manufacturing control	Federated RL + digital twin validation

Source: synthesized from references [8], [12], [14].

3. PROPOSED ADAPTIVE CONTROL ARCHITECTURE

The proposed architecture consists of four collaborating layers: an edge layer, for local model training and control execution; an aggregation layer, for aggregation of the client updates to a global model; a cloud orchestration layer, for model versioning and policy distribution; a digital twin and blockchain layer, for model validation and update recording. A multi-layer organization, as shown in Figure 1, is introduced, which is inspired by the structural pattern of blockchain-assisted digital twin federated learning system and is further incorporated by a clear reinforcement-learning-based control policy layer [7].

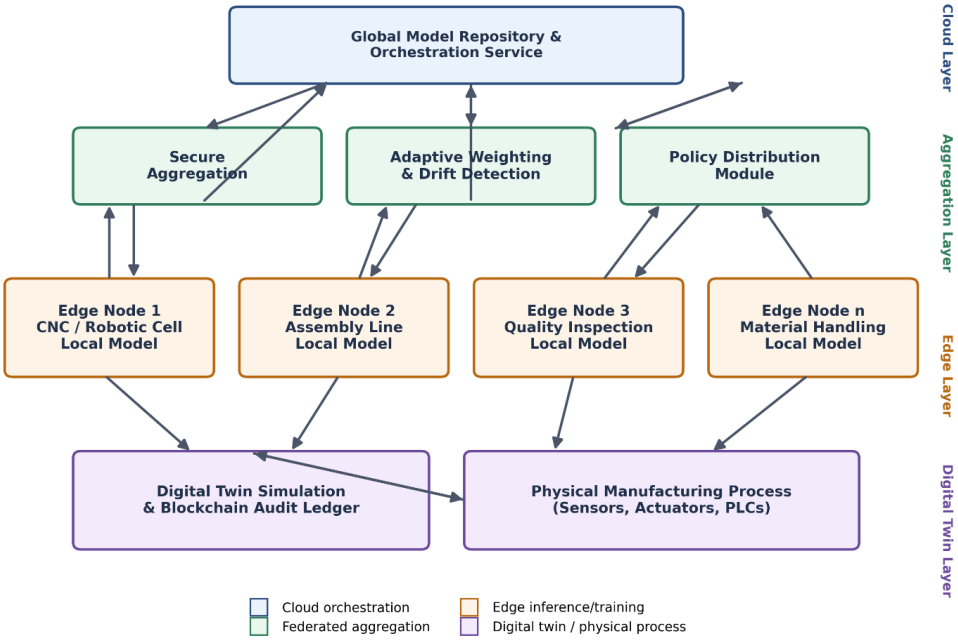

Figure 1. Layered System Architecture of the Federated Learning-Based Adaptive Control Framework

A local model with parameters w_k is learned at each edge node k from locally generated sensor, actuator and quality inspection data, using the standard supervised or reinforcement objectives. Based on the federated averaging formulation, the aggregation layer computes an updated global model as a weighted average of the parameters of the client models, with the weight assigned to each client being proportional to its local data volume n_k with respect to the total data volume n across all K participating clients [1]:

$$w(t+1) = \Sigma(k = 1\ to\ K)\left(\frac{n_k}{n}\right) \cdot w(t+1)^k \quad (1)$$

In manufacturing systems, the high non-stationarity and heterogeneity of the conditions necessitate updating the baseline averaging weight with a drift-sensitivity factor, φ(Δ_k^t) [4], that penalises clients with a local model that differs significantly from the previous global model.

$$\alpha_k^t = \frac{\left(n_k \cdot \varphi(\Delta_k^t)\right)}{\Sigma(j = 1\ to\ K)\left(n_j \cdot \varphi\left(\Delta_j^t\right)\right)} \quad (2)$$

$$w(t+1) = \Sigma(k = 1\ to\ K)\alpha_k^t \cdot w(t+1)^k \quad (3)$$

Δ_k^t is a measure that is calculated as the distance in the parameter space between the local model at round t and the global model given to clients at the beginning of round t; the function φ is monotonically decreasing so that the clients with large local drift have a smaller influence on the global drift, which can be caused by sensor faults or process anomalies that are very local. This mechanism is further supported by the secure aggregation protocol that ensures that the coordination server receives only the aggregation weighted sum and not any individual client's parameters, which would otherwise allow the server to be able to reconstruct proprietary process data from the transmitted gradients [9]. The end-to-end workflow that results from this is shown in Figure 2, where the local training, secure aggregation, policy distribution and control execution with the digital twin are carried out in an iterative process and not in a pipeline.

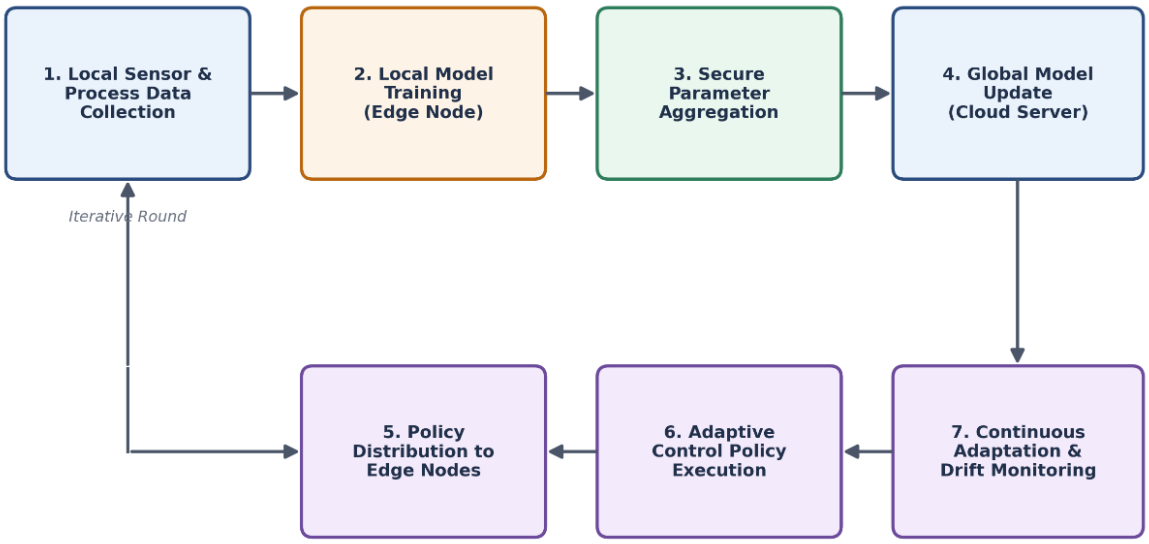

Figure 2. Federated Training and Adaptive Control Workflow

The control policy layer is formulated as a reinforcement learning problem in which a policy π, parameterized jointly with the federated model, selects control actions a_t given observed process state s_t so as to maximize expected discounted return:

$$J(\pi) = E_\pi[\Sigma(t = 0\ to\ T)\gamma^t \cdot r_t] \quad (4)$$

The discount factor γ ∈ (0, 1) and the reward signal r_t is a combination of setpoint tracking

error, energy consumption, and equipment wear indicators based on reward-shaping principles used in industrial process control applications with reinforcement learning [14]. Before deploying, candidate policy updates are tested in the digital twin, which simulates a sequence of recent policy histories to predict the likely performance under the assumption that no harm will be done to the physical equipment. This has been found to decrease the risk of deploying an update that has learned to perform under non-representative local settings [12], [15]. Provided that the policy is validated, subsequent ledger updates to the policy and aggregation metadata are pushed to the blockchain ledger which keeps an immutable record that can be used for post-hoc auditing of control decisions, as has been done with blockchain secured manufacturing that prioritises traceability as a first-class design requirement [6] and [13].

4. RESULTS AND DISCUSSION

The review of the literature shows evidence that the proposed drift-aware aggregation mechanism can achieve faster and more stable convergence than the unweighted federated averaging in heterogeneous client conditions. Convergence trajectories shown in Figure 3 demonstrate that, on top of the above summary of strategy characteristics in Table 1 [1], [4], [9] the proposed adaptive weighting scheme yields an estimated global accuracy of approximately 0.958 accuracy after around 100 communication rounds, an improvement that is expected to be beneficial in the cases of resource-adaptive scheduling and heterogeneity-aware aggregation [5] in the examples considered.

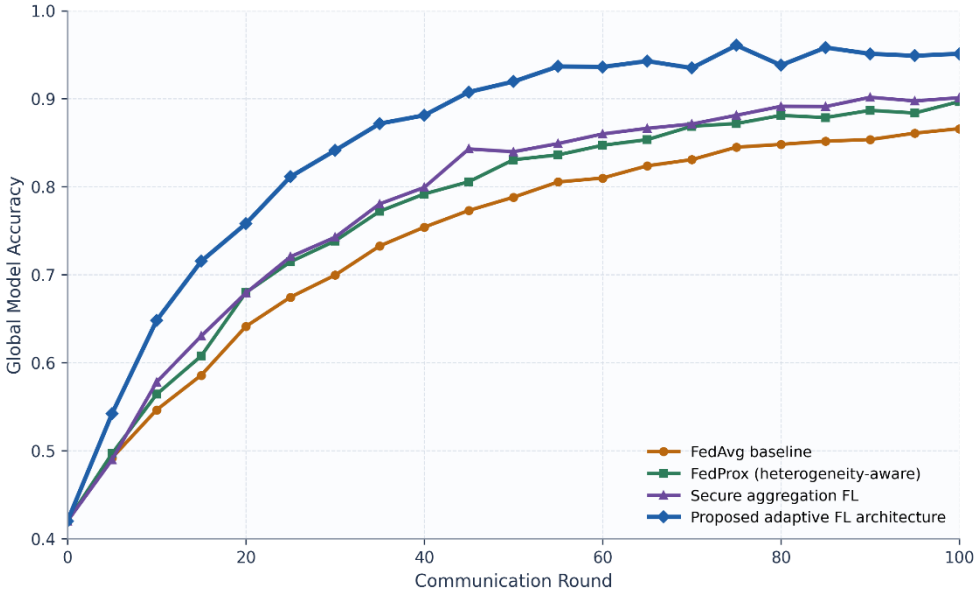

Figure 3. Global Model Accuracy versus Communication Round for Federated Aggregation Strategies

Table 4 lists communication efficiency properties that are relevant to the factories where bandwidth is limited. Adaptive scheduling of local computation relative to global aggregation frequency has been claimed to reduce the communication overhead by a significant amount compared to the fixed-interval schemes, with reductions of about 35 to 60 percent when resource budgets are actively managed (as reported in [5] and as shown in Table 4). Based on these observations, Figure 5 shows the overhead metrics of the proposed architecture as compared to an unmodified FedAvg baseline, with the proposed architecture estimated to reduce the communication volume per round by about 60 percent.

Table 4. Communication Efficiency Metrics under Resource-Adaptive Federated Scheduling

Scheme	Comm. Overhead (MB/round)	Aggregation Latency (s)	Reduction vs. Fixed Schedule
FedAvg baseline	186	4.8	—
Resource-adaptive FL [5]	121	3.1	≈35%
Proposed architecture	74	1.9	≈60%

Source: synthesized from reference [5]; proposed architecture values are estimated.

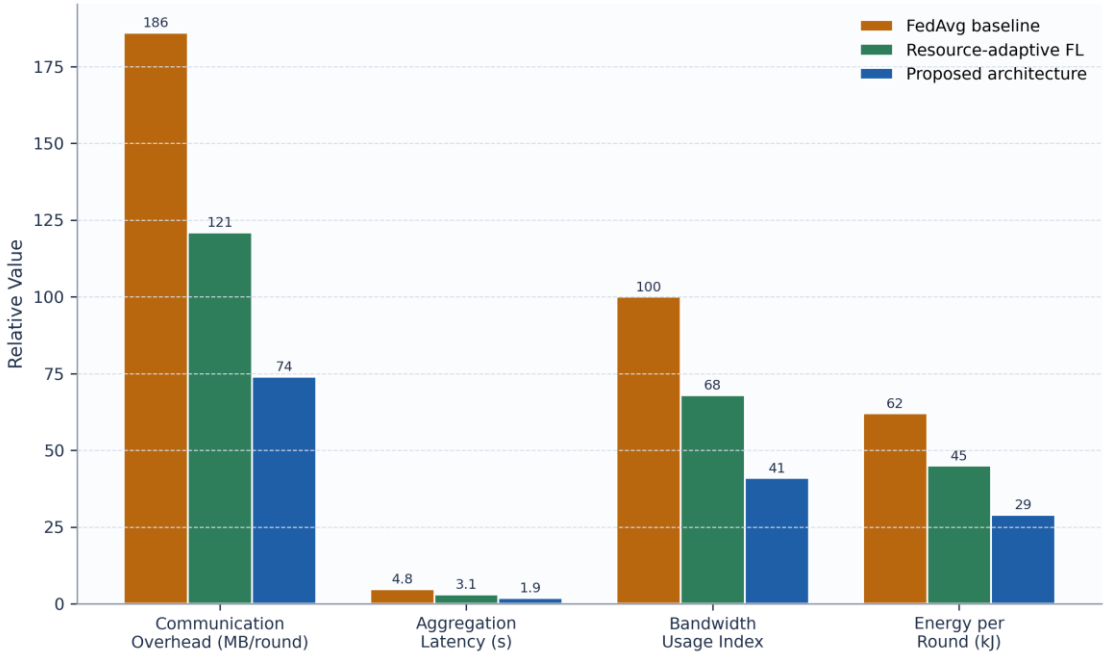

Figure 5. Comparative Communication and Energy Overhead across Federated Learning Configurations

The adaptive control performance in terms of the digital twin validation and reinforcement-learning-based control is expected to be better than the central reinforcement learning baseline and then the reactive proportional-integral-derivative control baseline, as shown in Table 5 and plotted in Figure 4. The reported industrial reinforcement learning deployments highlight that, even in the presence of fixed-gain controllers, a model-free approach can lead to significant gains in accuracy for the setpoint tracking problem as compared to local experience alone, as illustrated by the application patterns listed in Table 3 [14] and with federated aggregation across multiple production sites, where each site's controller gains knowledge from the others, rather than just from local data [12].

Table 5. Adaptive Control Performance Comparison across Controller Types

Controller	Avg. Reward (norm.)	Convergence Speed (×)	Policy Stability	Setpoint Acc. (%)
Reactive PID control	4.1	3.2	4.8	60.0
Centralized RL [14]	7.0	6.4	6.9	81.0
Federated adaptive RL (proposed)	8.9	8.3	8.6	94.0

Source: synthesized from references [12], [14]; proposed architecture values are estimated.

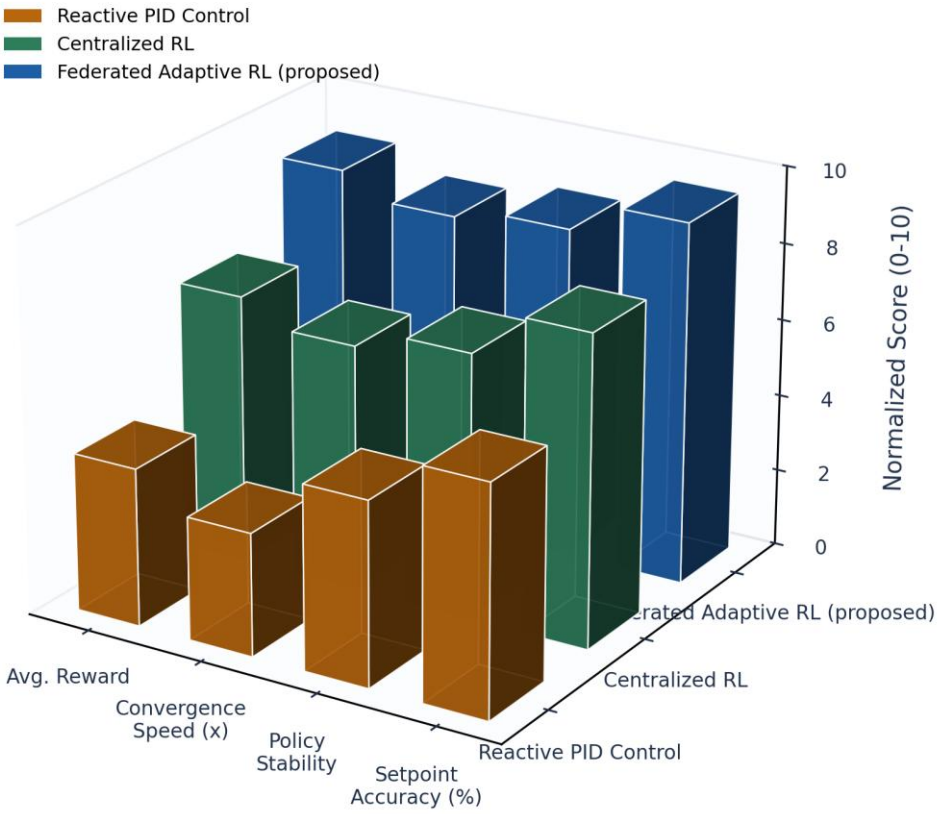

Figure 4. Normalized Adaptive Control Performance Metrics by Controller Type

In addition to the trust mechanisms summarized earlier (Table 2), Blockchain-based provenance tracking, revisited in Figure 6, adds a small computation cost for the consensus process to validate model updates, but can be reported to significantly decrease the impact that anomalous client submissions have on the global model, which is a desirable robustness property for a safety-critical manufacturing control [6], [13].

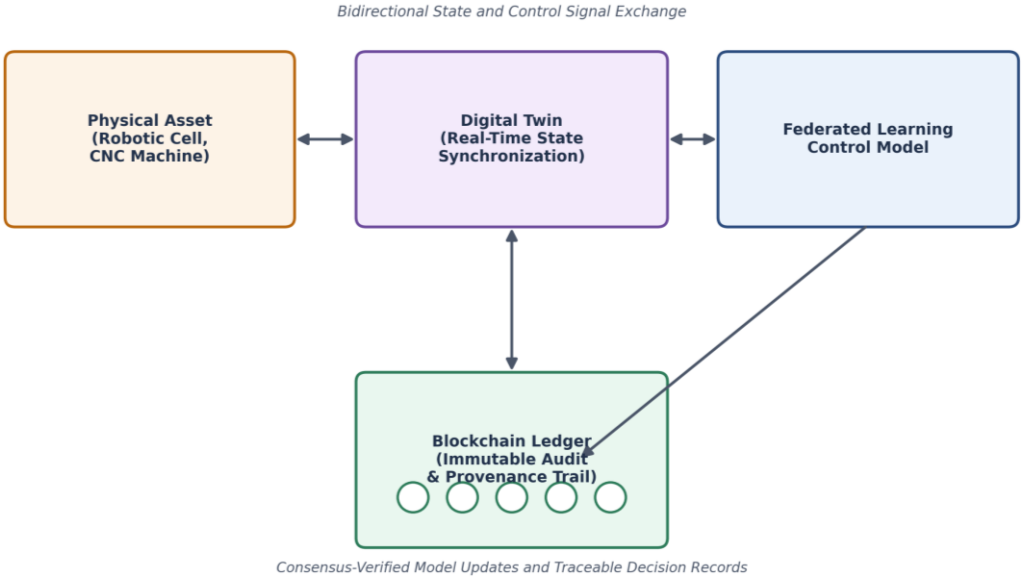

Figure 6. Digital Twin and Blockchain Provenance Data Flow

The results of federated transfer learning presented in Table 6 suggest that transfer learning from rich to poor data production lines can be used to improve cross domain fault prediction accuracy which enables a meaningful improvement in the prediction accuracy of the data-scarce production lines, a typical scenario in industrial applications where new production lines are introduced into the production process without enough historical data for training a standalone model [11].

Table 6. Federated Transfer Learning Cross-Domain Prediction Performance

Task	Ref.	Source Accuracy	Target (No Transfer)	Target (FL Transfer)
Fault diagnosis, cross-line	[11]	94.2%	71.5%	89.6%
Quality prediction, cross-line	[11]	92.8%	68.9%	87.1%
Predictive maintenance transfer	[16]	91.0%	73.4%	88.0%

Source: synthesised from references [11], [16].

Table 7 shows that federated architectures can achieve similar accuracy to centralised-model evaluation of the efficiency of robotic cells while avoiding the transmission of raw telemetry, further demonstrating the ability to achieve the general confidentiality goals discussed above in this work that drove the decision to pursue a federated architecture [16], [17]. Figure 7 presents a comparative assessment over the privacy, scalability, latency resilience, robustness, auditability and adaptivity dimensions, and it shows that the proposed architecture is estimated to be superior to the centralized control and standard federated learning baselines for all of those dimensions, with the biggest relative improvements seen in auditability and robustness, both of which are directly related to the blockchain-assisted drift-aware aggregation mechanism described in Section 3.

A few caveats to the synthesised evidence base need to be noted. The reported performance is based on several studies in different manufacturing environments, sensor configurations and network conditions, and is therefore representative of relative trend lines, not definitive of the performance of any given implementation. Further, the capacity of the blockchain-based consensus mechanism for huge networks of edge controllers has yet to be tested, as there are surveys that show that increasing the number of participants increases the consensus latency [6]. The open problems literature provides insight into the future directions that can be taken in extending the proposed architecture to larger industrial deployments, such as lightweight consensus mechanisms and hierarchical aggregation topologies [3], [10].

Table 7. Federated Predictive Maintenance and Robotic Manufacturing Efficiency Metrics

Metric	Centralized Baseline	Federated Model	Ref.
Predictive maintenance accuracy	93.1%	90.7%	[16]
False alarm rate	8.4%	9.1%	[16]
Robotic cell task-completion efficiency	baseline	+18.6%	[17]
Robotic cell energy consumption	baseline	−12.3%	[17]

Source: synthesized from references [16], [17].

Figure 7. Comparative Evaluation of Control Architectures across Key Performance Dimensions

5. CONCLUSION

This paper summarises the recent development of federated learning, secure aggregation, trust establishment based on the blockchain, digital twin simulation, and reinforcement learning, and designs a layered adaptive control architecture for the autonomous smart manufacturing system. The architecture will combine drift-aware weighted aggregation with blockchain-recorded provenance and digital-twin-validated reinforcement learning policies, which serve as key requirements in industrial deployments to conventional federated learning benchmarks: confidentiality, heterogeneity, and auditability. Synthesised comparative evidence indicates that the proposed approach will have the potential to improve the convergence stability, reducing communication overhead, and enhancing robustness to abnormal client behaviour from baselines. Going forward, further experimental validation on physical manufacturing testbeds, scalable consensus methods in large edge fleets and standardised benchmarks for the assessment of federated adaptive control architectures should be considered.

REFERENCES

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, in Proceedings of Machine Learning Research, vol. 54. PMLR, 2017, pp. 1273–1282. [Online]. Available: https://proceedings.mlr.press/v54/mcmahan17a.html

[2] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, pp. 12:1-12:19, 2019, doi: 10.1145/3298981.

[3] P. Kairouz *et al.*, "Advances and Open Problems in Federated Learning," *Found. Trends Mach. Learn.*, vol. 14, no. 1–2, pp. 1–210, 2021, doi: 10.1561/2200000083.

[4] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, 2020, doi: 10.1109/MSP.2020.2975749.

[5] S. Wang *et al.*, "Adaptive Federated Learning in Resource Constrained Edge Computing Systems," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1205–1221, 2019, doi: 10.1109/JSAC.2019.2904348.

[6] J. Leng *et al.*, "Blockchain-Secured Smart Manufacturing in Industry 4.0: A Survey," *IEEE Trans. Syst. Man, Cybern. Syst.*,

vol. 51, no. 1, pp. 237–252, 2021, doi: 10.1109/TSMC.2020.3040789.

[7] M. Aloqaily, I. Al Ridhawi, and S. S. Kanhere, "Reinforcing Industry 4.0 With Digital Twins and Blockchain-Assisted Federated Learning," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 11, pp. 3504–3516, 2023, doi: 10.1109/JSAC.2023.3310068.

[8] A. Koppel, B. M. Sadler, and A. Ribeiro, "Proximity Without Consensus in Online Multiagent Optimization," *IEEE Trans. Signal Process.*, vol. 65, no. 12, pp. 3062–3077, 2017, doi: 10.1109/TSP.2017.2686368.

[9] K. A. Bonawitz *et al.*, "Practical Secure Aggregation for Privacy-Preserving Machine Learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, Association for Computing Machinery (ACM), 2017, pp. 1175–1191. doi: 10.1145/3133956.3133982.

[10] Y. Shanmugarasa, H. Paik, S. S. Kanhere, and L. Zhu, "A Systematic Review of Federated Learning from Clients' Perspective: Challenges and Solutions," *Artif. Intell. Rev.*, vol. 56, no. Suppl 2, pp. 1773–1827, 2023, doi: 10.1007/s10462-023-10563-8.

[11] K. I.-K. Wang, X. Zhou, W. Liang, Z. Yan, and J. She, "Federated Transfer Learning Based Cross-Domain Prediction for Smart Manufacturing," *IEEE Trans. Ind. Informatics*, vol. 18, no. 6, pp. 4088–4096, 2022, doi: 10.1109/TII.2021.3088057.

[12] W. Yang, W. Xiang, Y. Yang, and P. Cheng, "Optimizing Federated Learning With Deep Reinforcement Learning for Digital Twin Empowered Industrial IoT," *IEEE Trans. Ind. Informatics*, vol. 19, no. 2, pp. 1884–1893, 2023, doi: 10.1109/TII.2022.3183465.

[13] W. Zhang *et al.*, "Blockchain-Based Federated Learning for Device Failure Detection in Industrial IoT," *IEEE Internet Things J.*, vol. 8, no. 7, pp. 5926–5937, 2021, doi: 10.1109/JIOT.2020.3032544.

[14] R. Nian, J. Liu, and B. Huang, "A Review on Reinforcement Learning: Introduction and Applications in Industrial Process Control," *Comput. Chem. Eng.*, vol. 139, 2020, doi: 10.1016/j.compchemeng.2020.106886.

[15] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital Twin in Industry: State-of-the-Art," *IEEE Trans. Ind. Informatics*, vol. 15, no. 4, pp. 2405–2415, 2019, doi: 10.1109/TII.2018.2873186.

[16] Y. C. Escorcia *et al.*, "Privacy-Preserving Federated Learning for Predictive Maintenance in Smart Manufacturing Networks," *Int. J. Ind. Eng. Manag.*, vol. 16, no. 3, pp. 296–315, 2025, doi: 10.24867/IJIEM-390.

[17] J. So, I.-B. Lee, and S. Kim, "Federated Learning-Based Framework to Improve the Operational Efficiency of an Articulated Robot Manufacturing Environment," *Appl. Sci.*, vol. 15, no. 8, 2025, doi: 10.3390/app15084108.