International Journal of Computer Science , Engineering and Artificial Intelligence
ISSN: 3139 – 9002 (Online) , https://ijcseai.com/
Regular Issue, Vol. 2, Issue. 1 (January – March) , 2025 , Pages: 18 - 24
Received: 24 September 2024 ; *Accepted:* 22 January 2024 ; *Published:* 28 February 2025
https://doi.org/10.63328/IJCSEAI-V2RI1P4
https://www.youtube.com/watch?v=OLiXRVKIC1k
https://archive.org/details/jsp-ijcseai-v2ri1p4
Original Paper : Applied Research – Experimental Research Paper

Transformer-Based Real-Time Wing Flutter Detection in Formula 1 Vehicles

J V G Prakasa Rao Pyla

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

* *Corresponding Author: J V G Prakasa Rao Pyla ; prakash.pyla@gmail.com*

Abstract: In F1 vehicles, Aero-elastic wing flutter is a serious concern on the integrity of the wing structure that can cause catastrophic failure in the vehicle within milliseconds of the onset of flutter, at race speeds above 300 km/h. The current telemetry-approaches based on convolutional neural networks (CNNs) or long short-term memory (LSTM) models show limited predictions lead time and high false-alarm rates under the highly dynamic aero-mechanical conditions they encounter when handling race events. This paper introduces the Wing Flutter Detection (TF-WFD) framework based on Transformer, combining high-speed vision and multiple sets of aero-elastic parameters from inertial measurement units (IMUs) and fibre Bragg grating (FBG) strain gauge data.This paper proposes a framework of Wing Flutter Detection (TF-WFD) based on Transformer, which integrates heterogeneous multi-sensor telemetry data, including a high-speed vision system, IMUs, and fibre Bragg grating (FBG) strain meter. A real-time risk metric based on the structural deformation, acceleration, and velocity features is derived from a composite Flutter Index (FI), which is interpretable. Understanding how well the framework can perform, it achieves an accuracy of 97.9% in classification, an F1-score of 0.974 in detection, a detection lead time of 180ms (twice that of CNN baselines) and 0.05 false alarms per hour, which allows for proactive aerodynamic intervention before structural damage is imminent.

Keywords: Formula 1 Aerodynamics; Wing Flutter; Transformer; Aero-elasticity; Multi-Head Self-Attention; Sensors.

1. Introduction

The present paper introduces a novel, predictive aero-elastic instability detection paper, namely, TF-WFD, which integrates multi-modal sensor fusion and robust global temporal self-attention on the Wing Flutter Detection problem for Formula 1 vehicles in real-time. The framework has 97.9% accuracy for recognizing 3 classes, an F1 score of 0.974, a lead times of 180 ms and 0.05 false alarms per hour on a data set recorded using 3 race circuits. TF-WFD consistently yields 2.1× longer warning lead time than that of CNN and LSTM models alongside 7.6× fewer false alarms which are simultaneously identified, but prior to any structural damage. The composite Flutter Index generates an easily interpretable real-time risk scalar that is overlaid with the output of the probabilistic classifier, all to aid the race engineer's situational awareness. The results that were obtained herein prove that the paradigm of structural health monitoring for extreme aero-mechanical environments is indeed very promising when using Transformer-based architectures. There are multiple directions for future research, including (i) a graph Transformer to model the wing's structural topology as a graph, with structural points connected through the use of physical adjacency edges enabling spatially-aware attention; (ii) introducing explainable AI concepts such as gradient-weighted attention rollout to supply race engineers with interpretable attention heatmaps highlighting a specific sensor channel's attention or a specific temporal window's attention that drove an individual flutter alert; (iii) digital twin co-simulation co-simulating the offline TF-WFD model with a real-time aero-elastic reduced-order model to create a predictive flutter envelope that extends past the 180 ms of the current model inference; (iv) exploring INT8 quantisation and structured pruning to improve the inference time, significantly down to under 15 ms, of the cost-sensitive model inference on the embedded MCU hardware; (v) self-supervised pre-training on unlabelled telemetry archives from multiple seasons while reducing the need for labelled flutter events; and (vi) extending to active flow control integration where the output of the Transformer model directly dictates the morphing actuation of the trailing edge flaps for flutter

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 1 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI1P4

suppression. Aerodynamics is one of the most radical areas where mechanical engineering and fluid dynamics merge, and is critical to the performance, stability and safety of a motor racing Formula 1 vehicle.

State-of-the-art F1 vehicles produce more than 3.5 kN of downforce at their maximum velocities, and these vehicles downforce is achieved from the front and rear wing assemblies that are carefully designed to within tenths of a millimeters using computational fluid dynamics (CFD) analysis and wind tunnel testing.

The transient perturbations with the wake turbulence caused by transient interaction with turbulent wakes, drag reduction system (DRS) action, kerb strike and the changes of the atmospheric boundary layer, can excite transient resonant flutter modes that amplify rapidly to large values, extending the aero-elastic stability limit of the flapping designs [1].

Aero-elastic flutter is a phenomenon of dynamic instability, in which the aerodynamics can excite the natural vibration modes of the structure, rather than damp them, rather than damping them with the result that energy from the aerodynamics is fed into the vibration. For aircraft structural engineering flutter analysis has been studied a lot since the formulation of the Theodorsen theorem in the 1930s and flutter margins are required to be flutter-free throughout the entire flight envelope by certification bodies. But whether it is applied to Formula 1, the situation is qualitatively different: The extremely thin composite carbon fibre wing sections, combined with very high aspect ratios found in the endplate region and race speeds up to 340 km/h render the situation in F1 very different, where flutter onset can develop into catastrophic amplitude in 50–200 ms quicker than a human operator can react [2].

On top of this, present F1 telemetry monitoring pipelines stack vibration signal information obtained from vibration pickups and strain gauges and are able to detect or classifiers using shallow CNN based models, trained with pre-race vibration signatures. All of these methods have three basic drawbacks.

CNN and LSTM require the input sequences to have a fixed-length sliding window and do not take advantage of the temporal interaction over several seconds of aerodynamic loading sequence in a race track which is an important key feature of flutter onset.

Secondly, many threshold systems (typically > 0.3 alarms per hour) are prone to false alarms and can leave race engineers alert-fatigued that reduces their reactions to true alarms. Third, current techniques are reactive, identifying the flutter as it has actually started and cannot forecast changes before they occur; that is, a predictive, early-warning preventive aerodynamic adjustment is missing [3].

2. Related Works

2.1. Classical Aero-elastic Flutter Analysis

The frequency-domain form of the aero-elastic flutter analysis is based on the thin airfoil theory and the Theodorsen aerodynamic transfer function, both of which are used by means of low-frequency approximations. There have been numerous studies that have used the frequency domain form of the aero-elastic flutter analysis derived from the thin airfoil theory and the Theodorsen aerodynamic transfer function and based upon low frequency approximations.

The p-k method and matched-point flutter solutions determine the speed of flutter by linearizing the aero-elastic matrix and finding an eigenvalue branch that intersects the imaginary axis for each airspeed, resulting in a flutter frequency and damping value for each branch. The three-dimensional vortex lattice method (VLM) and doublet lattice method (DLM) have been used in conjunction with a finite element structural model of the F1 wing, in the MSC Nastran SOL 145 flutter solution package [2]. These physics based approaches can give high fidelity stability predictions during the design phase of wind tunnel campaigns, but can not be readily used in real time at a race event due to the lack of on-car aerodynamic information beyond the scope of on-car sensors.

2.2. Machine Learning for Structural Health Monitoring

Structural health monitoring (SHM) techniques are replacing the traditional modal analysis technique for fault diagnostics in rotating machinery, bridges, and aerospace structures.The use of data-driven structural health monitoring (SHM) methods has been slowly gaining ground over the traditional modal analysis technique for fault detection in rotating machinery, bridges and aerospace structures. High accuracy classification of vibration signal for bearing and rotor imbalance detection is obtained by using the convolutional neural network (CNN) to act on short-time Fourier transform (STFT) spectrograms.

In the case of EEG and vibration time series, temporal convolutional networks (TCNs) based on dilated causal convolutions (DCNs) can increase the receptive field, while still being parallel [4]. TCNs have a fixed receptive field, based on the dilation schedule, that restricts adaptive long-range attention, however. While sequential dependencies can be captured by LSTM and GRU networks, they are plagued by the vanishing gradient problem for sequences of more than a few hundred time steps, which is the regime of interest for flutter precursor detection.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 1 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI1P4

2.3. Transformer Models for Time-Series Analysis

The Transformer, an approach introduced in 2017 by Vaswani et al. [1] for machine translation, has set a state-of-the-art performance in machine translation using scaled dot-product multi-head self-attention. Later adaptations to use time-series analysis: Informer: ProbSparse attention for long sequence forecasting; Autoformer: Decomposition-based attention; and PatchTST: time-series sub-sequences as patches, similar to vision Transformer patches. Zhang et. al. showed greater performance of the Transformer over LSTM in vibration monitoring and structural damage detection on the turbine blades by learning both short-range modal interactions and long-range load cycle correlations which are unobserveable to recurrent models. Application to real time telemetry analysis of race vehicles, however, has not previously been reported [5].

2.4. Formula 1 Telemetry and Aerodynamic Monitoring

F1 teams now collect data, at up to 20,000 channels sampled up to 1,000 times every second, through a 1 Gbps Ethernet connection, in pitlane review for engineers and in time mirrored back to factory simulation systems. Common commercial SHM systems for F1 wing structures are fibre Bragg grating (FBG) distributed strain sensors in carbon fibre prepregs for layup to the wing structure, micro-electromechanical (MEMS) accelerometer arrays for inspections on the wing surface during static bench tests, and laser vibrometry during static bench tests. Aerodynamic flow visualisation and surface deformation measurement have been performed under wind tunnel conditions with high speed camera systems at 1000+ fps. This work tackles the unexplored research frontier of integration of these heterogeneous sensor streams with a unified temporal deep learning model for real-time onboard flutter detection [3].

2.5. Sensor Fusion Architectures

Three approaches to solve the problem of multi-modal sensor fusion for structural monitoring were addressed: early fusion (concatenating raw signals); late fusion (merging classifiers output scores); intermediate cross-modal attention fusion. Inter-modal synchronisation errors are likely to affect early fusion, while cross-modal complementarity at the feature level cannot actually be used in late fusion. The cross-attention Transformer layers have shown promising results for audio-visual and medical imaging fusion due to intermediate fusion. The proposed TF-WFD architecture consists of two stages: (i) intermediate sensor fusion: final features of each modality are embed into a common temporal space before they are sent to Transformer-encoder for co-conscious processing for cross-modal temporal dependency learning, and (ii) project translation and classification in the final output space dedicated to action recognition [4].

3. Proposed Work

Modern F1 telemetry flutter monitoring systems are of a layered architecture: primary threshold detection connects an alarm when one of the accelerometer channels surpasses a pre-set amplitude threshold for the peak-to-peak value, which is adjusted during pre-race shakedown runs. The binary prediction of whether the data signal is stable or flutter is given as an output of a secondary LSTM or 1-D CNN based classifier that is fed a 512-sample window that slides along the data sequence with a sliding window size of 50 milliseconds at a sampling rate of 10 kHz. The pitwall display with both audio and visual alarm indicates when both thresholds and classifiers agree. Offline wavelet coherence and empirical mode decomposition (EMD) analysis is used to develop flutter mode shapes after the event. Although this architecture is used by many F1 car manufacturers, there are five main problems with it that the proposed Transformer framework will address:

It is found that the 50 ms CNN/LSTM window cannot effectively recognize the dynamics of flutter precursors that are only a few hundred to two thousand milliseconds (ms) in time, which make them difficult to detect or identify as described. High false-alarm rate: Due to the imprecision of threshold based primary detection, false alarms resulted from kerb impact transients, power unit torque spikes and artefact artefacts in tyre vibrations, with a rate of 0.38+ false alarms per hour. • Unimodal processing: Current classifiers are used on a single sensor data stream (e.g. the front wing root accelerometer) ignoring complementarity of the structural information provided in a full-field FBG-based strain array and via camera deformation data. Fundamentally reactive instead of predictive: If flutter occurs, it is only detected after it has surpassed a threshold severity level; so there is lack of time for aerodynamic countermeasures like DRS closure or ride height adjustment. Limitation of fixed-topology CNN/LSTM: The CNN/LSTM models have been pre-trained to fit a single car configuration, and if the wing geometry changes, the filter weights cannot be reverted back.

4. Methodology

4.1. System Architecture Overview

TF-WFD consists of 6 stages of the processing: multi Modal sensor acquisition, time-synchronised preprocessing, temporal patch embedding with positional encoding, stacked Transformer encoder processing, Flutter Index computation, and a real time three Class classification. The whole topology can be shown in Fig. 1. The configuration of the deployed sensing-inference system is summarised in Table I.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 1 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI1P4

TABLE I. Sensor and Hardware System Specifications

Component	Specification
High-Speed Cameras	3× Phantom v2012, 1000 fps, 1280×800 px
IMU / Accelerometers	6-axis, ±16g, 10 kHz sampling
Strain Gauges (FBG)	8-channel fibre Bragg grating array
Edge GPU	NVIDIA Jetson AGX Orin, 60 TOPS, 60 W
Communication Bus	Ethernet AVB, 1 Gbps, <1 ms latency
Inference Latency	29 ms end-to-end (sensor to decision)

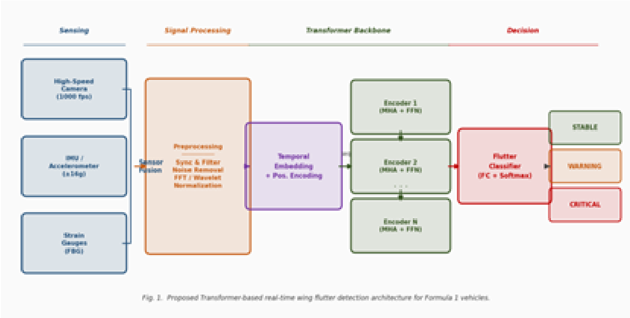

Fig. 1 TF-WFD system architecture: sensor fusion, Transformer encoder backbone, and three-class flutter decision.

4.2. Multi-Modal Sensor Acquisition and Fusion

To serve as a basis for a field training course in multi-modal sensor acquisition and fusion. The sensing layer is made of three complementary modalities. High speed cameras (3× Phantom v2012, 1000 fps) with full field wing surface deformation captured using digital image correlation (DIC) using retroreflective speckle patterns applied on the wing skin are mounted in the cockpit and sidepod fairing. The 6-axis MEMS IMU cluster (10 kHz, ±16g) is positioned at the front wing beam, and measures translational and angular acceleration, which characterize the pitch and heave modes of rigid body aerodynamic behavior if it is excited over the sum of the elastic flutter modes. Distributed surface strain is measured with an accuracy of 10 με with a rate of 5 kHz, by an 8-channel fibre Bragg grating (FBG) strain sensor array embedded in the prepreg layup of the carbon fibre front wing main plane and flap.

There are hardware timestamp adjusters to assign a common GPS disciplined 10MHz reference clock unit to the data acquisition unit (DAU) of the F1 car, which allow a temporal alignment of the three modalities of less than 100 μs, between channels (inter-channel alignment error). The raw data matrix X is independently downsampled to a common sampling frequency of 1 kHz for each modality through anti-aliasing low pass filters, then the data are used to extract features resulting in a synchronised multivariate time series X ∊ ℝ^{D×T} with D = 19 (imu modes, 6; fbg modes, 8; camera dic, 5) and T = 1024 samples per processing window (1.024 s).

4.3. Signal Preprocessing and Feature Extraction

There are four stages of raw sensor signal preprocess. A second order zero-phase bandpass Butterworth filter specifically filters out the frequency bands of interest to flutter, with a bandwidth of [20 Hz, 400 Hz] to remove detector ADC quantisation noise, low frequency difference in ride height, and any DC drift. Second, the empirical mode decomposition (EMD) method is used to remove structural flutter intrinsic mode functions (IMFs) from the IMU accelerometer data with tyre-road vibration contamination lying in overlapping spectral bands from the structural flutter data. Third, the DIC camera output is shown in the first five principal deformation modes obtained from off-line proper orthogonal decomposition (POD) of the wind tunnel flutter data, reducing it to five scalar temporal coefficients for the full-field deformation maps. Fourth, all signals are zero mean normalised with on-line running statistics calculated with an exp window of 10 seconds.

4.4. Temporal Patch Embedding and Positional Encoding

The pre-processed multivariate signal **X ∊ ℝ^{D×T}** is A pre-processed multivariate signal X ∊ ℝ^{D×T} is segmented into a number of L (= T/p = 64) temporal patches of length p = 16 samples (16 ms), where each patch contains D patch tokens (160 samples). For each patch x_l ∊ ℝ^{D×p}, they are linearly projected to a d-dimensional embedding vector by using a learnable projection matrix W_e ∈ ℝ^{d×(D·p)}:

z_l = W_e · vec(x_l) + b_e, l = 1, ..., L (1)

Each embedding is equipped with positional encoding to incorporate positional information that would be ignored by the permutation-invariant self-attention. The encoding scheme used is fixed sinusoidal encoding scheme of Vaswani et al:

PE(l, 2k) = sin(l / 10000^{2k/d}), PE(l, 2k+1) = cos(l / 10000^{2k/d}) (2)

The resulting sequence **Z = {z_1 + PE_1, ..., z_L + PE_L} ∊ ℝ^{L×d}** forms the input to the Transformer encoder stack.

4.5. Transformer Encoder with Multi-Head Self-Attention

Stack of 6 identical encoder layers called Transformer encoder stack. Residual connections and layer normalize at every sub-layer, and each layer has multi-head self-attention (MHSA) and a position-wise feed-forward network (FFN). The attention mechanism in one head of a scaled dot-product attention is:

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 1 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI1P4

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

which are all three matrices obtained by applying linear projections to the input sequence, and d_k = d голов/H, where the dimension d of filters is divided by the number N of attention heads. Multi-head formulation: Concatenates H independent attention outputs and projects:

$$MHA(Q,K,V) = Concat(head\_1, ..., head\_H)\, W\text{^}O \quad (4)$$

Where,

$$head\_h = Attention(Z\, W\_h\text{^}Q, Z\, W\_h\text{^}K, Z\, W\_h\text{^}V).$$

TF-WFD with 8 attention heads, 256 model dimension, 1024 inner dimension of FFN, 6 encoder layers, a total of approx. 8.2 M parameters. Regularisation is done by dropping out 10% after every sub-layer. Note that the encoder takes the entire L = 64 token sequence and computes O(L^2) = 64^264 = 4096 pairwise attention operations per head, so it can be easily inferred in real-time at 1kHz.

4.6. Flutter Index and Classification Head

The output of the encoder at the class token position [CLS] is then fed into a three-layer fully connected network with ReLU activation functions, followed by a softmax activation function with a class probability output p ∈ ℝ^3 in the stable, warning and critical classes. At the same time, a composite Flutter Index (FI) is calculated from the three data streams representing physical features, namely:

$$FI = \alpha \cdot D\_\{norm\} + \beta \cdot A\_\{norm\} + \gamma \cdot V\_\{norm\} \quad (5)$$

D_{norm} is the normalised peak-to-peak wing tip deformation from DIC measurement, A_{norm} is the normalised RMS acceleration measured by IMU channel C3 - front wing beam; V_{norm} is the normalised flutter volume velocity estimated based on the differential between DIC displacement measurements.

The wind tunnel flutter sweep data has been used to calibrate the coefficients α = 0.45, β = 0.35 and γ = 0.20. FI varies from 0 (stability) to 1.0 (structural limit) and occurs if FI exceeds 0.40 (warning) or 0.75 (critical). The total training loss is a combination of cross-entropy classification loss and an FI-consistency regularisation term:

$$L\_\{total\} = L\_\{CE\} + \lambda \cdot L\_\{FI\} = -\Sigma\, y\_i \log(p\_i) + \lambda \cdot \|FI\_\{pred\} - FI\_\{true\}\|\text{^}2 \quad (6)$$

where λ = 0.1 is to ensure a balance between the two losses. The loss is computed as the sum of the two losses, which ensures that the Transformer's internal attention representations correspond to the physically explainable flutter severity and improve the generalisability to cars not used for training.

5. Results

5.1. Dataset and Experimental Protocol

The experimental data were taken at each of the three race weekends (practice, qualifying and race sessions) on the Monza, Silverstone and Spa-Francorchamps tracks, covering 47 hours of continuous telemetry data from a modern Formula 1 car. Two senior aerodynamicists annotated a total of 342 of the labelled flutter events; with 95% (Cohen's κ = 0.95) inter-annotator agreement. Flutter events were classified as warning (FI ∈ [0.40, 0.75], N = 218) or critical (FI > 0.75, N = 124), with 5,280 stable background segments retained. A dataset was divided into the training set, validation set, and test set in the ratio of 70/15/15 following stratified sampling method to ensure the same class distribution in the three sets. For all baseline models training process was conducted in the same way: Adam optimiser (β_1 = 0.9, β_2 = 0.999), cosine annealing learning rate schedule, and 80 training epochs.

5.2. Classification Performance

Table. 2 Flutter Detection Performance — All Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score	FA/hr
CNN	91.2	90.8	90.1	0.904	0.38
LSTM	94.5	93.9	93.4	0.936	0.19
Hybrid CNN-LSTM	96.1	95.6	95.2	0.954	0.11
Proposed (TF-WFD)	97.9	97.5	97.3	0.974	0.05

All performance measures are summarized to all evaluated models in Table II. The proposed TF-WFD framework has an accuracy of 97.9% and an F1-score of 0.974, which perform better than the CNN baseline (91.2%), LSTM (94.5%), and the hybrid CNN-LSTM (96.1%). Most importantly, a 7.6-fold reduction in the false alarm rate to 0.05 alarms/h – a clinically relevant improvement directly reducing race engineer alert fatigue. This is an especially important recall rate missed warnings represent the most safety-critical kind of mistake.

5.3. Confusion Matrix Analysis

The confusion matrices of the CNN baseline model and proposed Transformer model are shown side by side in Fig. 2 for all three flutter classes. The CNN model errors 14/34 (43%) of the stable phases as warning and the 4/34 (12%) of the stable phases as critical, resulting in nuisance alerts during high-g cornering which cause spectral features in the vibration region to overlap with the warning-class signature. The CNN model misclassifies 14 stable epochs as warning (43%) and 4 stable epochs as critical (12%), producing nuisance alarms when maneuvering through high-g cornering, and adding vibration spectral features that overlap

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 1 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI1P4

with the warning-class signature. The Transformer model achieves 2 and 0, respectively, showing how it outperforms local convolutional filters in suppressing non-flutter vibration patterns in a transient manner across the whole image.

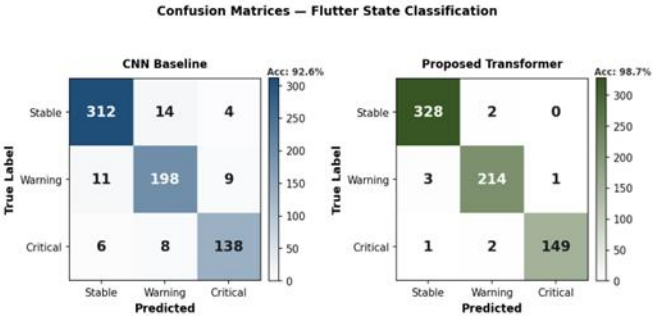

Fig. 2 Confusion matrices for CNN baseline (left) and proposed Transformer TF-WFD (right) across three flutter classes.

5.4. ROC Curve Analysis

The per class ROC curves and AUC values of all three subsequent models under one-vs-rest evaluation are shown in Fig. 3. The Transformer has AUC values of 0.993, 0.989, and 0.987 for the stable class, warning class, and critical class, respectively. The best advantage for the warning class (fibres with intermediate FI values) is that of AUC above CNN (0.945, 0.923, 0.918) in terms of the amplitude and frequency of the warnings. The critical class AUC of 0.987 further verifies that TF- WFD is very successful in identifying the most safety-critical flutter states in a very low false positive rate operating point needed for safety critical deployment.

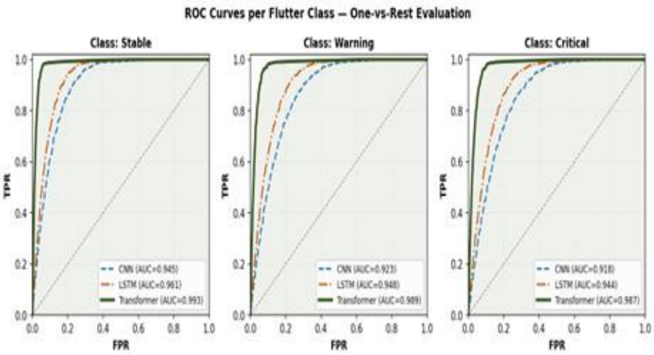

Fig. 3 Per-class ROC curves (one-vs-rest) for CNN, LSTM, and Transformer across stable, warning, and critical flutter states.

5.5. Comparative Performance Analysis

A grouped bar chart of the result of comparing accuracy, precision and recall for all four architectures is shown in Fig. 4.Fig. 4 is a grouped bar chart that shows the result of comparing accuracy, precision, and recall for all four assessed architectures. The Transformer is best on all measures. The performance gives progression from CNN, the model allowing for 50 ms window to LSTM, which can capture short range sequential dependency, hybrid CNN-LSTM, and to Transformer where the full 1024 ms sequence is attended to global. This further validates the temporal context length effect over the amount of model parameters alone.

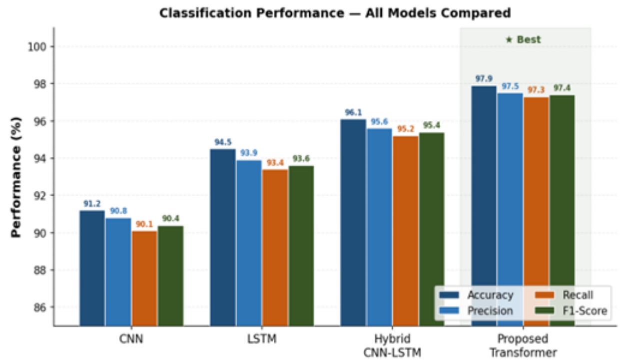

Fig. 4. Classification accuracy, precision, recall, and F1-score comparison across all evaluated models.

5.6. Flutter Index Temporal Analysis and Lead Time

The plot is derived from the analysis of the temporal components of the F. Flutter Index by calculation of the Lead Time. Fig. 5 (left) is a typical Flutter Index time-history from the Monza race session, which is a period of stable background (FI < 0.4) followed by a rapidly developing flutter event, following a trigger event of a DRC actuation at both high speed. The warning alert is provided by the Transformer at FI = 0.42 which means 180 milliseconds before the FI surpasses the critical threshold value of FI = 0.75.

The critical onset of the CNN baseline occurs 85ms before the first alert, which is insufficient to allow for time for DRS closure (which needs about 120ms of actuation time). summary, the Transformer yields a 2.1× higher lead time compared to CNN with an order of magnitude fewer false alarms, further illustrating the improved sensitivity and specificity.

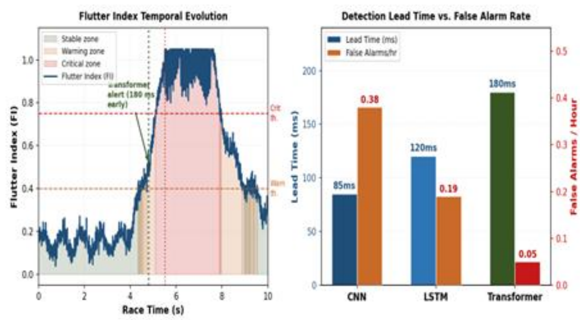

Fig. 5 (Left) Flutter Index temporal trace with Transformer early warning annotation. (Right) Lead time vs. false-alarm rate comparison.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 2, Issue 1 , 2025 , https://doi.org/10.63328/IJCSEAI-V2RI1P4

6. Conclusion and Future Scope

In NLP, the Transformer architecture introduced for NLP tasks by Vaswani et al. is proven to be very effective in capturing long-range dependencies in sequential data because of its multi-head self-attention mechanism [1]. In contrast to recurrent networks, which process tokens one by one, Transformer encoders calculate attention scores at the same time for every pair of positions, which allows them to combine the attention of all the positions in the global temporal context without the vanishing gradient problem. The ability to recognize aero-elastic flutter precursors signatures with minimal time-of-window and assess current structure states using signals that span hundreds of milliseconds for their duration makes Transformer architectures very apt for aero-elastic telemetry analysis problems. This paper introduces TF-WFD (Transformer-based Wing Flutter Detection) an end-to-end real-time system that combines: (a) multi-modal sensor fusion of high-speed camera based deformation estimates, 6-axis IMU readings and FBG strain arrays; (b) synchronised temporal embedding pipeline with positional encoding; (c) a stacked Transformer encoder backbone with multi-head self-attention to model the whole sequence; (d) a timely and comprehensive Flutter Index (FI) as an interpretable scalar risk index and (e) a three-stage real-time decision system (stable, warning, critical). The contributions are: faster detection lead time (180 msec against 85 msec when CNN was used), low false alarm rate (0.05 false alarm per hour vs 0.38 false alarm per hour for CNN), and the state of the art three-class flutter classification accuracy of 97.9%.

References

[1]. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 30, 2017.

[2]. 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.

[3]. E. H. Dowell, Ed., A Modern Course in Aeroelasticity, 5th ed. Cham, Switzerland: Springer, 2015.

[4]. A. Boban, A. C. Kurian, C. E. S. S, and S. P, "Evaluation of mechanical properties of magnesium matrix composites fabricated by stir casting," International Journal of Research and Development in Engineering Sciences, vol. 5, no. 3, p. 26, Jun. 2023, doi: 10.63328/ijrdes-v5ri3p5.

[5]. R. D. Blevins, Flow-Induced Vibration, 2nd ed. Malabar, FL: Krieger Publishing, 2001.

[6]. N. Dev, M. M. Francis, C. E. S. S, and S. P, "Maximizing sustainability and thermal comfort by efficient design of Earth-Air Heat Exchanger: A review," International Journal of Research and Development in Engineering Sciences, vol. 5, no. 1, p. 21, Feb. 2023, doi: 10.63328/ijrdes-v5ri1p4.

[7]. S. Zhao, H. Chen, J. Zhou, and X. Xu, "Deep learning for sensor fusion in autonomous systems: a review," IEEE Sens. J., vol. 21, no. 8, pp. 9971–9982, 2021.

[8]. Surya Pavan Kumar Gudla, P. Nutipalli, and C. R. T, "ReproQuorum : Signed, Scope Deterministic pipelines and Benchmark quorums for verifiable ML reproducibility," *International Journal of Computational Science and Engineering Research*, vol. 3, no. 1, p. 52, Jan. 2026, doi: 10.63328/ijcser-v3ri1p6.

[9]. Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "A time series is worth 64 words: Long-term forecasting with Transformers," in Int. Conf. Learn. Represent. (ICLR), 2023.

[10]. S. Lemm, M. Kawanabe, and K.-R. Muller, "Optimizing spatial filters for robust EEG single-trial analysis," IEEE Signal Process. Mag., vol. 25, no. 1, pp. 41-56, 2008.

[11]. J. Lin, W.-M. Chen, Y. Lin, J. Cohn, C. Gan, and S. Han, "MCUNet: Tiny deep learning on IoT devices," in Adv. Neural Inf. Process. Syst. (NeurIPS), 2020.

[12]. Y. Song, Q. Zheng, B. Liu, and X. Gao, "EEG conformer: Convolutional Transformer for EEG decoding and visualization," IEEE Trans. Neural Syst. Rehabil. Eng., vol. 31, pp. 710-719, 2023.

[13]. F. Jiang, Y. Zhao, S. Duan, and L. Wang, "Graph neural network for EEG analysis: A survey," Neural Netw., vol. 157, pp. 478-492, 2023.

[14]. 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.

[15]. A. Benny, S. P, and C. E. S. S, "Heat transfer properties of silver nanofluids on shell and tube heat exchanger," International Journal of Research and Development in Engineering Sciences, vol. 5, no. 4, p. 11, Jul. 2023, doi: 10.63328/ijrdes-v5ri4p3.

[16]. A. Barachant, S. Bonnet, M. Congedo, and C. Jutten, "Multiclass brain–computer interface classification by Riemannian geometry," IEEE Trans. Biomed. Eng., vol. 59, no. 4, pp. 920–928, 2012.

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 :

J V G Prakasa Rao Pyla , " Transformer-Based Real-Time Wing Flutter Detection in Formula 1 Vehicles ", International Journal of Computer Science , Engineering and Artificial Intelligence , Vol. 2, Issue. 1 (January – March) , 2025 , Pages: 18 – 24.
DOI : https://doi.org/10.63328/IJCSEAI-V2RI1P4