

A Neural ODE Approach to Aircraft Flight Dynamics Modelling

Gabriel Jarry*, Ramon Dalmau*, Xavier Olive†, Philippe Very*

*EUROCONTROL
Aviation Sustainability Unit (ASU)
Brétigny-sur-Orge, France

†ONERA – DTIS
Université de Toulouse
Toulouse, France

Abstract—Accurate aircraft trajectory prediction is critical for air traffic management, airline operations, and environmental assessment. This paper introduces NODE-FDM, a Neural Ordinary Differential Equations-based Flight Dynamics Model trained on Quick Access Recorder (QAR) data. By combining analytical kinematic relations with data-driven components, NODE-FDM achieves a more accurate reproduction of recorded trajectories than state-of-the-art models such as a BADA-based trajectory generation methodology (BADA4 performance model combined with trajectory control routines), particularly in the descent phase of the flight. The analysis demonstrates marked improvements across altitude, speed, and mass dynamics. Despite current limitations, including limited physical constraints and the limited availability of QAR data, the results demonstrate the potential of physics-informed neural ordinary differential equations as a high-fidelity, data-driven approach to aircraft performance modelling. Future work will extend the framework to incorporate a full modelling of the lateral dynamics of the aircraft.

Keywords—Aircraft performance modelling, Neural-ODE, trajectory prediction, physics-informed learning

I. INTRODUCTION

Accurate aircraft trajectory prediction is essential to air traffic management (ATM), airline operations, and aviation environmental assessment. Precise trajectory models are indispensable for conflict detection and resolution, sector capacity planning, and trajectory-based operations (TBO). Beyond operational safety and efficiency, high-fidelity predictions enable robust estimates of fuel consumption, emissions, and noise, thereby supporting the development of sustainable aviation strategies.

The most widely adopted reference for aircraft performance modelling is the Base of Aircraft Data (BADA) [1], [2], developed by EUROCONTROL. BADA provides a standardized set of equations and performance parameters to represent aircraft dynamics under nominal conditions. Its analytical formulation offers transparency and robustness, which explains its longstanding use in both operational and research contexts. Nevertheless, such models inevitably rely on simplifying assumptions and manufacturer-provided data, which can lead to divergences between simulated trajectories and those actually flown, particularly in speed profiles and mass dynamics as captured by Automatic Dependent Surveillance-Broadcast (ADS-B) measurements. Alternative open-source frameworks, such as OpenAP [3], have emerged to complement BADA by providing flexible, data-driven tools for trajectory modelling and validation.

This study introduces NODE-FDM, a novel data-driven Flight Dynamics Model that integrates large-scale Quick Access Recorder (QAR) data with physics-informed neural architectures. We focus first on the vertical and speed dimensions of the trajectory, including altitude profiles, flight path angle, and true airspeed. This paper does not address lateral trajectory management (e.g., route choice and turn dynamics), which we leave for future investigations. We apply the model to the Airbus A320-214 aircraft.

The proposed model combines analytical kinematic relations, regression-based propulsion and aircraft angle modelling, and a Neural Ordinary Differential Equations (ODE) formulation to predict the temporal evolution of altitude and speed with high fidelity. We evaluate its predictive performance against a BADA-based trajectory generation methodology (BADA4 performance model combined with trajectory control routines, hereafter referred to as BADA), focusing on trajectory reconstruction accuracy and suitability for environmental impact assessments. These assessments include developing CO₂ and non-CO₂ emissions inventories, which often require robust trajectory predictions even when only partial datasets or pre-flight planning information are available.

The remainder of this paper is organized as follows. Section II reviews the literature on aircraft trajectory prediction, covering model-based, data-driven, and hybrid physics-informed approaches, and discusses operational use-cases and open challenges. Section III describes the QAR datasets and preprocessing steps used to train the model. Section IV details the NODE-FDM methodology, including its modular architecture, state representation, neural and analytical layers, and training objectives. Section V presents quantitative results, benchmark comparisons with BADA, and analysis of trajectory reconstruction and fuel consumption estimation. Section VI discusses the limitations of the proposed approach, and Section VII concludes with a summary of findings and directions for future work.

II. LITERATURE REVIEW

Trajectory prediction and generation underpin a wide range of ATM and airline operations, from conflict detection and sequencing to environmental assessment and strategic planning. Two complementary paradigms dominate: model-based approaches grounded in aircraft performance and point-mass dynamics, and data-driven methods that learn spatio-temporal regularities from large-scale surveillance and weather datasets.



Physics-based point-mass formulations and performance databases remain the historical backbone of operational trajectory prediction. The EUROCONTROL BADA family provides aircraft-type parameters and procedures for 4D trajectory computation under forecast conditions [1], [2], while OpenAP extends these capabilities to the research community with an open-source framework [3], [4]. Beyond canonical models, the literature formalizes trajectory representation, compression, and optimization using splines, principal component analysis, and connections to optimal control and wavefront propagation under constraints [5]. Recent efforts by Poll and Schumann refine performance characteristics and cruise fuel modelling from aerodynamic theory and empirical data [6], [7]. These approaches provide interpretability and controllability, but rely on accurate state and intent estimates and may under-represent operational variability.

Data-driven methods have achieved notable gains in 4D trajectory prediction using supervised learning with surveillance and meteorological features. Early work by Ayhan and Samet demonstrated machine-learning-based prediction of arrival times and trajectories using historical and weather data [8]. Sequence models then became prevalent: Liu and Hansen proposed a seq2seq long short-term memory (LSTM) [9], Pang et al. combined convolutional encoders of convective weather “cubes” with recurrent decoders [10], and Ma and Tian introduced a convolutional neural networks (CNN)–LSTM on ADS-B data [11]. Online frameworks fuse point-mass physics, BADA parameters, and ADS-B conformance to update intent and improve estimated time of arrival precision [12]. Generative models address multi-modality and error accumulation: generative adversarial network (GAN)-based forecasters produce full-sequence 4D trajectories [13], variational autoencoders (VAE)/GAN generators improve realism in sparse regimes [14], privacy-preserving LSTM–GANs synthesize flows while mitigating re-identification risk [15], and imitation-learning methods such as TrajGAIL learn realistic route choices and flow-level statistics [16]. Data-driven models excel at capturing variability and airline-specific practices, though generalization off-distribution and physical guarantees remain challenging [17].

Hybrid strategies combine the adaptability of data-driven approaches with the consistency of physics-based models, either by calibrating physics-based predictors using data or by embedding known dynamics into neural architectures. Calibration work has improved climb predictions by inferring aircraft-specific parameters from radar data [18]. Physics-informed learning enforces equations of motion and aerodynamic limits via loss functions or architectural constraints. Neural ODEs formalize these approaches as continuous-time dynamical systems, trained efficiently with adjoint methods and evaluated using ODE solvers [19], [20]. Aerospace applications include learning residual dynamics under icing [21], hypersonic glide vehicle prediction [22], and differentiable predictive control for unmanned aerial vehicles [23]. Practical challenges remain regarding stiffness, training cost, stability, and hybrid dynam-

ics with discrete modes [24].

Beyond prediction, ATM applications also require the generation of realistic trajectories and flows with quantified realism and flyability. Statistical density models, such as Gaussian mixture models and vine copulas, combined with dimensionality reduction, reproduce flow statistics and go-around patterns while balancing realism and tractability [25], [26]. Evaluation frameworks assess operational realism, statistical coherence, similarity to observations, and flyability via simulator replay [27]. Interactive FPCA-based pipelines support clustering, deformation, and consistent generation for what-if analyses, for instance in noise footprint studies [28]. Aviation-focused GANs synthesize approach paths and identify atypical trajectories for safety management [29], while recent adversarial maximization–minimization methods assemble trajectories from mined manoeuvres under performance constraints, improving fidelity, diversity, and controller consistency [30].

Overall, model-based methods provide transparency and physical consistency but struggle with latent intent and operational diversity, data-driven predictors capture variability and multi-modality yet must address generalization, uncertainty, and physical plausibility, and physics-guided hybrids, including Neural ODE frameworks, offer a promising middle ground for robustness and realism. Within this spectrum, our approach falls into the physics-guided hybrid category, combining large-scale QAR data with analytical kinematics, regression-based propulsion modeling, and a Neural ODE formulation. Compared with purely model-based approaches such as BADA, it emphasizes accurate reconstruction of altitude and speed trajectories while remaining suitable for environmental impact assessments.

III. DATA SOURCES

The dataset used in this study is derived from Quick Access Recorder (QAR) data collected on nine different Airbus A320-214 aircraft operated by the airline Aigle Azur over two years. QAR systems provide high-frequency recordings, with the sampling rate depending on the parameter, typically ranging from 1 to 4 Hz. They capture a wide range of flight parameters, including aircraft state variables (e.g., altitude, true airspeed, Mach number), control inputs (e.g., thrust setting, flap configuration, landing gear status), and environmental conditions (e.g., temperature, wind components, runway elevation). Compared with publicly available surveillance data such as ADS-B, or radar data managed by air navigation service providers, QAR recordings offer higher temporal resolution and additional parameters critical for detailed aircraft performance modelling. However, QAR datasets are owned and maintained by airlines and are not routinely shared publicly, which limits access and restricts the ability to generalise models across operators.

The raw QAR data underwent several preprocessing steps to ensure consistency and suitability for machine learning applications. First, data were subsampled to 0.25 Hz (one point every four seconds). Second, all parameters were converted into standard SI units, and statistical measures (mean and standard deviation) were computed across the training set to



provide normalisation capabilities for the neural networks. Each flight was then divided into non-overlapping sequences of fixed length (60 time steps, equivalent to four minutes) to form the samples used for training and validation. Table I summarises the features extracted from the QAR data and used in the model.

The final dataset comprises 5,000 training flights, 500 validation flights, and 500 test flights, distributed across five airframes for training, two for validation, and two for testing, with no overlap of aircraft between the subsets. Each flight is divided into 4 minutes non-overlapping sequences and is represented by a multi-source tensor that groups three categories of variables:

- **state variables** $x(t)$ describe the instantaneous state of the aircraft and its motion. They include altitude, along-track distance, true airspeed, flight path angle, and estimated aircraft mass.
- **control variables** $u(t)$ reflect the pilot's or autopilot's inputs and mode selections. They include speed and altitude selected, flap setting, and landing gear status. When the speed target is set in Mach, the QAR also records the equivalent calibrated airspeed target (V_{sel}); in this study, only the V_{sel} representation is retained. Similarly, when the aircraft is commanded to follow a constant slope, the QAR does not log the corresponding vertical speed target explicitly; we compute the selected vertical speed ($V_{z,sel}$) from the recorded ground speed and target slope.
- **context and intermediate variables** $e(t)$ provide environmental conditions and derived parameters relevant to aircraft dynamics. They include static air temperature, horizontal wind components, engine parameters (e.g., N1 fan speed, thrust setting), and aerodynamic angles such as angle of attack.

IV. METHODOLOGY

The proposed NODE-FDM is organised as a modular, physics-informed architecture designed to estimate the temporal evolution of aircraft states, as illustrated in Figure 1. The model combines analytical equations, regression-based neural layers, and a Neural ODE formulation to ensure both physical consistency and flexibility, with a focus on the vertical profile and speed management. In this study, the scope is limited to the longitudinal dimensions of the trajectory: altitude, true airspeed, flight path angle, aircraft mass, and along-track distance. These variables play a central role in determining fuel consumption, emissions, and noise, making them particularly relevant for environmental assessment. Lateral dynamics (e.g. track angle and horizontal position) are not included at this stage and are left for future investigation.

The state vector is defined as:

$$x(t) = \{h(t), d(t), \gamma(t), V_{TAS}(t), m(t)\}, \quad (1)$$

where h denotes the altitude, d the along-track distance, γ the flight path angle, V_{TAS} the true airspeed, and m the

aircraft mass. These variables are either logged directly by the QAR from onboard instruments or derived from sensor-based computations available in the aircraft systems. This state vector corresponds to the standard point-mass formulation, widely applied in aircraft trajectory prediction and optimization.

In the implementation, a structured layer is defined and serves as a reusable building block within the model architecture. Each structured layer consists of three main components: an input normalizer, a backbone, and multiple output heads. The inputs are first normalized according to their empirical statistics and concatenated into a fixed-size representation. This representation is processed by the backbone, a two-layer fully connected network with 24 hidden neurons per layer, followed by ReLU activations. The extracted features are then passed to task-specific output heads, each implemented as a single-layer perceptron. This design allows the model to share a common latent representation while adapting flexibly to heterogeneous output variables (continuous or binary). Finally, the outputs of each head are denormalized to recover values on the original physical scale. We deliberately excluded dropout layers from our Neural ODE architecture. Unlike in discrete architectures (e.g., ResNets), applying standard dropout within continuous-time models is known to destabilize training, as the stochastic masking conflicts with the smooth dynamics required for stable ODE integration [31].

The complete model is composed of four main layers:

- a **trajectory layer** (analytical): deterministic computation of vertical speed, Mach number, and ground speed from current state and difference of selected vs real speed and altitude.
- an **angle layer** (neural network): a structured layer predicting angle-related parameters such as angle of attack and pitch from control inputs and aircraft states.
- an **engine layer** (neural network): a structured Layer predicting engine-related parameters such as fuel flow and fan speed (N_1) from control inputs and aircraft states.
- a **derivative layer** (neural network): a structured layer that estimates the derivatives dV_{TAS} , $d\gamma$ of the state vector, while the remaining derivatives are obtained through direct usage of existing variables: V_{GS} , $-m_{fuel}$, and V_z . Based on the concatenated state, control, context and intermediate variables, this block provides dV_{TAS} and $d\gamma$, which are then integrated with an ODE solver (Euler scheme in this work) to generate continuous-time trajectories.

The **trajectory layer** provides deterministic conversions from the state variables and context conditions to additional kinematic variables, ensuring consistency with physical laws. Specifically:

$$V_z = V_{TAS} \cdot \sin(\gamma), \quad (2)$$

$$M = \frac{V_{TAS}}{a}, \quad a = \sqrt{\gamma_{air} R T_{OAT}}, \quad (3)$$

$$V_{GS} = V_{TAS} - V_{\parallel}, \quad (4)$$



TABLE I: List of parameters used in the model with original QAR units and converted SI units.

| Feature | Symbol | QAR Unit | SI Unit | Description |
|---|---------------------|----------|---------|--|
| State variables $x(t)$ | | | | |
| altitude | h | ft | m | Standard altitude above mean sea level |
| distance along track | d | nm | m | Along-track distance from departure |
| flight path angle | γ | deg | rad | Angle between velocity vector and horizon |
| true airspeed | V_{TAS} | kt | m/s | Speed relative to the surrounding air mass |
| mass | m | kg | kg | Estimated aircraft mass |
| Control variables $u(t)$ | | | | |
| selected altitude | h_{sel} | ft | m | Altitude target set in the FMS/autopilot |
| selected speed | V_{sel} | kt | m/s | Speed target set in the FMS/autopilot |
| selected vertical speed | $V_{z,sel}$ | ft/min | m/s | Vertical speed target set in the FMS/autopilot (if set else 0.0) |
| flap setting | f | index | index | Current flap configuration 0,1,2,3,4 (FULL) |
| landing gear status | g | binary | binary | Gear up (0) / down (1) |
| speed brake command | SB | binary | binary | Speed brake retracted (0) / deployed (1) |
| Context and intermediate variables $e(t)$ | | | | |
| Context variables $e_0(t)$ | | | | |
| outside air temperature | T_{OAT} | °C | K | Ambient static air temperature |
| headwind component | V_{\parallel} | kt | m/s | Wind component along the trajectory |
| crosswind component | V_{\perp} | kt | m/s | Wind component orthogonal to the trajectory |
| Trajectory variables $e_1(t)$ | | | | |
| Mach number | M | – | – | Ratio of V_{TAS} to speed of sound |
| calibrated airspeed | V_{CAS} | kt | m/s | Airspeed derived from dynamic pressure |
| vertical speed | V_z | ft/min | m/s | Rate of climb or descent |
| ground speed | V_{GS} | kt | m/s | Speed relative to the ground |
| selected altitude difference | $h_{sel} - h$ | ft | m | Difference between selected and current altitude |
| selected speed difference | $V_{sel} - V_{CAS}$ | kt | m/s | Difference between selected speed and CAS |
| Aircraft angles $e_2(t)$ | | | | |
| angle of attack | α | deg | rad | Angle between chord line and relative wind |
| pitch angle | θ | deg | rad | Angle between aircraft longitudinal axis and horizon |
| Engine parameters $e_3(t)$ | | | | |
| N1 (fan speed) | N_1 | %RPM | %RPM | Low-pressure compressor rotation speed |
| fuel flow | m_{fuel} | kg/h | kg/s | Total fuel consumption rate |

where V_z is the vertical speed, M the Mach number, a the local speed of sound, γ_{air} the ratio of specific heats, R the perfect gas constant, T_{OAT} the outside air temperature, and V_{\parallel} the headwind component.

The evolution of the state vector is modelled as:

$$\frac{dx}{dt} = f_{\theta}(x(t), u(t), e(t)), \quad (5)$$

where $x(t)$ denotes the aircraft state vector at time t , $u(t)$ represents the control variables, and $e(t)$ corresponds to the context and intermediate variables. f_{θ} is represented by the derivative layer (a neural network). We compute the predicted trajectories through numerical integration of the ODE using an explicit Euler solver, implemented via the `odeint` function from the `torchdiffeq` library [19].


The model is trained by minimizing the difference between predicted and observed trajectories using a composite loss function. This loss combines errors across several groups of outputs, namely the state variables, aircraft angles, and engine parameters:

$$\mathcal{L} = \sum_{i \in \mathcal{C}} \alpha_i \ell_i(y_i^{\text{pred}}, y_i^{\text{true}}), \quad (6)$$

where ℓ_i is the error metric for feature i , α_i is a weighting coefficient, and $\mathcal{C} = \{x(t), e_2(t), e_3(t)\}$ (see Table I). In other words, the loss function directly compares the model's predicted variables with those recorded by the QAR. The resulting error signals are then used to adjust the weights of the neural ODE, gradually improving its prediction accuracy.

In this paper, ℓ_i corresponds to a mean squared error (MSE) and the weighting coefficients α_i are chosen inversely proportional to the empirical standard deviation of each variable, in order to mitigate differences in scale and absolute value. This formulation allows balancing between accuracy in integrated state values, and regressed engine and aircraft variables.

The model is trained on one GPU (RTX A6000 Ada Gen) using PyTorch for 1000 epochs with batch size of 500 and checkpoints on the validation set. Optimization was performed with the AdamW optimizer ($1r = 1 \times 10^{-4}$, `weight_decay` = 1×10^{-4} corresponding to L2 regularization).



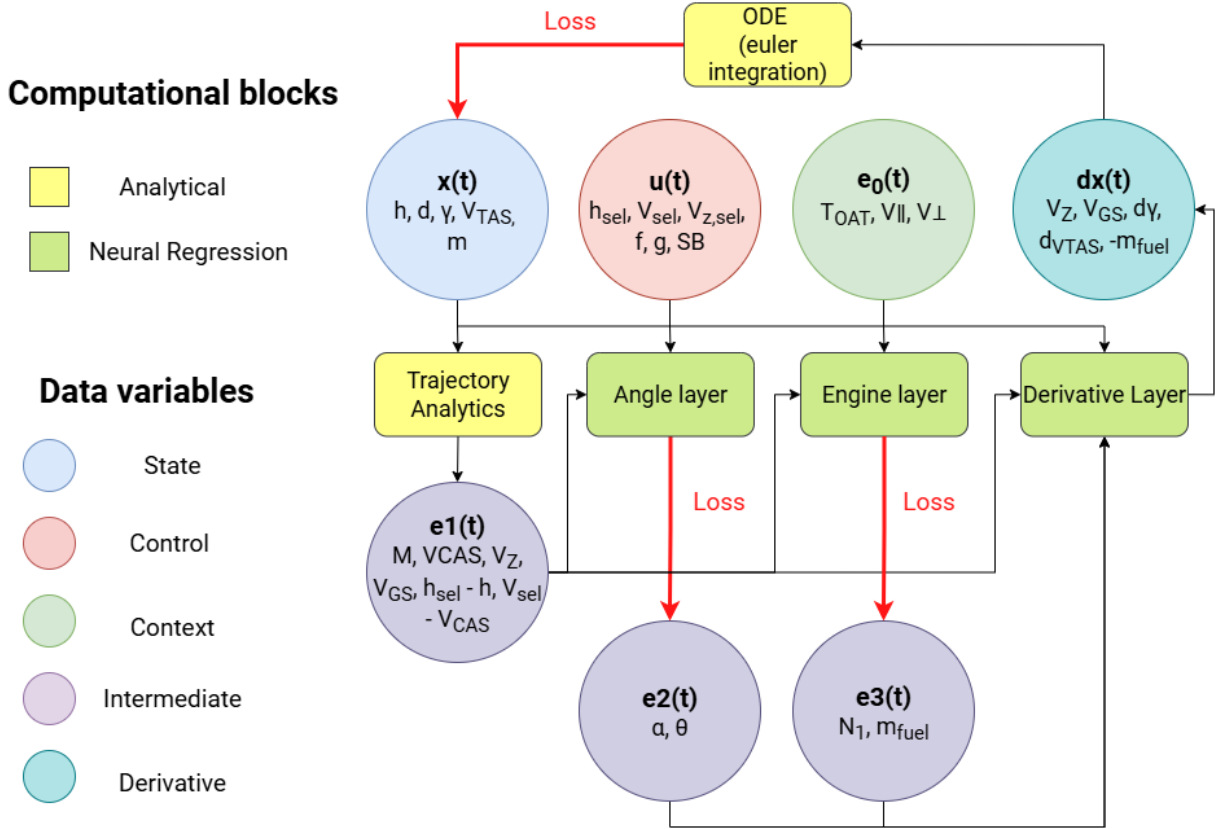


Figure 1: Architecture of the proposed model

V. RESULTS

To generate physically consistent benchmark trajectories from operational flight data, we used a BADA-based methodology combining the pyBADA 4.2 library with the TCL.py (Trajectory Calculation Library) module, together with control inputs mapped from QAR data. In this setup, the BADA4 aircraft performance model was instantiated for the A320-214, while trajectory propagation relied on the TCL routines `accDec_time`, `constantSpeedLevel`, `constantSpeedRating_time`, and `constantSpeedROCD_time`, applied over fixed 4-second intervals. For each QAR (Quick Access Recorder) record, the current aircraft state (altitude, true airspeed, Mach number, temperature deviation from ISA, configuration, and mass) was extracted and mapped to the corresponding BADA4 input variables. Calibrated airspeed and Mach number were derived from standard atmosphere relationships, while along-track wind and speed-brake deployment were also incorporated to ensure consistency with the recorded operational conditions. Importantly, the evaluation presented here concerns the entire BADA-based trajectory generation pipeline, not only the underlying BADA4 performance model.

Depending on the identified flight phase (climb, cruise, or descent) and the difference between commanded and actual speeds, the appropriate TCL routine was selected: `constantSpeedLevel` for cruise,

`constantSpeedRating_time` for climb or descent at constant speed, `accDec_time` when speed adjustments were required, and `constantSpeedROCD_time` when a vertical speed target had been engaged. At each integration step, the most recent simulated state, i.e., altitude, airspeed, rate of climb or descent, and mass, was fed back into the next computation, thereby ensuring a continuous and dynamically consistent simulation of the trajectory.

For testing purposes, complete trajectories (from take-off to landing) were generated in simulation by propagating the aircraft state from the first recorded point p_0 after take-off until the last point before landing. The propagation was performed over the same time horizon as the corresponding QAR trajectory, using the full sequence of recorded control inputs (speed targets, vertical mode selections, and configuration changes). This approach ensures that the simulated trajectory reproduces the temporal structure of the operational flight while remaining entirely model-driven, thereby providing a consistent benchmark for evaluating prediction performance.

The comparative analysis of NODE-FDM and BADA-based simulations, presented in Table II, highlights several clear patterns. Overall, NODE-FDM achieves lower mean absolute errors across altitude, airspeed, and mass compared with BADA, with the most significant gains observed during the climb and descent phases. For instance, the altitude error of NODE-FDM remains below 100 m in climb and around 200 m

TABLE II: Comparison of NODE-FDM and BADA performance across phases. Metrics are reported as mean (standard deviation), based on 849,490 trajectory points across 500 flights.

| parameter | phase | NODE-FDM | | | BADA | | |
|----------------------|--------------|-----------------|---------------|-----------------|-----------------|----------------|------------------|
| | | MAE | MAPE (%) | ME | MAE | MAPE (%) | ME |
| altitude [m] | All phases | 61.90 (145.54) | 2.49 (333.62) | 0.47 (158.15) | 167.47 (402.75) | 5.03 (315.60) | -75.45 (429.61) |
| | Climb | 93.03 (121.98) | 3.36 (569.52) | 57.64 (142.16) | 214.69 (180.19) | 6.14 (269.86) | 175.79 (218.31) |
| | Level Flight | 8.70 (33.50) | 0.26 (6.16) | -1.51 (34.57) | 12.32 (67.80) | 0.28 (9.20) | -8.26 (68.42) |
| | Descent | 182.04 (246.23) | 7.95 (467.70) | -49.49 (302.19) | 560.15 (713.73) | 17.39 (637.75) | -509.53 (750.69) |
| true air speed [m/s] | All phases | 1.35 (2.03) | 0.75 (1.27) | -0.04 (2.44) | 2.73 (6.16) | 1.64 (5.18) | -1.34 (6.61) |
| | Climb | 1.70 (1.93) | 0.95 (1.27) | 0.87 (2.42) | 3.31 (3.33) | 2.02 (2.66) | 1.86 (4.31) |
| | Level Flight | 0.70 (0.75) | 0.32 (0.43) | -0.04 (1.03) | 0.90 (1.62) | 0.45 (1.41) | -0.62 (1.75) |
| | Descent | 2.82 (3.33) | 1.76 (2.02) | -0.92 (4.27) | 7.32 (11.55) | 4.64 (10.15) | -6.51 (12.02) |
| γ [deg] | All phases | 0.24 (0.42) | N/A | 0.05 (0.48) | 0.46 (1.01) | N/A | 0.05 (1.11) |
| | Climb | 0.42 (0.51) | N/A | 0.04 (0.65) | 0.64 (1.11) | N/A | 0.03 (1.28) |
| | Level Flight | 0.08 (0.16) | N/A | 0.03 (0.18) | 0.07 (0.29) | N/A | 0.03 (0.30) |
| | Descent | 0.51 (0.60) | N/A | 0.09 (0.78) | 1.36 (1.49) | N/A | 0.09 (2.01) |
| mass [kg] | All phases | 64.89 (70.44) | 0.10 (0.11) | 7.26 (95.50) | 154.89 (122.70) | 0.25 (0.20) | 152.90 (125.17) |
| | Climb | 22.98 (31.34) | 0.04 (0.05) | 7.12 (38.21) | 60.22 (58.35) | 0.09 (0.09) | 59.65 (58.94) |
| | Level Flight | 74.91 (73.32) | 0.12 (0.12) | 6.06 (104.65) | 170.43 (119.16) | 0.27 (0.20) | 168.63 (121.69) |
| | Descent | 77.29 (74.27) | 0.13 (0.12) | 10.80 (106.64) | 202.96 (130.73) | 0.34 (0.22) | 199.05 (136.61) |

TABLE III: Comparison of consumption prediction performance (Predicted vs. BADA). Metrics are reported as mean (standard deviation), based on 849,490 trajectory points across 500 flights.

| parameter | NODE-FDM | | | BADA | | |
|-------------|---------------|-------------|---------------|-----------------|-------------|------------------|
| | MAE [kg] | MAPE (%) | ME [kg] | MAE [kg] | MAPE (%) | ME [kg] |
| consumption | 83.95 (79.65) | 1.54 (1.13) | 3.01 (115.74) | 163.52 (129.68) | 3.03 (1.75) | -149.67 (145.47) |

in descent, whereas BADA shows noticeably larger deviations, with errors of about 200 m in climb and more than 550 m in descent.

Similarly, NODE-FDM provides a closer reproduction of the recorded airspeed dynamics, with typical deviations around 1.35 m/s, while BADA shows larger discrepancies with more than 2.70 m/s mean absolute error.

When examining the flight path angle, NODE-FDM exhibits a smaller dispersion of errors, although both models encounter challenges in reproducing this variable accurately, reflecting the inherent sensitivity of γ to minor variations in altitude and airspeed. Regarding mass evolution, NODE-FDM closely follows the recorded values, with errors on the order of sixty kilograms, whereas BADA's estimates remain consistently biased by one hundred and fifty kilograms. This difference should be interpreted with caution, as BADA relies on performance coefficients calibrated on aircraft in factory new conditions rather than the specific QAR dataset employed here.


It is important to emphasise that BADA is a well-established and physically consistent performance model, widely adopted in operational and simulation contexts. Its higher errors in this comparison do not indicate a lack of validity; rather, they reflect methodological differences. Here, we are evaluating the entire BADA-based trajectory generation pipeline, which couples the BADA4 performance model with trajectory control routines using operational inputs, rather than the performance model in isolation. BADA is not intended to replicate individual recorded flights with exact control input fidelity, but to

provide a reliable and generalisable representation of aircraft performance across a broad range of conditions. NODE-FDM, by contrast, is trained directly on the same QAR data used for evaluation, which naturally gives it an advantage in reproducing the observed dynamics.

Taken together, these results demonstrate that NODE-FDM can achieve high-fidelity trajectory reproduction when provided with detailed flight data, while BADA remains a robust and reliable reference model, particularly well suited for large-scale studies and operational applications where generality, consistency, and transparency are essential.

The general trends observed in the quantitative analysis are further illustrated by the example trajectory shown in Figure 2. In this case, NODE-FDM more closely reproduces the recorded QAR dynamics across altitude, speed, and mass, whereas BADA exhibits larger deviations, particularly during descent. Nonetheless, both models capture the overall trajectory evolution and phase transitions, confirming the consistency between the aggregated statistical results and individual trajectory examples. The comparison presented in Table III illustrates the implications of different modelling approaches for estimating fuel consumption in the context of environmental impact assessments.

NODE-FDM provides a markedly more precise estimate of absolute consumption, with lower mean absolute and percentage errors compared to BADA. This improvement arises from the model's integration of aircraft- and trajectory-specific dynamics, as well as a direct loss on the fuel flow, which reduces biases. In particular, NODE-FDM accounts for the



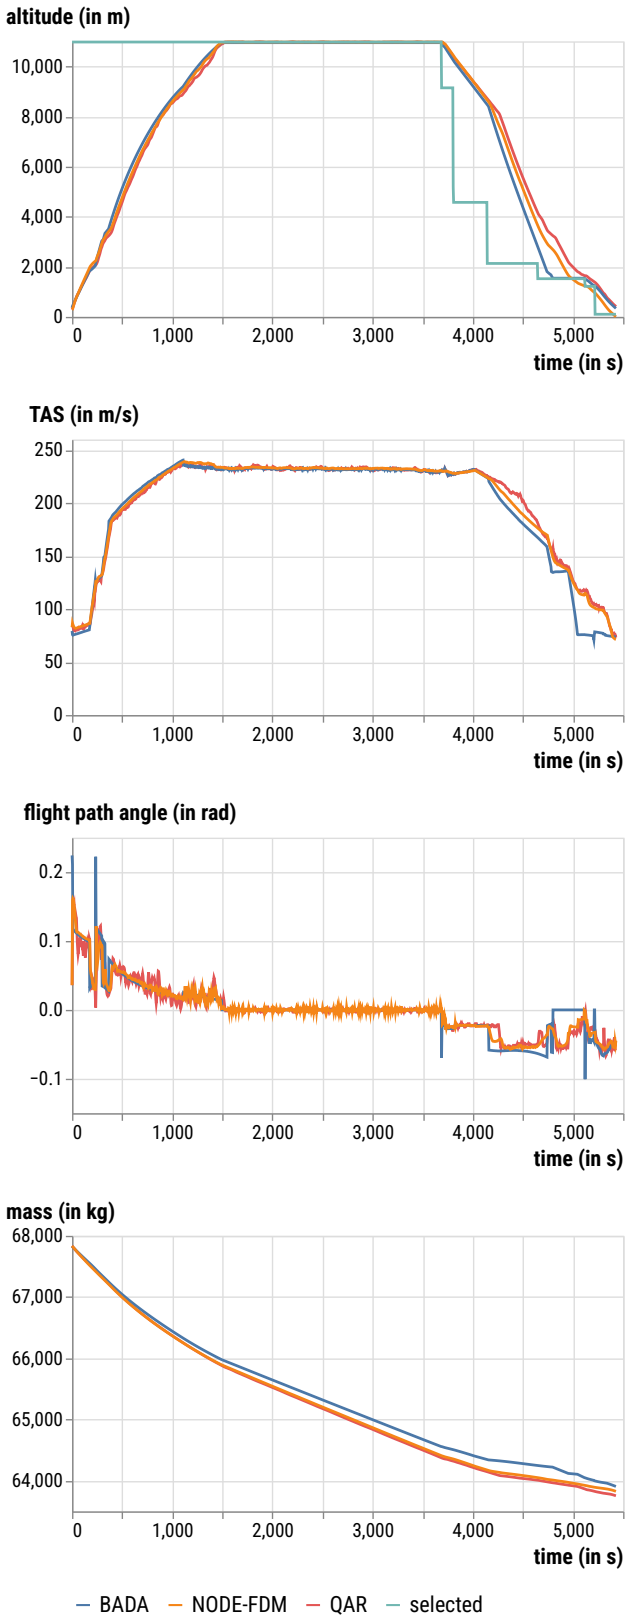


Figure 2: Example of trajectory reproduction from QAR data. The figure compares the recorded QAR trajectory (red) with the NODE-FDM simulation (orange) and the BADA-based benchmark (blue) across altitude, true airspeed, flight path angle, and mass. The selected altitude is shown as a blue-green line. Time is represented on the x-axis in seconds (one point every 4 s).

degradation of engine performance due to in-service fleet ageing, which BADA does not capture. Consequently, BADA tends to systematically under-estimate absolute fuel burn [32].

Nevertheless, despite this limitation, BADA successfully reproduces relative consumption trends across flights. The correlation between BADA-derived consumption and the actual reference remains strong, demonstrating its utility when the objective is to capture variability and patterns at a large scale, such as for trend analysis, relative comparisons, or traffic scenario studies.

VI. DISCUSSION AND LIMITATIONS

A first limitation of the present work is that NODE-FDM is restricted to the vertical-longitudinal dynamics of the trajectory, i.e., the evolution of altitude and speed along the aircraft's longitudinal axis, expressed through true airspeed. In this formulation, the aircraft is effectively represented in a body-fixed reference frame aligned with its nose, so only motion in the vertical plane is captured. Consequently, lateral behaviour, including variations in track angle, banking, and turn dynamics, is not represented. The framework therefore cannot yet support fully four-dimensional trajectory prediction, nor operational applications in which lateral manoeuvres and route choices play a decisive role. Extending the model to incorporate lateral dynamics represents a natural and necessary avenue for future research.

A second limitation is that the proposed approach is not strictly constrained by physical laws beyond the analytical kinematic relations embedded in the trajectory layer. While the Neural ODE formulation produces smooth and coherent trajectories, it does not guarantee adherence to aerodynamic or performance bounds. Unrealistic behaviours may therefore arise under unusual control inputs or when the model is applied outside the distribution of the training dataset. Introducing stronger physics-based constraints or regularisation terms could mitigate these risks and improve generalisation.

Finally, reliance on QAR data represents a structural limitation. QAR datasets provide rich, high-resolution information essential for model development and evaluation, but they are not widely accessible at scale. In the present work, the approach has been applied to a single aircraft type and model instance rather than across a full fleet, limiting the scope of the conclusions. This restricted availability poses challenges for developing models intended to generalise across fleets, operators, or aircraft types. Broader access to standardised, large-scale operational datasets would greatly facilitate such transferability.

VII. CONCLUSION

This paper has introduced NODE-FDM, a Neural ODE-based Flight Dynamics Model, illustrated using high-resolution Quick Access Recorder (QAR) data from Airbus A320-214 aircraft. By combining analytical kinematic relations with data-driven components, the model reproduces recorded trajectories more closely than the widely used Base of Aircraft Data (BADA) benchmark, particularly in altitude,

speed, and mass evolution. Comparative analysis shows that NODE-FDM substantially reduces mean errors during climb and descent phases while maintaining coherence in level flight, highlighting its potential as a high-fidelity trajectory reconstruction tool.

Beyond improvements in accuracy, these results demonstrate the benefits of hybrid modelling strategies that embed physical structure within learning architectures. Nevertheless, several limitations remain: the absence of lateral dynamics, the lack of strict physical constraints on the learned components, and the restricted availability of QAR data, which limits large-scale applicability across fleets. Addressing these challenges represents a promising avenue for future research, particularly by extending the framework to fully four-dimensional trajectories, integrating stronger physics-based constraints, and leveraging more accessible datasets. Future work will specifically investigate the use of automatic dependent surveillance-broadcast (ADS-B) data, coupled with a simplified version of NODE-FDM, to address scalability and enable deployment at fleet or network level while balancing fidelity and generalisability.

REFERENCES

- [1] A. Nuic and V. Mouillet, "User manual for the base of aircraft data (bada) revision 4," *EUROCONTROL Experimental Centre, EEC Technical/Scientific Report No. 12/11/22-58 v1.3*, 2016.
- [2] A. Nuic, D. Poles, and V. Mouillet, "Bada: An advanced aircraft performance model for present and future atm systems," *International journal of adaptive control and signal processing*, vol. 24, no. 10, pp. 850–866, 2010.
- [3] J. Sun, J. M. Hoekstra, and J. Ellerbroek, "Openap: An open-source aircraft performance model for air transportation studies and simulations," *Aerospace*, vol. 7, no. 8, p. 104, 2020.
- [4] J. Sun, "Openap. top: Open flight trajectory optimization for air transport and sustainability research," *Aerospace*, vol. 9, no. 7, p. 383, 2022.
- [5] D. Delahaye, S. Puechmorel, P. Tsiotras, and E. Féron, "Mathematical models for aircraft trajectory design: A survey," in *Air Traffic Management and Systems: Selected Papers of the 3rd ENRI International Workshop on ATM/CNS (EIWAC2013)*. Springer, 2014, pp. 205–247.
- [6] D. Poll and U. Schumann, "An estimation method for the fuel burn and other performance characteristics of civil transport aircraft in the cruise. part I fundamental quantities and governing relations for a general atmosphere," *The Aeronautical Journal*, vol. 125, no. 1284, pp. 257–295, 2021.
- [7] —, "An estimation method for the fuel burn and other performance characteristics of civil transport aircraft during cruise: part 2, determining the aircraft's characteristic parameters," *The Aeronautical Journal*, vol. 125, no. 1284, pp. 296–340, 2021.
- [8] S. Ayhan and H. Samet, "Aircraft trajectory prediction made easy with predictive analytics," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 21–30.
- [9] Y. Liu and M. Hansen, "Predicting aircraft trajectories: A deep generative convolutional recurrent neural networks approach," *arXiv preprint arXiv:1812.11670*, 2018.
- [10] Y. Pang, H. Yao, J. Hu, and Y. Liu, "A recurrent neural network approach for aircraft trajectory prediction with weather features from sherlock," in *AIAA Aviation 2019 Forum*, 2019, p. 3413.
- [11] L. Ma and S. Tian, "A hybrid cnn-lstm model for aircraft 4d trajectory prediction," *IEEE access*, vol. 8, pp. 134 668–134 680, 2020.
- [12] J. Zhang, J. Liu, R. Hu, and H. Zhu, "Online four dimensional trajectory prediction method based on aircraft intent updating," *Aerospace Science and Technology*, vol. 77, pp. 774–787, 2018.
- [13] X. Wu, H. Yang, H. Chen, Q. Hu, and H. Hu, "Long-term 4d trajectory prediction using generative adversarial networks," *Transportation Research Part C: Emerging Technologies*, vol. 136, p. 103554, 2022.
- [14] X. Chen, J. Xu, R. Zhou, W. Chen, J. Fang, and C. Liu, "Trajvae: A variational autoencoder model for trajectory generation," *Neurocomputing*, vol. 428, pp. 332–339, 2021.
- [15] J. Rao, S. Gao, Y. Kang, and Q. Huang, "Lstm-trajgan: A deep learning approach to trajectory privacy protection," *arXiv preprint arXiv:2006.10521*, 2020.
- [16] S. Choi, J. Kim, and H. Yeo, "Trajgail: Generating urban vehicle trajectories using generative adversarial imitation learning," *Transportation Research Part C: Emerging Technologies*, vol. 128, p. 103091, 2021.
- [17] G. Jarry, R. Dalmau, P. Very, and J. Sun, "On the generalization properties of deep learning for aircraft fuel flow estimation models," *Transportation Research Part C: Emerging Technologies*, vol. 176, p. 105143, 2025.
- [18] R. Alligier and D. Gianazza, "Learning aircraft operational factors to improve aircraft climb prediction: A large scale multi-airport study," *Transportation research part C: emerging technologies*, vol. 96, pp. 72–95, 2018.
- [19] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," *Advances in neural information processing systems*, vol. 31, 2018.
- [20] E. Haber and L. Ruthotto, "Stable architectures for deep neural networks," *Inverse problems*, vol. 34, no. 1, p. 014004, 2017.
- [21] J. Ma, Y. Li, J. Tu, Y. Zhang, J. Ai, and Y. Dong, "Development and implementation of physics-informed neural ode for dynamics modeling of a fixed-wing aircraft under icing/fault," *Guidance, Navigation and Control*, vol. 4, no. 01, p. 2350024, 2024.
- [22] S. Lu and Y. Qian, "Enhanced trajectory forecasting for hypersonic glide vehicle via physics-embedded neural ode," *Drones*, vol. 8, no. 8, p. 377, 2024.
- [23] J. Park and M. Kim, "Data-driven regulator of rotary-wing uav utilizing neural ordinary differential equations," *International Journal of Aeronautical and Space Sciences*, pp. 1–11, 2025.
- [24] L. Ruthotto and E. Haber, "Deep neural networks motivated by partial differential equations," *Journal of Mathematical Imaging and Vision*, vol. 62, no. 3, pp. 352–364, 2020.
- [25] T. Krauth, J. Morio, X. Olive, B. Figueat, and R. Monstein, "Synthetic aircraft trajectories generated with multivariate density models," *Engineering Proceedings*, vol. 13, no. 1, p. 7, 2021.
- [26] T. Krauth, A. Lafage, J. Morio, X. Olive, and M. Waltert, "Deep generative modelling of aircraft trajectories in terminal maneuvering areas," *Machine Learning with Applications*, vol. 11, p. 100446, 2023.
- [27] X. Olive, J. Sun, M. C. R. Murça, and T. Krauth, "A framework to evaluate aircraft trajectory generation methods," in *Proceedings of the 14th USA/Europe Air Traffic Management Research and Development Seminar, Saclay, France, 2021*, pp. 17–19.
- [28] G. Jarry, A. Hassoumi, D. Delahaye, and C. Hurter, "Interactive trajectory modification and generation with fpca," *CEAS Aeronautical Journal*, vol. 13, no. 2, pp. 371–383, 2022.
- [29] G. Jarry, N. Couellan, and D. Delahaye, "On the use of generative adversarial networks for aircraft trajectory generation and atypical approach detection," in *ENRI International Workshop on ATM/CNS*. Springer, 2019, pp. 227–243.
- [30] X. Gui, J. Zhang, X. Tang, D. Delahaye, and J. Bao, "A novel aircraft trajectory generation method embedded with data mining," *Aerospace*, vol. 11, no. 8, 2024.
- [31] J. Lee, Y. Oh, S. Kim, and D. Lim, "Neural regenerative stochastic differential equation: Dropout scheme for neural differential equations."
- [32] G. Jarry, R. Dalmau, P. Very, and J. Sun, "Aircraft fuel flow modelling with ageing effects: From parametric corrections to neural networks," 2025. [Online]. Available: <https://arxiv.org/abs/2509.15736>

