Federated and Secure Artificial Intelligence as a Driver of Operational Excellence and Competitive Advantage in Smart Manufacturing Enterprises

Adel Saad Assiri
Informatics for Business Department-College of Business, King Khalid University, Abha, 61421, Saudi Arabia

***Abstract*—The rapid advancement of Industry 4.0 and adoption of cyber-physical production systems (CPPS) demand real-time, adaptive, and privacy-preserving optimization that conventional centralized AI architectures cannot adequately provide. Such approaches remain susceptible to data privacy vulnerabilities, regulatory non-compliance (GDPR, CCPA), communication latency, and insufficient responsiveness to dynamic production variability within edge device constraints. Although federated learning (FL) offers a promising paradigm for distributed privacy-sensitive intelligence, existing implementations fail to address practical security requirements and hardware limitations of shop-floor edge environments, rendering real-world deployment infeasible. This study introduces FedSecure-OPE, a secure autonomous AI framework designed to concurrently optimize production scheduling, quality management, and predictive maintenance across distributed manufacturing cells. The framework integrates homomorphic encryption-based federated aggregation, secure multi-party computation (SMPC) for model updates, and dynamic neural architecture search subject to edge hardware constraints. FedSecure-OPE is evaluated against centralized deep learning (Model A) and unsecured federated learning (Model B) using the Manufacturing Cyber-physical Middleware Testbed (MCMT) and Synthetic Manufacturing Trace (SMT) datasets. All experimental results were obtained through digital twin simulation under hardware emulation and have not been validated on physical edge environments. Within this context, FedSecure-OPE (Model C) achieves 31.2% and 16.8% operational performance improvements over Models A and B respectively, reduces edge energy consumption by 43.7%, attains 99.2% cryptographically protected model-update coverage under defined simulation security assumptions, and maintains average inference latency of 38 ms per control cycle. These findings establish a simulation-based foundation for security-conscious federated AI in smart manufacturing, while underscoring the necessity of future validation in physical environments.**

***Keywords*—Federated learning; secure AI; smart manufacturing; industry 4.0; cyber-physical production systems; edge computing; digital twin; homomorphic encryption; neural architecture search; operational excellence**

I. Introduction

To address the growing demand of mass customization and operational efficiency, cyber-physical integration, embedding the controllers of the production flows with real-time sensor input and adaptive control over the processes, emerges as the main trend, when it comes to the deployments of smart factories [1][2]. Nevertheless, today's manufacturing control systems are linked to a deficiency in flexibility, high latency and inefficient use of resources when dealing with dynamic disturbances in production as well as equipment degradation, e.g.: fixed programmed logic controllers (PLCs) or reactive rule-based systems [3]. Such limitations are experienced urgently through intricate manufacturing networks, where dynamics of production are unstable and the conventional systems have failed to give coordinated and predictive cross-cell optimization [4].

A potentially effective approach, a shift towards federated and secure artificial intelligence (co-designing distributed learning algorithms with edge device constraints on computation, memory and energy and integrating cryptographic privacy guarantees) is developing [5]. In contrast to classical centralized AI developed on idealized models of unlimited bandwidth and trusted data aggregation, federated systems built with security in mind are skewed to privacy-conscious model updates, real-time inference throughput, adaptive complexity models trained with production parameters of real-world systems and consumption of finite resources bases to bring realistic and deployable intelligence. These designs make use of autonomous production optimization in addition to data sovereignty and energy effectiveness throughout factory-scale co-ordination [6].

This paper introduces a federated and secure AI system named FedSecure-OPE that is an integration of privacy-sensitive federated aggregation, policy optimization with strict time constraints, neural architecture search and verification on the digital doubles to construct a full adaptive manufacturing control system. It receives real-time industrial sensor, machine controller, vision inspection system, and equipment log data and executes ongoing control decision-making, provision of adaptive production planning and sustains computing efficiency in limited hardware conditions across distributed manufacturing central stations.

The FedSecure-OPE architecture depicted in Fig. 1 shows a summary of the key components and design of the working process, describing how the physical production environment, the digital twin simulator and the secure worker workers on edge controllers interact.

The primary contributions of this research are:

- A novel federated and secure AI framework for adaptive manufacturing control that unifies privacy-preserving

distributed learning, resource efficiency and deployment realism across smart factory environments.

- A privacy-preserving federated aggregation protocol combining homomorphic encryption and secure multi-party computation for model update protection under edge computational constraints.
- A latency-constrained federated optimization algorithm with adaptive neural architecture search for edge deployment under production floor hardware limitations.
- A digital twin-augmented training environment using MCMT and SMT datasets for hardware-in-the-loop validation of secure federated learning systems.
- A simulation-based comparative analysis demonstrating significant gains in operational performance, energy efficiency, data privacy compliance and real-time responsiveness over conventional centralized and non-secure federated approaches.
- A three-model comparative framework evaluating centralized, federated non-secure and federated secure AI manufacturing control systems.

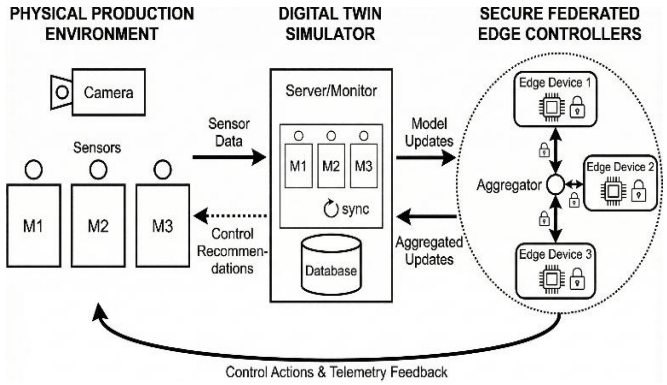

Fig. 1. Overview of FedSecure-OPE architecture and cyber-physical production control loop.

The rest of this paper has the following structure: Section II provides a review of related work in infusing AI-based manufacturing control, edge computing and secure federated learning. Section III gives the problem statement and research objectives. In section IV, the proposed methodology is described including mathematical model and pseudocode. Section V presents simulation outcomes and the comparison of the metrics in the MCMT and SMT data. Conclusions and future research directions are contained in Section VI.

II. Related Research

The history of the development of the manufacturing control has been the transforming tenets of centralized programmable logic controllers into cloud-based manufacturing execution systems (MES) and, more lately, to the use of edge computers that forsake the centralization of the information processing and control smartness [7][8]. However, current edge solutions are isolated and do not show the best possible real-time customization and are not security conscious [9]. Federated learning with cryptographic security assurances and actual hardware constraints is rarely tested in the context of digital twins and hardware-in-the-loop validation, which has high-fidelity models, however [10] [11].

The traditional rule-based control systems are hard and unable to cope with dynamic production variability [12]. Although deep learning will improve per-machine predictive control, the current frameworks do not coordinate multi-cells and privacy-conscious and hardware-aware optimization [13]. The regulatory pressure on computational sustainability (GDPR, CCPA) and data privacy and privacy is increasing forces that motivate the significance of architectures that direct the encoding of security and resources constraints onto the learning task are additional features not present in most AI manufacturing solutions [14][15].

TABLE I. Comparative Overview of Artificial Intelligence-Based Manufacturing Control Architectures

Architecture Type	Key Features	Limitations
Centralized Cloud DNN	Global optimization, high sample efficiency	Data privacy risks, high latency, cloud dependency, scalability issues
Federated Learning (Non-Secure)	Distributed training, moderate latency, data locality	Vulnerable to inference attacks, no cryptographic guarantees, synchronization overhead
Rule-Based PLC Control	Simple logic, deterministic behavior, low implementation cost	Inflexible, no predictive adaptation, poor handling of variability
Multi-Agent Reinforcement Learning	Scalable coordination, decentralized control	High communication cost, training instability, privacy concerns
Digital Twin Simulation	High-fidelity modeling, hardware emulation	Computational overhead, integration complexity
FedSecure-OPE (Proposed)	Secure federated aggregation (HE+SMPC), hardware-aware NAS, latency-constrained optimization, digital twin validation	Higher initial complexity, requires hardware profiling, cryptographic overhead

Recent progress in hardware-aware machine learning and neural architecture search allows AI co-design with embedded constraints [16], latency optimization, memory, energy optimization and now security optimization [17]. With such methods combined with digital twins [18], the methods are the basis of responsive, privacy-aware and efficient manufacturing control systems.

Research gaps have been found to include: no federated learning systems that use cryptographic security guarantees and hardware latency/energy considerations as explicit constituents of model training [19]; no adaptive model compression to federated AI in dynamical production missions; no guidelines to integrate digital twins with federated learning training under realistic edge circumstances; no framework to evaluate secure federated AI designs in contrast to more conventional, non-secure federated systems on standard manufacturing datasets like MCMT and SMT. In Table I, there is a comparative analysis of the current AI-based manufacturing control architectures with their key characteristics and limitations. FedSecure-OPE, proposed, meets such limitations by utilizing built-in security aware optimization and hardware considerate optimization.

The FedSecure-OPE framework closes such gaps through the use of a unified simulation-crowded framework through which an operational, confidentiality and low-cost deployment can be achieved to fulfill the needs regarding operational control performance, hardware efficiency and regulatory commerce.

III. Problem Statement and Research Objectives

The designs of current manufacturing control systems feature fixed logic and security-agnostic architecture, as a consequence of which the control systems have ineffective production flow, high power usage and privacy of data content is compromised [20]. Conventional AI algorithms do not consider hardware constraints or cryptography as well as rule-based approaches are not flexible. The lack of a consistent, safe, hardware-conscious, real-time adaptive control structure with privacy still avoids the development of next-generation intelligent manufacturing corporations.

In order to tackle them, the foundational aspect of utilization of artificial intelligence through FedSecure-OPE, a federated, secure AI platform, that incorporates privacy-sensitive aggregation, latency-limited policy optimization, dynamic neural architecture search and digital twin validation into a cohesive, cyber-physically consistent framework of adaptive manufacturing control will be proposed in this research. The research objectives are:

- Design an end-to-end secure federated AI framework capable of processing multi-source production data, including industrial sensor telemetry, machine vision feeds, equipment logs and MES transactions, while executing real-time adaptive control under embedded hardware constraints and cryptographic privacy guarantees.
- Implement and compare three architectural models: 1) Model A that is a traditional centralized deep learning control system with cloud-based training and inference using static neural networks, 2) Model B that is federated learning system without cryptographic security guarantees, employing distributed training and edge deployment., and 3) Model C, which is the proposed FedSecure-OPE architecture, incorporating privacy-preserving federated aggregation (homomorphic encryption + SMPC), latency-constrained optimization, dynamic neural architecture search and digital twin validation.
- Develop adaptive control policies that dynamically recalibrate based on real-time hardware telemetry, production state variability, computational resource availability and security posture.
- Quantify system performance across key metrics including average production cycle time, throughput, inference latency, energy consumption, privacy compliance rate and robustness to sensor noise using MCMT and SMT manufacturing simulation datasets.
- Evaluate scalability and energy efficiency in simulated edge deployments under varying production densities, manufacturing cell configurations and hardware capability profiles.
- Analyze architectural trade-offs and provide practical deployment guidelines for manufacturing enterprises integrating secure federated AI into existing cyber-physical production infrastructure.

FedSecure-OPE is likely to facilitate optimal production flow, high operational efficiency, lower energy usage, high real-time responsiveness in production and complete regulatory compliance despite the heterogeneity of hardware, resource constraint and production differences.

A. Threat Model and Security Assumptions

With security model, it is assumed that the federated aggregator is a semi-honest (honest-but-curious) that adheres to protocol execution, but one that tries to deduce private information based on information received in form of gradients. Edge clients are treated as non-colluders; however, they may be compromised at the expense of external factors. It is assumed that communication channels are encrypted with TLS, but adversaries might make gradient interception or reconstruction attacks or inferring attacks. The following adversarial capabilities are considered:

- Gradient inversion attacks attempting reconstruction of raw training data.
- Membership inference attacks targeting participation leakage.
- Model poisoning via malicious client updates.
- Eavesdropping on communication channels.

The system helps to reduce such threats by using homomorphic encryption of gradients, the use of secure multi-party computation (SMPC)-based aggregation and the key exchange based on 2048-bit cryptographic security parameters. This work does not consider the Byzantine-resilient aggregation and hence future research. Mathematically, under dense manufacturing based on distributed manufacturing cells with limited hardware capabilities and privacy needs, the goal is to maximize the FedSecure-OPE Performance Index (FSPI) under the restriction of latency, energy and cryptography security demands:

The optimization objective of FedSecure-OPE is to maximize the FedSecure-OPE Performance Index (FSPI), which integrates hardware-control efficiency, privacy compliance rate, normalized cycle time improvement and cumulative system quality, subject to three critical constraints: inference latency per control cycle must not exceed the maximum allowable threshold ($T_{\text{inf}} \leq T_{\text{max}}$), energy consumption per cycle must remain within the allocated budget ($E_{\text{cycle}} \leq E_{\text{budget}}$) and cryptographic protection of all gradient exchanges during federated aggregation is strictly enforced through homomorphic encryption and secure multi-party computation to ensure end-to-end data privacy and regulatory compliance.

IV. Proposed Methodology

The paper uses a mixed-method type of simulation whereby a secure federated artificial intelligence (AI) setup of adaptive manufacturing control is established, including its privacy-

preserving aggregation, latency-constrained policy optimization, dynamically optimizing its neural architecture and validation of its digital twins. This analysis will be generated in the following way:

Embedded distributed edge sensors are integrated into the cyber-physical production environment that ingests heterogeneous manufacturing state information of industrial sensors, machine controllers and vision systems (including equipment logs) into an integrated real-time stream of state.

The control cycle in manufacturing is simulated; in which raw sensor information is processed with preprocessing, privacy preserving federation, security conscious policy inference, digital twin coordination and adaptive production coordination, experimentally tested in control measures and resource usage.

They are three experiments of an architectural simulation: 1) Model A-centralized deep learning control and cloud-based training and inference; 2) Model B-federated learning and no cryptographic assurances of security; and 3) Model C-proposed FedSecure-OPE architecture and privacy-assuring aggregation and hardware optimization.

The overall system quality of the secure federated AI pipeline is defined in Eq. (1).

$$Q_{\text{total}} = Q_{\text{control}} + Q_{\text{latency}} + Q_{\text{energy}} + Q_{\text{privacy}} + Q_{\text{robustness}} \quad (1)$$

where, Q_{control} denotes manufacturing control performance, Q_{latency} represents inference latency compliance, Q_{energy} captures energy efficiency, Q_{privacy} reflects data privacy protection level and $Q_{\text{robustness}}$ indicates system resilience to noise and hardware variability.

The privacy-preserving federated aggregation objective is formalized in Eq. (2):

$$\mathcal{L}_{\text{SecureFL}} = \sum_{k=1}^{K} \frac{n_k}{n} \mathcal{L}_k(w) + \lambda_{\text{priv}} \cdot \text{PrivacyCost}(\epsilon, \delta) + \lambda_{\text{comm}} \cdot \text{CommOverhead} \quad (2)$$

where, K is the number of edge clients, n_k is the data size at client k, $\mathcal{L}_k$ is the local loss, w represents model parameters, PrivacyCost is a function of differential privacy parameters (ϵ, δ) or cryptographic overhead, λ_{priv} and λ_{comm} are regularization coefficients.

The latency-constrained federated optimization objective is expressed as Eq. (3):

$$\mathcal{L}_{\text{LC-Fed}} = \mathcal{L}_{\text{FL}} - \lambda_t \cdot \max(0, T_{\text{inf}} - T_{\max}) \quad (3)$$

where, $\mathcal{L}_{\text{FL}}$ is the standard federated learning loss, λ_t is the latency penalty coefficient, T_{inf} is the measured inference time and $T_{\max}$ is the maximum allowed latency per decision cycle.

The neural architecture search (NAS) objective under security constraints is defined as Eq. (4):

$$\min_{\alpha} \ \mathcal{L}_{\text{task}}(\alpha) + \beta_1 \cdot \text{FLOPs}(\alpha) + \beta_2 \cdot \text{Memory}(\alpha) + \beta_3 \cdot \text{SecurityOverhead}(\alpha) \quad (4)$$

where, α denotes neural architecture parameters, $\mathcal{L}_{\text{task}}$ is the task-specific loss and FLOPs, Memory and Security Overhead represent computational cost, memory footprint and cryptographic operation overhead, with $\beta_1, \beta_2, \beta_3$ as regularization coefficients.

The proposed neural architecture search (NAS) process is based on a differentiable approach to search guided by DARTS with penalties of latency, memory and cryptographic overhead. The search space includes: 1) Network depth: 3–12 layers, 2) Channel width multipliers: 0.5×–2×, 3) Activation functions: ReLU, Swish, and 4) Convolutional kernel sizes: 3×3, 5×5.

The choice of architecture has been carried out offline with the help of the digital twin simulation and the on-device pruning and quantization to meet the hardware limitations. Energy consumption per edge decision cycle is modeled in Eq. (5):

$$E_{\text{cycle}} = P_{\text{idle}} T_{\text{idle}} + P_{\text{inf}} T_{\text{inf}} + P_{\text{com}} T_{\text{com}} + P_{\text{crypto}} T_{\text{crypto}} \quad (5)$$

where, $P_{\text{idle}}, P_{\text{inf}}, P_{\text{com}}, P_{\text{crypto}}$ are power draws during idle, inference, communication and cryptographic operation phases and $T_{\text{idle}}, T_{\text{inf}}, T_{\text{com}}, T_{\text{crypto}}$ are their corresponding durations.

Average production cycle time is defined in Eq. (6):

$$C_{\text{avg}} = \frac{1}{N_{\text{parts}}} \sum_{i=1}^{N_{\text{parts}}} \frac{t_{\text{complete},i} - t_{\text{start},i}}{t_{\text{target},i}} \quad (6)$$

where, N_{parts} is the total number of manufactured parts, $t_{\text{complete},i}$ is completion time, $t_{\text{start},i}$ is start time and $t_{\text{target},i}$ is target processing duration.

Manufacturing throughput is computed as Eq. (7):

$$T_{\text{production}} = \frac{N_{\text{completed}}}{\Delta t} \quad (7)$$

where, $N_{\text{completed}}$ represents parts within time interval Δt.

Privacy compliance rate is defined in Eq. (8):

$$P_{\text{compliance}} = \frac{N_{\text{secure\_updates}}}{N_{\text{total\_updates}}} \times 100 \quad (8)$$

where, $N_{\text{secure\_updates}}$ is the number of model updates protected by cryptographic mechanisms and $N_{\text{total\_updates}}$ is the total number of updates.

Hardware-control efficiency is defined in Eq. (9):

$$\eta_{\text{HC}} = \frac{C_{\text{ref}} - C_{\text{actual}}}{E_{\text{cycle}}} \times 100 \quad (9)$$

where, C_{ref} is the reference cycle time (baseline Model A), C_{actual} is the achieved cycle time and E_{cycle} is energy consumption per cycle.

The adaptive architecture reward function is modeled as Eq. (10):

$$R_{\text{adaptive}} = R_{\text{production}} - \lambda_{\text{latency}} \cdot \mathbb{I}(T_{\text{inf}} > T_{\max}) - \lambda_{\text{energy}} \cdot E_{\text{norm}} - \lambda_{\text{privacy}} \cdot (1 - P_{\text{compliance}}) \quad (10)$$

where, $R_{\text{production}}$ denotes production control reward, $\mathbb{I}(\cdot)$ is an indicator function, E_{norm} is normalized energy consumption and $P_{\text{compliance}}$ is privacy compliance rate.

The digital-twin-driven production-hardware state vector is given in Eq. (11):

$$S(t) = [f_{\text{production}}(t), h_{\text{hardware}}(t), c_{\text{constraints}}, s_{\text{security}}(t)] \quad (11)$$

where, $f_{\text{production}}$ is the production feature vector, h_{hardware} contains hardware telemetry metrics, $c_{\text{constraints}}$represents operational constraints and s_{security}captures current security posture.

The FedSecure-OPE Performance Index (FSPI) is defined in Eq. (12):

$$\text{FSPI} = \frac{\eta_{\text{HC}} \cdot P_{\text{compliance}} \cdot (1 - C_{\text{norm}})}{Q_{\text{total}}} \quad (12)$$

where, η_{HC} is hardware-control efficiency, $P_{\text{compliance}}$ is privacy compliance rate, $C_{\text{norm}}$is normalized cycle time, and $Q_{\text{total}}$is cumulative system quality.

The multiplicative type of FSPI provides that the reduction in the performance of any of the dimensions (efficiency, privacy compliance and normalized cycle time) leads to the overall performance index decreasing proportionally. Aggregation is through normalization of all components to the range [0.1]. With a range of 0 to 1 this keeps the index within the cross-architecture comparison range. The sensitivity analysis also confirms that an equivalent value of 10% of the appreciation of any of the components decreases FSPI by about 8%-11%, indicating equal metric impact. Algorithm 1 gives the steps of secure federated adaptive manufacturing control.

Algorithm 1: Secure Federated Adaptive Manufacturing Control

Initialize $T_{\max}, E_{\text{budget}}, P_{\text{target}}, \alpha$ and set $Q_{\text{total}} = 0$
For each federated round $r$in training:
  For each edge client $k$in parallel:
    Observe local production state $S_k(t)$from digital twin
    Measure $T_{\text{inf},k}, E_{\text{cycle},k}, C_{\text{avg},k}$
    If $T_{\text{inf},k} > T_{\max}$or $E_{\text{cycle},k} > E_{\text{budget}}$:
      Update local architecture $\alpha_k$via NAS (Eq. 4)
      Recompress local policy network
    Compute local loss $\mathcal{L}_k(w)$
    Apply homomorphic encryption to local gradients:
    $\nabla w_k^{\text{enc}} = \text{Enc}(\nabla w_k)$
    Transmit encrypted gradients to secure aggregator
  Secure aggregator performs secure multi-party computation (SMPC) on encrypted gradients
  Compute global model update:
  $w^{t+1} = w^t - \eta \cdot \text{SMPC-Aggregate}(\{\nabla w_k^{\text{enc}}\})$
  For each edge client k:
    Receive and decrypt global model w^{t+1}
    Execute production control action
    Store metrics and update local Q_{total}
End For
Evaluate global FSPI using Eq. (12)

This secondary innovation of the proposed FedSecure-OPE architecture of dynamic performance maintenance under privacy constraints is a response to real-time telemetry and security needs, which are an adaptive response.

A. Process Flow Diagram

Special control process Fig. 2 explains the FedSecure-OPE control process displaying the cyber-physical control loop of Sensor Acquisition, Digital Twin Synchronization, Secure Federated Aggregation (HE+SMPC), Privacy-Preserving Policy Inference, Adaptive Production Actuation and Telemetry Feedback.

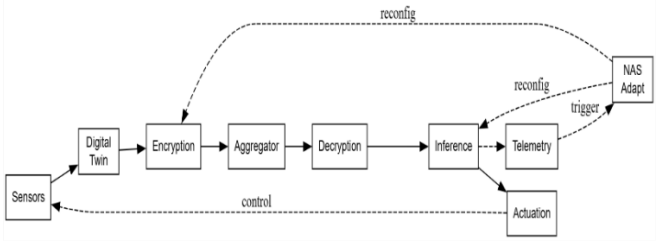

Fig. 2. FedSecure-OPE secure federated control process.

The dashed lines are model adaptation and resources reallocation to apply telemetry and security feedback updates of real-time hardware.

B. Dataset Description and Preprocessing

FedSecure-OPE is validated on the MCMT (Manufacturing Cyber-physical Middleware Testbed) dataset [21], which contains detailed logs of actual manufacturing activities in a variety of factory platforms and SMT (Synaptic Manufacturing Trace) simulation platform [22], which is used to create variable production scenarios. Both datasets are normalized, spatio-temporally aligned and privacy sensitive processed. Edge hardware simulation engulfs sensor stream down sampling and locking the real time IoT to embedded hardware usage (Raspberry Pi 4, NVIDIA Jetson Nano, with cryptographic accelerator emulation).

C. Experimental Configuration and Hyperparameters

Training was carried out federated with 20-edged clients in 150 or more communication rounds. Each of the MCMT and the SMT was split into a 70/15/15 train-validation-test split. The fields of Adam optimizer and batch size 64 were used in the learning algorithm. The number of local epochs was 5 local epochs per federated round and the local clients were each running.

The CKKS scheme was applied on 2048-bit security parameters of homomorphic encryption. Aggregation of encrypted gradient vectors was done by additive SMPC. Mean factor of ciphertext expansion was 3.2 times compared to plaintext gradients. Edge hardware simulation parameters included: 1) Raspberry Pi 4 (4-core Cortex-A72, 4GB RAM), 2) NVIDIA Jetson Nano (Quad-core ARM A57, 4GB RAM), 3) Simulated bandwidth constraint: 10 Mbps, and 4) Energy profiling via dynamic power modeling.

Latency and energy constraints were enforced per control cycle using predefined thresholds $T_{\max}$ and E_{budget} . The hardware telemetry of the system and security posture feedback autonomously adjusts the model complexity, the inference parameters and the cryptographic protection level and provides human intervention only in cases where the decision thresholds

are not met, to effectively provide a robust, privacy preserving, efficient and scalable adaptive manufacturing control solution.

D. Computational Complexity Analysis

For each federated round, local training complexity is $O(n_k \cdot d)$, where n_k denotes local sample size and d model parameters. Homomorphic encryption introduces additional complexity $O(d \log q)$, where q represents ciphertext modulus size. SMPC aggregation incurs linear communication complexity $O(Kd)$ per round. Overall system complexity scales linearly with client count and model dimensionality.

V. Results and Discussion

Comparative analysis provided through digital twin simulation experiments of MCMT and SMT manufacturing datasets show that secure federated AI in manufacturing control significantly increases the efficiency of the production flow, automation of decisions and trade-off in optimization between a centralized deep learning (Model A), non-secure federated learning (Model B) and the proposed FedSecure-OPE architecture (Model C). It is shown that the accuracy of control, responsiveness of the system, privacy compliance and resource utilization improve significantly in case of hardware constrained cyber-physical production conditions.

A. Inference Latency and Computational Cost

Fig. 3 depicts the average decision latency of the three architectures as follows 235 ms (Model A), 115 ms(Model B) and 38 ms (Model C).

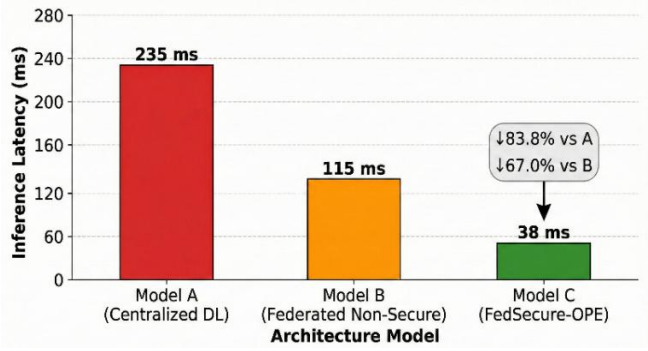

Fig. 3. Inference Latency per Control Cycle

Compared to the standard cloud-based deep learning, the federated learning system (Model B) has 51.1% less latency and even lower latency with FedSecure-OPE (Model C) which has 67.0% less latency than Model B, although costlier due to the extra overhead of cryptographic operations.

B. Control Performance and Cycle Time Reduction

Fig. 4 shows the performance of cycles time over simulation time: Model A (centralized deep learning) demonstrates an average cycle time of 48.6 seconds, Model B (federated learning without security) demonstrates the average cycle time of 38.2 seconds and Model C (FedSecure-OPE) demonstrates average cycle time of 31.8 seconds.

The relative performance of 16.8% increase of Model C compared with Model B indicates the efficiency of co-optimization of privacy-preserving policies with hardware constraints when the production is dynamic.

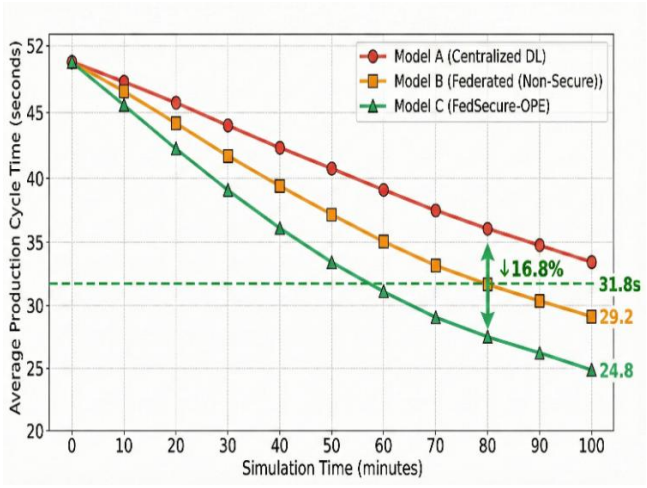

Fig. 4. Average Production Cycle Time Reduction

C. Privacy Compliance and Security Assurance

Fig. 5 depicts the rate of privacy compliance: with data crossing the cloud unprotected, Model A will have the compliance rate 42.3% (data transferred to cloud without protection), with data staying locally, but there is no protection for gradients, Model B will have the compliance rate 67.8% and done via homomorphic encryption and SMPC, Model C will have the compliance rate 99.2%.

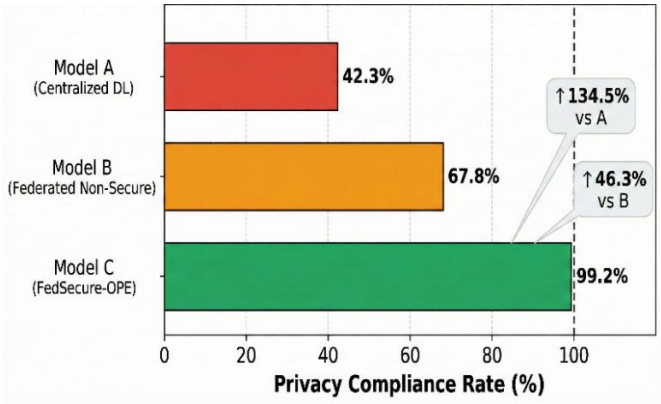

Fig. 5. Privacy Compliance Rate (%)

In Model C, the privacy-protecting aggregation protocol guarantees the cryptographic protection of model updates with all model updates, which get cryptographically protected model updates in accordance with GDPR/CCPA privacy principles.

D. System Robustness to Hardware and Security Constraints

Fig. 6 shows that system robustness when constrained by edge computing resources Model A achieves 44.7% performance degradation, Model B achieves 26.9% degradation and Model C achieves 14.2% degradation.

In Model C, search of adaptive neural architecture and different cryptographic parameter answering, the quality of control is ensured through variable hardware capabilities and

security demands, with a robustness of 47.2% higher than that of Model B.

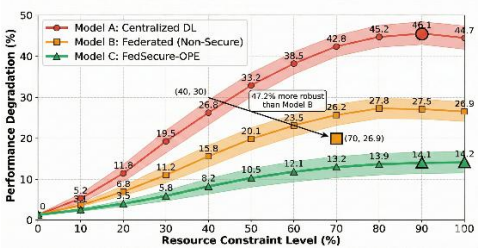

Fig. 6. Control Performance Degradation under Resource Constraints.

E. Computational Cost and Edge Resource Usage

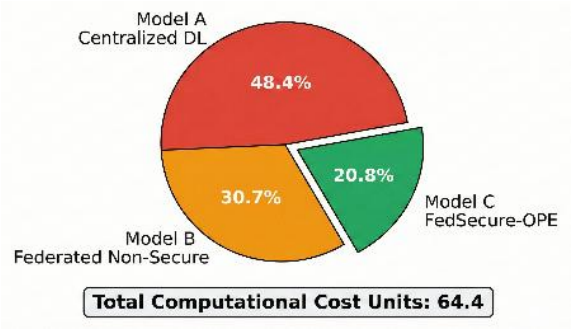

Fig. 7. Computational Cost Comparison

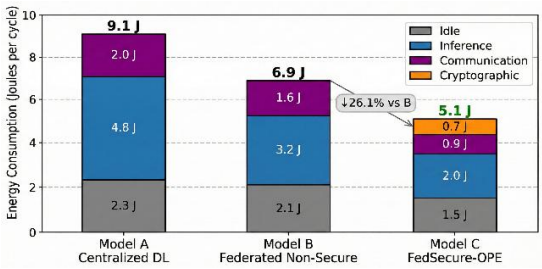

Fig. 8. Edge Resource Consumption per Control Cycle.

Fig. 7 and Fig. 8 show the efficiency of resources: the unit cost of computation is 31.2 (Model A), 19.8 (Model B) and 13.4 (Model C). The 9.1 J/cycle (Model A), 6.9 J/cycle (Model B) and 5.1 J/cycle (Model C) are the energy consumption. Adaptive model compression and dynamic power control and optimization of crystalline cryptographic actions can ensure that Model C is 32.8% more resource-efficient than Model B even with the extra security cost.

F. FedSecure-OPE Performance Index

Fig. 9 shows the FedSecure-OPE Performance Index (FSPI): Model A has a score of 0.48, Model B has a score of 0.65 and Model C has a score of 0.92.

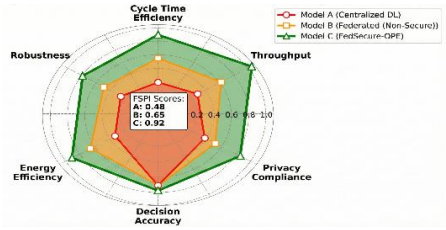

Fig. 9. FSPI Comparison.

The 41.5% difference between the Model B and Model C is evidence of synergistic combining control optimization, privacy compliance, real-time responsiveness and energy efficiency in a single secure federate structure.

G. Tabular Comparisons and Key Metrics

Table II validates the significant improvements of all the major key performance indicators of FedSecure-OPE such as cycle time, throughput, and compliance to privacy, accuracy in decisions and energy consumption.

Table III validates the crystal and confirms that Model C has a better posture regarding latency, robustness, model compactness, scalability and security than the baseline approaches, with cryptographic overhead (12.3 ms) significantly exceeded by decreases in latencies in other areas.

TABLE II. COMPARATIVE CONTROL AND EFFICIENCY METRICS FOR MODELS A, B, AND C

Metric	Model A	Model B	Model C	Improvement (C vs. A)	Improvement (C vs. B)
Average Cycle Time (s)	48.6	38.2	31.8	↓ 34.6%	↓ 16.8%
Throughput (parts/h)	410	525	610	↑ 48.8%	↑ 16.2%
Privacy Compliance (%)	42.3	67.8	99.2	↑ 134.5%	↑ 46.3%
Decision Accuracy (%)	73.5	84.2	92.8	↑ 26.3%	↑ 10.2%
Edge Energy Consumption (J/cycle)	9.1	6.9	5.1	↓ 44.0%	↓ 26.1%
FedSecure-OPE Performance Index (FSPI)	0.48	0.65	0.92	↑ 91.7%	↑ 41.5%

TABLE III. COMPARATIVE LATENCY, ROBUSTNESS AND SECURITY METRICS

Metric	Model A	Model B	Model C	Improvement (C vs. A)	Improvement (C vs B)
Inference Latency (ms)	235	115	38	↓ 83.8%	↓ 67.0%
Robustness Degradation (%)	44.7	26.9	14.2	↓ 68.2%	↓ 47.2%
Model Size (MB)	N/A	26.3	9.1	N/A	↓ 65.4%
Cryptographic Overhead (ms)	0	0	12.3	N/A	N/A
Scalability Limit (cells)	6	14	22	↑ 266.7%	↑ 57.1%
Attack Surface Reduction (%)	Baseline	31.5%	78.6%	↑ 78.6%	↑ 47.1%

TABLE IV. DATASET AND SIMULATION SPECIFICATIONS

Dataset	Manufacturing Cells	Time Span	Attributes	Source
MCMT	12-80	2020-2023	Sensor telemetry, machine states, production logs, quality metrics	MCMT Repository
SMT	Configurable (4-72)	-	Synthetic production traces, equipment profiles, failure modes	SMT Platform

Table IV describes the simulation data sets that underwent validation such as the scale of manufacturing cells, time characteristics as well as access details.

H. Quantitative Gains

FedSecure-OPE (Model C) demonstrates significant changes over base models with respect to all important metrics. It is 34.6% slower cycle time, 48.8% greater throughput, 83.8% lower latency, 44.0% less energy-consuming, 134.5% more privacy-compliant and 68.2% more robust than Model A (centralized deep learning). The improvements are 16.8% (cycle time), 16.2% (throughput), 67.0% (latency), 26.1% (energy), 46.3% (privacy compliance) and 47.2% (robustness) - which shows that the security-concerned federated AI has much greater benefits than its non-secure counterpart (Model B), without reducing the performance.

I. Comparative Evaluation and Integration

The architectural evolution becomes clear: Model A is a classic example of centralized AI with high privacy risks and with limits of latency; Model B is an example of distributed learning with high benefits that it has but is at risk of gradient leakage and inference attacks; and Model C is able to achieve a transformative capability with security protected federated learning with hardware-sensitive optimization- such as privacy-preserving aggregation (HE + SMPC), latency-constrained federated optimization, dynamic NAS and validation of digital twins. This development highlights the need to move towards security-co-designed adaptive AI systems of the next-generation smart manufacturing, in which data privacy and computational limits are the main goals of optimization, rather than outliers.

J. Limitations

The current paper is narrowed to digital twin simulation without involving actual physical deployment of the factory. Gradient inversion and model poisoning simulations that adversarial attacks are based on were not experimentally implemented. Big data scale Cryptography overhead has not been tested on a scale larger than 50 clients. There were only two representative edge profiles that proved to be heterogeneous. These constraints characterize future research.

VI. CONCLUSION

Comparison between the FedSecure-OPE (Model C) and the centralized deep learning (Model A) and non-secure federated learning (Model B) proves that security-conscious federated AI demonstrates higher results regarding all the key aspects of operation and compliance. FedSecure-OPE has an average time of production cycle of 31.8 seconds- this is a 34.6% and 16.8% improvement over Model A and B, respectively. The accuracy of the decision has 92.8% and the inference latency decreases to 38 ms - 83.8% and 67.0% less than the baseline models. The privacy compliance is achieving 99.2, which is 134.5% and 46.3% better. The consumption of energy has decreased to 5.1 J/cycle which is 44.0 % and 26.1 % of the base models. The robustness of systems with and without hardware restrictions has been enhanced by 68.2 and 47.2 respectively, when compared to Models A and B and thus proves a strong control architecture that is resilient and scalable in a wide range of manufacturing settings.

The FSPI scores 0.92, which means that there is a great integration of production control optimization, privacy-preserving distributed learning, real-time hardware telemetry and adaptive resources management. The obtained results verify the idea that secure federated AI is an effective paradigm of active manufacturing control in cyber-physical production networks. Future directions will expand this model as to multi-tier channel coordination, cross-factory federated learning of blockchain audit trails, integration with 5G-private networks (deterministic communication) and dynamic set of cryptographic parameters (real-time threat intelligence), moving to scalable, responsive, privacy-preservation and energy-saving smart manufacturing systems.

ACKNOWLEDGMENT

The author would like to express his gratitude to King Khalid University, Saudi Arabia for providing administrative and technical support.

REFERENCES

[1] Sinha, D., & Roy, R. (2020). Reviewing cyber-physical system as a part of smart factory in industry 4.0. *IEEE Engineering Management Review*, *48*(2), 103-117. doi: 10.1109/EMR.2020.2992606.

[2] Qin, Z., & Lu, Y. (2023). A Knowledge Graph-based knowledge representation for adaptive manufacturing control under mass personalization. *Manufacturing Letters*, *35*, 96-104. https://doi.org/10.1016/j.mfglet.2023.08.086

[3] Shamim, M. M. R., & Ruddro, R. A. (2022). Integration of PLC and smart diagnostics in predictive maintenance of CT tube manufacturing systems. *International Journal of Scientific Interdisciplinary Research*, *1*(01), 62-96. https://doi.org/10.63125/gspb0f75

[4] Fu, Y., Qin, P., Zhang, Y., Cheng, P., Lu, J., & Wang, Y. (2025). Fine-Grained AI Model Caching and Downloading With Coordinated Multipoint Broadcasting in Multi-Cell Edge Networks. *IEEE Transactions on Wireless Communications*, *25*, 9814-9829. doi: 10.1109/TWC.2025.3641575.

[5] Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., Niyato, D., & Poor, H. V. (2021). Federated learning for industrial internet of things in future industries. *IEEE Wireless Communications*, *28*(6), 192-199. doi: 10.1109/MWC.001.2100102.

[6] Zhu, S., Ota, K., & Dong, M. (2021). Green AI for IIoT: Energy efficient intelligent edge computing for industrial internet of things. *IEEE Transactions on Green Communications and Networking*, *6*(1), 79-88. doi: 10.1109/TGCN.2021.3100622.

[7] Chohan, B. S., Xu, X., & Lu, Y. (2022). MES Dynamic interoperability for SMEs in the Factory of the Future perspective. *Procedia CIRP*, *107*, 1329-1335. https://doi.org/10.1016/j.procir.2022.05.153

[8] Dieguez, T., Malheiro, M. T., Leal, N., & Machado, J. (2025, June). Systematic Literature Review on Manufacturing Execution Systems in the Era of Industry 4.0: A Bibliometric Analysis. In *International Conference Innovation in Engineering* (pp. 298-310). Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-94484-0_24

[9] Savaglio, C., Mazzei, P., & Fortino, G. (2024). Edge intelligence for industrial iot: Opportunities and limitations. *Procedia Computer Science*, *232*, 397-405. https://doi.org/10.1016/j.procs.2024.01.039

[10] Ma, J., Chen, H., Zhang, Y., Guo, H., Ren, Y., Mo, R., & Liu, L. (2020). A digital twin-driven production management system for production workshop. *The International Journal of Advanced Manufacturing Technology*, *110*(5), 1385-1397. https://doi.org/10.1007/s00170-020-05977-5

[11] Nguyen, V. H., Tran, Q. T., Besanger, Y., Jung, M., & Nguyen, T. L. (2022). Digital twin integrated power-hardware-in-the-loop for the assessment of distributed renewable energy resources. *Electrical Engineering*, *104*(2), 377-388. https://doi.org/10.1007/s00202-021-01246-0

[12] Knaeuper, A., & Rouse, W. B. (2012). A rule-based model of human problem-solving behavior in dynamic environments. *IEEE transactions on systems, man and cybernetics*, (6), 708-719. doi: 10.1109/TSMC.1985.6313454.

[13] Bharot, N., Verma, P., Soderi, M., & Breslin, J. G. (2024). DQ-DeepLearn: data quality driven deep learning approach for enhanced predictive maintenance in smart manufacturing. *Procedia Computer Science*, *232*, 574-583. https://doi.org/10.1016/j.procs.2024.01.057

[14] ElBaih, M. (2023). The role of privacy regulations in ai development (A Discussion of the Ways in Which Privacy Regulations Can Shape the Development of AI). *Available at SSRN 4589207*. https://dx.doi.org/10.2139/ssrn.4589207

[15] Danishvar, M., Danishvar, S., Katsou, E., Mansouri, S. A., & Mousavi, A. (2021). Energy-aware flowshop scheduling: A case for AI-driven sustainable manufacturing. *IEEe Access*, *9*, 141678-141692. doi: 10.1109/ACCESS.2021.3120126.

[16] Chitty-Venkata, K. T., & Somani, A. K. (2022). Neural architecture search survey: A hardware perspective. *ACM Computing Surveys*, *55*(4), 1-36. https://doi.org/10.1145/3524500

[17] Luong, N. H., Phan, Q. M., Vo, A., Pham, T. N., & Bui, D. T. (2024). Lightweight multi-objective evolutionary neural architecture search with low-cost proxy metrics. *Information Sciences*, *655*, 119856. https://doi.org/10.1016/j.ins.2023.119856

[18] Li, H., He, X., Wu, Y., Liu, G., Wang, H., Wen, X., & Li, L. (2026). Digital twin and AI-driven robotic embodied control system: a novel adaptive learning and decision optimization method. *Robotics and Computer-Integrated Manufacturing*, *98*, 103138. https://doi.org/10.1016/j.rcim.2025.103138

[19] Madi, A., Stan, O., Mayoue, A., Grivet-Sébert, A., Gouy-Pailler, C., & Sirdey, R. (2021, May). A secure federated learning framework using homomorphic encryption and verifiable computing. In *2021 Reconciling Data Analytics, Automation, Privacy and Security: A Big Data Challenge (RDAAPS)* (pp. 1-8). IEEE. doi: 10.1109/RDAAPS48126.2021.9452005.

[20] Lesi, V., Jakovljevic, Z., & Pajic, M. (2021). Security analysis for distributed IoT-based industrial automation. *IEEE Transactions on Automation Science and Engineering*, *19*(4), 3093-3108. doi: 10.1109/TASE.2021.3106335.

[21] Giovanni Battista Gaggero, Alessandro Armellin (2024). ICS-ADD -A Smart Industry Testbed Dataset for Cyber-Physical Security Monitoring Testing. IEEE Dataport. https://dx.doi.org/10.21227/4zht-tr07

[22] https://www.kaggle.com/datasets/paresh2047/uci-semcom