MANTECH Publications

Journal of Converging Technologies in Intelligent
Volume 1, Issue 1, January-April 2025

AI-Enabled Cyber-Physical Systems: Convergence of Smart Sensors, Predictive Analytics, and Autonomous Control — A Structured Review of Architectures, Benchmark Evidence, and Deployment Challenges

Arindam K. Sengupta[1], Sunita R. Dey[2]

ABSTRACT

Cyber-physical systems (CPS) tightly couple computation with physical processes, and the integration of artificial intelligence is transforming them from rule-driven automation into adaptive, self-optimising infrastructure. This review examines the convergence of three technology layers — smart sensors with edge intelligence, predictive analytics for prognostics and forecasting, and autonomous control based on model predictive control (MPC) and reinforcement learning (RL) — and evaluates how well the layers function as one closed loop. A structured, PRISMA-informed review of 32 peer-reviewed and benchmark-defining studies published between 2007 and 2021 was conducted, and a layered reference architecture was derived from it. Published remaining-useful-life (RUL) results on the NASA C-MAPSS benchmark show that deep convolutional models reduced root-mean-square error (RMSE) from 37.56 cycles for a multilayer perceptron to 12.61 cycles on subset FD001 (a 66% reduction), and by 9–22% relative to recurrent models across all four subsets. A qualitative appraisal of five control paradigms indicates that learning-based MPC and shielded safe RL offer the strongest balance of constraint satisfaction, adaptivity, and safety assurance. Persistent obstacles include edge resource limits, distribution shift, adversarial vulnerability, and the absence of certification-grade assurance. A research agenda built on digital twins, federated learning, and verifiable learning-enabled control is proposed.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

KEYWORDS: *Cyber-physical systems, Artificial intelligence, Smart sensors, Edge computing, Predictive maintenance, Remaining useful life, Model predictive control, Reinforcement learning, Digital twin, Safe autonomy*

INTRODUCTION

A cyber-physical system is an engineered system in which computation, communication, and control are integrated so tightly with physical dynamics that neither can be designed in isolation [1], [2]. Power grids, automated production lines, connected vehicles, and medical devices all belong to this class, and each must meet hard timing, safety, and reliability requirements while operating in uncertain environments [3]. Classical CPS rely on first-principles models, fixed thresholds, and hand-tuned controllers. These designs are dependable inside the envelope for which they were tuned, but they degrade when equipment ages, operating regimes shift, or sensor behaviour drifts.

Three developments now allow this limitation to be addressed. First, smart sensors with embedded processing convert raw physical signals into filtered, time-stamped, and often pre-interpreted data at the point of measurement [4]. Second, deep learning has made it feasible to extract degradation signatures and forecast future states from high-dimensional sensor streams [5], [6]. Third, advances in optimisation-based and learning-based control permit controllers to adapt their behaviour, rather than merely track a set-point [7], [8]. When joined, these layers form what the Industry 4.0 literature describes as a cognitive CPS, a system that senses, predicts, decides, and configures itself [9].

The problem addressed here is that these three layers are usually studied in isolation by separate research communities, so the evidence on how they perform as one integrated closed loop is scattered and hard to compare. Designers consequently lack consolidated guidance on which analytical and control techniques are mature enough for safety-relevant deployment, where the layers interact adversely, and what must be solved before certification. This review consolidates the evidence, derives a reference architecture, tabulates published benchmark results, and identifies the engineering obstacles that separate laboratory demonstrations from dependable field systems.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

LITERATURE REVIEW

Smart Sensing and Edge Intelligence

Advanced sensing has progressed from passive transducers to devices that combine MEMS elements, signal conditioning, and microcontroller-class inference [4]. Moving analytics towards the data source reduces latency and uplink traffic, which is the central argument of edge computing [10]. Tiny machine learning (TinyML) pushes this further by running compressed neural networks within kilobytes of memory and milliwatts of power, although benchmarking efforts show wide variation in accuracy, latency, and energy across hardware platforms [11]. For data that are unlabelled, as most industrial streams are, unsupervised anomaly detection remains the first line of defence, and its taxonomy of point, contextual, and collective anomalies still frames practical sensor-health monitoring [12].

Predictive Analytics and Prognostics

Data-driven prognostics estimate the remaining useful life (RUL) of a component from condition-monitoring signals. The NASA C-MAPSS turbofan degradation simulation became the standard benchmark for this task because it provides run-to-failure trajectories under varying operating conditions and fault modes [13]. Babu et al. showed that a convolutional network outperformed multilayer perceptron, support vector, and relevance vector regression baselines on this data [14]. Zheng et al. introduced stacked long short-term memory (LSTM) networks that exploit temporal dependence [15], and Li et al. later reported a deep convolutional network with time-window input that lowered error further on all four subsets [16]. A comprehensive review of deep learning for machine health monitoring concluded that representation learning removes much of the manual feature engineering that limited earlier approaches, but also noted the dependence of such models on large labelled run-to-failure datasets [6]. Because a point RUL estimate is of limited value to a downstream controller, uncertainty quantification matters; Monte Carlo dropout offers a practical approximation to Bayesian inference in deep networks [17].

Autonomous Control and Safe Learning

Model predictive control solves a constrained optimisation over a receding horizon and is the dominant advanced control method in process industries because it handles multivariable interaction and hard constraints explicitly [7]. Its weakness is dependence on model fidelity. Learning-based MPC addresses this by identifying or correcting the prediction model from data

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

while retaining constraint satisfaction guarantees [8], and a deployed example in data-centre cooling reported cooling-energy savings of about 9% over the incumbent controller [18]. Deep reinforcement learning offers a model-free alternative, with value-based agents reaching human-level performance in high-dimensional tasks [19] and actor-critic methods extending the approach to continuous actuation [20]. Unconstrained exploration, however, is unacceptable on physical plants, which motivated a family of safe-RL methods surveyed by Garcia and Fernández [21], including shielding, which monitors proposed actions and overrides those that violate formally specified safety properties [22]. Amodei et al. list reward misspecification, unsafe exploration, and distributional shift among the concrete safety problems of learning agents [23]. Surveys of autonomous-vehicle control confirm that these concerns dominate real deployments [24].

Integration Enablers and Threats

Digital twins, continuously synchronised virtual replicas of physical assets, provide the simulation substrate on which predictive and control policies can be trained and validated before field release [25], [26]. Federated learning allows distributed assets to train a shared model without exchanging raw data, and the FedAvg algorithm reduced the communication rounds needed by one to two orders of magnitude compared with synchronised stochastic gradient descent [27]. On the threat side, CPS inherit attack surfaces from both the cyber and physical domains [28], and deep networks are vulnerable to small adversarial perturbations that can flip a prediction [29].

RESEARCH GAP

The surveyed literature reveals four gaps. First, prognostics and control are typically evaluated on separate benchmarks, so there is no widely accepted protocol for assessing how prognostic error and uncertainty propagate into control performance and safety margins. Second, most benchmark results rely on simulated degradation data with clean sensor channels, whereas field data contain drift, missing values, and sensor faults. Third, comparisons of control paradigms are typically qualitative and application-specific, leaving designers without a common set of criteria covering constraint handling, adaptivity, edge-compute cost, and assurance. Fourth, security, interpretability, and certifiability are treated as add-ons rather than as design requirements of the integrated loop. This review addresses these gaps by proposing a unifying

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

architecture, a common appraisal framework, and a consolidated reading of benchmark evidence.

OBJECTIVES

- To derive a layered reference architecture that integrates smart sensing, predictive analytics, and autonomous control into a single closed-loop AI-enabled CPS.
- To synthesise published quantitative evidence on data-driven RUL prediction using the NASA C-MAPSS benchmark.
- To appraise five control paradigms against six engineering criteria relevant to safety-critical CPS deployment.
- To identify technical, security, and certification obstacles and propose a prioritised research agenda.

METHODOLOGY

Review Design and Search Strategy

This is a structured review with a quantitative synthesis component. The protocol followed the transparency principles of the PRISMA 2020 statement [30] and the guidelines for systematic reviews in engineering by Kitchenham and Charters [31]. Records were searched in IEEE Xplore, ScienceDirect, ACM Digital Library, SpringerLink, and Google Scholar using Boolean combinations of the terms cyber-physical system, smart sensor, edge intelligence, prognostics, remaining useful life, model predictive control, reinforcement learning, digital twin, and safe learning. Selection was purposive rather than exhaustive: priority was given to seminal definitions, high-citation methodological papers, and studies reporting reproducible results on public benchmarks. Because of this design, screening counts are not reported, and the 32 included works should be read as a curated evidence base. The full protocol is summarised in Table 1.

Table 1. Review Protocol and Selection Criteria

Protocol Element	Specification
Databases	IEEE Xplore, ScienceDirect, ACM Digital Library, SpringerLink, Google Scholar
Time window	2007–2021, with seminal earlier work

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

Protocol Element	Specification
Inclusion	Peer-reviewed or benchmark-defining; CPS, prognostics, or control focus; reproducible results
Exclusion	Non-English, opinion pieces, no methodological detail
Included works	32 purposively selected studies
Appraisal	Layer mapping, RMSE transcription, six-criterion scoring

Analytical Framework

Each included study was assigned to one of three technology layers (sensing and edge, predictive analytics, autonomous control) or to a cross-cutting enabler (digital twin, federated learning, security). A layered reference architecture was then constructed by identifying the data and command flows that link the layers (Figure 1). For the quantitative component, RMSE values for RUL estimation were transcribed from the original publications for the C-MAPSS subsets FD001 to FD004, which differ in the number of operating conditions (one versus six) and fault modes (one versus two) [13]. For the control component, five paradigms were scored from 1 (weak) to 5 (strong) on six criteria by the two authors independently, with disagreements resolved by discussion. These scores are an expert appraisal grounded in the cited literature, not measured outcomes. Because original studies used different data preprocessing and training protocols, cross-paper comparisons indicate relative standing rather than controlled experimental differences.

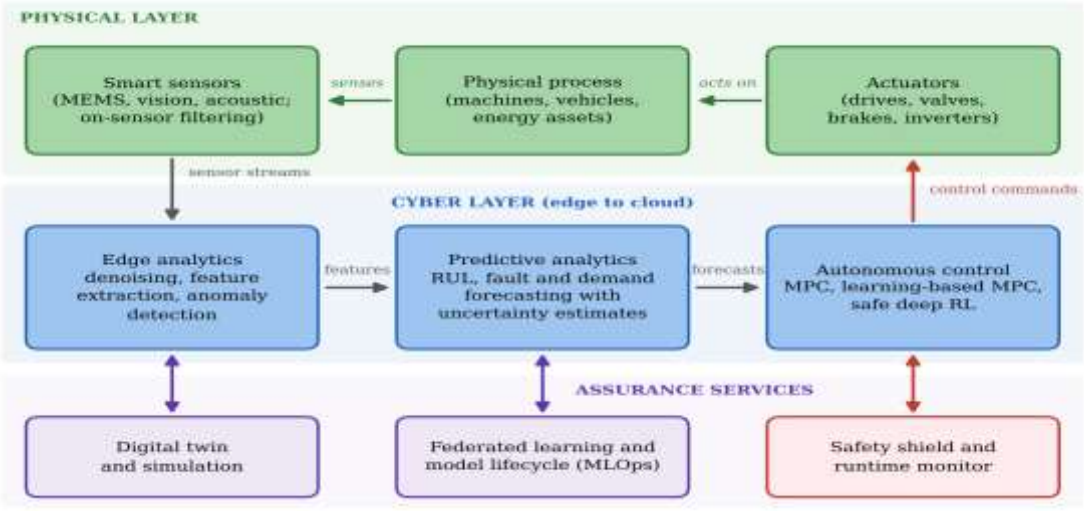


Figure 1. Layered reference architecture of an AI-enabled cyber-physical system showing the physical layer, the cyber layer.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

RESULTS AND FINDINGS

Architectural Synthesis

The reference architecture in Figure 1 organises the evidence into closed-loop data flow. Smart sensors feed edge analytics, which performs denoising, feature extraction, and anomaly detection close to the source. Edge outputs feed predictive models that issue forecasts with uncertainty estimates, and these forecasts parameterise an autonomous controller whose commands act on the physical process. Three assurance services operate alongside the loop: a digital twin for offline policy validation, a federated model lifecycle for continual learning across assets, and a runtime safety shield that can override unsafe commands. Table 2 summarises the evidence attached to each layer.

Table 2. Evidence Synthesis by Technology Layer

Layer	Representative Techniques	Reported Benefit	Key Constraint
Sensing and edge	MEMS, TinyML, anomaly detection [4], [11], [12]	Lower latency and uplink traffic	Memory and energy limits
Predictive analytics	CNN, LSTM, MC dropout [15], [16], [17]	RMSE 12.61 on FD001	Labelled data; calibration
Autonomous control	MPC, learning-based MPC, safe RL [8], [18], [22]	About 9% cooling-energy saving	Compute cost; verification
Enablers	Digital twin, federated learning [25], [27]	Safe training; fleet learning	Fidelity; communication

Prognostic Accuracy on the C-MAPSS Benchmark

Table 3 and Figure 2 consolidate the published RMSE values. On subset FD001, the multilayer perceptron recorded 37.56 cycles, support vector regression 20.96, relevance vector regression 23.80, and the shallow convolutional network of Babu et al. 18.45 [14]. The LSTM of Zheng et al. reached 16.14 [15], and the deep convolutional network of Li et al. reached 12.61 [16], a 66% reduction relative to the perceptron and a 40% reduction relative to support vector regression. Across all four subsets, the deep convolutional model reduced RMSE relative to the LSTM by 22% (FD001), 9% (FD002), 22% (FD003), and 17% (FD004). The smallest gain occurred on FD002, and absolute errors on the six-condition subsets (FD002 and FD004)

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

remained roughly 1.5 to 1.9 times those on the single-condition subsets, showing that multiple operating regimes remain the hardest case for purely data-driven prognostics.

Table 3. Published RMSE (Engine Cycles) on NASA C-MAPSS Subsets

Model	FD001	FD002	FD003	FD004	Source
MLP	37.56	n/r	n/r	n/r	[14]
SVR	20.96	n/r	n/r	n/r	[14]
RVR	23.80	n/r	n/r	n/r	[14]
CNN	18.45	n/r	n/r	n/r	[14]
LSTM	16.14	24.49	16.18	28.17	[15]
Deep CNN	12.61	22.36	12.64	23.31	[16]

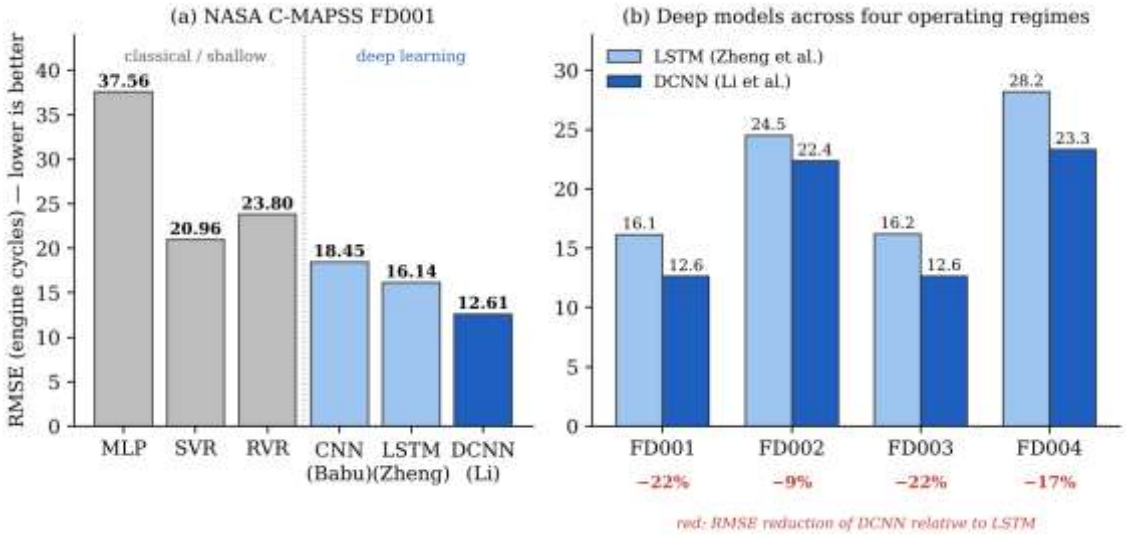


Figure 2. Published RUL prediction error on the NASA C-MAPSS benchmark: (a) RMSE on subset FD001 for classical and deep models; (b) LSTM versus deep convolutional network across subsets FD001–FD004. Red labels show relative RMSE reduction.

Appraisal of Control Paradigms

Figure 3 presents the qualitative appraisal. Conventional PID control scored highest for edge-compute efficiency and interpretability (5 each) but lowest for constraint handling and adaptivity (1 each). Linear MPC scored 5 for constraint handling, yet only 2 for robustness to model error and adaptivity to drift. Learning-based MPC scored 4 on robustness, adaptivity, and safety assurance while retaining a score of 5 for constraint handling, at the cost of a low

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

edge-efficiency score (2) because online optimisation and model updating are computationally demanding. Unconstrained deep RL scored the highest adaptivity (5) but the lowest safety assurance and interpretability (1 each). Shielding raised the safety assurance score of RL from 1 to 4 and constraint handling from 2 to 4, although its interpretability score (2) still depends on how transparent the shield specification is. Mean scores across the six criteria were 2.8 for PID, 3.3 for linear MPC, 3.7 for learning-based MPC, 2.5 for unconstrained RL, and 3.5 for shielded RL.

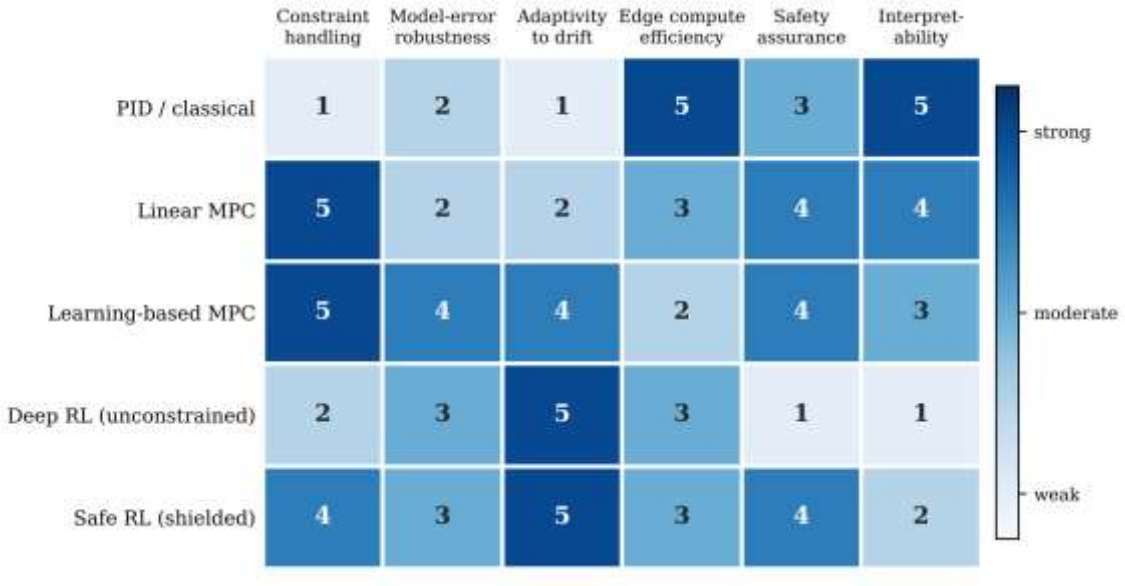


Figure 3. Qualitative appraisal (1 = weak, 5 = strong) of five control paradigms against six criteria relevant to AI-enabled CPS. Scores are the authors' literature-grounded expert judgement.

DISCUSSION

The principal finding is that the technology layers are individually mature enough to be useful but collectively under-integrated. Deep prognostic models now deliver errors that are small relative to typical engine lifetimes, yet they were evaluated as stand-alone regressors. A controller or maintenance scheduler that consumes an RUL estimate needs calibrated uncertainty, and none of the three RUL benchmark studies tabulated here reports it. Approximate Bayesian techniques [17] are a pragmatic route, but their calibration under distribution shift has not been systematically validated for CPS data. The larger errors on six-condition subsets reinforce this point: regime changes of exactly this kind occur in the field.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

The control appraisal explains why hybrid designs dominate current practice. MPC contributes explicit constraints and a principled safety argument, while learning contributes adaptivity. Learning-based MPC [8] and shielded RL [22] both place a verifiable envelope around a learned component, which is consistent with the engineering principle that learning should be permitted to optimise performance but not to define safety. The data-centre result [18] shows that such gains are realisable on real infrastructure rather than only in simulation.

The engineering significance is threefold. First, edge placement of anomaly detection and feature extraction lowers latency and bandwidth demand [10], [11], which is decisive for loops with millisecond deadlines. Second, digital twins [25], [26] allow aggressive policy exploration without risking physical assets, which is the only credible route for training RL agents on expensive equipment. Third, federated learning [27] enables fleet-level learning in sectors where operational data cannot leave the owner's premises. Practical applications include predictive maintenance of rotating machinery and aero-engines, adaptive thermal management in data centres, smart-grid load balancing, and driver-assistance functions in connected vehicles [32], [24]. In Indian manufacturing and power-distribution settings, where legacy equipment and intermittent connectivity are common, edge-first architectures with lightweight models are particularly relevant.

Security remains the least mature dimension. A compromised sensor can feed manipulated data to an otherwise correct predictor, and adversarial perturbations can degrade deep models without being obvious to operators [28], [29]. Anomaly detectors [12] offer partial protection, but a defence-in-depth strategy must also include authenticated sensing and cross-validation against physics-based models in the digital twin.

LIMITATIONS

- The evidence base is a purposive sample of 32 works rather than an exhaustive census, and publication bias towards positive results cannot be excluded.
- RMSE values were taken from different original papers with different preprocessing and training protocols, and C-MAPSS is a simulated benchmark, so the figures indicate relative standing rather than field performance.
- The control-paradigm scores in Figure 3 are expert judgements by two reviewers and have not been validated against a controlled experimental comparison on a common plant.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

FUTURE SCOPE

- Joint benchmarks that evaluate prognostic uncertainty, controller behaviour, and safety violations on the same closed-loop testbed, ideally using hardware-in-the-loop digital twins.
- Robust and drift-aware prognostics validated on real, noisy, and partially labelled fleet data, including self-supervised pre-training to reduce dependence on run-to-failure labels.
- Resource-aware learning-based MPC and shielded RL that fit within microcontroller-class edge budgets, together with formal verification of the shield specifications.
- Standardised assurance cases, adversarial-robustness testing, and explainability tools that make learning-enabled CPS certifiable in regulated domains.

CONCLUSION

AI-enabled CPS combine smart sensing, predictive analytics, and autonomous control into loops that can perceive degradation, anticipate failure, and adapt actuation. Published evidence shows large gains in prognostic accuracy, with deep convolutional models cutting RUL error on C-MAPSS FD001 by 66% relative to a multilayer perceptron, and shows that hybrid controllers combining optimisation or shielding with learning give the best balance of performance and safety. The integrated loop is nevertheless constrained by uncalibrated uncertainty, limited edge resources, distribution shift, and security exposure. Closing these gaps through joint benchmarks, digital-twin validation, federated lifecycle management, and verifiable learning-enabled control is the most direct route from promising demonstrations to certifiable, dependable autonomous infrastructure.

REFERENCES

1. Lee EA. Cyber physical systems: Design challenges. *Proceedings of the 11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC)*. IEEE. 2008; 363–369.
2. Baheti R, Gill H. Cyber-physical systems. *The Impact of Control Technology*. IEEE Control Systems Society. 2011; 12(1): 161–166.
3. Rajkumar R, Lee I, Sha L, Stankovic J. Cyber-physical systems: The next computing revolution. *Proceedings of the 47th Design Automation Conference*. ACM. 2010; 731–736.
4. Hancke GP, Silva BC, Hancke GP Jr. The role of advanced sensing in smart cities. *Sensors.* 2013; 13(1): 393–425.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

5. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature.* 2015; 521: 436–444.
6. Zhao R, Yan R, Chen Z, Mao K, Wang P, Gao RX. Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing.* 2019; 115: 213–237.
7. Mayne DQ. Model predictive control: Recent developments and future promise. *Automatica.* 2014; 50(12): 2967–2986.
8. Hewing L, Wabersich KP, Menner M, Zeilinger MN. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems.* 2020; 3: 269–296.
9. Lee J, Bagheri B, Kao HA. A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. *Manufacturing Letters.* 2015; 3: 18–23.
10. Shi W, Cao J, Zhang Q, Li Y, Xu L. Edge computing: Vision and challenges. *IEEE Internet of Things Journal.* 2016; 3(5): 637–646.
11. Banbury CR, Reddi VJ, Lam M, et al. Benchmarking TinyML systems: Challenges and direction. *arXiv preprint arXiv:2003.04821.* 2020.
12. Chandola V, Banerjee A, Kumar V. Anomaly detection: A survey. *ACM Computing Surveys.* 2009; 41(3): Article 15.
13. Saxena A, Goebel K, Simon D, Eklund N. Damage propagation modeling for aircraft engine run-to-failure simulation. *Proceedings of the International Conference on Prognostics and Health Management (PHM).* IEEE. 2008; 1–9.
14. Babu GS, Zhao P, Li XL. Deep convolutional neural network based regression approach for estimation of remaining useful life. *Database Systems for Advanced Applications (DASFAA), Lecture Notes in Computer Science.* Springer. 2016; 9642: 214–228.
15. Zheng S, Ristovski K, Farahat A, Gupta C. Long short-term memory network for remaining useful life estimation. *Proceedings of the IEEE International Conference on Prognostics and Health Management (ICPHM).* 2017; 88–95.
16. Li X, Ding Q, Sun JQ. Remaining useful life estimation in prognostics using deep convolution neural networks. *Reliability Engineering and System Safety.* 2018; 172: 1–11.
17. Gal Y, Ghahramani Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *Proceedings of the 33rd International Conference on Machine Learning (PMLR).* 2016; 48: 1050–1059.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

18. Lazic N, Boutilier C, Lu T, Wong E, Roy B, Ryu MK, Imwalle G. Data center cooling using model-predictive control. *Advances in Neural Information Processing Systems (NeurIPS).* 2018; 31.
19. Mnih V, Kavukcuoglu K, Silver D, et al. Human-level control through deep reinforcement learning. *Nature.* 2015; 518: 529–533.
20. Lillicrap TP, Hunt JJ, Pritzel A, et al. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971.* 2015.
21. Garcia J, Fernández F. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research.* 2015; 16: 1437–1480.
22. Alshiekh M, Bloem R, Ehlers R, Könighofer B, Niekum S, Topcu U. Safe reinforcement learning via shielding. *Proceedings of the AAAI Conference on Artificial Intelligence.* 2018; 32(1).
23. Amodei D, Olah C, Steinhardt J, Christiano P, Schulman J, Mané D. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565.* 2016.
24. Kuutti S, Bowden R, Jin Y, Barber P, Fallah S. A survey of deep learning applications to autonomous vehicle control. *IEEE Transactions on Intelligent Transportation Systems.* 2021; 22(2): 712–733.
25. Tao F, Zhang H, Liu A, Nee AYC. Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics.* 2019; 15(4): 2405–2415.
26. Rasheed A, San O, Kvamsdal T. Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access.* 2020; 8: 21980–22012.
27. McMahan HB, Moore E, Ramage D, Hampson S, Arcas BA. Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (PMLR).* 2017; 54: 1273–1282.
28. Humayed A, Lin J, Li F, Luo B. Cyber-physical systems security — A survey. *IEEE Internet of Things Journal.* 2017; 4(6): 1802–1831.
29. Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. *Proceedings of the International Conference on Learning Representations (ICLR).* 2015.
30. Page MJ, McKenzie JE, Bossuyt PM, et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ.* 2021; 372: n71.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved

31. Kitchenham B, Charters S. *Guidelines for performing Systematic Literature Reviews in Software Engineering (EBSE Technical Report).* Keele University and University of Durham. 2007.
32. Monostori L, Kádár B, Bauernhansl T, et al. Cyber-physical systems in manufacturing. *CIRP Annals.* 2016; 65(2): 621–641.

***Author for Correspondence:**
Arindam K. Sengupta
E-mail: arindam.sengupta78 @yahoo.com

[1]Assistant Professor, Department of Computer Science and Engineering, Narula Institute of Technology, Agarpara, Kolkata, West Bengal
[2]Associate Professor, Department of Computer Science and Engineering, Narula Institute of Technology, Agarpara, Kolkata, West Bengal

Received Date: January 09, 2025
Accepted Date: January 23, 2025
Published Date: February 06, 2025

Citation: Arindam K. Sengupta, Sunita R. Dey. AI-Enabled Cyber-Physical Systems: Convergence of Smart Sensors, Predictive Analytics, and Autonomous Control — A Structured Review of Architectures, Benchmark Evidence, and Deployment Challenges. Journal of Converging Technologies in Intelligent. 2025; 1(1): 30-43p.

Page 30-43 © MANTECH PUBLICATIONS 2025. All Rights Reserved