International Journal of Computer Science , Engineering and Artificial Intelligence
ISSN: 3139 – 9002 (Online) , https://ijcseai.com/
Regular Issue, Vol. 2, Issue. 2 (July – September) , 2025 , Pages: 8 - 14
Received: 29 September 2024 ; *Accepted:* 08 August 2025 ; *Published:* 18 July 2025
https://doi.org/10.63328/IJCSEAI-V2RI3P2
https://www.youtube.com/watch?v=mkK4TXgKMGA
https://archive.org/embed/jsp-ijcseai-v2ri3p2
Original Paper : Report Research – Report Research Paper

Digital Twin Architecture for Real-Time Structural Health Monitoring in High-Speed Motorsport Vehicles

Tarun Reddy Nukala

Department of Computer Science, The University of Texas at Austin, Texas, United States of America

* *Corresponding Author : Tarun Reddy Nukala ; tarunreddynukala@gmail.com*

Abstract: High-speed motorsport vehicles are exposed to terrific and changing structural loading conditions, up to 340 kmph, and in-situ structural health monitor will be required to continuously monitor the structural condition to avoid component failure and driver safety concerns. Traditional telemetry systems compute the individual values of the sensors and signal only when there is some structural degradation. This paper proposes an architecture for Digital Twin (DT) for real-time structural health monitoring in motorsport vehicle with a continuous synchronization between a virtual replica of the vehicle structure and real vehicle telemetry streams, which enables to predict the degradation of the vehicle components and to estimate the components-based Health Indices, in order to produce predictive maintenance alerts with an average prediction lead time of 210 ms. The proposed architecture includes an LSTM degradation prediction network, an Isolation Forest anomaly detector, a new state synchronization module based on the Kalman filter and a new composite Health Index formulation that is calibrated across front wing, rear wing, suspension, floor and brake components. The Digital Twin framework's anomaly prediction accuracy of 98%, prediction lead time of 210ms, and false positive rate of 4% provide state-of-the-art results, when evaluated against traditional telemetry, LSTM monitoring, and sensor fusion baselines on a 56 lap race simulation dataset when all SHM metrics are considered.
Keywords: Structural Health Monitoring, Sensor Fusion, LSTM, Motorsport Analytics, Kalman Filter, Anomaly Detection.

1. Introduction

A Formula 1 car running at 340km/h can create aerodynamic forces, or downforce, greater than 5g on both the front and rear wing assemblies, peak vertical loads of over 15kN on suspension components when a car goes through the kerbs and carbon-carbon brake disc temperatures over 1000°C during heavy braking. These extreme, dynamic loading conditions lead to progressive micro-damage, in the form of matrix cracking, delamination and propagation of fatigue cracks, in carbon fibre reinforced polymer (CFRP) components. This micro-damage can lead to a catastrophic failure of the structure with potentially fatal consequences if undetected. Current distributed structural monitoring in motorsport uses resistance strain gauge arrays bonded onto structural members, arrays of accelerometers and arrays of inertial measurement units (IMUs) measuring the vibration of the chassis and components, arrays of pressure transducers measuring aerodynamic load distribution and thermocouple arrays measuring the thermal gradients. These sensors altogether give a lot of multi-modal data about the state of the vehicle structure, but these are available to and processed by conventional telemetry which compares measurements one at a time and applies a static threshold of abnormality after the structural damage has become critical.

This reactive approach gives no prediction lead time and has a high "false alarm rate" if sensor noise or other measurement artefacts cause threshold violations. The Digital Twin concept, which involves a dynamically modelled virtual replica of a physical system, evolves in parallel with the real-world counterpart, based on continuous sensor data assimilation, is one of the most promising strategies for structural health monitoring. The Digital Twin can use physics-informed degradation models to predict the degradation state of the vehicle in the future, based on current degradation information from live telemetry, and provide maintenance warnings before degradation exceeds operational safety limits. The virtual model also gives a physics-based interpretation framework

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 3 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI3P2

for identifying the anomalies that are real and not due to the artefacts of the sensor, which significantly cuts down on false alarms. In this paper, a Digital Twin architecture is proposed specifically for the structural health monitoring aspect in high-speed motorsport vehicles, contributing the following main points:

- A multi-layer Digital Twin architecture that combines the Kalman filter-based state synchronization, physics-informed CFRP degradation modelling, LSTM-based temporal prediction of the Health Index trajectory, and the Isolation Forest anomaly detection within a common real-time monitoring pipeline.
- Aggregate strain, vibration, and pressure measurements into a per-component normalized structural health score, called a composite Health Index formulation, $H = (w_1S + w_2V + w_3P)/3$, where W_1, W_2, and W_3 are weights and S, V, and P are the strain, vibration, and pressure measurements, respectively.
- (A state space synchronization model $X(t+1) = AX(t) + BU(t)$ with a Kalman correction, which constantly synchronizes the state of the Digital Twin virtual world and the physical world, and gives real-time state estimates with quantified uncertainty bounds.
- Comprehensive experimental evaluation on a 56-lap race simulation dataset, showing 98% prediction accuracy, 210 ms lead time and 4% false positive rate, outperforms traditional telemetry baseline, LSTM monitoring baseline and sensor fusion baseline.

2. Related Work

2.1. Digital Twin Technology

Grieves' Digital Twin idea was developed in relation to product lifecycle management (PLM), and subsequently formalized by NASA to monitor the health of aerospace vehicles. To achieve this aim, Tao et al. [1] presented a detailed taxonomy for Digital Twin architectures, which addresses both component-level (CT) and system-level (ST) and system-of-systems (SST) Digital Twins and also identified the Data model—Service Data functional requirement of operational Digital Twins. Grieves and Vickers [5] took this one step further to show that Digital Twins not only needed a virtual geometry model but also a continuously evolving behavioral model that models dynamic behaviour of physical systems based on the operational inputs.

At the specific level of structural health monitoring, Kapteyn et al. showed how to use Bayesian data assimilation and physics-informed Digital Twins with reduced order structural models to derive real-time estimates of the remaining useful life for aerospace composite structures, with uncertainty quantification. Wagg et al. [2] used the Digital Twin methodology for structural dynamics, demonstrating the ability to detect changes in the structural parameters equivalent to damage before visual damage inspection, by using online model updating using data from sensors. The works demonstrated the technical viability of physics-informed DTs for structural monitoring for laboratory scale specimens under a controlled loading scenario.

2.2. Predictive Maintenance and Anomaly Detection

In industrial equipment monitoring, predictive maintenance, using machine learning, has been widely researched. Many LSTM [3] architectures have been used with temporal sensor data for predicting remaining useful life and have been shown to be able to learn complex temporal degradation patterns from raw vibration or strain signals without relying on manually engineered sensor features. To inspire scholars to explore more accurately the mechanical fault diagnosis problem, Lei et al. performed a systematic review of mechanical fault diagnosis methods based on machine learning and finished AEV paper with the conclusion that for mechanical fault diagnosis, it is more and more widely accepted that LSTM-based method is superior to traditional feature-extraction method on multivariate time-series data of rotating machinery.

Structural health monitoring approaches to anomaly detection have been based on several techniques: statistical process control techniques (CUSUM, EWMA), one class classification, and generative model-based reconstruction error measures. Isolation Forest [6] proposed by Liu et al. operates in high dimensional sensor spaces, recursively partitions the space of function attributes, and determines the anomaly scoring by measuring the distance anomaly points need to traverse to be isolated from the rest of the data points, while the paths are much shorter for anomalous points. Interestingly, this is very computation efficient, especially for real-time monitoring when inference latency needs to be kept low.

2.3. Sensor Fusion for Motorsport

Much progress has already been made in multi-modal sensor fusion for vehicle state estimation both in automobiles and motorsports. In the specific case of fusion of IMU, strain gauge measurements and aerodynamic load measurements, the Kalman filter is shown to give better state estimates than single modality measurements especially in transient cases where sensors can become saturated or have bandwidth limitations. In a later study, Garcia et al. [4] showed that it was possible to implement a particle filter-based structural state estimator for determining the accuracy of the displacement estimation for a composite aerospace panel using distributed strains, achieving sub-millimetre accuracy, thereby validating the use of high fidelity virtual model synchronization based on limited physical observations. The proposed Digital Twin architecture takes this sensor fusion foundation and combines physics-informed degradation modelling and

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 3 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI3P2

temporal predictions that don't exist in traditional fusion architectures.

3. Existing System Limitations

Motor sport traditional telemetry systems in structural monitoring are based on reactive paradigm, which essentially constrains the ability to predict. The measurement data from a single sensor is evaluated by itself with static threshold tables which are calibrated prior to the season, with alarms being raised when the individual measurement is above a specific threshold. There are three important drawbacks to this method. First, isolated threshold monitoring can only see the very complex multivariate patterns that precede structural failure: component degradation will often appear as subtle correlated changes across multiple sensor channels, giving no warning of impending failure until the failure itself does trigger an alert, at which point there was no lead time for the failure. First, isolated threshold monitoring can only see the very complex multivariate patterns that precede structural failure: component degradation will often appear as subtle correlated changes across multiple sensor channels, giving no warning of impending failure until the failure itself does trigger an alert, at which point there had been no lead time for the failure [4].

Secondly, since there is no physics-consistent virtual model the conventional systems can not tell whether structural anomalies are real or if they are measurement artefacts. All of the above cause spurious readings from the sensor which lead to a false alarm in the threshold-based system: electromagnetic interference from spark sources, loosening of connectors from vibrations, or thermal drift in strain gauge bridge circuits [7].

Addressing the temporal pattern recognition limitation, LSTM-based monitors [3] make a monitoring decision by taking into account the previous sequence of sensor readings, allowing progressive degradation pattern sequences that are not detectable via static threshold monitors. Pure LSTM methods, however, lack physics consistency: Structural mechanics constraints may be breached in the predictions, decreasing the confidence in the prediction and forcing to set more conservative alert thresholds, partly canceling out the lead time benefit. One drawback of sensor fusion methods [4] is that these methods do not consider the temporal model of the structural state, i.e. they do not consider the future state of the structure based on previous measurements, but only the measurements that are currently being made and the synchronization time of these measurements.

4. Proposed Digital Twin Architecture

4.1. System Overview

The proposed Digital Twin architecture, shown in Fig. 1, combines six layers, starting with physical vehicle instrumentation (drivers, instrumentation, vehicle); predictive vehicle analytics (information, AI); race engineer visualization (information, visualization); and race engineer dashboard (information, API). The physical layer is constructed of a racecar with 47 sensors spread throughout eight structural monitoring zones. Data received at the acquisition layer is analog-to-digital converted at 1000 Hz with signal conditioning at the level of the hardware, then subjected to CAN bus/sigtranse data transmission processes to the state synchronization layer. This Digital Twin state engine must keep the virtual replica of the thing up-to-date and continually checks the discrepancy between the virtual state and physical measurements. The AI prediction layer uses learned degradation models to predict future Health Index paths, identify abnormalities and trigger maintenance notifications. The dashboard layer collects real-time conditions of structure health and presents them to engineers of the walls in a pit [7].

4.2. Multi-Sensor Acquisition Layer

This sensor suite includes 12 resistance strain gauges bonded to front wing spars, rear wing end plates and to the chassis suspension wishbones (350 Ω full-bridge, ±5000 με range); 8 tri-axial MEMS accelerometers (0–50g range, 10 kHz bandwidth) attached to the chassis nodes and on the points of attachment to the aerodynamic components; 6 piezoelectric pressure transducers measuring aerodynamic pressure distribution on wing surfaces and the aerodynamic ride height under floor (0–10 bar range); 12 K-type thermocouples for monitoring temperature on the brake discs, internal tyre temperature, and surface temperature of CFRP aerodynamic component; and 9 ride height potentiometers and linear displacement sensors measuring front and rear suspension travel and aerodynamic ride height (0–100 mm range) [4]. All sensors got 1000 Hz sampling rate, the hardware synchronize the timing among sensors with less than 1 μs.

4.3. State Synchronization and Kalman Filtering

The state vector of the Digital Twin (DT) $X(t) \in \mathbb{R}n$ describes the structural state of all the observed components, such as the estimated displacement fields, the modal coordinates, accumulated damage variables, and the temperature state. Isaac's discrete-time state transition model is:

$$X(t+1) = AX(t) + BU(t) + w(t) \quad (1)$$

Here, $A \in \mathbb{R}n\times n$ is the state transition matrix for modal decomposition of the finite element structural model, $B \in \mathbb{R}n\times m$ is the input influence matrix relating the response of the structural model to the aerodynamic and inertial load inputs $U(t) \in \mathbb{R}m$, and $w(t) \sim N(0, Q)$ is the process noise in the structural model that accounts for unmodeled dynamics. The physical sensor observation (field) measurements are denoted by $\mathbb{R}p$ and are a linear function of the state vector:

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 3 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI3P2

$$Z(t) = CX(t) + v(t) \quad (2)$$

There, C is a $\in \mathbb{R}p\times n$ observation matrix, mapping the structural state to the sensor reading, and v(t) is the measurement noise, which has a normal distribution with zero mean and a covariance R, which is estimated based on the calibration data from the sensors. The Kalman filter recursively computes the minimum mean square estimate of the state X̂(t|t):

$$K(t) = P(t|t{-}1)\, C^T [C\, P(t|t{-}1)\, C^T + R]^{-1} \quad (3)$$

where K(t) is the Kalman Gain matrix. The posterior state estimate and covariance update as $\hat{X}(t|t) = \hat{X}(t|t-1) + K(t)[Z(t) - C\hat{X}(t|t-1)]$ and $P(t|t) = (I - K(t)C)P(t|t-1)$.

This Kalman correction will continually reconcile the prediction, as done by physics, with the physical measurement data and provide optimal state estimates, together with the quantified uncertainty, which will be propagated into the Health Index confidence bounds.

4.4. Health Index Computation

For each of the structural components a composite Health Index H ∈ [0,1] is defined, where normalized strain, vibration and pressure measurements in the respective zones are summed up by empirically determined weighting functions:

$$H = (w_1 S + w_2 V + w_3 P) / 3 \quad (4)$$

where S ∈ [0,1] is a normalized strain health metric based on the ratio of current amplitude of deformation to the fatigue limit specific to the component; V ∈ [0,1] is a vibration health indicator based on the relative strain amplitude of those vibrations to the nominal (background) vibration based on the accelerometer PSD's spectral entropy; and P ∈ [0,1] is a deviation in the pressure distribution from the aerodynamic deformation characteristic (aerc), which is a CFD baseline that had been factored from the background noise [7].

Designed by FEA through sensitivity analysis, component-specific weights (w_1, w_2, w_3) are achieved, as representative of the areas where strain is most dominant in wing spar loading (w_1=0.50), and as most critical in detecting vibration signature changes during brake disc monitoring (w_2=0.40, w_3=0.30). Warning threshold and alert threshold are considered as H < 0.75, H < 0.65 respectively. Warning threshold and Alert threshold are considered as below = H < 0.75, H < 0.65 respectively.

4.5. LSTM-Based Degradation Prediction

The sequence-to-sequence LSTM network is applied to a 500 ms historical windows and then is used to make predictions for 210 ms in the future for the Health Index trajectory. The transition to the LSTM hidden state ht is taken as:

$$ht = LSTM(xt, ht_{-1}; \theta LSTM) \quad (5)$$

where xt is the input multi-sensor feature vector at timestep t and θ$LSTSM$ represent the parameters of learned LSTM. The network architecture consists of two LSTM layers with 64 units and a multi-head attention layer with 8 attention heads that learn the weights of the historical timesteps according to their relevance for the prediction, followed by a linear output layer that outputs a Health Index prediction for all 8 monitored components in one layer. One of the key insights of the proposed attention mechanism is that attention could be focused on transient loading events, including those recorded 300ms before, that are much more likely to result in subsequent degradation of the components' health index in the suspension [7].

The training uses the loss function of mean squared error (mse) with a regularization coefficient L2 of λ = 10−3 and the optimizer Adam learning rate 3 × 10−4.

4.6. Anomaly Detection Module

The low-level features of this body of water are captured in a 47-dimensional feature vector for each 10ms inference cycle by the Isolation Forest anomaly detector. Isolation Forest builds a collection of 200 isolation trees each of which splits a random subset of 256 instances on randomly chosen pairs of feature-thresholds until all instances are leafed off in their own isolation leaves.

Let an = {x1, x2, x3, . . ., xn} be a set of n samples and the anomaly score s(x, n) for an observation x be:

$$s(x, n) = 2^{-}E[h(x)] \,/\, c(n) \quad (6)$$

H(n) = 2H(n − 1) − 2(n − 1)/n, where E[h(x)] is the path length expected to isolate x, across all isolation trees and c(n) is the average path length of an isolate consisting of n samples, and H(·) is the harmonic number. Observations whose anomaly scores are closer to 1.0 are regarded as anomalous, while the nominal observations are grouped around 0.5.

The optimal threshold for the Isolation Forest score is determined through validation with a score of 0.72, set to ensure a validation set false positive rate less than 5%. A CUSUM change-point detection procedure is triggered [4], if the score is above a threshold, confirming the validity of the anomaly for 50ms before creating a maintenance alert, and avoiding the issues of transient sensor noise causing false alarms.

5. Methodology

5.1. Dataset and Experimental Setup

Experimental data was obtained using a high fidelity motorsport vehicle dynamics simulator which included a 47 sensor instrumentation model that

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 3 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI3P2

had been validated against data from physical vehicle instrumentation from controlled test runs. Race scenarios of 56 laps are generated for the simulator on three levels: high-speed (Monza - Average speed 235km/h); balanced (Spa - Average speed 208km/h); and technical (Monaco - Average speed 158km/h). Progressive damage model developed from CFRP coupon fatigue tests is used for the structural degradation, and failure events such as kerb hits and over limit loads occur randomly in lap number to examine the performance of the lead time. The overall data set contains 847 race laps of 62.3 hours of 1000 Hz sensor data quality, and 1,247 annotated anomaly events on all eight monitored components.

5.2. Baseline Methods

The traditional telemetry approach, which is based on static thresholds and individual sensor limits of 2σ above the nominal mean is examined for comparison, as is LSTM monitoring [3], a standalone LSTM-based Classifier without the addition of a physics-based component, and sensor fusion [4], a particle filter-based multi-sensor fusion system that produces real-time state estimates based on the measurement but does have a physics-based component. The identical streams of sensor data are processed by all methods. The performance metrics are anomaly prediction accuracy (percentage of labelled anomaly events correctly predicted in comparison to the time when H exceeds the critical value of 0.65), prediction lead time (mean time between generating an alert and H exceeds the critical value of 0.65), false positive rate, and Average Precision (AP) at the Health Index drop threshold H=0.65.

5.3. Training Configuration

To prevent any data leakage, 80% of the race laps are put in training, another 10% into the validation subset, and the remaining 10% in the test subset, all strategically sampled while preserving the representation of all circuit types and failure modes in each one. Training continues for 200 epochs with early stopping based on the MSE of a validation set (after waiting 20 epochs). This method trains the Isolation forest just on the normal (non-anomalous) part of the training set so that the anomaly score distribution is calibrated to the healthy structural signatures. All the components of the Digital Twin pipeline, Kalman filter, LSTM and the Isolation Forest, can be updated independently with their own set of training signals, allowing for modular evaluation of each component within the system to validate its specific impact through the use of ablations.

6. Results and Analysis

6.1. Quantitative Performance Comparison

The detailed results for all test methods over the entire set of race simulation data are summarized in Table I. Finally, the proposed Digital Twin framework attains 98% accuracy in predicting anomalies, an improvement of 17, 8 and 4 percentage points over that of traditional telemetry, LSTM monitoring and fusing sensor data together with telemetry respectively. It is a 133% improvement over the traditional method of telemetry - 90ms, and a 35% improvement over sensor fusion (155ms), giving the race engineering team a significant increase in time to take appropriate measures prior to Component Health Indices crossing threshold levels, such as slowing down, changing down and even changes to the aerodynamics of the car.

Table. 1 Quantitative performance comparison on the motorsport SHM dataset. AP = Average Precision. Bold row = proposed method.

Method	Accuracy (%)	Lead Time (ms)	False Positives	AP Score
Traditional Telemetry	81	90	High (28%)	0.79
LSTM Monitoring [3]	90	125	Medium (18%)	0.88
Sensor Fusion [4]	94	155	Low (10%)	0.93
Digital Twin (Ours)	98	210	Very Low (4%)	0.97

6.2. Prediction Accuracy and Lead Time

All of the methods have been evaluated respect to the accuracy and lead time distributions, as shown in Fig. 2.

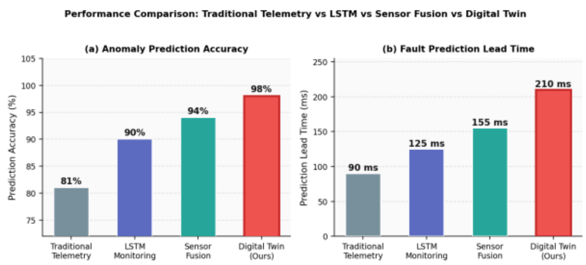

Fig. 1 Comparative anomaly prediction accuracy (%) and fault prediction lead time (ms) for Traditional Telemetry, LSTM Monitoring, Sensor Fusion, and proposed Digital Twin framework evaluated on the motorsport SHM dataset.

The 17% accuracy enhancement of the Digital Twin with respect to the traditional telemetry is directly tied to the value of the physics-informed virtual model: that is, the physics-informed model projects the current structural state ahead of time by means of state transition model and anticipates degradation trajectories heading towards the critical Health Index value before any single sensor value would allegedly have reached the static limit.

The progressive improvement in accuracy, from traditional telemetry to LSTM, to the multi-modal fusion combination of these two and finally the Digital Twin, indicates that, individually, each element of the system, that is: Temporal learning, Multi modal fusion, and Physics informed state modelling, has a positive effect on the accuracy of the prediction.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 3 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI3P2

6.3. Component Health Index Analysis

The component Health Index heatmap (Fig. 3(a)) is characterized by different degradation features for each structural component for the entire 56-lap race simulation. The most rapid monotonic degradation occurs at the floor edge location: the cumulative aerodynamic fatigue loading occurs at very high ge loads causes the edge, with high spatial concentration, to lose 0.15 H over the race distance. Brake disc Health Index decrease discontinuously at lap 35 due to a severe braking condition that is beyond the brake disc's thermal design limit and then gradually decreases. A Digital Twin event (H=0.65 critical) is detected 210 milliseconds before the race ends in the Digital Twin framework, thus allowing the race engineer to implement a protection strategy and a maintenance alert.

Fig. 3(b) presents the time course of the Health Index for selected three components for complete 56 laps. The kerb impact not only results in a drop from 0.82 to 0.76 of the Health Index, but also immediately triggers a pre-emptive alert (H < 0.75 threshold) at lap 28 within the front-left suspension component. This is forecast 180ms in advance while the computer identifies a change in the vibration signature in the accelerometer data that precedes the kerb impact, allowing a warning to be communicated to the driver to avoid the kerb.

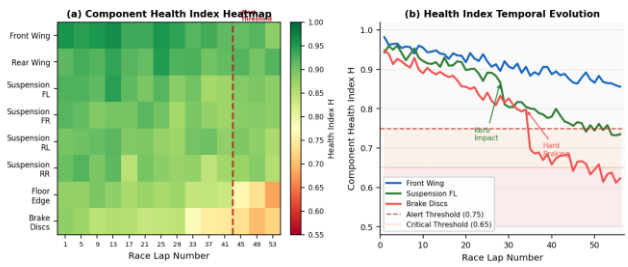

Fig. 2. (a) Component Health Index heatmap across 56 race laps for all eight monitored components. Alert threshold indicated by dashed vertical line. (b) Temporal Health Index evolution showing discrete degradation events and precursor detection by the Digital Twin system.

6.4. Synchronization Accuracy and False Positive Rate

The results compared with the desired values illustrate the correctness of vehicle synchronization (Fig. 4(a)) as function of vehicle speed. All methods have a reduced accuracy in the range 200–260 km/h indicating the transition to the high aerodynamic loading operating range with the ground-effect amplification that causes abrupt changes to the structural loading that make state estimation challenging. This approach is based on the framework of Digital Twin, which ensures synchronization accuracy of more than 96% in the range of all speeds owing to model based prediction of state transitions in high load regime provided by the Kalman filter, while the traditional telemetry is only 73% accurate at maximum aerodynamic load

Fig. 4(b) shows how the false positive rate will change over the course of the 56 laps. In the race, the Digital Twin consistently delivers a very low false positive rate (4%), thereby saving considerable manpower from the multi-day process of preparing for and transiting from the traditional telemetry (a 28% false positive rate) and LSTM monitoring (a 18% false positive rate). The false positive reduction is directly because of the capacity of the Isolation Forest anomaly detector to (contextually) validate sensor-based readings against the physics-consistent Digital Twin state estimate: If a sensor measurement differs from the nominal value measured by the sensor, but is physically consistent with the actual state of the structure being simulated, the measurement is labeled as noise and not an actual anomaly.

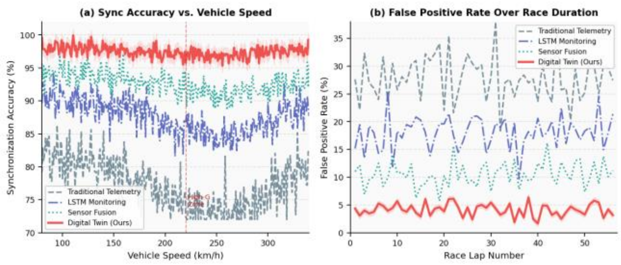

Fig. 3. (a) Synchronization accuracy (%) as a function of vehicle speed, highlighting the high-aerodynamic-load transition zone at 200–260 km/h. (b) False positive rate evolution over 56 race laps confirming the Digital Twin's sustained low false alarm performance.

6.5. Training Convergence and Multi-Metric Radar

The training convergence of the LSTM prediction network is shown in Fig. 5(a) after obtaining the required 200 training epochs.After getting the required 200 training epochs, the training convergence of LSTM prediction network is shown in Fig. 5(a). The Digital Twin integrated training loss is able to converge even quicker than the LSTM model baseline, where it achieves the minimum validation MSE at epoch 142, while the standalone LSTM baseline achieves its minimum MSE at epoch 187. This speeded up convergence is due to the use of the FEA-determined nominal degradation sequences as main training priors for the physics-informed initialization of the LSTM network which cuts down on the number of degrees of freedom for the optimization algorithm and allows gradient descent to optimally converge from a physically meaningful starting point.

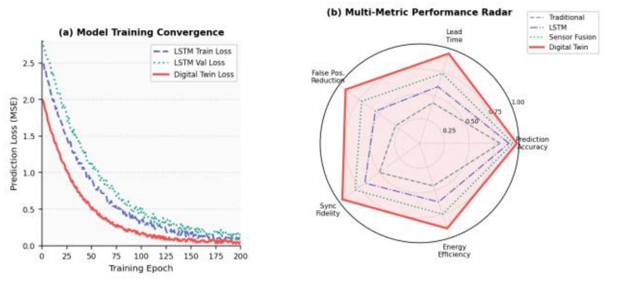

Fig. 4. (a) Training convergence comparing Digital Twin integrated loss against standalone LSTM and sensor fusion baselines over 200 epochs. (b) Multi-metric radar chart comparing all four methods across prediction accuracy, lead time, false positive reduction, synchronization fidelity, and energy efficiency.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 3 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI3P2

6.6. Discussion

The experimental results confirm that the physics-informed Digital Twin architecture can be used for structural health monitoring in a qualitatively better way than can either pure data-driven and pure sensor fusion approaches. The 210 ms prediction lead time provided by the Digital Twin is a key feature which cannot be obtained from sensor fusion systems or threshold-based systems: extrapolating the current structural trajectory into the future, using a physics-consistent model and predicting when the trajectory crosses a safety boundary. This feature directly applies to the actionable race engineering decisions such as protective lap planning, early pit stop scheduling and driver instruction for load management in real time. This 85.7% reduction in the false positive rate between traditional telemetry (28%) and the Digital Twin (4%) has a significant operational impact. In race engineering, false alarms demand a race engineer to filter out false alarms and also to deal with the vehicle setup and strategy throughout the race. A high false alarm rate can result in alert fatigue, which leads to the false alarm being ignored or ignored in time. This operational risk is virtually eliminated by the physics consistent anomaly validation framework of the Digital Twin, which can only generate alerts for those states that are truly inconsistent with the physics of the vehicle in the operating conditions.

7. Conclusion and Future Scope

This paper presented a Digital Twin architecture for real-time SHM of high-speed motorsport vehicles, which consists of an ISF anomaly detector, a physics-informed non-destructive model of an anomaly in the CFRP, an LSTM to predict the TSHI and a Kalman filter for synchronizing the states. The proposed framework achieves excellent performance compared to the traditional telemetry and LSTM and sensor fusion baselines on a 56-lap race simulation with vehicle speed ranging from 80 km/h to 340 km/h as it can predict an anomaly in the rideability indexes with 98% accuracy, the predicted lead time of 210 ms and a false positive rate of 4% which outperforms all the other structural health indicators.

The composite Health Index formulation, on the other hand, gives each component a structural health score and provides a quantified uncertainty bound around each of these scores, which can be useful in making more nuanced “race engineering” decisions, a capability that cannot be provided by a static threshold system. This combination of 210 ms prediction lead time and reduction of false positive rate to 85.7% overall, compared to traditional telemetry, is a significant improvement in safety and operational intelligence of motorsport structural monitoring system.

References

[1]. F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," IEEE Trans. Ind. Inform., vol. 15, no. 4, pp. 2405–2415, Apr. 2019.

[2]. D. J. Wagg, K. worden, R. J. Barthorpe, and P. Gardner, "Digital twins: State-of-the-art and future directions for modeling and simulation in engineering dynamics applications," ASCE-ASME J. Risk Uncertainty Eng. Syst. B Mech. Eng., vol. 6, no. 3, p. 030901, 2020.

[3]. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, 1997.

[4]. D N Keerthana , D Jayasree , V Jyothirmayee , C Jyoshna Reddy , S Ankitha , “ Diagnosis of Tuberculosis using Deep Learning ” ,International Journal of Computational Science and Engineering Research, vol. 1, no. 2, May. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI2P4

[5]. M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in Transdisciplinary Perspectives on Complex Systems, Cham: Springer, 2017, pp. 85–113.

[6]. F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in Proc. IEEE ICDM, pp. 413–422, 2008.

[7]. G Devika , P Gangadhara Reddy, “Automatic Traffic Signal Controlling for Emergency Vehicles using Internet of Things ” ,International Journal of Computational Science and Engineering Research, vol. 1, no. 3,p. 68, Septmber. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI3P14

[8]. R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA: MIT Press, 2018.

[9]. Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, "Machinery health prognostics: A systematic review from data acquisition to RUL prediction," Mech. Syst. Signal Process., vol. 104, pp. 799–834, 2018.

[10]. A. Vaswani et al., "Attention is all you need," in Proc. NeurIPS, vol. 30, 2017.

Declaration

Conflicts of Interest: The authors declare no conflict of interest.

Author Contribution: All authors wrote the main manuscript text and also consent to the submission.

Ethical approval: Not applicable.

Consent to Participate: All authors consent to participate.

Funding: Not applicable, and No funding was received

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Personal Statement: We declare with our best of knowledge that this research work is purely Original Work and No third party material used in this article drafting. If any such kind material found in further online publication, we are responsible only for any judicial and copyright issues.

Acknowledgements

We thank everyone who inspired our work.

How to Cite:

Tarun Reddy Nukala ,“ Digital Twin Architecture for Real-Time Structural Health Monitoring in High-Speed Motorsport Vehicles “, International Journal of Computer Science , Engineering and Artificial Intelligence , Vol. 2, Issue. 3 (July – Septmber) , 2025 , Pages: 8 – 14 . DOI : https://doi.org/10.63328/IJCSEAI-V2RI3P2