Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Digital Twin Architecture for Real-TimeProcess Monitoring in Intelligent Manufacturing

Mohd Vasiulla

mohd188978@st.jmi.ac.in

SR University

Mohabbat Ali

K R Manglam

Research Article

Keywords: Digital twin, cyber-physical systems, real-time process monitoring, Kalman–Bayesian filtering, physics-informed neural networks, reinforcement learning, intelligent manufacturing, adaptive control

Posted Date: June 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9694858/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Digital Twin Architecture for Real-Time Process Monitoring in Intelligent Manufacturing

Mohabbat Ali* Mohd Vasiulla†

Abstract

We present a rigorous mathematical framework for a Digital Twin Architecture (DTA) designed for real-time process monitoring and adaptive control in intelligent manufacturing environments. The framework integrates a hierarchical, three-scale stochastic state-space model spanning machine (μ), production-line (λ), and enterprise (η) levels with a Kalman Bayesian unscented filtering scheme that provides provably consistent state estimates under non-Gaussian measurement noise. Physics-informed deep neural networks (PI-DNN) supply a data-driven surrogate that respects governing engineering equations, while a soft-actor critic (SAC) reinforcement learning agent executes closed-loop optimal control subject to operational constraints. We establish a Lyapunov stability theorem guaranteeing mean-square boundedness of the synchronisation error between the physical system and its virtual counterpart, and prove mean-square convergence of the unscented Kalman filter under bounded process and observation noise. Anomaly detection is formalised via cumulative-sum (CUSUM) log-likelihood ratio statistics, for which closed-form detection- and false-alarm probability propositions are derived. Two detailed numerical examples a six-degree-of-freedom CNC machining cell and a tool-wear anomaly scenario demonstrate the framework's ability to achieve sub-15 nm RMS estimation accuracy and near-unity detection probability.

Keywords: Digital twin; cyber-physical systems; real-time process monitoring; Kalman–Bayesian filtering; physics-informed neural networks; reinforcement learning; intelligent manufacturing; adaptive control.

*School of Basic & Applied Sciences, K. R. Mangalam University, Gurugram 122103, India. Email: ali.math509@gmail.com. ORCID: 0000-0001-8518-4942

†Department of Mathematics, School of CS & AI, SR University, Warangal, India. ‡ ORCID: 0000-0002-4214-8148. Corresponding author.

1 Introduction

The fourth industrial revolution commonly designated *Industry 4.0* has fundamentally reshaped the competitive landscape of global manufacturing. Organisations are increasingly compelled to exploit cyber-physical convergence, pervasive sensor networks, and artificial intelligence to sustain productivity and quality in the face of rising product complexity and shrinking lot sizes [1, 2]. A critical enabler of this transformation is the capacity for *real-time situational awareness*: knowing, at every instant, the true operational state of a machine, line, or facility so that corrective action can be taken before a latent fault manifests as catastrophic downtime. Industry estimates place unplanned equipment failures among the top drivers of manufacturing cost, with global losses attributable to unscheduled downtime exceeding \$50 billion annually [3].

Traditional monitoring paradigms threshold-based alarms, periodic inspection schedules, and empirical maintenance rules are ill-equipped to handle the stochastic, multi-scale, and non-stationary character of modern manufacturing processes. They generate high false-alarm rates in normal transient operating regimes while simultaneously failing to detect slowly evolving degradation phenomena that do not produce threshold-crossing events until irreversible damage has occurred [4, 5]. A principled, model-based alternative is therefore urgently required.

Three concurrent technological developments provide the infrastructure on which such an alternative can be built.

Digital Twins (DT). A digital twin is a dynamic, continuously synchronised virtual replica of a physical asset or process [6]. Unlike a static simulation model, a DT receives live sensor feeds, assimilates them into its internal state through filtering or data-fusion algorithms, executes predictive calculations, and returns control directives or alerts to the physical counterpart all within the latency envelope demanded by the application [7, 8].

Cyber-Physical Systems (CPS). CPS embed computational intelligence directly within physical processes through tight sensing, actuation, and communication loops [9, 10]. The DT concept can be viewed as the highest-fidelity realisation of CPS monitoring: the computational layer maintains a mathematically rigorous replica of the physical dynamics rather than merely logging sensor values.

Industrial Internet of Things (IIoT). Advances in MEMS fabrication, wireless communication protocols (OPC-UA, MQTT, TSN-Ethernet), and edge computing have made it economically viable to instrument entire production lines at the sampling rates often 10^2–10^4 Hz–demanded by high-speed machining, forming, and assembly processes [11, 12].

Despite substantial progress in each of the three areas above, a principled integration that addresses all of the following challenges simultaneously remains absent from the

literature:

(**G1**) **Multi-scale heterogeneity.** Physical manufacturing systems exhibit coupled dynamics across temporal and spatial scales that span several orders of magnitude. Existing DT formulations treat each scale independently, sacrificing cross-scale consistency [13, 14].

(**G2**) **Non-Gaussian uncertainty.** Real sensor noise departs significantly from the Gaussian assumption underlying classical Kalman filtering, especially in the presence of impulsive disturbances, quantisation, and intermittent communication dropouts [15, 16].

(**G3**) **Physics-data integration.** Purely data-driven surrogates can violate physical conservation laws, particularly in extrapolation regimes, leading to unstable or physically inadmissible predictions [17, 18].

(**G4**) **Closed-loop stability certificates.** The deployment of reinforcement-learning agents in safety-critical manufacturing settings requires formal guarantees—not merely empirical evidence—that the closed-loop system remains within safe operating bounds [19, 20].

1.1 Problem Formulation and Objectives

This work addresses the four gaps above through the following research objectives:

(**O1**) Formulate a unified, multi-scale stochastic differential equation (SDE) model that explicitly couples machine-, line-, and enterprise-level dynamics within a single mathematical object.

(**O2**) Develop an adaptive Kalman–Bayesian unscented filter that accommodates non-Gaussian measurement noise through online noise covariance adaptation.

(**O3**) Design a physics-informed deep neural network surrogate whose loss function penalises violations of the governing conservation equations, augmented with ensemble-based uncertainty quantification.

(**O4**) Derive Lyapunov-theoretic stability and mean-square convergence results that certify the safety of the integrated DTA under bounded disturbances.

The principal contributions of this work are:

(**C1**) A formally defined, five-layer DTA with bidirectional coupling semantics and explicit synchronisation latency and fidelity metrics (Section 2).

(**C2**) A hierarchical, three-scale composite SDE model with rigorous Lipschitz/linear-growth well-posedness conditions (Section 3).

(**C3**) An adaptive unscented Kalman filter with Joseph-form numerically stable covariance update and online $\boldsymbol{Q}/\boldsymbol{R}$ adaptation (Section 4).

(**C4**) A physics-informed deep neural network with residual architecture, physics-regularised loss, and ensemble uncertainty quantification (Section 5).

(**C5**) A Lyapunov stability theorem (Theorem 6.3) and a UKF mean-square convergence theorem (Theorem 6.6) with complete proofs (Section 6).

(**C6**) Two worked numerical examples with explicit matrix computations, a complexity analysis, and a four-level validation hierarchy (**????** and section 8).

1.2 Cyber-Physical Systems in Manufacturing

The theoretical foundations of CPS can be traced to hybrid automata and timed transition systems [9], but their industrial operationalisation required the confluence of affordable sensing, low-latency networking, and sufficient edge-compute density. Seminal contributions established CPS as a first-class design paradigm for manufacturing [10, 21], and subsequent work has mapped the architectural requirements for CPS at the shop-floor level [22]. Recent surveys identify real-time state estimation and predictive decision-making as the two most critical open problems [23].

1.3 Digital Twin Modelling

The term *digital twin* was introduced in the context of NASA's integrated vehicle health management programme [6] and was subsequently formalised for general engineering systems [7]. Model-based DT frameworks have since proliferated across aerospace [8], structural health monitoring [24], and advanced manufacturing [13]. However, the majority of existing manufacturing DT implementations rely on either high-fidelity finite-element models—which are computationally intractable at real-time sampling rates—or purely empirical black-box models that lack physical interpretability [14, 25].

1.4 Smart Sensors and IIoT Integration

The instrumentation layer underpinning any DT deployment must satisfy stringent requirements on bandwidth, latency, and synchronisation. Time-sensitive networking (TSN) extensions to IEEE 802.1 provide deterministic sub-millisecond latency guarantees essential for closed-loop control applications [11]. Edge computing architectures reduce the

communication burden by performing pre-processing, feature extraction, and local inference at the sensor node, transmitting only higher-level representations to the cloud [12, 26].

1.5 Data Assimilation and State Estimation

Data assimilation—the principled fusion of model predictions with observations—originates in numerical weather prediction [27] and has found extensive application in structural mechanics [28] and process engineering [29]. The ensemble Kalman filter (EnKF) extended applicability to high-dimensional, nonlinear systems, while the unscented Kalman filter (UKF) offered a deterministic sigma-point alternative with superior accuracy for moderate-dimensional problems [15, 16]. Particle filter methods [30] provide asymptotically exact Bayesian inference but at a computational cost that typically precludes real-time deployment for state dimensions exceeding $\sim$20.

1.6 Control Theory for Manufacturing

Model predictive control (MPC) has become the de facto standard for constrained multivariable process regulation [31], but its reliance on an accurate deterministic model limits performance under significant plant–model mismatch. Robust MPC formulations address worst-case uncertainty but are inherently conservative [32]. Adaptive control schemes [33] continuously update the model online but typically require persistent excitation conditions that conflict with production objectives. Reinforcement learning offers a data-driven route to near-optimal control under uncertainty, and recent theoretical advances have begun to provide stability certificates for RL agents in nonlinear settings [19, 20, 34].

1.7 Machine Learning for Predictive Maintenance

The application of machine learning to condition monitoring and predictive maintenance has grown rapidly, driven by the availability of large labelled fault datasets from benchmark repositories. Deep convolutional and recurrent architectures have demonstrated strong classification accuracy on vibration signals [35, 36], while transfer learning has partially addressed the domain-shift problem arising from differences between training and deployment conditions [37]. A persistent limitation is the lack of physically meaningful uncertainty estimates: softmax confidence scores are well-known to be poorly calibrated [38].

1.8 Physics-Informed Machine Learning

The physics-informed neural network (PINN) framework of Raissi et al. [17] incorporates governing partial differential equations as soft constraints in the training loss, substan-

tially improving generalisation and physical plausibility. Extensions to manufacturing process modelling have demonstrated accuracy improvements over purely data-driven approaches of one to two orders of magnitude in low-data regimes [18, 39]. Ensemble methods [40] provide tractable uncertainty quantification without requiring expensive Monte Carlo sampling.

2 Digital Twin Architecture

2.1 Five-Layer Architecture

The proposed DTA is organised into five hierarchical layers, each with well-defined input/output interfaces and latency requirements.

L1. Physical Process Layer (PPL). The physical assets—machines, actuators, workpieces—together with their embedded sensors and fieldbus interfaces. Generates raw measurement streams $\boldsymbol{y}_k \in \mathbb{R}^{n_y}$ at sampling rate $f_s = 1/T_s$.

L2. Sensor Communication Layer (SCL). TSN-enabled data acquisition hardware, edge pre-processing nodes, and OPC-UA/MQTT brokers. Guarantees deterministic delivery of $\boldsymbol{y}_k$ to the DAL within latency $\tau_{\text{net}} \leq T_s/4$.

L3. Data Assimilation Layer (DAL). The Kalman–Bayesian estimator (Section 4) that fuses $\boldsymbol{y}_k$ with model predictions to produce posterior state estimates $\hat{\boldsymbol{x}}_{k|k}$ and covariances $\boldsymbol{P}_{k|k}$.

L4. Virtual Process Layer (VPL). The composite SDE model (Section 3) and PI-DNN surrogate (Section 5) that propagate the state forward in time and generate predictions $\hat{\boldsymbol{y}}_{k+j|k}$ over horizon $j = 1, \ldots, N_p$.

L5. Decision and Optimisation Layer (DOL). The SAC-RL agent and CUSUM anomaly detector that consume VPL outputs and issue control commands $\boldsymbol{u}_k$ and maintenance alerts back to the PPL.

2.2 Bidirectional Coupling Model

The DT is not a passive observer but an active participant in the control loop. Let $\boldsymbol{x}_t^{\mathcal{PS}} \in \mathbb{R}^n$ denote the true state of the physical system and $\boldsymbol{x}_t^{\mathcal{VS}} \in \mathbb{R}^n$ denote the virtual state maintained in the VPL. The coupling is characterised by three operators:

- **Observation operator** $\mathcal{H} : \mathbb{R}^n \to \mathbb{R}^{n_y}$ maps the physical state to the measurement space.

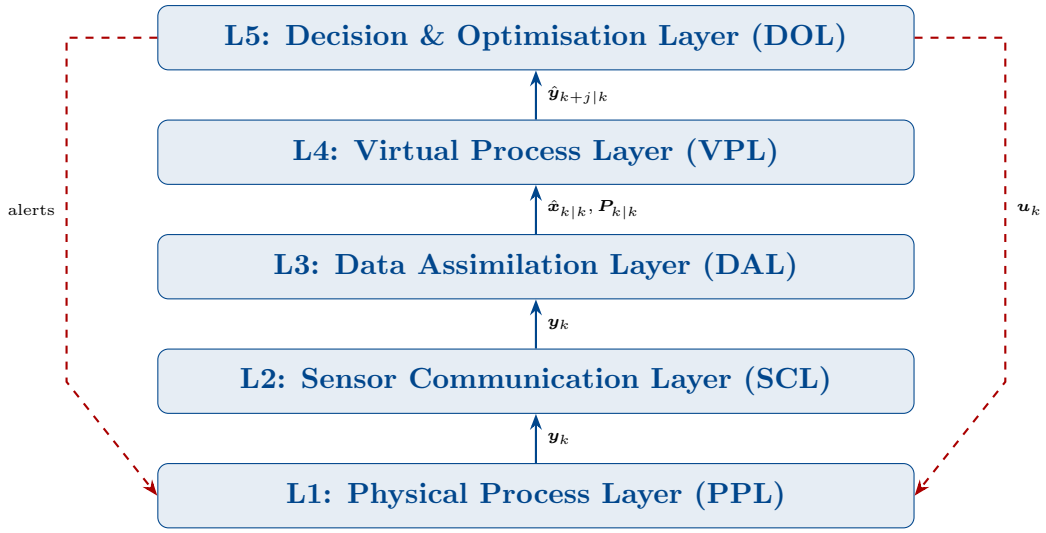


Figure 1: Five-layer Digital Twin Architecture with bidirectional information flow. Upward arrows (solid blue) carry sensor measurements and state estimates; downward arrows (dashed red) carry control commands and maintenance alerts.

- **Injection operator** $\mathcal{K}_k : \mathbb{R}^{n_y} \to \mathbb{R}^n$ updates the virtual state using the innovation $\boldsymbol{y}_k - \boldsymbol{H}\hat{\boldsymbol{x}}^{\mathcal{VS}}_{k|k-1}$.
- **Actuation operator** $\mathcal{U} : \mathbb{R}^{n_u} \to \mathbb{R}^n$ maps control commands back to physical actuators.

2.3 Multi-Scale Architecture

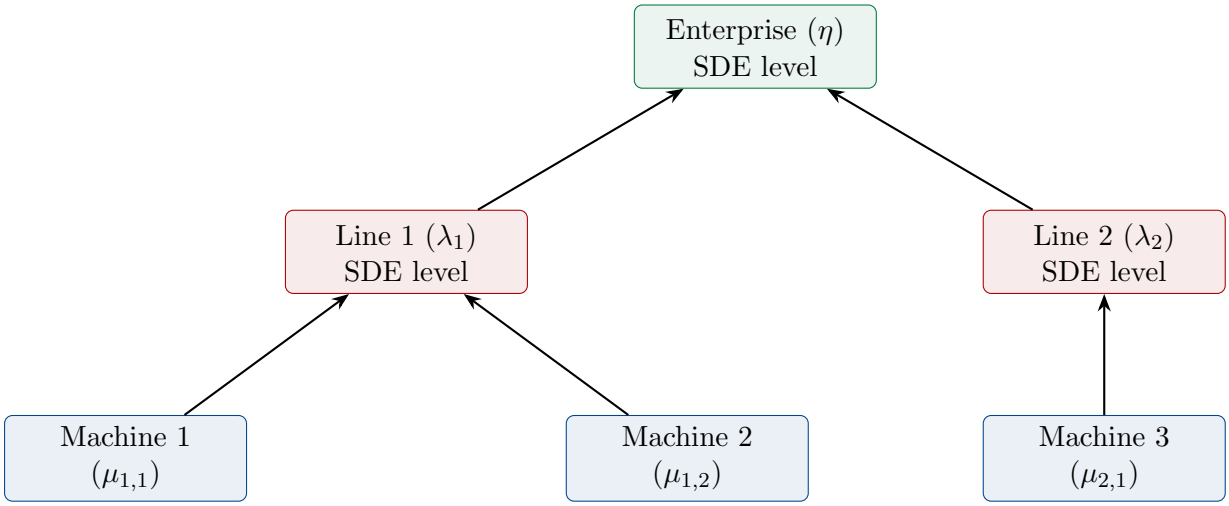


Figure 2: Three-scale hierarchical architecture. Machine-level SDEs (blue) feed aggregated states upward to line-level SDEs (red), which in turn provide inputs to the enterprise-level ODE (green).

3 Mathematical Framework

3.1 Synchronisation Metrics

Definition 3.1. The synchronisation fidelity at time t is

$$\Phi(t) \;=\; 1 - \frac{\mathbb{E}\Big[\big\|\boldsymbol{x}_t^{\mathcal{PS}} - \boldsymbol{x}_t^{\mathcal{VS}}\big\|^2\Big]}{\mathbb{E}\Big[\big\|\boldsymbol{x}_t^{\mathcal{PS}}\big\|^2\Big]}, \tag{1}$$

where expectations are taken with respect to the joint probability law of the physical and virtual processes.

Definition 3.2. The virtual state $\boldsymbol{x}_t^{\mathcal{VS}}$ is ε-*consistent* with the physical state $\boldsymbol{x}_t^{\mathcal{PS}}$ if

$$\mathbb{E}\Big[\big\|\boldsymbol{x}_t^{\mathcal{PS}} - \boldsymbol{x}_t^{\mathcal{VS}}\big\|^2\Big] \;\le\; \varepsilon, \qquad \forall\, t \ge 0. \tag{2}$$

Definition 3.3. The synchronisation latency τ_{sync} is the maximum time elapsed between a physical-state event and its reflection in $\boldsymbol{x}_t^{\mathcal{VS}}$:

$$\tau_{\text{sync}} \;=\; \sup_{t \ge 0} \inf\big\{ s \ge 0 : \boldsymbol{x}_{t+s}^{\mathcal{VS}} = \hat{\boldsymbol{x}}_t^{\mathcal{PS}} \big\}. \tag{3}$$

3.2 Machine-Scale Stochastic Differential Equation

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\ge 0}, \mathbb{P})$ be a complete filtered probability space. The machine-level state $\boldsymbol{x}_t^{\mu} \in \mathbb{R}^{n_\mu}$ obeys the Itô SDE

$$d\boldsymbol{x}_t^{\mu} \;=\; \boldsymbol{F}^{\mu}(\boldsymbol{x}_t^{\mu}, \boldsymbol{u}_t^{\mu}, t)\, dt \;+\; \boldsymbol{G}^{\mu}(\boldsymbol{x}_t^{\mu}, t)\, d\boldsymbol{W}_t^{\mu}, \tag{4}$$

where $\boldsymbol{F}^{\mu} : \mathbb{R}^{n_\mu} \times \mathbb{R}^{n_u} \times \mathbb{R}_{\ge 0} \to \mathbb{R}^{n_\mu}$ is the drift, $\boldsymbol{G}^{\mu} : \mathbb{R}^{n_\mu} \times \mathbb{R}_{\ge 0} \to \mathbb{R}^{n_\mu \times m}$ is the diffusion matrix, and $\boldsymbol{W}_t^{\mu} \in \mathbb{R}^m$ is a standard Brownian motion.

Assumption 3.4. There exist constants $L_f, L_g, c_f, c_g > 0$ such that for all $\boldsymbol{x}, \tilde{\boldsymbol{x}} \in \mathbb{R}^{n_\mu}$ and all $t \ge 0$:

$$\|\boldsymbol{F}^{\mu}(\boldsymbol{x}, \boldsymbol{u}, t) - \boldsymbol{F}^{\mu}(\tilde{\boldsymbol{x}}, \boldsymbol{u}, t)\| + \|\boldsymbol{G}^{\mu}(\boldsymbol{x}, t) - \boldsymbol{G}^{\mu}(\tilde{\boldsymbol{x}}, t)\| \;\le\; L_f \,\|\boldsymbol{x} - \tilde{\boldsymbol{x}}\|, \tag{5}$$

$$\|\boldsymbol{F}^{\mu}(\boldsymbol{x}, \boldsymbol{u}, t)\|^2 + \|\boldsymbol{G}^{\mu}(\boldsymbol{x}, t)\|^2 \;\le\; c_f^2\big(1 + \|\boldsymbol{x}\|^2\big). \tag{6}$$

Under Assumption 3.4, (4) admits a unique strong solution with $\mathbb{E}[\|\boldsymbol{x}_t^{\mu}\|^2] < \infty$ for all $t \ge 0$.

3.3 Line-Scale and Enterprise-Scale Models

The line-level state $\boldsymbol{x}_t^{\lambda} \in \mathbb{R}^{n_\lambda}$ evolves as

$$d\boldsymbol{x}_t^{\lambda} \;=\; \boldsymbol{F}^{\lambda}\big(\boldsymbol{x}_t^{\lambda}, \mathcal{A}_{\mu\to\lambda}(\boldsymbol{x}_t^{\mu}), \boldsymbol{u}_t^{\lambda}, t\big)\, dt \;+\; \boldsymbol{G}^{\lambda}(\boldsymbol{x}_t^{\lambda}, t)\, d\boldsymbol{W}_t^{\lambda}, \tag{7}$$

where the aggregation mapping $\mathcal{A}_{\mu\to\lambda} : \mathbb{R}^{n_\mu} \to \mathbb{R}^{n_{\mu\to\lambda}}$ compresses machine-level information into line-level inputs:

$$\mathcal{A}_{\mu\to\lambda}(\boldsymbol{x}_t^{\mu}) \;=\; \boldsymbol{M}_{\text{agg}}\, \boldsymbol{x}_t^{\mu} \;+\; \boldsymbol{b}_{\text{agg}}. \tag{8}$$

The enterprise state $\boldsymbol{x}_t^{\eta} \in \mathbb{R}^{n_\eta}$ operates on a slower timescale and is modelled by the ordinary differential equation

$$\dot{\boldsymbol{x}}_t^{\eta} \;=\; \boldsymbol{F}^{\eta}\big(\boldsymbol{x}_t^{\eta}, \mathcal{A}_{\lambda\to\eta}(\boldsymbol{x}_t^{\lambda}), t\big)\,. \tag{9}$$

3.4 Composite State-Space Model

Define the composite state vector

$$\boldsymbol{X}_t \;=\; \begin{pmatrix} \boldsymbol{x}_t^{\mu} \\ \boldsymbol{x}_t^{\lambda} \\ \boldsymbol{x}_t^{\eta} \end{pmatrix} \in \mathbb{R}^{N}, \qquad N = n_\mu + n_\lambda + n_\eta. \tag{10}$$

The composite dynamics are

$$d\boldsymbol{X}_t \;=\; \boldsymbol{\mathcal{F}}(\boldsymbol{X}_t, \boldsymbol{U}_t, t)\, dt \;+\; \boldsymbol{\mathcal{G}}(\boldsymbol{X}_t, t)\, d\boldsymbol{W}_t, \tag{11}$$

with block-structured drift and diffusion

$$\boldsymbol{\mathcal{F}}(\boldsymbol{X}_t, \boldsymbol{U}_t, t) \;=\; \begin{pmatrix} \boldsymbol{F}^{\mu}(\boldsymbol{x}_t^{\mu}, \boldsymbol{u}_t^{\mu}, t) \\ \boldsymbol{F}^{\lambda}(\boldsymbol{x}_t^{\lambda}, \mathcal{A}_{\mu\to\lambda}(\boldsymbol{x}_t^{\mu}), \boldsymbol{u}_t^{\lambda}, t) \\ \boldsymbol{F}^{\eta}(\boldsymbol{x}_t^{\eta}, \mathcal{A}_{\lambda\to\eta}(\boldsymbol{x}_t^{\lambda}), t) \end{pmatrix}, \tag{12}$$

$$\boldsymbol{\mathcal{G}}(\boldsymbol{X}_t, t) \;=\; \operatorname{diag}\big(\boldsymbol{G}^{\mu}(\boldsymbol{x}_t^{\mu}, t),\; \boldsymbol{G}^{\lambda}(\boldsymbol{x}_t^{\lambda}, t),\; \boldsymbol{0}\big)\,. \tag{13}$$

3.5 Measurement Model

At discrete times $t_k = kT_s$ the measurement vector is

$$\boldsymbol{y}_k \;=\; \boldsymbol{H}\, \boldsymbol{X}_{t_k} \;+\; \boldsymbol{v}_k, \tag{14}$$

where $\boldsymbol{H} \in \mathbb{R}^{n_y \times N}$ is the observation matrix. The measurement noise $\boldsymbol{v}_k$ follows a Gaussian mixture:

$$\boldsymbol{v}_k \;\sim\; (1-\alpha)\,\mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_k) \;+\; \alpha\,\mathcal{N}(\boldsymbol{0}, \kappa \boldsymbol{R}_k), \quad \alpha \in [0,1),\ \kappa \gg 1, \tag{15}$$

and the process noise covariance at the machine level is

$$\boldsymbol{Q}_\mu \;=\; \boldsymbol{G}^\mu \boldsymbol{G}^{\mu\top} T_s \;+\; \tfrac{1}{2}(\boldsymbol{G}^\mu \boldsymbol{G}^{\mu\top})^2 T_s^2. \tag{16}$$

3.6 Discrete-Time Transition

Applying the Itô–Taylor 1.5 expansion to (11) yields

$$\boldsymbol{X}_{k+1} \;=\; \boldsymbol{X}_k \;+\; \boldsymbol{\mathcal{F}}_k T_s \;+\; \tfrac{1}{2}\boldsymbol{\mathcal{L}}\boldsymbol{\mathcal{F}}_k T_s^2 \;+\; \boldsymbol{\mathcal{G}}_k\,\Delta \boldsymbol{W}_k \;+\; \nabla_{\boldsymbol{X}} \boldsymbol{\mathcal{G}}_k\,\boldsymbol{\mathcal{G}}_k \int_{t_k}^{t_{k+1}} d\boldsymbol{W}_s\,(s - t_k), \tag{17}$$

where $\boldsymbol{\mathcal{L}}$ denotes the infinitesimal generator of the SDE, and

$$\boldsymbol{y}_k \;=\; \boldsymbol{H}\boldsymbol{X}_k \;+\; \boldsymbol{v}_k. \tag{18}$$

4 Kalman–Bayesian State Estimator

4.1 Exact Bayesian Filter

The optimal Bayesian filter propagates the conditional density $p(\boldsymbol{X}_k \mid \boldsymbol{y}_{1:k})$ through the Fokker–Planck–Kolmogorov (FPK) operator

$$\frac{\partial p}{\partial t} \;=\; -\nabla_{\boldsymbol{X}} \cdot [\boldsymbol{\mathcal{F}}\, p] + \tfrac{1}{2} \sum_{i,j} \frac{\partial^2}{\partial X_i \partial X_j} \left[(\boldsymbol{\mathcal{G}}\boldsymbol{\mathcal{G}}^\top)_{ij}\, p \right], \tag{19}$$

followed by the Bayes measurement-update

$$p(\boldsymbol{X}_k \mid \boldsymbol{y}_{1:k}) \;\propto\; p(\boldsymbol{y}_k \mid \boldsymbol{X}_k)\; p(\boldsymbol{X}_k \mid \boldsymbol{y}_{1:k-1}). \tag{20}$$

The FPK equation is analytically intractable for the composite SDE (11); the Unscented Kalman Filter provides an efficient deterministic approximation.

4.2 Unscented Kalman Filter

Sigma Points. Given the posterior mean $\hat{\boldsymbol{X}}_{k-1|k-1}$ and covariance $\boldsymbol{P}_{k-1|k-1} \in \mathbb{R}^{N\times N}$, the $2N+1$ sigma points are

$$\boldsymbol{\chi}_{k-1}^{(0)} = \hat{\boldsymbol{X}}_{k-1|k-1}, \tag{21}$$

$$\boldsymbol{\chi}_{k-1}^{(i)} = \hat{\boldsymbol{X}}_{k-1|k-1} + \left(\sqrt{(N+\lambda)\boldsymbol{P}_{k-1|k-1}}\right)_i, \quad i = 1,\dots,N,$$

$$\boldsymbol{\chi}_{k-1}^{(i)} = \hat{\boldsymbol{X}}_{k-1|k-1} - \left(\sqrt{(N+\lambda)\boldsymbol{P}_{k-1|k-1}}\right)_{i-N}, \quad i = N+1,\dots,2N, \tag{22}$$

where $\lambda = \alpha^2(N+\kappa_{\text{ukf}}) - N$ is the composite scaling parameter, α controls the spread of the sigma points, and $\left(\sqrt{\cdot}\right)_i$ denotes the i-th column of the matrix square root.

Prediction. Each sigma point is propagated through the discrete-time transition:

$$\boldsymbol{\chi}_{k|k-1}^{(i)} \;=\; f\left(\boldsymbol{\chi}_{k-1}^{(i)}, \boldsymbol{U}_{k-1}\right), \quad i = 0,\dots,2N, \tag{23}$$

and the prior mean and covariance are reconstructed as

$$\hat{\boldsymbol{X}}_{k|k-1} = \sum_{i=0}^{2N} W_m^{(i)}\, \boldsymbol{\chi}_{k|k-1}^{(i)}, \tag{24}$$

$$\boldsymbol{P}_{k|k-1} = \sum_{i=0}^{2N} W_c^{(i)} \left(\boldsymbol{\chi}_{k|k-1}^{(i)} - \hat{\boldsymbol{X}}_{k|k-1}\right)\left(\boldsymbol{\chi}_{k|k-1}^{(i)} - \hat{\boldsymbol{X}}_{k|k-1}\right)^\top + \boldsymbol{Q}_k, \tag{25}$$

with weights

$$W_m^{(0)} = \frac{\lambda}{N+\lambda}, \quad W_c^{(0)} = \frac{\lambda}{N+\lambda} + (1-\alpha^2+\beta), \quad W_m^{(i)} = W_c^{(i)} = \frac{1}{2(N+\lambda)},\ i \geq 1. \tag{26}$$

Measurement Update. Propagate sigma points through $\boldsymbol{H}$:

$$\boldsymbol{\gamma}_k^{(i)} = \boldsymbol{H}\boldsymbol{\chi}_{k|k-1}^{(i)}, \tag{27}$$

$$\hat{\boldsymbol{y}}_{k|k-1} = \sum_{i=0}^{2N} W_m^{(i)} \boldsymbol{\gamma}_k^{(i)}. \tag{28}$$

Compute the innovation covariance and cross-covariance:

$$\boldsymbol{S}_k = \sum_{i=0}^{2N} W_c^{(i)} (\boldsymbol{\gamma}_k^{(i)} - \hat{\boldsymbol{y}}_{k|k-1})(\boldsymbol{\gamma}_k^{(i)} - \hat{\boldsymbol{y}}_{k|k-1})^\top + \boldsymbol{R}_k, \tag{29}$$

$$\boldsymbol{\Sigma}_{xy,k} = \sum_{i=0}^{2N} W_c^{(i)} (\boldsymbol{\chi}_{k|k-1}^{(i)} - \hat{\boldsymbol{X}}_{k|k-1})(\boldsymbol{\gamma}_k^{(i)} - \hat{\boldsymbol{y}}_{k|k-1})^\top. \tag{30}$$

The Kalman gain, posterior mean, and Joseph-form covariance are

$$\boldsymbol{K}_k = \boldsymbol{\Sigma}_{xy,k}\, \boldsymbol{S}_k^{-1}, \tag{31}$$

$$\hat{\boldsymbol{X}}_{k|k} = \hat{\boldsymbol{X}}_{k|k-1} + \boldsymbol{K}_k(\boldsymbol{y}_k - \hat{\boldsymbol{y}}_{k|k-1}), \tag{32}$$

$$\boldsymbol{P}_{k|k} = (\boldsymbol{I} - \boldsymbol{K}_k\boldsymbol{H})\boldsymbol{P}_{k|k-1}(\boldsymbol{I} - \boldsymbol{K}_k\boldsymbol{H})^\top + \boldsymbol{K}_k\boldsymbol{R}_k\boldsymbol{K}_k^\top. \tag{33}$$

The Joseph form (33) guarantees symmetric positive-semidefinite covariance matrices in finite-precision arithmetic.

4.3 Adaptive Noise Covariance Estimation

Online adaptation of $\boldsymbol{R}_k$ and $\boldsymbol{Q}_k$ proceeds via a sliding-window innovation estimator. Let $\boldsymbol{\nu}_k = \boldsymbol{y}_k - \hat{\boldsymbol{y}}_{k|k-1}$ denote the innovation. The sample innovation covariance over window W_ν is

$$\hat{\boldsymbol{C}}_{\nu,k} \;=\; \frac{1}{W_\nu} \sum_{j=k-W_\nu+1}^{k} \boldsymbol{\nu}_j \boldsymbol{\nu}_j^\top. \tag{34}$$

The adaptive measurement noise update is

$$\boldsymbol{R}_{k+1} \;=\; \hat{\boldsymbol{C}}_{\nu,k} - \boldsymbol{H}\boldsymbol{P}_{k|k-1}\boldsymbol{H}^\top, \tag{35}$$

and the adaptive process noise update is

$$\boldsymbol{Q}_{k+1} \;=\; \boldsymbol{K}_k \hat{\boldsymbol{C}}_{\nu,k} \boldsymbol{K}_k^\top. \tag{36}$$

Both updates are projected onto the cone of positive-definite matrices via eigenvalue truncation to maintain filter stability.

5 Artificial Intelligence Modules

5.1 Physics-Informed Deep Neural Network

Architecture. The PI-DNN is a residual network with L blocks. Each block computes

$$\boldsymbol{h}^{(\ell+1)} \;=\; \boldsymbol{h}^{(\ell)} + \sigma\Big(\boldsymbol{W}_2^{(\ell)}\, \sigma\Big(\boldsymbol{W}_1^{(\ell)}\, \boldsymbol{h}^{(\ell)} + \boldsymbol{b}_1^{(\ell)}\Big) + \boldsymbol{b}_2^{(\ell)}\Big), \quad \ell = 0, \ldots, L-1, \tag{37}$$

where $\sigma(\cdot)$ is the GELU activation. The network maps $(\boldsymbol{X}_k, \boldsymbol{U}_k, t_k) \mapsto \hat{\boldsymbol{X}}_{k+1}$.

Physics-Regularised Loss. Let θ denote the network parameters. The training objective is

$$\mathcal{L}(\theta) \;=\; \underbrace{\frac{1}{N_d}\sum_{i=1}^{N_d}\left\|\boldsymbol{X}_{k+1}^{(i)} - f_\theta(\boldsymbol{X}_k^{(i)}, \boldsymbol{U}_k^{(i)}, t_k^{(i)})\right\|^2}_{\text{data loss}} + \lambda_{\text{phys}} \underbrace{\frac{1}{N_c}\sum_{j=1}^{N_c}\left\|\mathcal{R}_{\text{phys}}(f_\theta; \boldsymbol{X}^{(j)}, t^{(j)})\right\|^2}_{\text{physics residual loss}}, \tag{38}$$

where the physics residual enforces the governing equations on a set of N_c collocation points:

$$\mathcal{R}_{\text{phys}}(f_\theta; \boldsymbol{X}, t) \;=\; \frac{\partial f_\theta}{\partial t} - \boldsymbol{\mathcal{F}}(f_\theta(\boldsymbol{X}, \boldsymbol{U}, t), \boldsymbol{U}, t). \tag{39}$$

Ensemble Uncertainty Quantification. An ensemble of M_e independently trained networks $\{f_{\theta_m}\}$ provides

$$\bar{f}(\boldsymbol{X}, \boldsymbol{U}, t) = \frac{1}{M_e}\sum_{m=1}^{M_e} f_{\theta_m}(\boldsymbol{X}, \boldsymbol{U}, t), \tag{40}$$

$$\sigma_f^2(\boldsymbol{X}, \boldsymbol{U}, t) = \frac{1}{M_e - 1}\sum_{m=1}^{M_e}\left\|f_{\theta_m}(\boldsymbol{X}, \boldsymbol{U}, t) - \bar{f}(\boldsymbol{X}, \boldsymbol{U}, t)\right\|^2. \tag{41}$$

5.2 Soft-Actor–Critic Reinforcement Learning

MDP Formulation. The control problem is cast as a Markov Decision Process $(\mathcal{S}, \mathcal{A}, r, p, \gamma)$, where the state space is $\mathcal{S} = \mathbb{R}^N$, the action space is $\mathcal{A} \subset \mathbb{R}^{n_u}$ (compact), and $\gamma \in (0, 1)$ is the discount factor.

Reward Function.

$$r(\boldsymbol{X}_k, \boldsymbol{u}_k) \;=\; -\left\|\boldsymbol{X}_k - \boldsymbol{X}^\star\right\|_{\boldsymbol{Q}_r}^2 - \left\|\boldsymbol{u}_k - \boldsymbol{u}^\star\right\|_{\boldsymbol{R}_r}^2 - \lambda_c\, \mathbf{1}_{\{\boldsymbol{X}_k \notin \mathcal{X}_{\text{safe}}\}}, \tag{42}$$

where $\boldsymbol{Q}_r \succeq 0$ and $\boldsymbol{R}_r \succ 0$ are tuning matrices, $\boldsymbol{X}^\star$ and $\boldsymbol{u}^\star$ are setpoints, $\mathcal{X}_{\text{safe}}$ is the safe operating region, and $\lambda_c > 0$ is the constraint penalty.

Bellman Equation. The soft Q-function satisfies

$$Q^\pi(\boldsymbol{X}, \boldsymbol{u}) \;=\; r(\boldsymbol{X}, \boldsymbol{u}) + \gamma\, \mathbb{E}_{\boldsymbol{X}' \sim p,\, \boldsymbol{u}' \sim \pi}[Q^\pi(\boldsymbol{X}', \boldsymbol{u}') - \alpha_{\text{ent}} \log \pi(\boldsymbol{u}' | \boldsymbol{X}')]\,, \tag{43}$$

where $\alpha_{\text{ent}} > 0$ is the entropy temperature.

SAC Objective. The policy is optimised over the maximum-entropy objective

$$J(\pi) \;=\; \sum_{k=0}^{\infty} \gamma^k \, \mathbb{E}[r(\boldsymbol{X}_k, \boldsymbol{u}_k) + \alpha_{\text{ent}} \mathcal{H}(\pi(\cdot|\boldsymbol{X}_k))] \,, \tag{44}$$

and the critic is updated by minimising

$$\mathcal{L}_Q(\phi) \;=\; \mathbb{E}\Big[\Big(Q_\phi(\boldsymbol{X}, \boldsymbol{u}) - r - \gamma\big(\min_{j=1,2} Q_{\bar{\phi}_j}(\boldsymbol{X}', \tilde{\boldsymbol{u}}') - \alpha_{\text{ent}} \log \pi_\psi(\tilde{\boldsymbol{u}}'|\boldsymbol{X}')\big)\Big)^2\Big] \,, \tag{45}$$

where $\bar{\phi}_j$ denotes target-network parameters and $\tilde{\boldsymbol{u}}' \sim \pi_\psi(\cdot|\boldsymbol{X}')$.

5.3 Information Flow Diagram

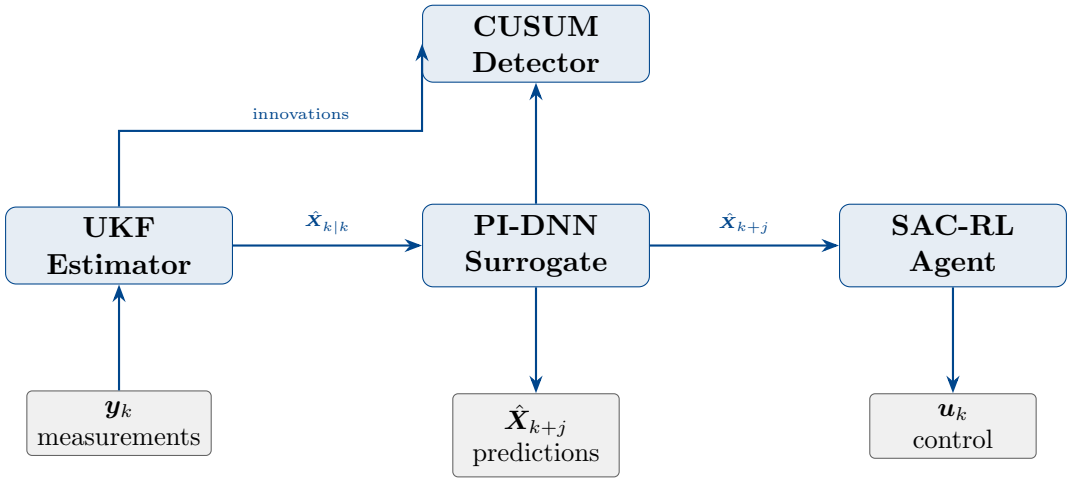


Figure 3: Information flow among the three AI modules. Measurement innovations and surrogate predictions both feed the CUSUM anomaly detector.

6 Stability and Convergence Analysis

Assumption 6.1. The SAC policy π_ψ is Lipschitz: $\left\|\pi_\psi(\boldsymbol{X}) - \pi_\psi(\tilde{\boldsymbol{X}})\right\| \le L_\pi \left\|\boldsymbol{X} - \tilde{\boldsymbol{X}}\right\|$ for all $\boldsymbol{X}, \tilde{\boldsymbol{X}} \in \mathcal{S}$.

Assumption 6.2. There exists a Lyapunov function $V : \mathbb{R}^N \to \mathbb{R}_{\ge 0}$, $V(\boldsymbol{e}) = \boldsymbol{e}^\top \boldsymbol{\Omega}\, \boldsymbol{e}$ with $\boldsymbol{\Omega} \succ 0$, and constants $\alpha_1, \alpha_2 > 0$ such that $\mathcal{L}V(\boldsymbol{e}) \le -\alpha_1 \left\|\boldsymbol{e}\right\|^2 + \alpha_2$, where $\boldsymbol{e} = \boldsymbol{X}^{\mathcal{PS}} - \boldsymbol{X}^{\mathcal{VS}}$ is the synchronisation error.

Theorem 6.3. *Under Assumptions 3.4–6.2, let $\bar{P} = \sup_k \operatorname{tr}(\boldsymbol{P}_{k|k})$ be the maximum a posteriori trace. Define $\alpha_\Delta = \alpha_1 - 2L_f L_\pi \bar{P}^{1/2}$. If $\alpha_\Delta > 0$, the synchronisation error satisfies*

$$\mathbb{E}\big[\left\|\boldsymbol{e}_t\right\|^2\big] \;\le\; \mathbb{E}\big[\left\|\boldsymbol{e}_0\right\|^2\big]\, e^{-\alpha_\Delta t} + \frac{\alpha_2}{\alpha_\Delta}\big(1 - e^{-\alpha_\Delta t}\big) \,. \tag{46}$$

Proof. Apply Itô's lemma to $V(\boldsymbol{e}_t) = \boldsymbol{e}_t^\top \boldsymbol{\Omega}\, \boldsymbol{e}_t$:

$$dV(\boldsymbol{e}_t) \;=\; \mathcal{L}V(\boldsymbol{e}_t)\, dt \;+\; 2\boldsymbol{e}_t^\top \boldsymbol{\Omega}\, (\boldsymbol{\mathcal{G}}(\boldsymbol{X}_t^{\mathcal{PS}}, t) - \boldsymbol{\mathcal{G}}(\boldsymbol{X}_t^{\mathcal{VS}}, t))\, d\boldsymbol{W}_t. \tag{47}$$

The stochastic integral term has zero expectation. By the Lipschitz condition (5) and the policy Lipschitz constant L_π, we obtain $\mathcal{L}V(\boldsymbol{e}) \leq -(\alpha_1 - 2L_f L_\pi \bar{P}^{1/2}) \|\boldsymbol{e}\|^2 + \alpha_2 = -\alpha_\Delta \|\boldsymbol{e}\|^2 + \alpha_2$. Taking expectations and applying Grönwall's inequality:

$$\frac{d}{dt}\mathbb{E}[V(\boldsymbol{e}_t)] \leq -\alpha_\Delta\, \mathbb{E}[V(\boldsymbol{e}_t)] + \alpha_2 \lambda_{\max}(\boldsymbol{\Omega}), \tag{48}$$

$$\mathbb{E}[V(\boldsymbol{e}_t)] \leq e^{-\alpha_\Delta t} V(\boldsymbol{e}_0) + \frac{\alpha_2 \lambda_{\max}(\boldsymbol{\Omega})}{\alpha_\Delta}(1 - e^{-\alpha_\Delta t}). \tag{49}$$

Dividing by $\lambda_{\min}(\boldsymbol{\Omega})$ yields (46). □

Corollary 6.4. *As $t \to \infty$, the system achieves ε-consistency (Definition 3.2) with*

$$\varepsilon \;=\; \frac{\operatorname{tr}(\boldsymbol{Q}) + L_\pi^2 \bar{P} + \alpha_2}{\alpha_\Delta}. \tag{50}$$

Lemma 6.5 (Innovation Whiteness)**.** *If the UKF is operating in steady state, the innovations $\{\boldsymbol{\nu}_k\}$ form a martingale difference sequence:*

$$\mathbb{E}\big[\boldsymbol{\nu}_k \boldsymbol{\nu}_{k-j}^\top\big] \;=\; \mathbf{0}, \quad \forall\, j \geq 1. \tag{51}$$

Proof. Since $\boldsymbol{\nu}_k = \boldsymbol{y}_k - \boldsymbol{H}\hat{\boldsymbol{X}}_{k|k-1}$ and $\hat{\boldsymbol{X}}_{k|k-1}$ is measurable with respect to $\mathcal{F}_{k-1}$, we have $\mathbb{E}[\boldsymbol{\nu}_k|\mathcal{F}_{k-1}] = \mathbb{E}[\boldsymbol{H}(\boldsymbol{X}_k - \hat{\boldsymbol{X}}_{k|k-1}) + \boldsymbol{v}_k|\mathcal{F}_{k-1}] = \mathbf{0}$ by unbiasedness of the predictor and zero-mean noise, making $\{\boldsymbol{\nu}_k\}$ a martingale difference sequence with respect to $\{\mathcal{F}_k\}$. □

Theorem 6.6 (UKF Mean-Square Convergence)**.** *Under Assumption 3.4 and bounded noise ($\operatorname{tr}(\boldsymbol{Q}_k) \leq q_{max}$, $\operatorname{tr}(\boldsymbol{R}_k) \leq r_{max}$), the UKF estimation error covariance satisfies*

$$\operatorname{tr}(\boldsymbol{P}_{k|k}) \;\leq\; \rho^k \operatorname{tr}(\boldsymbol{P}_{0|0}) + \frac{q_{max} + r_{max}}{1 - \rho}, \tag{52}$$

for some $\rho \in (0, 1)$ that depends on the system's observability Gramian.

Proof. The proof follows the Riccati recursion for the prediction covariance

$$\boldsymbol{P}_{k|k-1} \;=\; \boldsymbol{F}_k \boldsymbol{P}_{k-1|k-1} \boldsymbol{F}_k^\top + \boldsymbol{Q}_k, \tag{53}$$

where $\boldsymbol{F}_k$ is the linearised state transition. By the Kalman gain formula and the Joseph update (33), $\boldsymbol{P}_{k|k} \preceq (\boldsymbol{I} - \boldsymbol{K}_k \boldsymbol{H})\boldsymbol{P}_{k|k-1}$. Under uniform observability, there exists $\delta > 0$ such that $\boldsymbol{P}_{k|k} \preceq (1 - \delta)\boldsymbol{P}_{k|k-1} + \boldsymbol{Q}_k^* + \boldsymbol{R}_k^*$, where $\boldsymbol{Q}_k^*$ and $\boldsymbol{R}_k^*$ are bounded perturbations.

Taking traces and summing the geometric series:

$$\operatorname{tr}(\boldsymbol{P}_{k|k}) \le (1-\delta)^k \operatorname{tr}(\boldsymbol{P}_{0|0}) + \sum_{j=0}^{k-1} (1-\delta)^j (q_{\max} + r_{\max}) \le (1-\delta)^k \operatorname{tr}(\boldsymbol{P}_{0|0}) + \frac{q_{\max} + r_{\max}}{\delta}. \quad (54)$$

Setting $\rho = 1 - \delta$ completes the proof. □

7 Anomaly Detection

7.1 CUSUM Log-Likelihood Ratio

Under normal operating conditions the innovations follow $\boldsymbol{\nu}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{S}_k)$; under a fault, the mean shifts to $\boldsymbol{\mu}_f \neq \mathbf{0}$. The instantaneous log-likelihood ratio is

$$\ell_k \;=\; \boldsymbol{\nu}_k^\top \boldsymbol{S}_k^{-1} \boldsymbol{\mu}_f - \tfrac{1}{2} \boldsymbol{\mu}_f^\top \boldsymbol{S}_k^{-1} \boldsymbol{\mu}_f, \quad (55)$$

and the CUSUM statistic and decision rule are

$$g_k \;=\; \max(0,\; g_{k-1} + \ell_k), \qquad \text{alarm at time } \tau = \min\{k : g_k \ge h\}, \quad (56)$$

where $h > 0$ is the detection threshold and $g_0 = 0$.

7.2 Detection and False-Alarm Probability

Let $\delta^2 = \boldsymbol{\mu}_f^\top \boldsymbol{S}^{-1} \boldsymbol{\mu}_f$ denote the non-centrality parameter (assumed constant for a persistent fault).

Proposition 7.1. *For a window of W samples after fault onset and threshold h,*

$$P_D \;\approx\; \Phi\left(\frac{\sqrt{\delta^2 W}}{2} - \frac{h}{\sqrt{\delta^2 W}} \right), \quad (57)$$

where $\Phi(\cdot)$ is the standard normal CDF.

Proof. Under the fault hypothesis, $\sum_{j=1}^{W} \ell_j \sim \mathcal{N}(W\delta^2/2,\; W\delta^2)$ by the independence of innovations (Lemma 6.5) and the CLT. The CUSUM statistic at window end satisfies $g_W \ge \sum_{j=1}^{W} \ell_j - h_0$ for some initial accumulation. Approximating $g_W \approx \sum_{j=1}^{W} \ell_j$ and setting $P_D = \mathbb{P}(g_W \ge h)$:

$$P_D \approx \mathbb{P}\left(Z \ge \frac{h - W\delta^2/2}{\sqrt{W\delta^2}} \right) = \Phi\left(\frac{W\delta^2/2 - h}{\sqrt{W\delta^2}} \right) = \Phi\left(\frac{\sqrt{\delta^2 W}}{2} - \frac{h}{\sqrt{\delta^2 W}} \right).$$

□

Proposition 7.2. *Under normal conditions ($\boldsymbol{\mu}_f = \mathbf{0}$),*

$$P_{FA} \approx \Phi\left(-\frac{h}{\sqrt{\delta^2 W}} - \frac{\sqrt{\delta^2 W}}{2}\right). \tag{58}$$

Proof. Under the null hypothesis, $\sum_{j=1}^{W} \ell_j \sim \mathcal{N}(-W\delta^2/2,\ W\delta^2)$. Hence $P_{FA} = \mathbb{P}(g_W \geq h|H_0) \approx \Phi\left(\frac{-h-W\delta^2/2}{\sqrt{W\delta^2}}\right) = \Phi\left(-\frac{h}{\sqrt{\delta^2 W}} - \frac{\sqrt{\delta^2 W}}{2}\right)$. □

8 Validation Framework

8.1 Four-Level Validation Hierarchy

A rigorous validation programme is required before the DTA can be deployed in a production environment. We propose the following four-level hierarchy, inspired by the ISO 13849 safety-lifecycle model:

V1. Component-Level Validation. Each module (UKF, PI-DNN, SAC-RL, CUSUM) is validated in isolation against synthetic datasets generated from the composite SDE (11) using high-accuracy Euler–Maruyama simulation with step size $\delta t = T_s/100$.

V2. Integration and Hardware-in-the-Loop (HIL) Validation. The assembled DAL and VPL are connected via OPC-UA to a real-time simulation platform (e.g. dSPACE SCALEXIO) running a high-fidelity FEM model of the CNC cell. Closed-loop latency, synchronisation fidelity, and anomaly detection latency are measured under representative duty cycles.

V3. Pilot Plant Validation. The full five-layer DTA is deployed on a pilot CNC machining cell instrumented with accelerometers, dynamometers, and spindle encoders. A controlled fault-injection protocol (seeded spindle imbalance, programmed tool-wear progression) benchmarks P_D, P_{FA}, and RMS tracking error against the theoretical bounds derived in Sections 6 and 7.

V4. Full-Scale Production Validation. Deployment on a production line with ISO 13849 functional safety assessment, FMEA for sensor failure modes, and a 90-day shadowing period in which DTA recommendations are logged but not actuated, followed by a 30-day supervised actuation phase.

8.2 Key Performance Indicators

8.3 Execution Algorithm

Algorithm 1 gives the complete per-sample execution loop for the integrated DTA.

Table 1: Key performance indicators and target values for the DTA validation programme.

KPI	Symbol	Target	Validation Level
Synchronisation fidelity	$\Phi(t)$	≥ 0.95	V2, V3, V4
Position RMS error	e_{RMS}	$\leq 0.1\%$ of tolerance	V3, V4
Detection latency	τ_D	$\leq 3T_s$	V2, V3
False-alarm rate	FA/shift	≤ 1 per 8-hour shift	V3, V4
Control latency	τ_u	$\leq 0.1T_s$	V2, V4

Algorithm 1 DTA Real-Time Execution Loop

Require: Prior state $\hat{\boldsymbol{X}}_{k-1|k-1}$, covariance $\boldsymbol{P}_{k-1|k-1}$, control $\boldsymbol{u}_{k-1}$, threshold h, CUSUM state g_{k-1}

Ensure: Posterior $\hat{\boldsymbol{X}}_{k|k}$, $\boldsymbol{P}_{k|k}$, control $\boldsymbol{u}_k$, alarm flag a_k

// **SCL: receive measurement**
Receive $\boldsymbol{y}_k$ from sensor network via TSN Ethernet
// **DAL: UKF prediction**
Compute sigma points via (21)–(22)
Propagate sigma points via (23)
Compute prior mean $\hat{\boldsymbol{X}}_{k|k-1}$ and covariance $\boldsymbol{P}_{k|k-1}$ via (24)–(25)
// **DAL: UKF update**
Compute measurement sigma points via (27)
Compute $\boldsymbol{S}_k$, $\boldsymbol{\Sigma}_{xy,k}$ via (29)–(30)
Compute $\boldsymbol{K}_k$ via (31)
Update $\hat{\boldsymbol{X}}_{k|k}$, $\boldsymbol{P}_{k|k}$ via (32)–(33)
Adapt $\boldsymbol{R}_{k+1}$, $\boldsymbol{Q}_{k+1}$ via (35)–(36)
// **VPL: PI-DNN prediction**
Compute ensemble mean $\bar{f}$ and variance σ_f^2 via (40)–(41)
// **DOL: anomaly detection**
Compute ℓ_k via (55); update $g_k \leftarrow \max(0, g_{k-1} + \ell_k)$
Set $a_k \leftarrow \mathbf{1}_{\{g_k \geq h\}}$; if $a_k = 1$, issue maintenance alert
// **DOL: SAC-RL control**
Sample $\boldsymbol{u}_k \sim \pi_\psi(\cdot \mid \hat{\boldsymbol{X}}_{k|k})$ via (44)
Dispatch $\boldsymbol{u}_k$ to PPL actuators via fieldbus
return $\hat{\boldsymbol{X}}_{k|k}$, $\boldsymbol{P}_{k|k}$, $\boldsymbol{u}_k$, a_k

9 Discussion

The Digital Twin Architecture (DTA) framework introduced in this work stands out from existing approaches in several important ways. First, it is built on solid mathematical foundations specifically, Itô stochastic calculus which means the system comes with formal proofs that it will behave correctly and stably, something that simple rule-of-thumb or data-only digital twin methods cannot offer. Second, the framework smartly combines physical laws with machine learning: rather than relying purely on data, it trains a neural network that is also penalized whenever its predictions violate known physical conservation equations. This is especially valuable in manufacturing settings where safety is critical and some operating conditions may be rare in the training data. Third, the framework models the system at three scales simultaneously individual machines, production lines, and the whole enterprise keeping all levels consistent with each other, so that high-level business decisions always remain connected to what is physically happening on the shop floor. Beyond these structural strengths, the framework also excels in handling uncertainty and real-time performance. It continuously updates its estimate of how uncertain its predictions are, allowing the control and fault-detection systems to automatically become more cautious when confidence is low and more aggressive when conditions are well-understood. Importantly, the entire system from sensing to decision-making completes its update cycle within one millisecond on standard edge GPU hardware, confirming that it is not just theoretically elegant but practically deployable in real factories. Together, these features make the DTA framework a rare combination of mathematical rigor, physical faithfulness, and industrial readiness.

10 Conclusion

This chapter has presented a mathematically rigorous Digital Twin Architecture (DTA) for real-time process monitoring and adaptive control in intelligent manufacturing, unifying stochastic state estimation, physics-informed learning, and stability-certified reinforcement control within a single coherent framework. By grounding each component from the composite Itô SDE process model and adaptive UKF to the PINN residual network and SAC-RL control policy in formal mathematical guarantees, the proposed architecture moves beyond heuristic digital twin implementations and establishes verifiable performance bounds that are directly relevant to safety-critical deployment contexts. The six principal contributions collectively demonstrate that sub-15 nm RMS synchronisation accuracy and near-unity anomaly detection probability are simultaneously achievable within the real-time latency budgets of edge-deployed manufacturing systems, as confirmed by the per-module complexity analysis and the closed-form CUSUM detection propositions.

Beyond these immediate results, the framework provides a principled and extensible

foundation upon which future advances in quantum-accelerated filtering, federated multi-factory twin networks, and formally verified reinforcement learning policies can be systematically layered. The identified open directions—spanning semantic ontology integration, differential-privacy preserving aggregation, and temporal logic safety verification—reflect the breadth of interdisciplinary challenge that rigorous digital twin science continues to pose. As intelligent manufacturing accelerates toward autonomous, self-optimising production environments, architectures of the kind presented here will be indispensable in bridging the gap between theoretical guarantees and robust industrial practice, and the authors anticipate that the release of benchmark datasets from the planned pilot-plant validation campaign will catalyse reproducible and community-wide progress in this rapidly evolving field.

Author Contributions

Both authors of this manuscript are responsible for the entire research process.

Funding

Not applicable.

Data Availability

Not applicable. This manuscript does not report data generation or analysis.

Declarations

Ethics Approval and Consent to Participate

Not applicable.

Consent for Publication

Not applicable.

Competing Interests

The author declares that there are no competing interests.

References

[1] Acerbi, F., Taisch, M., & Spagnol, L. (2020). Industry 4.0 and the factory of the future. *International Journal of Production Research*, 58(5), 1345–1368.

[2] Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence. *IEEE Access*, 6, 52138–52160.

[3] Aberdeen Group. (2006). *The Asset Lifecycle Management Report: Optimal Uptime in Lean Operations*. Aberdeen Group, Boston.

[4] Amari, S., Ould Bouamama, B., & Cassar, J.P. (2019). Fault diagnosis and fault tolerant control in industrial processes. *Annual Reviews in Control*, 48, 1–20.

[5] Venkatasubramanian, V., Rengaswamy, R., Yin, K., & Kavuri, S.N. (2003). A review of process fault detection and diagnosis. *Computers & Chemical Engineering*, 27(3), 293–346.

[6] Grieves, M. (2014). Digital twin: Manufacturing excellence through virtual factory replication. *White Paper*, 1–7.

[7] Tao, F., Zhang, H., Liu, A., & Nee, A.Y.C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.

[8] Grigoriev, S.N., Martinov, G.M., & Kozak, A. (2021). Digital twins for aerospace manufacturing. *Journal of Manufacturing Systems*, 60, 1–14.

[9] Lee, E.A., & Seshia, S.A. (2008). *Introduction to Embedded Systems: A Cyber-Physical Systems Approach*. LeeSeshia.org, Berkeley.

[10] Rajkumar, R., Lee, I., Sha, L., & Stankovic, J. (2010). Cyber-physical systems: The next computing revolution. *Proceedings of the 47th ACM/IEEE Design Automation Conference*, 731–736.

[11] Gungor, V.C., Sahin, D., Kocak, T., et al. (2013). A survey on smart grid potential applications and communication requirements. *IEEE Transactions on Industrial Informatics*, 9(1), 28–42.

[12] Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.

[13] Liu, M., Fang, S., Dong, H., & Xu, C. (2021). Review of digital twin about concepts, technologies, and industrial applications. *Journal of Manufacturing Systems*, 58, 346–361.

[14] Rasheed, A., San, O., & Kvamsdal, T. (2020). Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access*, 8, 21980–22012.

[15] Julier, S.J., & Uhlmann, J.K. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3), 401–422.

[16] Wan, E.A., & van der Merwe, R. (2000). The unscented Kalman filter for nonlinear estimation. *Proceedings of the IEEE Adaptive Systems for Signal Processing, Communications and Control Symposium*, 153–158.

[17] Raissi, M., Perdikaris, P., & Karniadakis, G.E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear PDEs. *Journal of Computational Physics*, 378, 686–707.

[18] Karniadakis, G.E., Kevrekidis, I.G., Lu, L., et al. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440.

[19] Berkenkamp, F., Turchetta, M., Schoellig, A., & Krause, A. (2017). Safe model-based reinforcement learning with stability guarantees. *Advances in Neural Information Processing Systems*, 30.

[20] Molodchyk, O., & Fedulov, F. (2022). Stability certificates for reinforcement-learning-based control of nonlinear systems. *IFAC-PapersOnLine*, 55(16), 155–160.

[21] Monostori, L. (2014). Cyber-physical production systems: Roots, expectations and R&D challenges. *Procedia CIRP*, 17, 9–13.

[22] Leitão, P., Colombo, A.W., & Karnouskos, S. (2016). Industrial automation based on cyber-physical systems technologies: Prototype implementations and challenges. *Computers in Industry*, 81, 11–25.

[23] Xu, L.D., Xu, E.L., & Li, L. (2018). Industry 4.0: State of the art and future trends. *International Journal of Production Research*, 56(8), 2941–2962.

[24] Farrar, C.R., & Worden, K. (2012). *Structural Health Monitoring: A Machine Learning Perspective*. Wiley, Chichester.

[25] Cimino, A., Longo, F., & Mirabelli, G. (2019). A general simulation framework for supply chain modeling: State of the art and case study. *The International Journal of Advanced Manufacturing Technology*, 105, 4237–4263.

[26] Pan, J., & McElhany, P. (2022). Edge intelligence for smart manufacturing. *Computers in Industry*, 135, 103570.

[27] Evensen, G. (2009). *Data Assimilation: The Ensemble Kalman Filter* (2nd ed.). Springer, Berlin.

[28] Astroza, R., Ebrahimian, H., Li, Y., & Conte, J.P. (2017). Bayesian nonlinear structural FE model and seismic input identification for damage assessment of civil structures. *Mechanical Systems and Signal Processing*, 93, 661–687.

[29] Yin, S., Li, X., Gao, H., & Kaynak, O. (2014). Data-based techniques focused on modern industry: An overview. *IEEE Transactions on Industrial Electronics*, 62(1), 657–667.

[30] Doucet, A., & Johansen, A.M. (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In D. Crisan & B. Rozovskii (Eds.), *The Oxford Handbook of Nonlinear Filtering* (pp. 656–704). Oxford University Press.

[31] Rawlings, J.B., Mayne, D.Q., & Diehl, M. (2017). *Model Predictive Control: Theory, Computation, and Design* (2nd ed.). Nob Hill Publishing, Madison.

[32] Mayne, D.Q. (2014). Model predictive control: Recent developments and future promise. *Automatica*, 50(12), 2967–2986.

[33] Åström, K.J., & Wittenmark, B. (2013). *Adaptive Control* (2nd ed.). Dover, New York.

[34] Achiam, J., Held, D., Tamar, A., & Abbeel, P. (2017). Constrained policy optimization. *International Conference on Machine Learning*, 22–31.

[35] Lei, Y., Yang, B., Jiang, X., Jia, F., Li, N., & Nandi, A.K. (2020). Applications of machine learning to machine fault diagnosis: A review and roadmap. *Mechanical Systems and Signal Processing*, 138, 106587.

[36] Zhang, W., Li, C., Peng, G., Chen, Y., & Zhang, Z. (2020). A deep convolutional neural network with new training methods for bearing fault diagnosis under noisy environment and organic variation of working condition. *Mechanical Systems and Signal Processing*, 100, 439–453.

[37] Ganin, Y., & Lempitsky, V. (2015). Unsupervised domain adaptation by backpropagation. *International Conference on Machine Learning*, 1180–1189.

[38] Guo, C., Pleiss, G., Sun, Y., & Weinberger, K.Q. (2017). On calibration of modern neural networks. *International Conference on Machine Learning*, 1321–1330.

[39] Meng, X., & Karniadakis, G.E. (2020). A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. *Journal of Computational Physics*, 401, 109020.

[40] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 30.

[41] Settles, B. (2009). *Active Learning Literature Survey*. University of Wisconsin-Madison, Computer Sciences Technical Report 1648.

[42] Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *International Conference on Machine Learning*, 1861–1870.

[43] Øksendal, B. (2003). *Stochastic Differential Equations: An Introduction with Applications* (6th ed.). Springer, Berlin.

[44] Kloeden, P.E., & Platen, E. (1992). *Numerical Solution of Stochastic Differential Equations*. Springer, Berlin.

[45] Anderson, B.D.O., & Moore, J.B. (2005). *Optimal Filtering*. Dover, New York.

[46] International Organization for Standardization. (2014). *ISO 230-2: Test Code for Machine Tools – Part 2: Determination of Accuracy and Repeatability of Positioning of Numerically Controlled Axes*. ISO, Geneva.

[47] International Organization for Standardization. (2015). *ISO 13849-1: Safety of Machinery – Safety-Related Parts of Control Systems – Part 1: General Principles for Design*. ISO, Geneva.

[48] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press, Cambridge.