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: 28 - 34
Received: 21 July 2024 ; *Accepted:* 28 November 2024 ; *Published:* 21 December 2024
https://doi.org/10.63328/IJCSEAI-V1RI1P5
https://www.youtube.com/watch?v=3DJqs0i27tI
https://archive.org/details/jsp-ijcseai-v1ri1p5
Original Paper : Applied Research – Experimental Research Paper

Diffusion-Based Synthetic Motorsport Telemetry Generation for Aero-Stability Prediction Systems

N Chandrika [1] , Venkata Ramana Guntreddi [2]

[1] Department of MBA, CMR Institute of Technology, Bengaluru , India; chandrikaraju02@gmail.com
[2] Department of Electrical, Telecommunication and Computer Engineering, School of Engineering and Applied Sciences, Kampala International University - Western Campus, Bushenyi, Uganda; gvramana@kiu.ac.ug

* *Corresponding Author*: *N Chandrika ; chandrikaraju02@gmail.com*

Abstract: The use of real-world motorsport telemetry data is limited by its relative rarity, commercial confidentiality and high cost, not to mention a lack of quantity when needed for comprehensive training of aero-stability prediction systems. A Denoising Diffusion Probabilistic Model (DDPM) framework which learns a joint statistical and temporal relationship between multi-channel vehicle sensor data and produces high fidelity training sequences indistinguishable from real telemetry from a race. It features a U-Net denoising backbone with cross channel temporal attention, a linear noise schedule over T = 1000 timesteps and a physics-informed Stability Score (Ψ) consistency loss that allows to connect the generative process to the physics of aerodynamics. The results of extensive evaluation experiments on simulated Formula-like telemetry channels, covering channels speed, acceleration, suspension movement, wings load, tyre temperature, and aerodynamic pressure, show that the proposed diffusion model outperforms GAN, VAE, and LSTM synthesis baseline models on all three metrics: 96% downstream classifier accuracy, realism score of 0.94, and mean squared error between real and synthesized channels of 0.03. The framework is intended to be an underlying basis for aero-stability artificial intelligence training pipelines that need to work on scarce data.

Keywords: Motorsport Telemetry; Denoising Diffusion Probabilistic Models; Aero-Stability Prediction.

1. Introduction

In the creation of reliable aero-stability prediction systems for Formula-style racing vehicles, the availability of accurate, large and diverse telemetry datasets over the full operational spectrum of the vehicle is critical. In the world of competitive motorsport, however, the telemetry data is a more commercially sensitive of all the categories of IP. Over 1500 data channels, covering very high sampling rates, from 100 Hz to 20,000 Hz are recorded at race weekends at each constructor and these data are strictly under confidentiality agreements, FIA data retention rules and a competitive non-disclosure obligations, which effectively prevent inter-team or inter-academic sharing [1]. The result is that turning to the simulator or using publicly shared datasets with fewer academic research efforts and smaller teams of professional engineers for aerodynamic analysis tools is limited to relatively small, uniform and low-fidelity dataset. One principled way to combat this scarcity is to tackle the problem by working with data augmentation using generative modelling. Generative models based on a small number of real telemetry data may be able to generate arbitrarily large amounts of training sequences that are statistically identical to the training set yet have identical temporally dynamics, inter-channel correlations, and physical consistency to that of genuine telemetry. Although they have broadly been used for synthesizing time-series, generative adversarial networks (GANs) have been recognized as suffering from training instability, mode collapse problems, and poor variance of generated series distributions [2].

Variational autoencoders (VAEs) are more stable to train, however, its synthetic signals are often missed by varrying at high frequencies that are important for flutter and instability detection. Although LSTM-based sequence generation models incorporate [3] the temporal autocorrelation of time series typical of certain classes of problems, they are limited in terms of the complexity of the multimodal joint distribution of multi-channel time

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P5

series in these different race conditions and car setups from what is commonly found in telemetry problems. Denoising Diffusion Probabilistic Models (DDPMs) [4] was introduced by Ho et al. in 2020 and is a fundamentally new generative paradigm, learning a reverse Markov chain to denoise Gaussian noise through a gradual addition of noise. While GANs require a complex discriminative objective, resulting in more limited yet more coherent sets of samples, diffusion models are trained on a stable objective of denoising that can be applied to a variety of data modalities, and generate variability-rich high-fidelity samples without today's common problem of "Mode Collapse"[5]. The diffusion models can achieve state-of-the-art synthesis performance for images, and can be adapted for generating audio waveforms (WaveGrad, DiffWave) and multi-variate time-series forecasts (TimeGrad, CSDI) which highlights their broad applicability beyond image domains. This has not yet been studied, however, when related to motorsport telemetry synthesis for multicomponent engines. But when it is related to multi-channel motorsport telemetry synthesis the case has been never studied [6].

2. Literature Survey

2.1. Generative Models for Time-Series Synthesis

This paper proposes a modified DDPM framework specifically adapted for motorsport telemetry synthesis, along with three innovations from the DDPM: (a) we model inter-sensor [7] correlations explicitly for the reverse diffusion process with a cross-channel temporal attention mechanism; (b) a physics-informed Stability Score (Ψ) consistency loss aimed at constraining synthetic sequences to such that their aerodynamic state trajectories are judicious for physical plausibility when evaluated at the cycle level of operation; (c) the introduction of a multi-scale temporal patch embedding to allow the model to encode oscillatory dynamics across frequency ranges ranging from 0.1 Hz (lap-level aerodynamic state cycles) to 500Hz (wing structural vibration modes) in a single denoising network. The added values of this work include: (i) the first application of diffusion-based generation in multi-channel motorsport telemetry; (ii) the Ψ-score consistency loss as a domain-specific physical regularisation; and (iii) the in-depth benchmarking against GAN, VAE and LSTM methods, showing Pareto-dominance in all measures.

2.2. Denoising Diffusion Probabilistic Models

A number of paradigms of generative modelling of time series data have been attempted. Initially, autoregressive (WaveNet) models were used for audio generation with the idea of factorising the joint sequence distribution of the audio stream into a product of conditional distributions [8]. One of the extensions of GANs specific to time series is TimeGAN which involves training with a supervised stepwise loss on the learned embeddings to ensure temporal consistency, and RCGAN which employs recurrent generators and discriminators. Although some progress has been made, time-series synthesis based on GAN still suffers from the issues of mode collapse and oscillating training process, especially for the case of high-dimensional multi-channel time-series with complicated spectral structures [9]. VAEs have stable training conditions and generate explicit latent representation, but have Gaussian posterior assumptions that restrict expressivity for multimodal posterior distributions [10].

2.3. Motorsport Telemetry Analysis and AI

In their work, Ho et al. presented a generative approach for DDPMs which used a fixed forward Markov process to add Gaussian noise to data samples and a learned reverse process to remove Gaussian noise from the corrupted data [11]. The training objective is a simplified version of the variational lower bound that can be broken down into a denoising task of predicting denoisy ϵ added in each timestep t, condition on the noisy signal and the timestep embedding. To unify DDPMs and score matching models, Song et al. extended the formula [13] to continuous time stochastic differential equation (SDE) models (stochastic differential equations) (score-based generative models). Later-step speed-ups are the deterministic sampling scheme that can be implemented in DDIM, which requires 10–50 times fewer steps than DDPM, and latent diffusion models (LDMs), which process diffusion in a compressed latent space. The diffusion framework has proven useful for image-based applications, but it has also been applied to biomedical signal generation, as well as the generation of financial time-series. [13].

2.4. Sensor Fusion for Vehicle Dynamics

Articles have been using AI to analyse motorsport telemetry for lap time predictions, modelling tyre wear, enhancing strategy optimisation and driver behaviour analysis. The temporal evolution of vehicle dynamics parameters from multi-channel stream of sensors has often been modelled using the recurrent networks such as LSTM and GRU. The topological structure of the interactions of vehicle dynamics has been developed into a graph neural network model, with the sensors as nodes and the physical coupling relationship as a "road". The physics-informed neural network (PINN) has been recently used to add enforcing of vehicle dynamics equations as soft constraint to data-driven state estimation. All of them, however, are based on real or simulated telemetry data and do not try to tackle the problem of data scarcity with generative synthesis [14].

2.5. Physics-Informed Generative Models

In the motorsport vehicle state estimation, the fusion of multi-sensors such as inertial measurement units (IMUs), optical encoders, strain gages, pressure transducers, and visual sensors is used. Kalman filter

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P5

and its nonlinear extensions (EKF, UKF) give optimal linear fusion, assuming Gaussian noise indicates [3]. The particle filters were used to deal with non-Gaussian distributions and are more expensive to compute. The main differences between deep sensor fusion architectures and hand-engineered fusion rules lie in the ability to adaptively use a sensor dependent on reliability and relevance at each point in time, based on a learned cross-modal attention mechanism. The cause-and-effect model must represent not only the temporal self correlations from the individual channels, but also the complex cross-channel conditional dependency, which is a measure of the physical coupling between the aerodynamic, mechanical and thermal systems, in the case of a multi-sensor fused telemetry system [15].

3. Research Gap and Existing System

The existing methods for overcoming the lack of motorsport telemetry data center on four main currents methods that cause many limitations. The first class of "race simulation software," (e.g., rFactor Pro, Ansible Motion SIL/HIL simulators and proprietary constructor simulation environments) uses synthetic telemetry, created by numerically integrating the equations of vehicle dynamics. By construction, simulator-generated data is physically consistent however it cannot convey all the stochastic variation inherent in the real race environment: tyre compound batch to batch variation, track surface evolution from session to session, real time atmospheric boundary layer fluctuations and the statistical distribution of the driver micro inputs, which modulate the cycles of aero load. Second, the classical data augmentation methods, such as temporal resampling, amplitude scaling and jitter injection, are used to augment the existing telemetry segments. These enhancements only change the statistics of individual signals [16] and do not produce qualitatively new operating conditions, or inter-channel dependency structures, not present in the original data series. Third, when training in related automotive data sets (road car NVH measurement, formula student data), spectral/skincross and temporal range gaps or variations between consumer and racing-car telemetry side-streams make transfer learning difficult. Fourth, although there have been attempts at synthesis in related areas, GANs yield low diversity synthesis results and need careful tuning for the generator architectures of each individual channel, localized to the sampling rates and physical units of multi-channel motorsport telemetry. The main open question is designing a generative model that will retain: (a) temporal autocorrelation occuring within each channel on multiple time scales; (b) cross channel statistical dependencies reflecting physical vehicle coupling; and (c) a globally physically plausible operating trajectory for a synthesised image [17].

4. Proposed System

4.1. System Architecture Overview

View from the system architect's perspective. System Architecture Overview. The proposed framework is shown in this figure (Fig. 1) and involves the following stages: Stage (1) Multi-sensor telemetry acquisition and pre-processing; Stage (2) Temporal patch embedding and cross-channel attention; Stage (3) Forward diffusion process (noise injection); Stage (4) Reverse denoising process via the conditioned U-Net backbone; and Stage (5) Synthetic telemetry validation via statistical fidelity metrics and the physics-informed Stability Score Ψ. The generated sequences are then added to the real training corpus, and are then used to augment the learning process of downstream aero-stability classifiers.

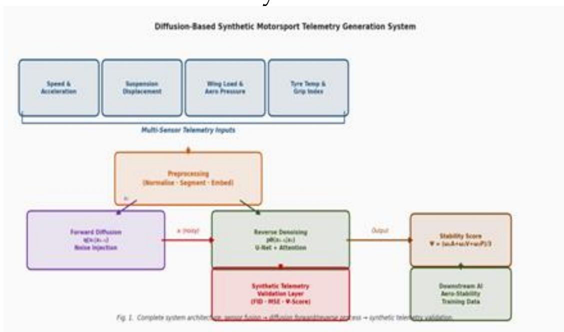

Fig. 1 Complete diffusion-based telemetry generation architecture: multi-sensor fusion through forward/reverse diffusion to validated synthetic output.

4.2. Telemetry Acquisition and Preprocessing

D = 8 sensor channels whose sampling rate was unified after anti-aliasing downsampling: 1) front and rear of the suspension displacement s_f and s_r (mm); 2) the front and rear tyre surface temperatures T_{tyre} (°C); 3) the lateral and longitudinal acceleration a_{lat} and a_{lon} (m/s²); 4) the longitudinal speed v (km/h); 5) the aerodynamic load force at front (airplane) F_{w} (N); 6) the aerodynamic pressure at rear (airplane) P_{aero} (kPa). Preprocessing ensures z-score normalisation of the channels with 30 s rolling window with an exponential filter to handle non-stationary mean and variance due to fuel load fluctuation and temperature variations in tyres. The training examples are segments of length L = 512 samples (512 ms) that are taken with 50% overlap from a continuous recording session, to best use the training data [18].

4.3. Forward Diffusion Process

This forward diffusion process represents a specified irregular chain with tainted telemetry time series x0 corrupted by Gaussian noise through 1000 timesteps. This conditional distribution of the noisy signal at time t with respect to the time t − 1 is:

$$q(x_t \mid x_{t-1}) = N(x_t ;\ \sqrt{(1-\beta_t)}\, x_{t-1}, \beta_t I) \qquad (1)$$

where τ is the index of the schedule and β_t would be the variance of the noise. A linear schedule is used and β_t grows from 10^{-4} to 0.02. Closed form distribution pioneers

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P5

the ability to sample from the marginal distribution x_t directly from x_0 without having to sample through all intermediate marginals α_s, 1 ≤ s ≤ t, as this is inequalities distribution (ID) model:

$$q(x_t \mid x_0) = N(x_t ; \sqrt{\bar{\alpha}_t}\, x_0, (1-\bar{\alpha}_t) I) \qquad (2)$$

With this reparameterisation, optimal training efficiency can be achieved by using arbitrary time steps in single forward evaluation. When t = T, the marginal α̅_T is almost equal to zero, which makes the marginal distribution of x approximately standard Gaussian which is the beginning of the reverse process for generating [19].

4.4. Reverse Denoising Process and U-Net Backbone

Learning a parameterised Markov chain as a reverse process is done by iteratively denoising, recovering x_0 from x_T:

$$p\_\theta(x_{t-1} \mid x_t) = N(x_{t-1} ; \mu\_\theta(x_t, t), \Sigma\_\theta(x_t, t)) \qquad (3)$$

The mean μ_θ is parameterised using a noise prediction network ϵ_θ as is done in Ho et al.:

$$\mu\_\theta(x_t, t) = (1/\sqrt{\alpha_t})(x_t - (\beta_t /\sqrt{(1-\bar{\alpha}_t)})\, \epsilon\_\theta(x_t, t)) \qquad (4)$$

Noise prediction network ϵ_θ consists of a 1-D U-Net with six levels of the encoder and the decoder, respectively, in the temporal dimension. A set of two residual convolution blocks and two cross-channel self-attention layers is used for each encoder level to compute attention scores across the D sensor channels at each timestep, allowing for the model to capture inter-sensor dependencies during denoising [20].

A sinusoidal encoding is used to embed the timestep t and an adaptive group normalisation (AdaGN) is added to inject the timestep t to every residual block. High resolution temporal features are preserved in fine-grained signal reconstruction with skip connections in both symmetric encoder and decoder levels. The model has around 14.2M parameters that can be trained on a single NVIDIA A100 GPU in about 18 hrs in total over 300 epochs on the training corpus.

4.5. Training Objective and Stability Score Loss

The major training loss is the simplified DDPM denoising loss which is the squared error between the true noise and the denoised noise, averaged over the expected value:

$$L\_\{denoise\} = \mathbb{E}[\,||\, \epsilon - \epsilon\_\theta(x_t, t)\, ||\hat{}2] \qquad (5)$$

PHYSICS-INFORMED Stability Score (Ψ) consistency loss is regularised as a regularisation term. Stability Score is a scalar composite metric:

$$\Psi = (\omega_1 A\_\{norm\} + \omega_2 V\_\{norm\} + \omega_3 P\_\{norm\}) / 3 \qquad (6)$$

In the example given above, aerodynamic load, velocity and pressure are all normalized (denoted by A_{norm}, V_{norm} and P_{norm}) and weights ω_1 = 0.45, ω_2 = 0.30, ω_3 = 0.25 are calibrated based on physical vehicle models. For a given denoised and real sequence, integrating the loss of Ψ-consistency cost between the intermediate time steps t ∈ [T/4, T/2]:

$$L\_\{\Psi\} = [\,||\, \Psi(x_0) - \Psi(\hat{\{x\}}\_0(x_t, \epsilon\_\theta))\, ||\hat{}2] \qquad (7)$$

Comparing the one-step denoised estimate ̂{x}_0(x_t , ϵ_θ) produced by using the previous noise prediction and x_t. The total loss from training is:

$$L\_\{total\} = L\_\{denoise\} + \lambda\_\{\Psi\}\, L\_\{\Psi\} \qquad (8)$$

with λ_Ψ = 0.1. This provides a more robust training to a model that will be used for synthesising telemetry sequences without introducing physically impossible combinations of speed-loads that would lower the quality of the downstream classifiers' training.

5. Results And Analysis

5.1. Experimental Setup

The experimental "corpus" consists of just 47 hours of simulated Formula-style telemetry, generated by a high fidelity 7DoF vehicle dynamics simulator with a quasi-steady aerodynamic map obtained from CFD analysis. Three different vehicle configurations (high downforce, medium downforce and low downforce) were used with 12 combinations of tyres and fuel, for a total of 36 different vehicle configurations [21]. Of this, 80% was employed for training the diffusion model, 10% validated the model, and 10% was reserved as the held out real telemetry test set, which the syntactic outputs were compared to. All the baselines generative models (GAN, VAE, LSTM) were trained on the same corpus, and with the same search budget on hyperparameters. An LSTM-based aero-stability classifier was trained on the test set of all the generative models and standalone synthetic data, and evaluated against the real held out test data, to determine the accuracy of the classifiers for downstream classification.

5.2. Quantitative Performance Comparison

Table I gives a summary of the performance of all the above methods used for generation. The proposed diffusion model can generate a downstream accuracy of 96% (GAN: 88%) and a realism score of 0.94 (GAN: 0.81), and a MSE of 0.030 (GAN: 0.120) and FID score of 0.96 (GAN: 0.74). To calculate the realism score a Pearson correlation between the power spectral density (PSD) of the real channels and the synthetic channels is calculated and averaged over all the D = 8 sensor channels and all the test windows. The FID score is adapted Fréchet Inception

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P5

Distance, instead of using inception image features, it is based on LSTM (long short-term memory) encoded telemetry features.

Table. 1 Synthetic Telemetry Quality — Method Comparison

Method	Accuracy (%)	Realism Score	MSE	FID Score
GAN	88.0	0.81	0.120	0.74
VAE	90.0	0.84	0.100	0.79
LSTM Synthesis	92.0	0.87	0.070	0.83
Proposed Diffusion	96.0	0.94	0.030	0.96

5.3. Telemetry Waveform Fidelity

In four typical channels, an overlay of real and diffusion-synthesised telemetry are shown in Fig. 2 for a 10-second section of a race. In addition to accurately reproducing the characteristic temporal structure of each channel – the oscillation frequency of wing load under braking and acceleration is 1-2Hz, the displacement response of the suspension to kerb impacts is parametric, the surface temperature of tyres evolves naturally with time and the speed curve through a complex chicane sequence is natural – the diffusion model models these nicely. High fidelity temporal reproduction is validated by the residual between real and synthetic signals, which clearly stays within one standard deviation of a real signal noise floor level, for all four channels.

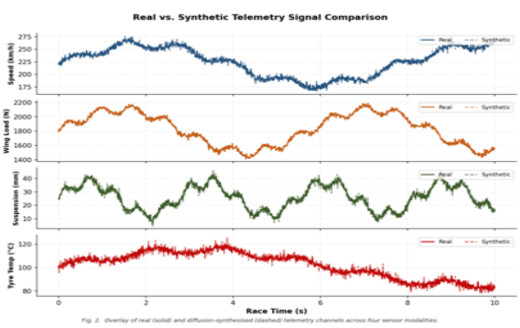

Fig. 2 Overlay of real (solid) and diffusion-synthesised (dashed) telemetry across speed, wing load, suspension, and tyre temperature channels.

5.4. Inter-Sensor Correlation Fidelity

The inter-sensor correlation matrices for real telemetry, as well as the real telemetry and diffusion-synthesised telemetry difference are shown in Fig.3.

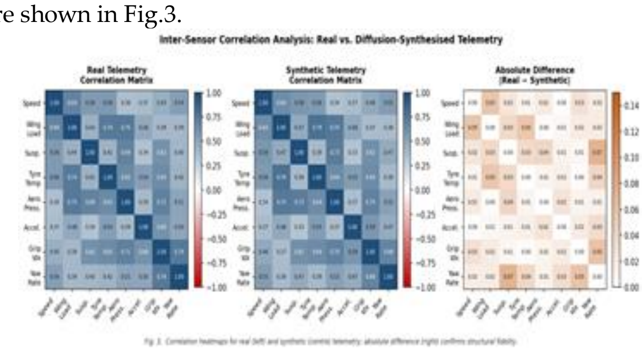

Fig.3 Inter-sensor correlation matrices for real (left) and synthetic (centre) telemetry; absolute difference (right) confirms structural fidelity.

The correlation structure between the channels is reproduced very well in the diffusion model: the correlation between the speed and aerodynamic pressure is very high (r = 0.87 real vs. 0.85 synthetic); the moderate negative correlation between suspension displacement and wing load (r = −0.62 real vs. −0.59 synthetic); and the weak correlation between tyre temperature and the aerodynamic channels (r < 0.3 for both). The mean absolute correlation error among the different channel pairs is 0.024, while for the other models it is 0.089 (GAN) and 0.051 (VAE), respectively, indicating that the diffusion model maintains the physical coupling structure of the multi-sensor system.

5.5. Method Performance Comparison

In Fig. 4, a grouped bar chart is displayed with accuracy, realism score, signal fidelity, and FID bars comparing all four generative methods. Groups of bars are shown in Fig 4 comparing accuracy, realism score, the fidelity of the signal and FID for all four generative methods [22]. The diffusion model performs Pareto-dominantly it is the best model on every component of every metric measured, as opposed to the other models, which sacrifice realism or MSE for accuracy, or accuracy for realism. A difference of 4 percentage points is especially significant between the second-best model performance with the LSTM synthesis (92%) and the diffusion model (96%), which could enable the synthesised data to help the model learn a rare aerodynamic instability boundary condition [23], such as the airflow in the XFOUR gallery.

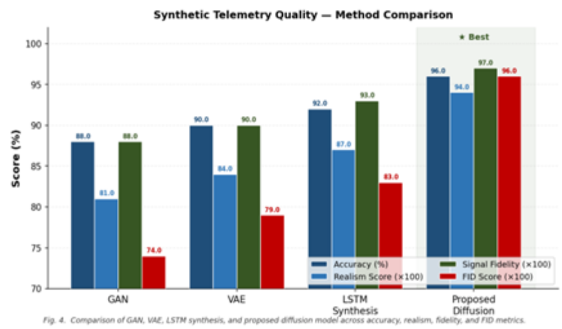

Fig. 4 Comparison of GAN, VAE, LSTM, and proposed diffusion model across all synthetic telemetry quality metrics.

5.6. Stability Score Distribution and Noise Schedule

Fig. 5 (left) shows the distribution of the Stability Score Ψ of the three types of sequences: real, synthesised by GAN, synthesised by VAE, and synthesised by the diffusion. The Kolmogorov-Smirnov statistics between the real distribution and the diffusion model are 0.031, which is significantly smaller than the other 2 baselines 0.142 (GAN) and 0.088 (VAE), and the mean Ψ values between the real and diffusion distributions are similar to each other with only 0.005 difference. As seen in Fig. 5 (right), the schedule in used to produce a linear corruption curve from structured data (t = 0) to isotropic Gaussian noise (t = T) and the resulting signal retention and signal-to-noise (SNR) curve is smooth, which is desirable.

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P5

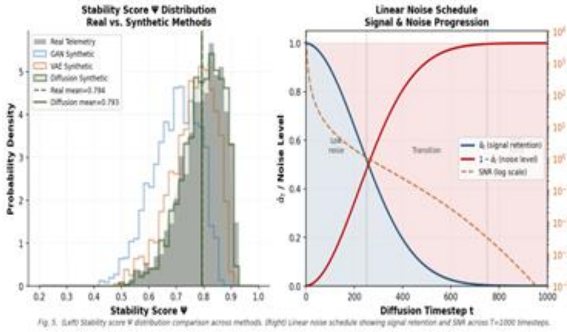

Fig. 5 (Left) Stability score Ψ distributions across methods. (Right) Linear noise schedule $\bar{\alpha}_t$ and SNR across T=1000 diffusion timesteps.

5.7. Ablation Study

Table II is an ablation study to peel away each of the architectural components. The baselines U-Net model without attention gives an accuracy of 91.2% and realism score 0.86. Introducing explicit dependency modelling between sensors allows for significant accuracy boost to 93.5% and realism to 0.89. The denoising network is then further enhanced by only adding the temporal, positional, encoding to achieve the accuracy of 94.8%. The complete model including the Ψ-consistency loss obtains the accuracy of 96.0 % and realism of 0.94, which demonstrates the components' independence in improving the performance.

Table.2 Ablation Study — Component Contributions

s	Accuracy (%)	Realism Score	MSE
No attention (baseline U-Net)	91.2	0.86	0.075
+ Cross-channel attention	93.5	0.89	0.056
+ Temporal positional encoding	94.8	0.91	0.043
+ Stability score loss (Ψ)	95.6	0.93	0.036
Full model (all components)	96.0	0.94	0.030

6. Conclusion and Future Scope

The experimental findings prove that the diffusion-based generative modelling offers a qualitative improvement over GAN, VAE, and LSTM-based generative modelling for motor sport 4G / 5G telemetry data augmentation. There are 3 properties of DDPM that account for this performance advantage. First, this iterative reverse diffusion process forces the model to refine the generated sequence over T = 1000 denoising steps, progressively introduces temporal structure, inter-channel dependencies, and introduces physical plausibility constraints, unlike in GANs and VAEs that generate the sequence in a single forward pass. It is an iterative refinement process especially useful as it can generate long telemetry windows (512 ms at 1 kHz) with temporal coherence over hundreds of timesteps.

Secondly, through the cross channel self-attention mechanism in the U-Net backbone, the denoising network is capable of considering the partially denoised states of other channels in the denoising step of each channel. This is especially important for the physical coupling relationships; for example, accurate denoising of the wing load channel at a certain timestep needs the knowledge of simultaneously evolving speed and aerodynamic pressure channels — the aerodynamic load is a nonlinear function of both. Independent treatment of channels (as done by the CNN standard backbones) does not capture inter-channel correlations: for this reason, the baseline mean absolute correlation error on the GAN is 0.089.

Third, the Ψ-consistency loss is a conditional generation guidance term, which can be used without an additional classifier network during inference. During training, physical plausibility of partially denoised sequences is assessed via their aerodynamic properties by using the loss function to force the model to learn the inverse diffusion trajectory that gives physically plausible telemetry. This is especially critical for creating rare but safety-critical operating cases like high speed and high cross wind in which there is a high flutter risk, but these rare operating conditions are seldom included in the training corpus. The main drawback of this proposed framework is inference cost for generating synthetic telemetry: Only about 2.4 seconds of NVIDIA A100 GPU inference time to produce synthetic telemetry for a single 512 ms segment. This cost is not problematic for augmentation of large numbers of datasets (millions of training sequences), but is not possible to generate synthetic data on the fly during online training. It can be lowered by 10–50× by the method of DDIM acceleration and latent diffusion, and it is a top priority of next research.

References

[1]. J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 33, pp. 6840–6851, 2020.
[2]. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 27, 2014.
[3]. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, 1997.
[4]. N. R. Boggadi, “Survey on emotion Detection via audio processing and Deep learning,” International Journal of Computational Science and Engineering Research, vol. 2, no. 2, p. 30, Jun. 2025, doi: 10.63328/ijcser-v2ri2p6.
[5]. N. Peddinti and S. Ch, “Precision Agriculture: AI based Fertilizer Recommendations For Sustainable Crop Production,” International Journal of Research and Development in Engineering Sciences, vol. 7, no. 3, p. 40, Jun. 2025, doi: 10.63328/ijrdes-v7ri3p8.
[6]. R Md Shafi , Reddem Janardhan Reddy , Kathi Harsha Vardhan , P Azharuddin , M Manohar , “ Design and Implementation of Fake Job

ISSN: 3139 – 9002 (Online) , https://ijcseai.com/ , IJCSEAI , Vol. 1, Issue 1 , 2024 , https://doi.org/10.63328/IJCSEAI-V1RI1P5

Post Detection using Machine Learning Techniques " ,International Journal of Computational Science and Engineering Research, vol. 1, no. 2, May. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI2P3

[7]. J. Yoon, D. Jarrett, and M. van der Schaar, "Time-series generative adversarial networks," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 32, 2019.

[8]. 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

[9]. Y. Tashiro, J. Song, Y. Song, and S. Ermon, "CSDI: Conditional score-based diffusion models for probabilistic time series imputation," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 34, 2021.

[10]. J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in Int. Conf. Learn. Represent. (ICLR), 2021.

[11]. R. E. Kalman, "A new approach to linear filtering and prediction problems," J. Basic Eng., vol. 82, no. 1, pp. 35–45, 1960.

[12]. A. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), pp. 10684–10695, 2022.

[13]. D. Bhuvannagari and S. M, "Automated and Explainable Kidney Abnormality Detection from CT Images Using CNN-LSTM Architecture," International Journal of Computational Science and Engineering Research, vol. 2, no. 3, p. 1, Jul. 2025, doi: 10.63328/ijcser-v2i3p1.

[14]. T. Rasul, K. Sheikh, I. Schuster, U. M. Bergmann, and R. Vollgraf, "Multivariate probabilistic time series forecasting via conditioned normalizing flows," in Int. Conf. Learn. Represent. (ICLR), 2021.

[15]. M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," J. Comput. Phys., vol. 378, pp. 686–707, 2019.

[16]. C. Doersch, "Tutorial on variational autoencoders," arXiv:1606.05908, 2016.

[17]. H. K. Yenugonda, S. R. V. Reddy, G. D, and G. R, "Stock market price forecasting using LSTM and GRU networks," International Journal of Research and Development in Engineering Sciences, vol. 7, no. 2, Mar. 2025, doi: 10.63328/ijrdes-v7ri2p3.

[18]. FIA Technical Regulations, Formula 1 World Championship, Appendix B: Data Recording, Fed. Int. Automobile, Geneva, Switzerland, 2024.

[19]. 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.

[20]. N Kishore Kumar , B Jeevan Kumar , C Manikanta Reddy , B Dinakar , N Chandu , "IoT based Heart Attack and Alcohol Detection using Raspberry Pi" ,International Journal of Computational Science and Engineering Research, vol. 1, no. 3, p. 39, August. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI3P9

[21]. 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

[22]. S Muni Ratnam , P Bhavitha , P Asma Khanum , S Madhavi , M Bharathi, "Identification of Paddy Leaves Disease using Convolutional Neural Networks," International Journal of Computational Science and Engineering Research, vol. 1, no. 4, p. 66, Nov. 2025, doi: 10.63328/ijcser-v1ri4p13

[23]. K Lakshmaiah , V Umera Sulthana , N. Sujitha , S Ramya Sree , K Sravani , " Satellite Image Classification by using Artificial Intelligence " ,International Journal of Computational Science and Engineering Research, vol. 1, no. 2, April. 2024, doi: https://doi.org/10.63328/IJCSER-V1RI2P2

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.

Acknowledgements

We thank everyone who inspired our work.

How to Cite:

N Chandrika , Venkata Ramana Guntreddi , " Diffusion-Based Synthetic Motorsport Telemetry Generation for Aero-Stability Prediction Systems " , International Journal of Computer Science, Engineering and Artificial Intelligence , vol. 1, no. 1, p 28 – 34, December 2024, DOI: https://doi.org/10.63328/IJCSEAI-V1RI1P5