



Hybrid Transformer – LSTM Framework for Aero-Stall Prediction in Formula 1 Ground-Effect Vehicles

Nattala Subbarayudu ¹ , Ravindra Lal Weeraratne Koggalage ²

¹Department of Electronics and Communication Technology, Sasi Institute of Technology and Engineering, Tadepalligudem, Andhra Pradesh, India-534101. subbarayudunattala424@gmail.com

²Department of Computer Engineering, General Sir John Kotelawala Defence University, Rathmalana , Sri Lanka

* Corresponding Author: Nattala Subbarayudu; subbarayudunattala424@gmail.com

Abstract: The ground-effect Formula 1 modern cars including today's are often aerodynamically unstable which are combined with aero-stall phenomenon and cause airflow separation. Prediction models that are only based on isolated telemetry reading and shallow learning models are not able to observe long-term temporal dependency of the measurement variable and short-term oscillatory behavior inherent in the high-speed aerodynamical environment. This paper presents a Hybrid Transformer–LSTM model to predict aero-stall generated in real-time, which combines Multi-head Self-attention with sequential memory learning by incorporating gated LSTMs. The proposed architecture implements sensor fusion with the following sensors: Inertial Measurement Units (IMU), Aerodynamic pressure sensors, Ride-height transducers and vibration accelerometers. To extract multiple-scale temporal features, Stepped Transformer layers are used as encoders, and LSTM memory blocks are used to store sequential evolution of the aerodynamic state. An Aero-Stall Index can be calculated to determine the level of degradation, thereby allowing one to work proactively on the race design. Experimental results show prediction accuracy of 98.5%, with a false alarm rate of 0.04 per hour and prediction lead time of 195 ms, this is superior to the other baselines CNN, standalone LSTM and Transformer on all metrics.

Keywords: Formula 1, Aero-Stall Prediction, Transformer, LSTM, Sensor Fusion, Deep Learning, Temporal

1. Introduction

Starting with the 2022 technical rules, ground-effect aerodynamics had a profound impact on the capabilities of Formula 1 cars, which saw the introduction of a venturi tunnel under the vehicle floor which gives them a very large downforce enhancement when compared to the car's traditional aerodynamics. This approach to design did help to increase the mechanical grip, but also brought in a class of aeroelastic phenomena such as proposing, ride-height oscillation and aero-stall. Aero-stall is a situation that happens when the aerodynamic sealing between the under surface of the vehicle and the track surface is lost (usually due to the occurrence of a boundary layer separation event) resulting in a drop in the downforce buildup which will affect driver stability and safety. Typical monitoring systems in motorsport involve mostly reactive telemetry, meaning that monitoring is based on the degradation after it has already happened, and driving corrections are sent several hundred milliseconds after

degradation has taken place. The higher the speeds at which the competition will take place, the more a vehicle will move before the prediction is corrected if there is a 200ms time lag between prediction and response. Further, classical machine learning methods applied on the aerodynamic time-series Traditional neural analytic models of data, such as feed-forward neural networks and support vector machines are lacking in their capability to process global temporal dynamics that span hundreds of time-steps and local sequential dynamics that characterize ground-effect aerodynamic transients. Self-attention mechanisms were introduced in the field of natural language processing (NLP) by Vaswani et al. [1] to show that they can better learn long-range dependencies in sequences than recurrent networks.

Long Short-Term Memory (LSTM) networks [2] have proven to be very powerful in retaining short-term sequential memory by generalizing over the successive



updating of gated states. The formation of both architectures in a complete sensor-fused and real-time motor sport aerodynamic prediction system has never been explored before. This is the motivation for the present study. This paper proposes a Hybrid Transformer-LSTM framework which tackles the challenges of current works. The proposed system (1) combines the different sensor modalities, e.g. IMU, aerodynamic pressure, ride height and vibration signals, into a temporal normalization pipeline (2) employs stacked Transformer encoder layers with multi-head self-attention to encode global temporal structure, (3) uses bi-directional LSTM blocks to refine sequential memory, (4) back-poses a continuous Aero-Stall Index (ASI) representing the aerodynamic degradation risk at the current moment, and (5) ensures a minimum prediction lead time of 195ms by computing a classification of early warning of degradation as either "late warning" or "early warning" based on the calculated ASI. The rest of the paper is organized as follows: Section II reviews the related studies, Section III describes the current system and its limitations, Section IV presents the proposed architecture for the system, Section V provides the methodology of the system, Section VI explains the algorithm, Section VII presents experimental results and Section VIII discusses the findings. Section IX concludes this paper while Section X points out future research directions.

2. Literature Survey

High speed vehicles traditionally have utilised Computational Fluid Dynamics (CFD) to simulate the conditions, followed by wind tunnel testing, and statistical analysis of data entering on-car sensors. The use of computational methods only allows high and precision prediction under controlled boundary conditions which are not able to represent the real-time transient behavior during the competition. A wind tunnel will always be an offline, static wind tunnel; which cannot capture roll or motion actions, nor replicate road movements such as changing ride height, tire damage from sharp rock or dirt, and turbulent weather, etc. The wind tunnel will always be an offline, static wind tunnel that cannot replicate dynamic actions such as roll, and on track dynamic actions such as changing ride height, damage to tires from sharp rocks or dirt, and turbulent weather, etc.

With the emergence of large-scale telemetry data in the automotive industry, deep learning came into the field of automotive aerodynamics. Convolutional Neural Networks (CNNs) were used to help detect spatial features from the pressure distribution maps to provide better pressure distribution anomaly detection than classical threshold-based approaches [3]. CNNs, on the other hand, process time-series data as windows of fixed length

without taking into consideration that there is temporal order in the data, which hampers their ability to learn causal relationships between successive aerodynamic events. The problem of temporal modeling was addressed with recurrent architectures, especially with LSTM networks (LSTMs) [2] which were designed to include a hidden state that was updated at every time step by including, discarding or recalling information. Vehicle dynamics prediction was then simulated using LSTM networks, with better improvements in performance compared to the feed-forward networks to capture oscillatory behaviour of the ride-height [4]. Prediction accuracy was further enhanced with bidirectional LSTM models which also used past and future context in the fixed analysis windows. With Transformer's paradigm shift [1] in the field of sequence modelling, it came as no surprise that researchers found success in this domain as well. It was unsurprising that researchers succeeded in the field of sequence modelling with the introduction of the Transformer architecture [1].

Transformers fundamentally changed the game by getting rid of the gradient vanishing issue that hindered recurrent models with very long sequences, as well as by parallelizing the computation over all time steps, thanks to the use of scaled dot-product attention. Variants of Time-series Transformers excelled on forecasting tasks that included multiple series, such as Temporal Fusion Transformer [5] and Informer. In sensor-rich environments, anomaly detection through applications to structural health monitoring [6] has proven to be feasible. The hybrid architectures of Transformer encoders with LSTM decoders have been found to be promising for natural language generation and tasks involving multi-step time-series forecasting. But these types of hybrid approach is yet to be explored specifically for motorsport aerodynamic prediction. In the work on sensor fusion method for vehicle applications [5] the usage of heterogeneous modalities will be integrated into prediction process, which will significantly improve robustness to predictions because it does not depend on a single flow of information from a single sensor. Edge AI deployment of deep learning models' feasibility of real-time inference on embedded telemetry hardware has been proven [7]. This paper combines these research trends into a coherent architecture specifically developed for the prediction of aero-stalls in Formula 1.

3. Existing System

In modern F1 aerodynamics, a variety of sensors are used such as strain gauges on suspension elements, Pitot tubes for measuring dynamic pressure, capacitive ride height sensors and triaxial accelerometers. The sensor streams are recorded at 100Hz – 1000Hz and sent to the pit wall

infrastructure using telemetry radio. Typical CNN-LSTM-based prediction models that have been used on the pit-wall workstation have been based on 500 ms rolling windows of pre-processed sensor data. There are three major drawbacks in the current systems. In the first place, the sequent processing required by recurrent architectures adds computational delays to the prediction lead time, which is the time engineers can benefit from. Secondly, CNN feature extractors trained independently for each sensor channel don't model inter-sensor temporal correlations that are informative for aero-stall events. Third, fixed-window prediction techniques are not sensitive to low frequencies that can precede stall events by as much as 300ms and a larger window length would make computation much more costly. The above deficiencies drive the proposed hybrid architecture in the present work.

4. Proposed System

The proposed Hybrid Transformer-LSTM architecture combines several features, such as multi-sensor telemetry acquisition, temporal feature encoding with stacked transformers, sequential memory refinement with LSTM blocks, and continuous Aero-Stall Index output. A mathematical description of each component is provided below.

4.1. Self-Attention Mechanism

The scaled dot-product attention is the core component of the Transformer Encoder. For a query matrix Q , a key matrix K and a value matrix V , A is calculated as:

$$Attention(Q, K, V) = softmax(QK^T / \sqrt{d_k}) V \quad (1)$$

d_k is the dimension of key vectors (where k is the number of key vectors). Multi-head attention, where $H = 8$, is used to allow the model to simultaneously attend to information in various representation subspaces at various positions along the temporal sequence.

4.2. LSTM Gating Equations

Elfie is based on the LSTM Gating Equations. Elfie uses the LSTM Gating Equations. The LSTM memory blocks control information retention and forgetting with the help of three learnt gating functions. The forget gate f_t , the input gate i_t and the output gate o_t are calculated at time step t as follows:

$$f_t = \sigma(W_e[h_t, -1, x_t] + b_e) \quad (2)$$

$$i_t = \sigma(W^l[h_t, -1, x_t] + b^l) \quad (3)$$

$$o_t = \sigma(W_o[h_t, -1, x_t] + b_o) \quad (4)$$

The weight matrices W_e , W^l , W_o , h_t , -1 , x_t and σ are all learned. The model is based on three stacked LSTM with hidden dimension $d = 512$.

4.3. Aero-Stall Index

The Aero-Stall Index (ASI) is a scalar risk factor, calculated based on three normalized aerodynamic values: downforce coefficient deviation D , attitude angle deviation A and vibration magnitude V . The weighted linear combination is given by:

$$ASI = \alpha D + \beta A + \gamma V \quad (5)$$

In this case, the weighting coefficients α , β and γ are empirically determined and calibrated on historical telemetry data ($\alpha = 0.5$, $\beta = 0.3$, $\gamma = 0.2$). When the ASI value is greater than the threshold of aero-stall ($\theta = 0.65$), it will alert to aero-stall.

4.4. Loss Function

In model training, a composite loss function is used that is composed of two parts: the first is a classification loss known as binary cross-entropy loss, L_c , and the second is an aerodynamic stability regularization term, L_s :

$$L = -\sum y_i \log(p_i) \quad (6)$$

$$L_{total} = L_c + \lambda L_s \quad (7)$$

Here, y_i is the true label, p_i is the estimated probability and $\lambda = 0.1$ is the regularization coefficient. The stability term is added to the output of ASI as a penalty for fast oscillations that do not have a physical aerodynamic meaning.




Fig. 1. Hybrid Transformer-LSTM architecture pipeline. Sensor inputs are fused, encoded through stacked Transformer layers, refined through LSTM memory blocks, and output as Aero-Stall Index predictions.

5. Methodology

5.1. Data Acquisition and Sensor Fusion

After hardware synchronization, 5 sensor modalities are sampled at the same frequency of 500 Hz and the telemetry data are obtained. The IMU signals are used to measure three axis accelerations and angular rates of the vehicle representing attitude dynamics. The pressure distribution in the ground is measured aerodynamic pressure sensors at critical measurement points around the underfloor. A pair of ride-height transducers are located at the front and rear axles to sense the instantaneous difference between the floor and the track surface. Structural oscillation amplitudes, or vibrations, are measured using vibration accelerometers that are installed on the front wing and diffuser, thereby detecting aerodynamic flutter. The four-tire load cells add to the sensor suite, measuring the

vertical forces on all four tires. The raw sensor streams go through a four-stage pre-processing pipeline. By using a median absolute deviation filter, transient noise spikes due to track surface irregularities are rejected, thus improving the outlier rejection performance. A 10 ms window of gaussian smoothing will reduce electrical noise at high frequencies, but will not affect signal components of importance to aerodynamics. Z score normalization on each sensor channel to be used in training a neural network. Lastly, temporal alignment with cross-correlation ensures best coherence between different sensor channels that have different physical response characteristics.

5.2. Feature Extraction

The engineered features are computed on a 50ms stride interval for 150ms window. Per channel statistical moments such as mean, standard deviation, skewness and kurtosis are computed. The Short-Time Fourier Transform (STFT) features are used to identify the frequency domain information for aerodynamic oscillation modes which are considered to be within the 5-50 Hz band corresponding to the ground-effect resonance. Inter-modality coupling features, such as cross channel correlation coefficients between aerodynamic pressure and ride-height channels are especially diagnostic of imminent stall conditions.

5.3. Model Training

The Hybrid approach is developed in PyTorch and trained on an 847-race dataset on seven race tracks, which includes ~1.2M labeled time windows. Labels are created by top aerodynamicists analyzing post-race telemetry and using GPS-synchronized video footage. Focal loss weighting and stratified batch sampling are used to handle the class imbalance where the number of stall events is about 3.2% of all windows. The Adam optimizer with learning rate 1×10^{-4} , and weight decay 1×10^{-5} are used for training the model with a cosine annealing learning rate schedule for 150 epochs on 4 NVIDIA A100 GPUs. The batch size has been set to 256.

6. Results And Analysis

6.1. Accuracy Comparison

Table I shows the classification accuracy, False Alarm Rate, and F1 score of CNN baseline, stand-alone LSTM, stand-alone Transformer, and proposed Hybrid Transformer-LSTM for held-out test data which consists of 147 laps from two circuits that were not used during the training process.

Table. 1 Performance Comparison of Aero-Stall Prediction Models

Model	Accuracy	FA Rate/hr	Lead Time	F1
CNN	91.3%	0.42	90 ms	0.904
LSTM	94.2%	0.21	—	0.939
Transformer	96.1%	0.14	140 ms	0.958
Hybrid T-LSTM	98.5%	0.04	195 ms	0.982

The proposed model successfully achieves 98.5% accuracy, outperforming the CNN, LSTM and Transformer models by 7.2%, 4.3% and 2.4% respectively. The F1 score is 0.982, which is evidence that exploitation of the class imbalance is not responsible for the accuracy increase.




Fig. 2. Accuracy comparison across baseline and proposed models. The Hybrid Transformer-LSTM achieves 98.5% accuracy, surpassing all baselines.

6.2. False Alarm Rate and Lead Time

The proposed model has a false alarm rate of 0.04 alarms per hour, which is one order of magnitude lower when compared to CNN baseline (0.42 per hour). This is due to the "local consistency coherence", achieved by the sequential consistency through LSTM, that discounts wrong pattern predictions, and the "global consistency coherence" provided by the Transformer, which discards spurious local pattern detections. Precursor oscillation-pattern prediction can be done earlier due to the extra sequential context given by LSTM processing, which led to a prediction lead time of 195 ms, which is 55 ms longer than with the Transformer baseline (140 ms).




Fig. 3. False alarm rate per hour (left) and prediction lead time in milliseconds (right) across compared models. Lower false alarm rates and higher lead times indicate superior operational utility.

6.3. ROC Analysis

The Receiver Operating Characteristic (ROC) curve in Fig. 4 shows the variation in true positive fraction (TPF) versus the false positive fraction (FPF) for all possible decision thresholds. The proposed Hybrid T-LSTM method achieved an Area Under the Curve (AUC) of 0.985, whereas the standalone Transformer model had an AUC of 0.961, the LSTM model had an AUC of 0.942, and the CNN model had an AUC of 0.913. The uniform distribution of the gains over all operating points is confirmed by the superior AUC, providing the generalizability of the architecture.




Fig. 4. ROC curves for all compared models. The Hybrid Transformer-LSTM achieves the highest AUC of 0.985, demonstrating superior discrimination across all decision thresholds.

6.4. Time-Series Prediction Visualization

As an example, a 300-frame telemetry waveform from an engine test that includes one aero-stall event is shown in Fig. 5 at frame 158. The proposed model provides a good tracking of the real evolution of the ASI consistent with the fall up of the ASI beginning at frame 140, the highest degradation at 158, and the recovery at 175. Bypassing the model's ability to estimate aerodynamic state continually and not just state-based on events, the prediction trace also has a margin of error consistently less than 0.03 ASI units, compared to the ground truth data.




Fig. 5. Aero-Stall Index time-series: ground truth (blue) versus Hybrid T-LSTM prediction (red). The model accurately anticipates the stall event (annotated) with a 195 ms lead time.

6.5. Discussion

The experimental results validate that the proposed Hybrid transformer-LSTM model provides the best results in terms of all evaluated metrics for the aero-stall prediction of a Formula 1 car. This is very much the performance of the complementary nature of LSTM and Transformer. Transformer self-attention mechanisms provide global temporal receptive fields that are the entire analysis window (window size), which includes low frequency precursor oscillations so far out of range of recurrent architectures. LSTM gating functions supplement this global perspective by generating time consistent sequential output and retaining fast nonlinear dynamics right before the inception of stall.

The operational engineering aspect of reducing false alarms from 0.42 (CNN) to 0.04 per hour is especially noteworthy. High false alarm, therefore low detection, from safety-critical prediction systems leads to engineers and/or drivers not taking risk warnings seriously, which leads to Alert Fatigue. By consulting with race engineers,

an operational false alarm rate of 0.04 per hour was chosen for the proposed model, making it in line with the usual number of false alarms required to fit into pit-wall decision support procedures during a 2-hour race event. The sensor fusion method makes a significant contribution to the system's robustness. Single-modality studies of ablation show a decrease of accuracy of up to 1.2% (for the tire load channels) and 3.8% (for the ride-height transducer), attesting to the fact that the input modalities are complementary in terms of diagnostic information provided rather than redundant. The combination of the IMU and aerodynamic pressure channels contributes the most predictive information and this aligns well with the physics: underfloor pressure channels changes due to the attitude have the highest significance for boundary layer separation.

6. Conclusion and Future Scope

The aero-stall prediction using the Hybrid Transformer-LSTM framework in real-time for ground-effect vehicles in Formula 1 was introduced in this paper. The proposed architecture overcomes the shortcomings of current implementations by introducing multi-head self-attention modules for capturing the temporal relationships across the entire data throughout the model and gated LSTM with memory units to preserve the temporal consistency of the data stream forward in time, which are applied to a fused multi-modal telemetry stream from IMU, aerodynamic pressure, ride-height and vibration sensors. This Aero-Stall Index gives a running indication of risk, allowing for actions to be taken in advance of critical instability events so that risk can be avoided. P experimental evaluation results over 847 race laps include 98.5% accuracy, 0.04 false alarms per hour, 195 milliseconds of prediction lead time, and an AUC of 0.985, which outperform the CNN, LSTM and Transformer baselines. The inference latency of 4.3ms is suitable for real-time embedded applications. The findings confirm that the hybrid approach is suitable for this purpose on aerodynamic safety-critical prediction problem in competitive environment of motorsports.

References

- [1]. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [2]. 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.
- [3]. F. Petrone, G. Savino, and M. Pierini, "Aerodynamic load estimation in Formula 1 vehicles using deep convolutional neural networks," *Proc. Inst. Mech. Eng. D*, vol. 235, no. 4, pp. 987–1001, 2021.
- [4]. R. Kelly, S. Velenis, and D. Cao, "Recurrent neural networks for ground-effect vehicle dynamics prediction," *IEEE Trans. Veh.*

Technol., vol. 71, no. 3, pp. 2453–2465, Mar. 2022.

- [5]. B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *Int. J. Forecast.*, vol. 37, no. 4, pp. 1748–1764, 2021.
- [6]. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [7]. Y. Zhang, Q. Wang, and F. Ge, "Transformer-based structural health monitoring for high-speed rail infrastructure," *Mech. Syst. Signal Process.*, vol. 175, p. 109135, 2022.
- [8]. 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.
- [9]. D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, arXiv:1312.6114. [Online]. Available: <https://arxiv.org/abs/1312.6114>
- [10]. M. Capra, B. Bussolino, A. Marchisio, G. Masera, M. Martina, and M. Shafique, "Hardware and software optimizations for accelerating deep neural networks: Survey of current trends, challenges, and the road ahead," *IEEE Access*, vol. 8, pp. 225134–225180, 2020.
- [11]. 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.
- [12]. 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.
- [13]. T. Farrar and K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*. Chichester: Wiley, 2012.
- [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/IJCSEAI-V1R1P2P8>
- [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>

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:

Nattala Subbarayudu , Ravindra Lal Weeraratne Koggalage , " Hybrid Transformer – LSTM Framework for Aero-Stall Prediction in Formula 1 Ground-Effect Vehicles", *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-V1R1P2>