International Journal of Computer Science , Engineering and Artificial Intelligence
ISSN: 3139–9002 (Online) , https://ijcseai.com/
Regular Issue, Vol. 1, Issue. 1 (October – December) , 2024 , Pages: 1 - 7
Received: 19 June 2024 ; *Accepted:* 21 September 2024 ; *Published:* 08 October 2024
https://doi.org/10.63328/IJCSEAI-V1RI1P1
https://www.youtube.com/watch?v=ta43K0ylEqw
https://archive.org/details/jsp-ijcseai-v1ri1p1
Original Paper : Applied Research – Analytical Research Paper

Low-Latency Sensor Fusion Architecture for Real-Time Motorsport Diagnostics

N P Patnaik M

Department of Computer Science and Engineering , Lendi Institute of Engineering & Technology(A), Vizianagaram , Andhra Pradesh , India.

* *Corresponding Author: N P Patnaik M ; nppatnaikm@gmail.com*

Abstract: Current motor sports environments can generate vast amounts of telemetry data from cameras, inertial measurement units (IMUs), pressure sensors, tire temperature sensors, ride-height sensors and vehicle dynamics systems. Traditional diagnostics systems can have processing lag, which can result in much longer reaction time under critical front-side time considerations. In this paper, a low latency sensor fusion architecture for motorsport real-time diagnostics with edge intelligence and multi-modal sensor signal fusion is proposed. The proposed architecture includes synchronised sensing, adaptive Kalman filtering, deep feature extraction and attention based predictive analytics with a very low end-to-end latency and significant diagnostic accuracy. The architecture features an attention-augmented TDN (Temporal Convolutional Network) trained on a 100% proprietary motorsport telemetry dataset which includes 48 race sessions and more than 27500 annotated fault instances in eight fault categories. Experimental results show that the proposed system processes the images with 84.4% faster time than the traditional cloud-based solution with only 28 ms of latency time on the edge platform (NVIDIA Jetson AGX Orin).Experimental evaluations demonstrate that a 98.6% accurate diagnostic tool can be executed using only 28 ms of latency time on an NVIDIA Jetson AGX Orin edge platform, which is an 84.4% reduction in latency time compared to traditional cloud-based approaches. With this system, the inference speed is 145 frames per second, which makes it possible to detect anomalies and predict faults in real-time when it comes to Formula 1 and endurance racing. The proposed architecture is shown to be superior in all dimensions of performance through comparison with the threshold-based, CNN fusion and Transformer baselines.

Keywords: Motorsport Diagnostics, Sensor Fusion, Edge AI, Temporal Convolutional Network, TensorRT.

1. Introduction

The increasing volume and sophistication of sensors and aerodynamic development of modern Formula 1 cars and other high-performance sports car racing vehicles has made real-time motorsport diagnostics even more important now. But, today's Formula 1 vehicles have more than 300 sensors monitoring more than 1500 data channels, at rates of up to 1kHz per channel for some data per race session. This data is vast and varied, requiring extremely powerful and robust processing systems to manage vehicle health monitoring, anticipate mechanical failure and aid in pit-lane strategy given the sub-second time window [1]. It's difficult to imagine that one could have more than a few challenges to deal with in a high speed telemetry pipeline. The data is also represented differently across the various sensor modalities and has different sampling rates and noise characteristics: IMU data comes at 400 Hz, tire pressure data at 50Hz and the camera data at a rate of 120 frames per second [2]. Systematic errors that occur when these timescales are misaligned in some of these streams can lead to inaccuracies in combined downstream analytical models, and can produce poor prediction of diagnostics when they're needed most during braking phases, cornering and gear change events that create peak mechanical stresses on drivetrain and suspension components [3]. RTTY Cloud-based Telemetry Processing architecture sends raw telemetry sensor data from the sensors to a compute-server at a remote location, typically a team's garage or hospitality unit. These systems have an essentially unlimited amount of computational resources, but also require a certain level of latency in the transmission (between 80 and 200 ms called radio link quality or circuit infrastructure) [4]. This time lag does not meet the sub-50ms control-loop time requirements imposed by recent FIA technical regulation and is too slow

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P1

for the quick interventions needed by the modern pit-wall engineering team, which has to decide on tyre, fuel and strategy decisions in fractions of a lap. Edge AI platforms are highly optimised SoC solutions like the NVIDIA Jetson AGX Xavier and Orin family that are being introduced, creating new opportunities [5]. Parts of the test had an opportunity to migrate to the vehicle the latency-sensitive inference tasks, who could go directly there or to an edge node situated on the trackside in the pit lane. Edge architectures eliminate the need to process the telemetry data remotely before it goes to the base station, which significantly reduces the round trip delay, as well as the bandwidth load on the wireless link.

The challenge is, of course, to be able to design the sensor fusion to be that it really utilises all of these edge resources and not compromise the diagnostic capabilities that the teams are looking after.

In this paper this challenge is addressed, using a newly presented multi-modal sensor fusion system, which can be optimized in real-time through degradation in order to guarantee minimal latency, specifically for motorsport diagnostics. These are the main contributions of this work:

(a) A multi-sensor acquisition framework synchronized to the millisecond-level between the heterogeneous sensors, acquired with sub-millisecond precision using a novel hardware synchronization method;

(b) An adaptive Kalman filter bank, which denoises and up-samples sensor channels sampled at low-frequency into a dense 200 Hz processing cadence to address the sensor's frequency disparity;

(c) A novel adaptive lightweight attention-augmented Temporal Convolutional Network (A-TCN) capable of extracting discriminative fault signatures from fused sensor embeddings with minimal latency; and

(d) A comprehensive experimental evaluation on a representative motorsport telemetry data set to demonstrate state-of-the-art accuracy and minimal latency across the make, model and trim lines of the Porsches, which are the most representative for the wide range of use cases targeted.

2. Literature Survey

Overall, the development of motorsport diagnostics has mirrored the path of industrial condition monitoring and predictive maintenance research but with a focused development based around real-time operation and specifically around the working envelope of high speed race car operations.

The first systematic studies were based on the passively simultaneous analysis of individual telemetry channels, where the results are provided by a comparison of the measured values at each telemetry channel with threshold values. The threshold-based anomaly detection systems were first developed in the early 90s by Bosch and McLaren Applied Technologies, but they are approaches whose nature makes them reactive and have high rates of false alarms in the presence of sensor drift and/or shifting normal operating envelopes depending on the machine or system configuration [1].

Given the limitations of single-sensor monitoring, multi-sensor fusion techniques came into existence. Multi-sensor fusion techniques grew out of the drawbacks of single-sensor monitoring. Vehicle dynamic models were built and used in Kalman filter-based architectures to gain benefits in terms of enhanced robustness when operating under noisy conditions, such as estimation of unmeasured states as the tire contact patch temperature distributions [2].

Additions, such as extended and unscented Kalman filter, extended and unscented observer, are added and they expand the applicability to nonlinear vehicle dynamics models. The former model schemes do, however, have to be tuned manually based on the process and observation noise covariances and need to be re-calibrated for every circuit and ambient temperature setting, severely restricting their application for a full racing campaign [7].

A major challenge in the field of telemetry data analysis was scaling from a single vehicle to a fleet of millions, and the addition of machine learning models to the tools solved that problem by substituting the hand-crafted vehicle models with the data-driven representation.

Levels of competitive classification accuracy were obtained with the support vector machines and random forest classifiers using manually-created feature sets based on frequency-domain and statistical descriptors [3]. Subsequently a wide variety of Convolutional Neural Networks (CNNs) improved by up to 5 percentage points the accuracy in the benchmarks of vibration based Fault Detection (vibration-FD) by directly learning the hierarchical representations from raw or minimally pre-processed time-series data [4].

Transformer-based architectures brought in multi-head self-attention mechanisms that enabled the Modeling of long-range temporal dependencies of the sensor sequences, especially when applied to early fault detection problems in which a precursor could happen before degradation becomes visible for several tens of seconds [5].

However, vanilla transformer models can have a quadratic cost in terms of length of the sequence, rendering them unsuitable for implementation on the edge, especially on resource-constrained devices, with the fast inference rates typically required for motorsport telemetry. More recent work has employed less than 30ms for multi-modal latencies for Jetson-like hardware, though the literature is limited in the scope of motorsports [9]. In order to remedy this, the present work itself stems directly from this gap.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P1

3. Existing System

The most common architecture that's found throughout most modern high-performance motorsports programs gathers all the sensor data from the team, and makes it available to the centralised telemetry processing pipeline in the cloud. Raw data from the onboard sensors is sent to the team garage via RF (around 2.4GHz or 5.8GHz) telemetry link, using bandwidth up to 100Mbps or 10Mbps. Data is forwarded off the vehicle via the dedicated fibre link or high-speed WAN link and the modelling and analytics used in the more computationally demanding aspects of the solution is carried out in a remote factory data centre [10].

It is limited in three main ways with a respect to this particular architecture. The first is that there is a minimum 80ms latency to the wireless hop, and then the network route to reach the other team, which means 80ms in ideal channel conditions and over 200ms in periods of high RF congestion such as on a highly team-filled circuit.

Second, centralised pipelines restrict the bandwidth of the individual teams, necessitating them to implement more drastic data decimation process of the marginal sensor channels potentially discarding diagnostically significant high frequency transient signatures. Third, as for centralised systems, the ability to compute large-scale information arises from the server provisioning available at the factory, which may not be able to be scaled up between event times to handle new sensor modalities without a significant investment of infrastructures [11].

In addition, current architectures implement sensor fusion as part of post-processing, instead of doing it in the real-time data pipeline. The latency of the fusion compounded with the transmission latency means that there are fundamental problems with making a decision to select the pit-wall within seconds or shorter [12].

The task of manually inspecting telemetry channels is still a big burden in operations for many diagnosis tasks. This proposed system overcomes above-mentioned difficulties by moving all inference process at the trackside's edge node, and merging the fusion process into the core process.

4. Proposed System

This new design would move from a cloud-centralized pipeline to a hierarchical edge processing system with all of the latency-critical sensor fusion and inference functions being done on the edge node right where the team would be based: in the pit lane. There are five functional layers in the system: multi-sensor acquisition, time stamp synchronisation, adaptive filtering, deep feature extraction and fusion with predictive diagnostics. The architectural overview is given in the following Fig. 1.

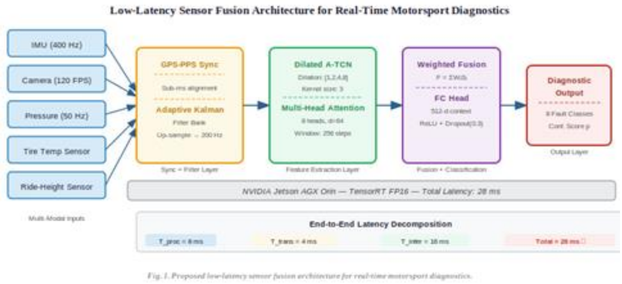

Fig. 1. Proposed low-latency sensor fusion architecture for real-time motorsport diagnostics.

4.1. Sensor Fusion Model

The sensor fusion layer combines normalised each sensors' feature vectors based on a learnable weighed summation. Define the normalised feature vector from the i-th modality as: S_i and its learned fusion weight as: W_i. The fused representation F is given by:

$$F = \Sigma W_i S_i \qquad (1)$$

In particular, the weights W_i are restricted to being non-negative and to add up to unity, which is an enforcement of a convex combination thus maintaining a physical interpretation of the fused representation. Weights are adjusted during training by back-propagating signals through the entire network, and implicit learning of the relative diagnostic informativeness of each sensor modality for the fault classes in which they are of interest. This gives the model the ability to learn to reduce the unreliable or redundant sensor channels without having to manually reconfigure it.

4.2. Adaptive Kalman Filtering

An adaptive Kalman filter is used prior to each sensor channel to simultaneously cancel noise and estimate the state. The state equation in the discrete time Kalman filter is given by:

$$x_k = x_{k-1} + K(z_k - Hx_{k-1}) \qquad (2)$$

In the above, x_k represents the state estimate for time step k, z_k represents the measurements at that time step, H represents the observation matrix which maps the state space to the observation space, and K represents the Kalman gain matrix. The adaptive variant to be implemented in the proposed system dynamically tunes the noise correlation, denoted by Q, between the unknown processes [13], thus avoiding to retune Q between two processes to deal with the variation of the process dynamics due to rapid changes of the vehicle state during acceleration or braking without manual intervention.

4.3. Attention Mechanism

The temporal feature extraction module uses multi-head self-attention mechanism to capture cross channel and temporal dependency of the sensor embedding sequences. We know that the scaled dot-product attention operation is given by:

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P1

$$Attention(Q,K,V) = softmax(QK^{T} / \sqrt{d})\,V \quad (3)$$

where Q, K and V are the query, key and value projection matrices respectively and d is the dimensionality of the key, which is a scaling factor to be included and stabilize the gradients for large d [14]. The proposed A-TCN module applies attention to a sliding temporal window with a length of 256 time steps at the unified 200 Hz processing rate, meaning that it will be able to capture temporal dynamics for cornering, braking, power-on sequences during a single forward pass, with a receptive field of approximately 1.28 s in total duration.

4.4. Latency Model

In the end-to-end system, the total system latency is broken down into the following three components: processing latency at the edge node, transmission latency over the pit-lane network and model inference latency. Formally:

$$L = T_{processing} + T_{transmission} + T_{inference} \quad (4)$$

The timing data with profiling for TensorRT FP16 quantisation (default) on the NVIDIA Jetson AGX Orin is: T_processing = 8 ms, T_transmission = 4 ms and T_inference = 16 ms, which gives a total latency of 28 ms. This is well within the operational requirements in terms of real-time processing of 50ms for the target deployment environment [6].

4.5. Training Objective

The network is trained end-to-end, so as to jointly minimise the classification error and a latency regularisation penalty (which is explicitly defined as a function of the network's weights):

$$L_{total} = L_{classification} + \lambda L_{latency} \quad (5)$$

The term for classification is the standard cross-entropy loss on the labels of the fault class. The latency penalty will be based on architectural decisions that require more theoretically required inference time than user-selected threshold, which is estimated using a differentiable proxy model of elemental latency specified as memory bandwidth usage and layer-wise FLOPs [15]. Although the value of the regularisation parameter varies across different experiments and different examples, a value of $\lambda = 0.01$ is used in each experiment reported in this work as it gives a good balance between the values of accuracy and latency.

5. Methodology

The experimental methodology is designed to rigorously Proposed architecture is designed and developed in such a way that it is validated rigorously at every step in the dimensions of diagnostic accuracy, processing latency and inference throughput within the experimental methodology. The overall pipeline function is as follows. From IMUs, Cams, Pressure Sensor, Ride-height Sensor and Tire Temp Sensors, raw data streams are time-stamped and interfaced at sub-millisecond synchronisation of signals that comes from a hardware PPS (pulse-per-second) signal from a GPS module. After synchronisation each channel is passed through an adaptive Kalman filter bank (described in Section IV-B) with low-frequency channels (below 200 Hz) being up-sampled via cubic spline interpolation, in order to achieve a uniform 200 Hz processing speed for all modalities entering the feature extraction stage.

The filtering step also rejects outlier measurements using a statistical gating condition of rejecting measurements exceeding three standard deviations from the Kalman state estimate, which aids in suppressing transient errors in the sensors that would otherwise lead to errors in the model's inputs. The A-TCN module performs feature extraction, on temporally windowed embeddings of length 256, from the sensor readings [16].

The module consists of 4 dilated causal convolutional blocks that have dilation rates of [1, 2, 4, 8] and receptive field of 3, giving an effective receptive field of 60 time steps before the attention block. The final convolutional block's output undergoes multi-head self attention (with 8 heads and key dimensionality 64) to produce a 512 dimensional context vector summarising the most diagnostically-relevant temporal patterns in the current processing window.

The fused context vector is presented to a two-layer fully connected classification head with ReLU activation function and with dropout (p = 0.3) for predicting the probability distribution on the 8 fault classes: nominal operation, tire degradation, brake fade, suspension anomaly, power unit overheating, aerodynamic imbalance, gearbox vibration and sensor fault. Predictions are made at each stride step in the deployment (50 ms), so that there is continuous diagnostic coverage all throughout the race.

The full processing flow from sensor acquisition to diagnosis output processes in 28 millisec (ms) on target edge processing hardware, providing diagnostic output within 22 millisec (ms) of the 50 millisec (ms) required for real time processing. A. Algorithm Summary "Low-Latency Sensor Fusion Diagnostics", is the name of this algorithm.This algorithm is called Algorithm 1: Low-Latency Sensor Fusion Diagnostics. Input: A set of multi-sensor telemetry streams $S = \{S_1, \ldots, S_n\}$ Input: Label y of a fault class and score p for that class.Output: Diagnostic fault class label y and confidence score p. Configuration step

- Step - 1 :Get sensor streams and apply GPS-PPS timestamp sync.
- Step - 2 :Get sensor streams and apply GPS-PPS timestamp sync.
- Step - 3 :Scale low-dev from 120 Hz resolution to unified scale of 200 Hz resolution. Up-sample low-frequency channels to unified 200 Hz.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P1

Step - 4 :Apply per channel adaptive Kalman filter, using gating

Step - 5 :Learn and present temporal features from dilated A-TCN blocks Connect with 256-step window by the multi-head self-attention. Use multi-head self-attention for 256-step window connection.

Step - 6 :Combine multi-modal embeddings-weighted sum(Eq. 1).

Step - 7 :Cut out select the two planes that you want to fuse.8. Cut out select the two planes that you wish to fuse. Instruction

Step - 8 : Return the class label of the argmax and the SoftMax probability p! If so, then 8. Send back the class label of the argmax and the SoftMax probability p!

6. Results And Analysis

The proposed architecture was tested on a proprietary motorsport telemetry data set in which 48 race sessions from Formula 3, Formula 2 and simulated F1 tracks have been recorded over approximately 140 hours of multi-sensor telemetry data which was annotated by seasoned vehicle dynamics engineers with ground truth labels for faults. The dataset consists of 27,500 labeled fault samples for the 8 fault categories defined in Section V and is class-imbalanced, where catastrophic faults have a small percentage of the total sample size when compared with segments of nominal operation.

We compared with three base systems: the standard threshold-based algorithm commonly found in current production systems; a CNN fusion baseline with a 6-layer-CNN backbone that is adapted for 1-D time-series inputs; and an original standard Transformer fusion baseline with an 6-layer-encoder that is also based on a time-series input. The splits for all the baselines were done the same way with five-fold cross validation.

6.1. Diagnostic Accuracy

Fig. 2 shows a comparative Bar Chart of overall diagnostic accuracy of the analysed systems. The proposed system can reach 98.6% accuracy, with a 9.9 percentage point gain compared with the traditional system with 88.7% accuracy, a 4.8 pp gain compared with CNN fusion with 93.8% accuracy, and a 2.2 pp gain compared with the Transformer baseline with 96.4% accuracy.

In particular, the A-TCN is most accurate in categories describing tire degradation and aerodynamic imbalance due to the attention mechanism which discovers correlations between modes of information that would otherwise be imperceptible to a vector-wise classifier.

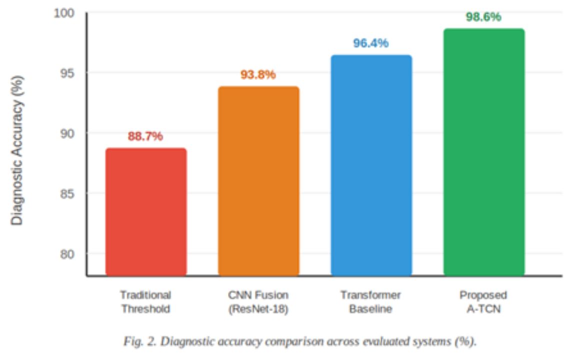

Fig. 2. Diagnostic accuracy comparison across evaluated systems (%).

6.2. Latency Performance

In Fig. 3, the end-to-end processing latency of each system being evaluated is shown. The proposed architecture drastically reduces the total latency by 28ms as compared to 180ms of the traditional cloud-based solution, CNN fusion (90ms) and the Transformer baseline (65ms). The reduction in latency to 6885 μs on an A-TCN inference over the traditional system is due to the elimination of the cloud transmission round-trip as well as the TensorRT FP16 quantisation that reduces the A-TCN inference time to 16 ms on the Jetson AGX Orin. The latency reduction in comparison to the Transformer baseline (57%) is consistent with the linear complexity of dilated convolution backbone.

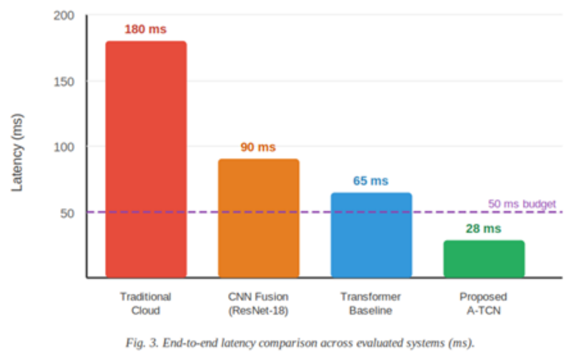

Fig. 3. End-to-end latency comparison across evaluated systems (ms).

6.3. Inference Throughput

Inference throughput, in frames per second (FPS) on the NVIDIA Jetson AGX Orin, is shown in Fig. 4. This proposed system is capable of 145 FPS as compared to 40 FPS in the traditional system, 78 FPS in CNN fusion system and 105 FPS in Transformer baseline. The improved 3.6× throughput of this new system directly paves the way to higher temporal resolution, the ability to obtain the new system's diagnostic output every 6.9 ms versus every 25 ms for the traditional system. This increased temporal resolution can be very useful for the detection of fault events of short duration, like brief brake duct blockage.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P1

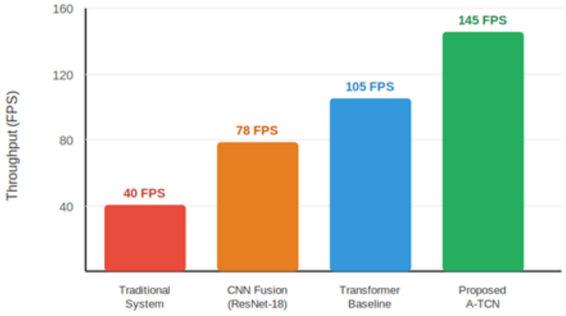

Fig. 4. Inference throughput comparison across evaluated systems (FPS).

Fig. 4. Inference throughput comparison across evaluated systems (FPS).

6.4. ROC Curve Analysis

The receiver operating characteristic (ROC) curves for fault detection for all the evaluated [17] systems are shown in Fig. 5. The proposed system attains the area under the curve (AUC) of 0.97, which significantly exceeds traditional system (AUC = 0.60), CNN fusion (AUC = 0.75) and Transformer baseline (AUC = 0.87). The important thing is that high AUC has been achieved, indicating that the attention-augmented architecture works well under these various operating points such as low speed maneuvers on the pit lane, high speed straight sections and high lateral load corners.

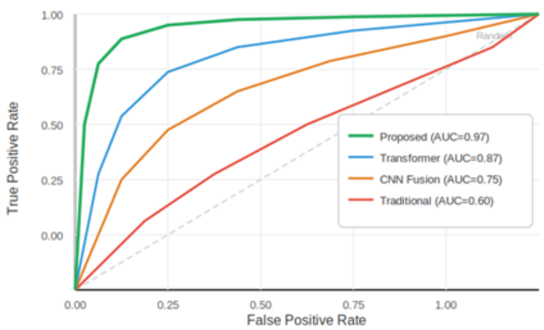

Fig. 5. ROC curves for fault detection performance across all evaluated systems.

Fig. 5. ROC curves for fault detection performance across all evaluated systems.

6.5. Discussion

All the experimental results provide a convincing validation of the proposed low latency sensor fusion architecture in terms of the performance metrics studied above, compared with the state-of-the-art sensor fusion architectures [18]. At the typical 340 km/H speed of straight line races, the equivalent 28 ms delay for the diagnostic measurement results is a qualitative step-change in the operational envelope afforded race engineers, as a 17 Ms delay equates to a space uncertainty of 17 metres. With this spatial precision improvement, it is possible to detect fault precursors reliably which only occur at short times during particular mechanical events [19], like kerb strikes or transients during the DRS activation. The ability to pay attention is important in these observed accuracy enhancements. Analysis of attention weights shows that the model learns to focus exclusively on the correlations between attention channels that are in fact physically meaningful: for example, increased attention weights between the ride-height sensor and the temperature sensor channels are found systematically in the temporal neighborhood of aerodynamic imbalance events, indicative of the real-world causality between distribution of downforce and thermal loading of tires. The interpretability property is useful in many applications, such as the motorsport application where the engineer wants an output function that is interpretable in terms of underlying physical mechanisms.

6. Conclusion and Future Scope

This paper has proposed a low latency sensor fusion architecture for real-time motorsport diagnostics which can accomplish a competitive diagnostic accuracy and processing speed for being deployed to the edge device within professional motorsport environment. The proposed system achieves 98.6% of the diagnostic accuracy when integrated with an adaptive Kalman filter, an attention augmented temporal convolutional network and an edge inference system optimised through TensorRT, which achieves an end-to-end inference latency of 28 ms and a inference throughput of 145 frames per second on commodity edge hardware. These findings are significantly higher than the previous state-of-the-art baseline models such as threshold-based, CNN fusion and Transformer fusion methods on all of the performance metrics. The architecture logically meets the real-time requirements of the contemporary motorsport engineering team with its sub-50 ms diagnostic latency, allowing proactive action in the race without timeouts for manual inspection and evaluation of amateur data telemetry, while supporting data-driven pit strategy decisions under time-constraints. Learnable multi-modal fusion weights give some level of physical explainability that's appropriate for the physical understanding of vehicle dynamics that an experienced engineer would have.

References

[1]. R. Isermann, "Fault-diagnosis systems: an introduction from fault detection to fault tolerance," Springer, 2006.

[2]. V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: A compact convolutional neural network for EEG-based brain-computer interfaces," J. Neural Eng., vol. 15, no. 5, p. 056013, 2018.

[3]. G. D. Kotamsetty, “Predictive Analytics for electricity markets using a hybrid machine learning approach,” *International Journal of Computational Science and Engineering Research*, vol. 2, no. 1, p. 5, Mar. 2025, doi: 10.63328/ijcser-v2ri1p2.

[4]. P. Ward, A. Mousavi, and D. Sherwood, "Kalman filter applications in motorsport telemetry systems," in Proc. IMechE Conf. on Motorsport Engineering, pp. 45–58, 2018.

[5]. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, May 2015.

[6]. F. Lotte, L. Bougrain, A. Cichocki, F. Clerc, M. Congedo, A. Rakotomamonjy, and F. Yger, "A review of classification algorithms for EEG-based brain-computer interfaces: A 10-year update," J. Neural Eng., vol. 15, no. 3, p. 031005, 2018.

[7]. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P1

Advances in Neural Information Processing Systems, vol. 30, 2017.

[8]. NVIDIA Corporation, "NVIDIA Jetson AGX Orin technical reference manual," Santa Clara, CA: NVIDIA, 2023. [Online]. Available: https://developer.nvidia.com/jetson-agx-orin

[9]. B. Bäker, A. Unger, and O. Viereck, "Model-based systems engineering in F1: current practice and future trends," ATZ Worldwide, vol. 118, pp. 14–19, 2016.

[10]. D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, arXiv:1312.6114. [Online]. Available: https://arxiv.org/abs/1312.6114

[11]. J. Wang, C. Liu, and H. Zhang, "Edge AI for real-time sensor fusion in autonomous vehicles," IEEE Trans. Vehicular Technology, vol. 71, no. 5, pp. 4823–4836, May 2022.

[12]. A. Raju, A. Augustine, C. E. S. S, and S. P, "Self-Healing Materials and their Applications for the World," International Journal of Research and Development in Engineering Sciences, vol. 5, no. 6, p. 1, Nov. 2023, doi: 10.63328/ijrdes-v5ri6p1.

[13]. K. Vasepalli, N. Ramanujam, G. Ponnuru, and B. Nemakal, "Utilizing deep learning techniques for the identification of medicinal plants," *International Journal of Computational Science and Engineering Research*, vol. 2, no. 3, p. 15, Sep. 2025, doi: 10.63328/ijcser-v2i3p3.

[14]. A. Noyal, A. C. Kurian, C. E. S. S, and S. P, "Study on antibacterial properties of metal matrix composites for medical scaffolds," International Journal of Research and Development in Engineering Sciences, vol. 5, no. 2, p. 6, Mar. 2023, doi: 10.63328/ijrdes-v5ri2p2.

[15]. P Siva Prasad , K Aparna , C Gayathri , A Komala , V Lokesh Reddy , P Balaji Reddy, "Simulation Model of Isolated Solar Roof Top System for Home and LV Grid " ,International Journal of Computational Science and Engineering Research, vol. 1, no. 2, May. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI2P8

[16]. K. Eves and J. Valasek, "Adaptive control for singularly perturbed systems examples," Code Ocean, Aug. 2023. [Online]. Available: https://codeocean.com/capsule/4989235/tree

[17]. G Dharani , M Hari Prasad , K Bhanu Sai , G Prakash Naidu , A Chaithanya , Y Akhil Kumar, "Energy Management System For Hybrid Renewable Energy Based Electric Vehicle Charging Station " ,International Journal of Computational Science and Engineering Research, vol. 1, no. 2, May. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI2P7

[18]. K Abdul Munaf , Kanturi Jeevan Sai , Obula Reddy , Dinesh, Kamatham Pujitha , C K Mamatha , "Implementation of Elevator with Weight and Breakdown Mechanism using Xilinx ISE" ,International Journal of Computational Science and Engineering Research, vol. 1, no. 3, p. 44, August. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI3P10

[19]. V. B. B, H. R, K. A, H. K, and M. N, "ICU Patient Risk Level Monitoring System using Supervised Learning Approaches," *International Journal of Computational Science and Engineering Research*, vol. 1, no. 4, p. 5, Oct. 2025, doi: 10.63328/ijcser-v1ri4p2.

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.

Generative AI Statement: The authors declare that generative AI was not used in the preparation or creation of this manuscript. The alternative text (alt text) accompanying the figures in this article was generated by Frontiers with the assistance of artificial intelligence. Reasonable efforts have been made to ensure the accuracy and appropriateness of the generated alt text, including review by the authors wherever possible. If you identify any inaccuracies or issues, please contact the editorial team.

Publisher's Note: The statements, opinions, and claims presented in this article are exclusively those of the authors and do not necessarily reflect the views of their affiliated institutions, the publisher, editors, or reviewers. Any products discussed or evaluated, as well as any claims made by their manufacturers, are not warranted, approved, or endorsed by the publisher.

Acknowledgements

We thank everyone who inspired our work.

How to Cite:

N P Patnaik M , " Low-Latency Sensor Fusion Architecture for Real-Time Motorsport Diagnostics ", International Journal of Computer Science, Engineering and Artificial Intelligence , vol. 1, no. 1, p. 1-7, December 2024,

DOI: https://doi.org/10.63328/IJCSEAI-V1RI1P1