The Aeronautical Journal (2026), 1–31
doi:10.1017/aer.2026.10199

RESEARCH ARTICLE

Development of an icing-resilient autopilot for automated egress of small UAVs from dangerous icing conditions

Zoey Riss and Christoph Deiler

Institute of Flight Systems, German Aerospace Center Brunswick Site: Deutsches Zentrum fur Luft- und Raumfahrt, Braunschweig, Germany
Corresponding author: Christoph Deiler; Email: christoph.deiler@dlr.de

Received: 10 December 2025 UTC; **Revised:** 9 June 2026 UTC; **Accepted:** 10 June 2026 UTC

Keywords: UAV; icing; simulation; flight control; system identification

Abstract
Small fixed-wing unmanned air vehicles (UAVs) weighting under 20 kg are typically not protected against icing. Yet, commercial and military users find interest in operating this class of platform under cold climates where encounters with icing conditions are likely. This paper introduces a new concept named sense-and-egress that aims at safeguarding small UAVs in case of an inadvertent icing encounter. A short overview of the system's operating principles and a plausible architecture candidate are given. Then, the benefits and drawbacks of the proposed solution are assessed with respect to both unprotected (no ice detection or deicing devices active) and fully ice-protected vehicles. Enabling technologies and knowledge gaps are briefly explored in order to gauge the maturity of the concept. The non-linear flight dynamics model of an existing UAV is built up using results from numerical aerodynamics software. The model is then improved and validated using system identification from existing flight test data. A bespoke path-following autopilot is designed to control the UAV during the escape manoeuvre. The selected control system structure and the tuning procedures for both the attitude controller and the total energy control system are detailed. The ability of the baseline autopilot to perform a simple egress manoeuvre with both clean and iced airframes is assessed in simulation. Simple yet effective solutions are implemented to improve the autopilot performance in icing while preserving its gust resistance capabilities.

Nomenclature

a_X, a_Y, a_Z	translational accelerations in the body coordinate system
BVLOS	beyond visual line of sight
CFD	computational fluid dynamics
$C_{L,0}, C_{D,0}$	lift and drag coefficients at zero angle-of-attack
$Cl_{\delta L}$	aileron effectiveness coefficient
$Cm_{\delta M}$	elevator effectiveness coefficient
CG	centre of gravity
CS	certification specifications
EASA	European Union Aviation Safety Agency
EKF	extended Kalman filter
FCS	flight control system
g	acceleration due to gravity
GNSS	Global Navigation Satellite System
h	geodetic altitude
HTP	horizontal tail plane

IMU	inertial measurement unit
IPS	icing protection system
lat, lon, h	latitude, longitude and altitude in the geographic coordinate system
LTI	linear time-invariant
MVD	median volume diameter
PID	proportional integral derivative
SAE	Society of American Engineers
TAS	true airspeed
TECS	total energy control system
u, v, w	translational velocities in the body coordinate system
u_A, v_A, w_A	translational velocities in the wind coordinate system
UAV	unmanned air vehicle
p, q, r	roll, pitch and yaw rates
$\text{pos}_N, \text{pos}_E$	position in the local-level coordinate system
V_N, V_E, V_D	translational velocities in the local-level coordinate system
x_{CG}	centre of gravity longitudignal position

Greek symbol

α	angle-of-attack
β	angle of sideslip
γ	flight path angle
$\delta_L, \delta_M, \delta_N$	ailerons, elevator and rudder deflection angles
θ	pitch angle
τ	throttle fraction
ϕ	bank angle
χ	course angle
ψ	heading angle

Subscripts

SIM	simulation model
FIT	fitted model
r	reference signal
d	disturbance signal

1.0 Introduction

Small fixed-wing unmanned aerial vehicles (UAVs) have become instrumental in collecting affordable aerial images for military or civilian purposes. However, the threat of icing severely restricts their employment in regions subject to cold climates [1]. Platform-specific effects typically include higher droplet collection efficiency compared to manned aircraft and more severe aerodynamic penalties due to ice accretions at low Reynolds numbers [2]. Although comprehensive icing protection systems (IPS) are being developed [3] to allow flight into known icing, their size, weight and power footprint may not be justified for operations under moderate icing risk. Hence, there is a need for a compromise solution that restores some level of operational capability while reducing the penalties associated with an extensive IPS.

This work introduces the concept of a sense-and-egress module designed to trigger and execute an escape manoeuvre, allowing the UAV to exit the icing conditions before the performance and flying qualities of the platform degrade to a critical level. The first part details the operating principles of the system and proposes a suitable architecture. The benefits and drawbacks of the proposed solution are assessed with respect to alternative solutions. Enabling technologies and knowledge gaps are briefly

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

explored in order to gauge the maturity of the concept. In the second part, the flight dynamics model of the PX-31 Falk UAV is derived from numerical simulations and experimental data. An emphasis is put on tailoring and validating the model structure using system identification from flight tests. In the third section, a synthesis and simulation framework is developed to tune the platform's guidance and control laws. The selected autopilot architecture is detailed, as well as the frequency- and time-domain analysis of the two main controllers. An assessment of the effects of icing on the autopilot trajectory is provided in the case of a simple egress manoeuvre. A permanent elevator limiter and an adaptive bank angle limiter are added to address the shortcomings of the baseline autopilot. Special attention is given to the compromise between reduction of trajectory excursions and minimisation of manoeuvre completion time, as well as the preservation of gust resistance capabilities.

2.0 The sense-and-egress concept

The purpose of the sense-and-egress module is to fill the capability gap between unprotected and fully ice-protected UAVs by increasing the chance of platform survival in case of an inadvertent icing encounter. In that case, the module initiates and conducts an escape manoeuvre to attempt to save the platform. This is a key difference with respect to IPS, which are designed to ensure safe flight into and through known icing conditions, offering greater capability at the cost of extensive hardware integration. IPS have been adopted by most commercial or military aircraft and are currently being developed for small UAVs. However, there is a need for more frugal solutions suitable for environments presenting a limited icing risk or platforms unable to bear the performance penalties due to IPS. Compared to legacy UAVs without ice protection, the platforms equipped with sense-and-egress technology should benefit from greatly improved survivability in icing conditions thanks to timely ice detection, protection of critical components and adequate mitigation strategies. This would enable a shift in operational perspective and risk assessment: The no-go criteria associated with atmospheric conditions could be significantly relaxed, with a consequent increase in availability and profitability.

2.1 Operating principles

The sequence of actions performed by the sense-and-egress may be more or less extensive, as a result of the trade-off between performance and complexity. The most elaborate sequence consists of the following steps:

1. Detection of ice accretion
2. Estimation of current and future flight envelopes
3. Generation of a suitable egress manoeuvre (trajectory and constraints on flight variables)
4. Manoeuvre execution: tracking of trajectory under flight variable constraints

Note that the flight performance of the platform is to be continuously monitored during the execution of the manoeuvre. If, at any instant within the flight envelope prediction horizon, the trajectory tracking problem becomes unfeasible, the remote pilot may be prompted to relax constraints on the egress manoeuvre. Conversely, the manoeuvre may be terminated if monitoring shows that the flight performance of the platform has recovered.

2.2 Technical description

The main function of the sense-and-egress system can be performed by a set of interconnected modules, each of them with a specific purpose. The chart shown in Fig. 1 suggests a decomposition of the system containing four main modules and two optional ones. The function of each main module can be summarised as follows:

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

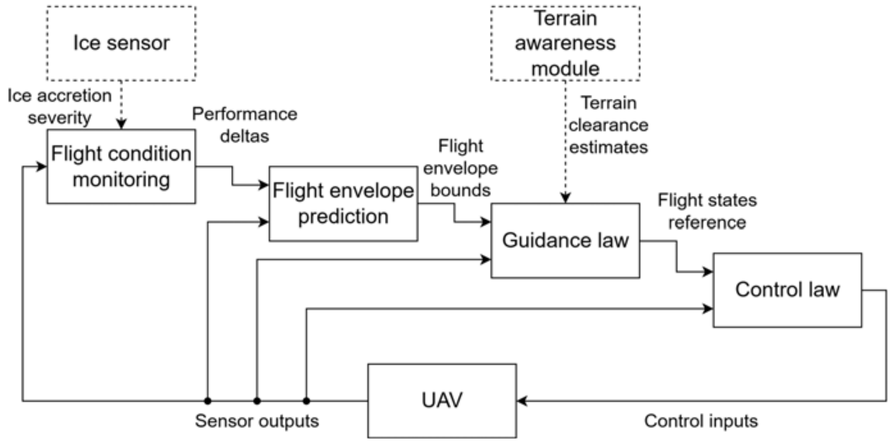

Figure 1. *Architecture of a sense-and-egress system.*

- The flight condition monitoring tracks the evolution of the flight variables to quantify performance degradation and detect ice accretion. The performance indicators to be monitored and fed to the flight envelope prediction algorithm include the power imbalance (icing-induced change of total energy) proposed by Deiler and Fezans [4]. This quantity is used to implement an indirect ice detection method, which can be supplemented by external sensor data to improve sensitivity and reliability [5].
- The flight envelope prediction module estimates the current operating limits based on the performance deltas provided by the monitoring routine. The purpose of this module is to allow the guidance and control laws to operate closer to the platform lift and power limits in order to speed up its exit from hazardous conditions. The evolution of those limits is forecasted over a relevant time horizon to anticipate further performance degradation while the egress manoeuvre is already in progress.
- The guidance law commands the flight variables to perform a safe evasive manoeuvre, accounting for degrading operating limits. Its goal is to minimise exit time under performance and control constraints. Terrain data provided by an additional module could be used to better exploit the potential energy of the platform.
- The control law drives the flight variables to follow guidance commands and reject disturbances in spite of model uncertainties. The performance of the control loop conditions the quality of the trajectory tracking.

2.3 Complementarity with other icing protection concepts

This section addresses the key differences and similarities between sense-and-egress and IPS. From an implementation standpoint, the main difference between the two systems lies in the extent of the protected components. The sense-and-egress only requires protection of the critical components, which are the most vulnerable to icing (i.e. the air data probes and propellers). The IPS, however, needs to provide deicing or anti-icing action on the lifting surfaces such as the wings and horizontal tail. This can be accomplished through a variety of thermal, mechanical or chemical means. The actions taken also differ according to the purpose of the system and its installed hardware: In the IPS case, the availability of extensive deicing or anti-icing hardware allows for the planned flight path and focus on mission completion. In the sense-and-egress case, the lack of comprehensive icing countermeasures prompts the system to abort the flight plan and perform an evasive manoeuvre in order to focus on platform survival. Hence, compared to IPS, sense-and-egress may offer the following advantages:

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

- Reduced performance penalties outside of icing conditions: The weight penalty associated with wing and tail de-icing systems is avoided, which may be significant when the aspect ratio of those lifting surfaces is high.
- Reduced unit cost: The above-mentioned systems may represent a significant fraction of the platform cost or require a more complex manufacturing process.
- Improved reliability: A lower number of non-redundant systems implies a lower probability of failure. De-icing systems, which are integral to the airframe, may be less serviceable than other components.
- Availability and scalability: As the few ice-protected components required by the sense-and-egress are not part of the airframe, they may be retrofitted to legacy platforms or even shared between a fleet of UAVs of the same type.

For all its merits, it may, however, present the following drawbacks:

- Broader development effort: As sense-and-egress heavily relies on the prediction of the platform flight dynamics and includes a flight control module, the development costs associated with modelling and control may be higher than for IPS. The interface requirements with the legacy platform autopilots may also generate additional complexity.
- More restrictive operational requirements: As this system does not allow flight into known icing, the range of atmospheric conditions in which UAV operations remain possible is more limited than for IPS.
- Loss of mission capability on an icing encounter: Since the platform is required to manoeuvre in order to survive an icing encounter, adherence to the initial flight plan can no longer be guaranteed, which may be unacceptable for time-critical scenarios. The energy expenditure due to evasive manoeuvres can be prohibitive and may require aborting the mission.

2.4 Enabling technologies and knowledge gaps

The aim of this section is to propose a brief overview of the relevant previous works in each of the topics that are crucial to the development of a sense-and-egress module.

2.4.1 Flight dynamics modelling

The approach explored by Deiler [6] is to describe the dynamics of the iced aircraft as the sum of the dynamics of the clean aircraft and the contribution due to ice accretion. In the frame of this work, the base model of the clean UAV was obtained using computational fluid dynamics (CFD) [7], panel methods [8] and wind tunnel tests of the propulsion subsystem [9]. The perturbation component has been computed from CFD simulations on insulated wing and tailplane aerofoil sections [7]. The method detailed in Ref. [8] uses a multiplicative scheme instead of the additive contribution implemented by Ref. [6]. It must be noted that the applicability of this iced model is limited to a single case of atmospheric conditions with a given set of static temperature, liquid water content and droplet mean volume diameter values. The ice accumulation rate is assumed to be constant, with linear performance degradation up to the state corresponding to the maximum icing time.

As the topology of ice accretions and the severity of the resulting aerodynamic penalties strongly depend on the atmospheric conditions [10], which are not commonly monitored on board of small UAVs, the icing model should represent a worst-case scenario with regard to expected operating conditions. Local measurements of ice thickness provided by onboard ice sensors may be used to drive a multi-model strategy. which could allow to reduce the conservatism of this approach. The source data of the icing model could also be updated to include the recent icing wind tunnel campaign carried out with the wing aerofoil section [11]. Ultimately, the measurement of the forces and moments acting on the whole aircraft would yield an even more representative model of its flight dynamics in icing, but the associated costs and time expenditure would be significant.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

2.4.2 Performance-based ice detection

This ice detection technology is based on the estimation of the rate of change of the total energy of the aircraft [4]. The in-flight estimated value is compared to the expected value known from the nominal aircraft model. Ice is detected when the difference exceeds a specific threshold, which is chosen as a compromise between detection sensitivity and false alarm rate. This methodology has recently been validated on business jets and large transport aircraft [12]. However, its applicability to small UAVs remains to be proven as there may be less data available to build an accurate nominal aircraft model.

2.4.3 Flight envelope prediction

Two definitions of the flight envelope are proposed in the literature: the most extensive one is described as the manifold of the state space from the 6 degrees-of-freedom equations of motion [13] where the aircraft can remain controllable [14]. The more computationally tractable one, also called the trim flight envelope, only includes the equilibrium points of the system. Those points are usually constrained by thrust and angle-of-attack (α) limitations, which may both be affected by icing. Modelling of α-induced aerofoil flow separation (stall) and its effect on aircraft flight dynamics has been carried out via a variety of methods, including system identification from flight test data [15]. However, variants of this method compatible with real-time processing (i.e. online system identification) have failed to accurately capture stall phenomena observed on small delta-wing UAVs [16]. This begs the question of whether or not online system identification methods could estimate the effects of icing on the flight envelope.

Assuming that the icing phenomenon could be captured, the data collected would only represent its influence on the current operating condition, which itself may be far from the flight envelope bounds. Thus, one could either adjust the operating condition, ultimately putting the UAV at risk, or infer the performance delta at the bound from its value at a safer operating point, creating potentially large uncertainties in the process. Prediction of the evolution of the operating limits may also not be straightforward, as the severity of icing buildup could depend on some of the flight variables, such as α, altitude or airspeed, which vary during the manoeuvre. Rough estimates could be obtained by extrapolating the data trend, but physics-based evolution laws should be sought to consolidate the predictions.

All that being said, flight envelope predictions may be inaccurate if available at all. Thus, the guidance and control laws should be designed to be able to perform the egress manoeuvre in the absence of accurate flight envelope bounds. It may also be wise to postpone the development of such methods until a baseline autopilot is implemented.

2.4.4 Guidance and control laws

Both guidance and control aspects of the autopilot are tightly coupled, as the bandwidth of the control law conditions the performance of the trajectory tracking loop. This is because it may be desirable to enforce time-scale separation between both loops. Under this assumption, the control and guidance gains can be tuned sequentially, starting with the inner (control) loop [13].

Robust synthesis techniques can be employed to guarantee closed-loop stability in the presence of model uncertainties [17], which may be correlated with the performance deltas determined by the flight condition monitoring module. Structured $\mathcal{H}_\infty$ synthesis appears to be a suitable choice as it is able to balance performance, robustness and actuator usage objectives [18]. Moreover, it can be used on any prescribed controller structure, which simplifies its implementation both in simulation and on flight hardware.

Model predictive control is also of interest because of its ability to account for time-dependent constraints on the flight states [19]. However, solving the online optimisation problem is often computationally intensive and may not always be feasible. Moreover, obtaining stability guarantees could be a challenge for such classes of systems [20].

Figure 2. *The Falk UAV on its launch rail; DLR (CC BY-NC-ND 3.0).*

3.0 Flight dynamics model of the Falk UAV

3.1 Platform description

The PX-31 Falk, pictured in Fig. 2, is a small fixed-wing UAV manufactured by the Norwegian company Maritime Robotics. Its maximum take-off weight is 22 kg for a payload of 7 kg and a wing span of 3.2 m. It is powered by an electric motor driving a pusher propeller up to a ceiling of 20.000 ft and a maximum airspeed of 40 m/s. This platform is capable of Beyond Visual Line Of Sight (BVLOS) operation thanks to a long-range very high frequency (VHF) radio link. As current operational limitations [21] preclude BVLOS operations when static ground air temperature drops below a given threshold, the availability of this platform could be greatly improved by providing protection against inadvertent icing. Hence, the Falk is a prime candidate for the application of sense-and-egress algorithms.

3.2 Baseline mixed-fidelity model

Both linear and non-linear formulations of the UAV flight dynamics model are required to, respectively, synthesise the control laws and assess the sense-and-egress autopilot in icing conditions. The baseline representation of the aerodynamic forces and moments is derived from the numerical works of Skogan [8] including dynamic coefficients from the panel code XFLR5, as well as CFD-computed lift, drag and pitching moment polars originally published by Lindner et al. [7]. This CFD dataset covers both the main wing and horizontal tail plane (HTP) aerofoil s for various icing conditions. For this work, the worst-case scenario has been chosen according to the most severe continuous maximum icing case (as defined by EASA CS-25 Appendix C) identified by Lindner [7]. This scenario entails a static air temperature of -2° and a droplet median volumetric diameter (MVD) of 15 μm with an accretion time of 21.5 min.

The model structure proposed by Skogan and adopted in this work resolves the longitudinal forces and moments at both the main wing and HTP centres of pressure while the lateral directional actions are applied at the vehicle centre of mass. This decomposition models the effect of the downwash induced by the wing on the HTP as well as the pitch damping effect provided by the tail. Both lifting surfaces have individual non-linear lift, drag and pitching moment polars, allowing wing stall, tailplane stall or any combination of those effects to be captured. Numerous general aviation accidents [22, 23] show that this level of detail is crucial for flight dynamics investigations in icing conditions. Note that the wing incidence angle of the Falk is 4°, giving a stall angle of respectively 7° and 4° for clean and iced configurations.

The propulsion system model is given by the experimental findings of Løw-Hansen [9]. Note that the propeller is assumed to be ice-protected, hence remaining free from any performance degradation due to icing. Relevant terms of the inertia tensor are estimated from pendulum experiments performed by Løw-Hansen at the Norwegian University of Science and Technology. It is also assumed that the effect of ice accretion on the mass and inertia properties of the UAV is negligible. Finally, the formulation of the

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

6 degrees-of-freedom equations of motion is provided by Beard and McLain [13] with additions from Zipfel [24] regarding the quaternion formulation of the attitude angles and angular rates. The linear, time-invariant (LTI) forms of this model are obtained through a Jacobian linearisation performed under the assumption of steady-level flight.

3.3 Tailoring to flight test data

Even though the model structure has been selected to simulate the Falk dynamics in icing conditions with adequate fidelity, there may be significant deviations with respect to the behaviour of the actual platform. Hence, data from a functional check flight has been leveraged in an attempt to bridge this gap. The opportunistic nature of the test data leads to significant limitations in the analysis. The adverse factors detrimental to the accuracy of the system identification procedure include poor quality of the sensor package, lack of adequate excitation of all the dynamics of interest, as well as the unknown influence of external wind disturbances.

3.3.1 Data compatibility analysis

In order to consolidate the sensor data provided by the platform, measured linear accelerations and rotation rates from the inertial measurement unit (IMU) are fed into a motion reconstruction algorithm, which outputs the expected geographic position and velocities, as well as the Euler and airflow angles. The reconstructed measurements are compared to the values provided by the onboard extended Kalman filter (EKF), which processes the Global Navigation Satellite System (GNSS) data. As the EKF outputs v_A and β showed a bias incompatible with the directional stability characteristics of the platform, these signals had to be filtered before the estimation and are displayed in their processed states. In order to improve consistency, the accelerometer bias and constant geographic wind were estimated using the maximum likelihood criterion as implemented in DLR proprietary software Fitlab.

Figures 3 and 4 depict the fit of the reconstructed measurements to the EKF outputs (in blue), both before (in red) and after (in green) the identification procedure. The plots show that the integration of the IMU rate gyros appears to be reliable, as the reconstructed Euler angles come very close to the EKF values. The reconstructed position, however, diverges from its ground truth. The estimation of accelerometer bias allows for counteracting this phenomenon and promotes data consistency. Additionally, the estimation of constant geographic wind significantly improves the fit of the airflow angles. Hence, this data harmonisation procedure is deemed successful and is systematically and independently performed on both the identification and the validation time series. Tables 1 and 2 provide the numerical values of the estimated accelerometer bias and geodetic wind components. Note that the average amplitude of those factors is higher for the validation sequence (Val) than for the identification sequence (ID). This can be interpreted as the validation sequence measurements being more severely affected by wind disturbances and other unmodelled phenomena.

3.3.2 Identification of the revised model

The first step of the model adjustment is to add or select the parameters that are to be estimated in order to improve the fit to the actual flight data. To model the additional lift and drag provided by the fuselage, which has been neglected in the aerodynamic model build-up, offset terms are added to the lift and drag polars of the wing. Since the longitudinal centre of gravity (CG) position at which the UAV was flown was not documented, an offset term has also been added to its initial guess based on manufacturer data. Finally, the control effectiveness derivatives of both elevator $Cm_{\delta M}$ and ailerons $Cl_{\delta L}$ have been chosen as tuning parameters. Note that the rudder effectiveness derivative has been excluded from the estimation procedure due to the lack of adequate rudder excitation in the available flight data.

Once the model structure has been frozen, the identification dataset can be used to estimate the free parameters. The identification sequence is 30 s long and can be viewed in Figs 5, 6 and 7. The algorithm and software used are similar to what was used for conducting the data harmonisation procedure.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

Table 1. Estimated accelerometer bias

	a_X bias (m/s^2)	a_Y bias (m/s^2)	a_Z bias (m/s^2)
ID	$4.317e^{-2}$	$1.540e^{-1}$	$3.894e^{-2}$
Val	$-1.494e^{-1}$	$-3.066e^{-1}$	$2.134e^{-1}$

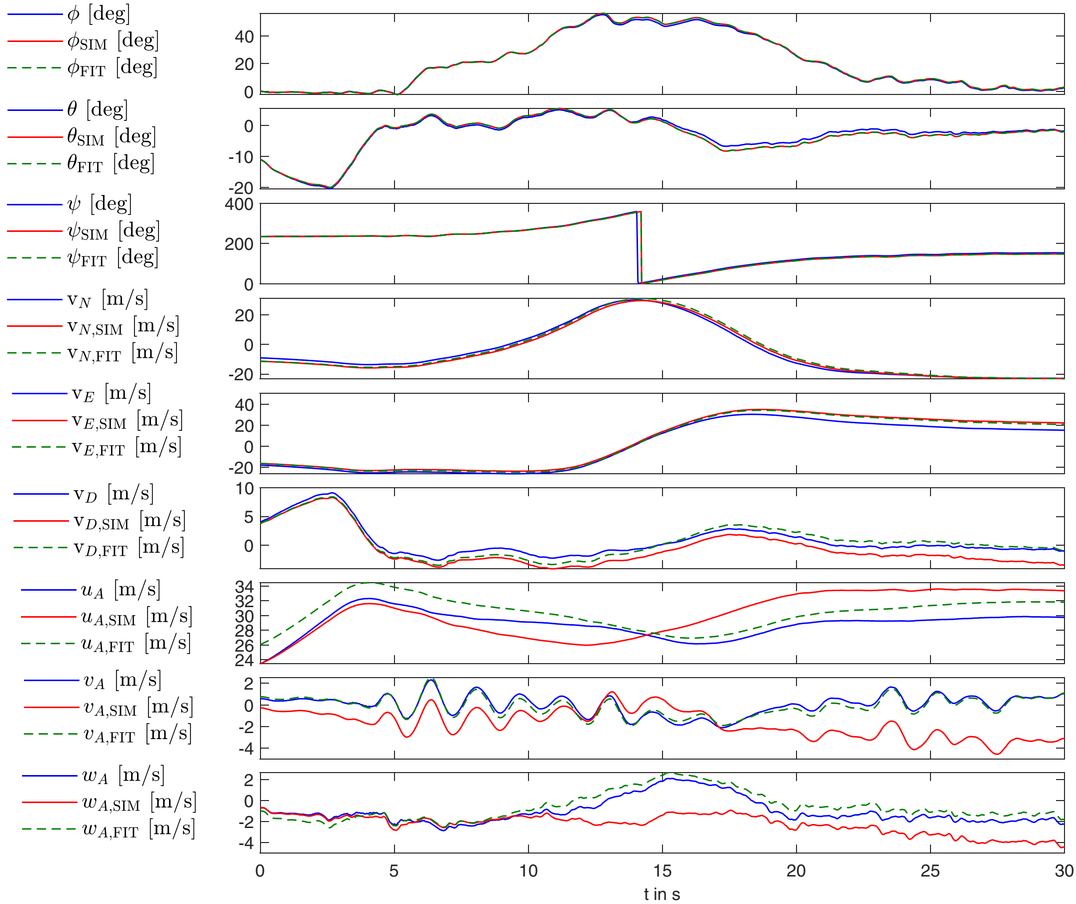

***Figure 3.** Fit of reconstructed flight data (1/2).*

Table 3 presents the estimated values of the chosen model parameters. The negative values for the lift and drag offsets reported in Table 3 do not match their initial interpretation as fuselage interference effects. Instead, one could see those values as biases that are tuned to compensate for the mismatch between the behaviour of modelled and actual components. This is especially true for the propulsion system, as the propeller and engine flown are of a different type than what has been tested in the wind tunnel to obtain the propulsion model. Hence, a negative drag bias could indicate an attempt to compensate for an underperforming propulsion model. Conversely, a negative lift bias corresponds to an overperforming lift model, i.e. due to the lack of modelling of wing sweep effects. The estimated CG offset is small and well within the manufacturer's limits. The aileron and elevator effectiveness have been, respectively, increased by 49% and 76% compared to their initial values provided by panel methods. This level of inaccuracy is not unusual for such preliminary analysis tools, as discrepancies of more than 40% on aileron effectiveness have been reported between XFLR5 and semiempirical methods [25].

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

Table 2. *Estimated constant wind*

	$V_{w,N}$ (m/s)	$V_{w,E}$ (m/s)	$V_{w,D}$ (m/s)
ID	$6.655e^{-1}$	2.741	$-1.667e^{-1}$
Val	$-3.357e^{-1}$	4.255	$-3.114e^{-1}$

Table 3. *Estimated model parameters*

$C_{L,0}$ bias	$C_{D,0}$ bias	x_{CG} bias	$Cl_{\delta L}$	$Cm_{\delta M}$
$-2.107e^{-1}$	$-1.169e^{-2}$	$2.839e^{-3}$	$3.526e^{-1}$	1.493

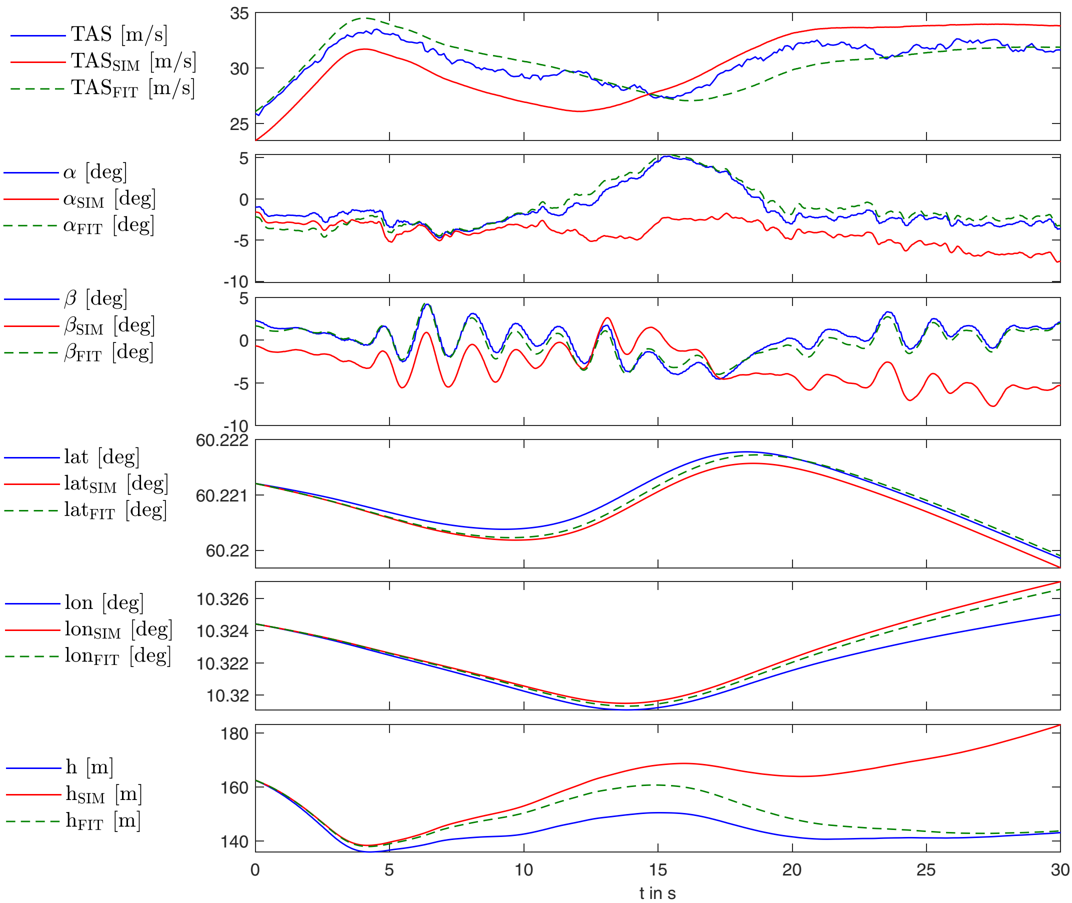

Figure 4. *Fit of reconstructed flight data (2/2).*

3.3.3 Model fit on identification and validation data

The fit of the revised dynamic model on the identification dataset is shown in Figs 5, 6 and 7. Linear accelerations and body rotation rates are well matched, and so are the Euler angles. Body frame inertial velocities and airflow angle predictions are less accurate, with significant deviations or time lags. In particular, the revised model is not able to replicate the variations of the angle-of-attack. The model errors on inertial velocities explain the divergence of the measured and predicted geographic positions towards the end of the identification time series.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

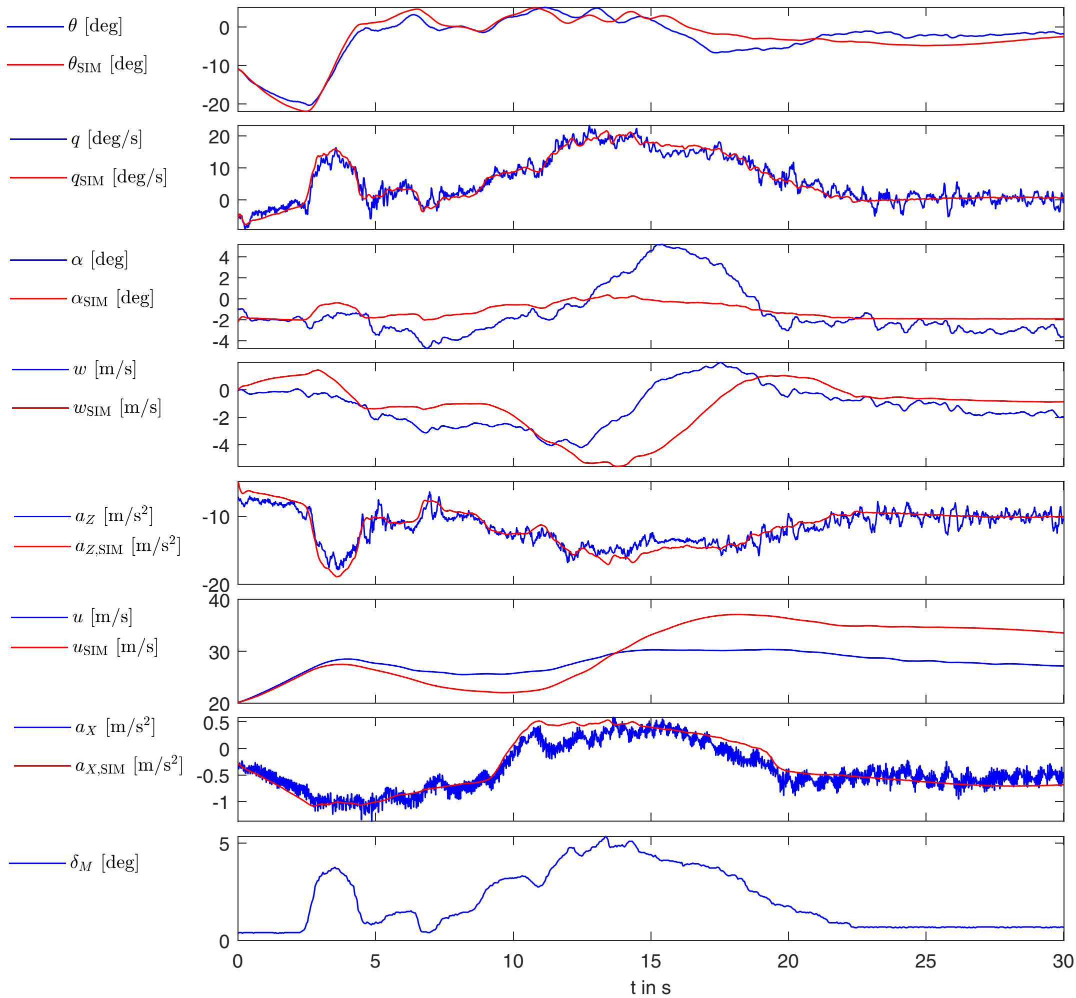

***Figure 5.** Model fit on identification data: longitudinal dynamics.*

The validation time series is 50 s long and is displayed in Figs 8, 9 and 10. Here, the divergence between the measured and predicted positions and attitudes is more apparent, especially on the longitudinal axis. Large pitch oscillations develop in the vertical plane, which indicate that the phugoid mode of the simulated airframe has been inadvertently excited. The pitch divergence starts around 20 to 25 s, when the model is getting into a steep left bank, whereas the actual aircraft returns to wings level. This bank angle error creates an apparent lift discrepancy that disturbs the longitudinal equilibrium of the aircraft, exciting its phugoid mode.

3.3.4 Discussion on model adequacy

As uncovered in the previous section, the revised model shows some serious limitations when predicting the long-term behaviour of the aircraft, but appears to adequately capture the initial response of the UAV. Since the fast dynamic modes are the most relevant for flight control design [26], this model was deemed suitable for this purpose and for simulation of the closed-loop dynamics. Moreover, the deficiencies of the model are most prevalent at lower frequencies, which are covered by the bandwidth of the control system. Hence, it is assumed that the controller can be designed to compensate for the disturbances due to model mismatch, for instance, by constraining its output sensitivity function (see

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

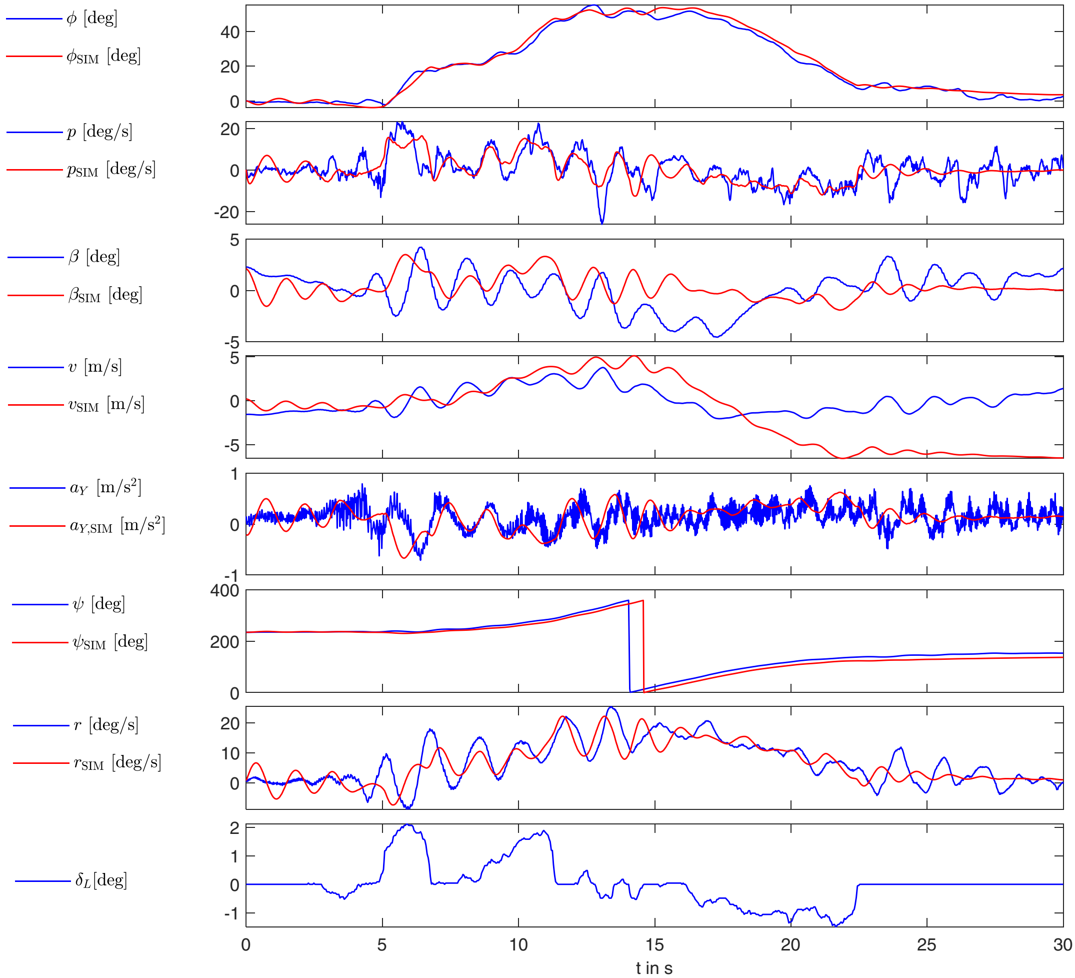

Figure 6. *Model fit on identification data: lateral-directional dynamics.*

section 4.1.1). The fitness for purpose of the model-controller combination should be demonstrated through hardware-in-the-loop or flight testing in order to verify this hypothesis. This perspective is left for future work.

4.0 Baseline autopilot design

The purpose of the autopilot is to safely conduct the escape manoeuvre in the presence of ice contamination on the lifting surfaces. The autopilot implementation pictured in Fig. 11 includes both control and guidance laws. In practice, the horizontal guidance algorithm provides a course target that the flight control system (FCS) attempts to follow. The FCS itself is composed of three controllers: the course controller tracks the course angle χ by commanding a bank angle ϕ, while the total energy control system (TECS) is tasked with following both airspeed and altitude targets using the throttle and a commanded pitch attitude θ. Finally, the attitude controller takes in the sideslip angle β prescribed by the guidance law as well as the bank and pitch references obtained, respectively, from the course controller and TECS. It then manoeuvres the airframe by acting on its ailerons, elevator and rudders. In the following section, the design of all three controllers is conducted using dynamic models linearised at the operating point (TAS = 30 m/s, h = 100 m).

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

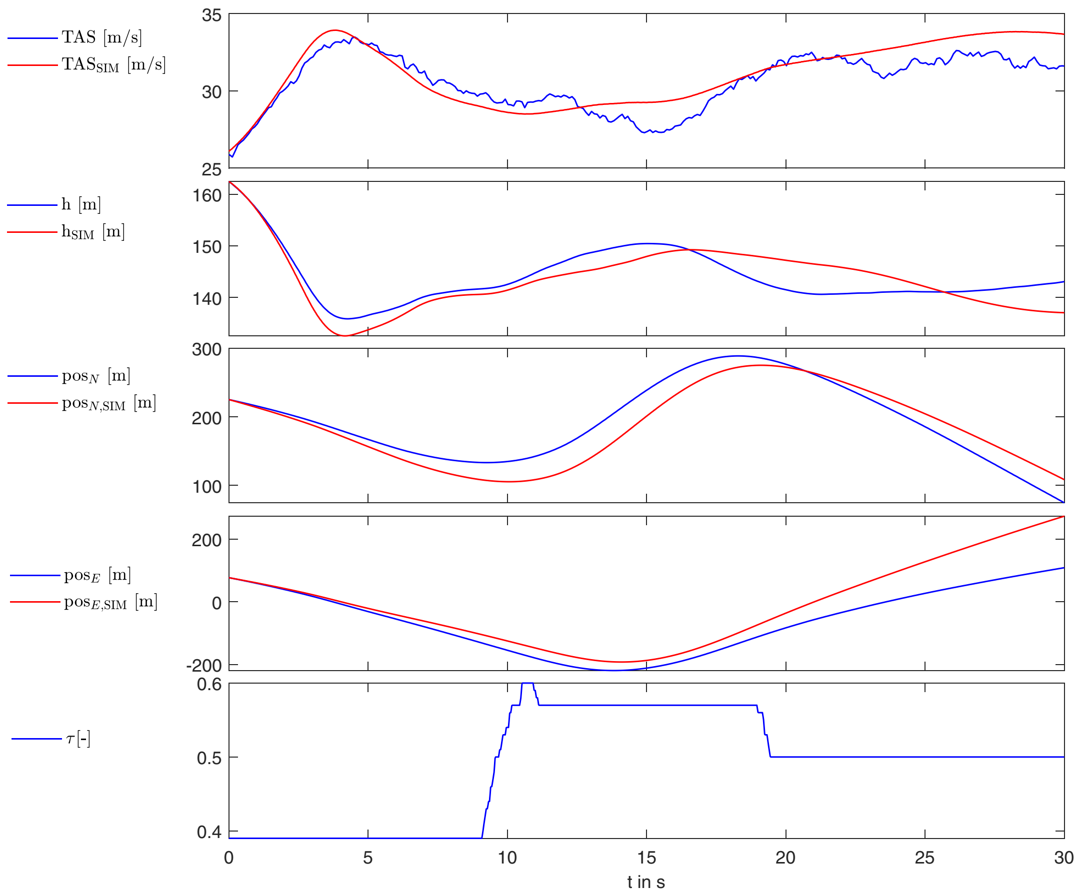

Figure 7. *Model fit on identification data: airspeed and position.*

4.1 Attitude controller

The FCS design process is carried out from the innermost controller (i.e. acting directly on the plant) outwards. The layout of the attitude controller, pictured in Fig. 12, features a dual-loop architecture with additional feedforward gains, which has been commonly used in aerospace applications [27, 28]. The inner angular rate regulation loop allows for the dampening of the natural dynamics of the plant G while avoiding the need to differentiate potentially noisy attitude angle signals. The regulator block K_{reg} provides a proportional action on the angular rates p, q and r. The outer 'servo' loop tracks the reference signals through a proportional-integral action K_{serv} on the bank, pitch and sideslip angles. The feedforward allows improved reference tracking performance and is injected through the K_{FF} block, which implements a proportional action on β, ϕ and θ. The controller implements an anti-windup scheme (not pictured) to protect against the effects of control surface deflection limits on the integrator. The choice of sideslip angle instead of heading as one of the channels stems from the fact that the UAV course is steered using a bank-to-turn scheme; hence, there is no need to control the heading with the rudder. Instead, this surface is dedicated to sideslip regulation, ensuring coordinated turns and reducing the likelihood of a spin at angles of attack close to the stall limit. Even though small UAVs are rarely fitted with sideslip sensors, estimates stemming from the EKF [29] are available on most common flight software such as Ardupilot.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

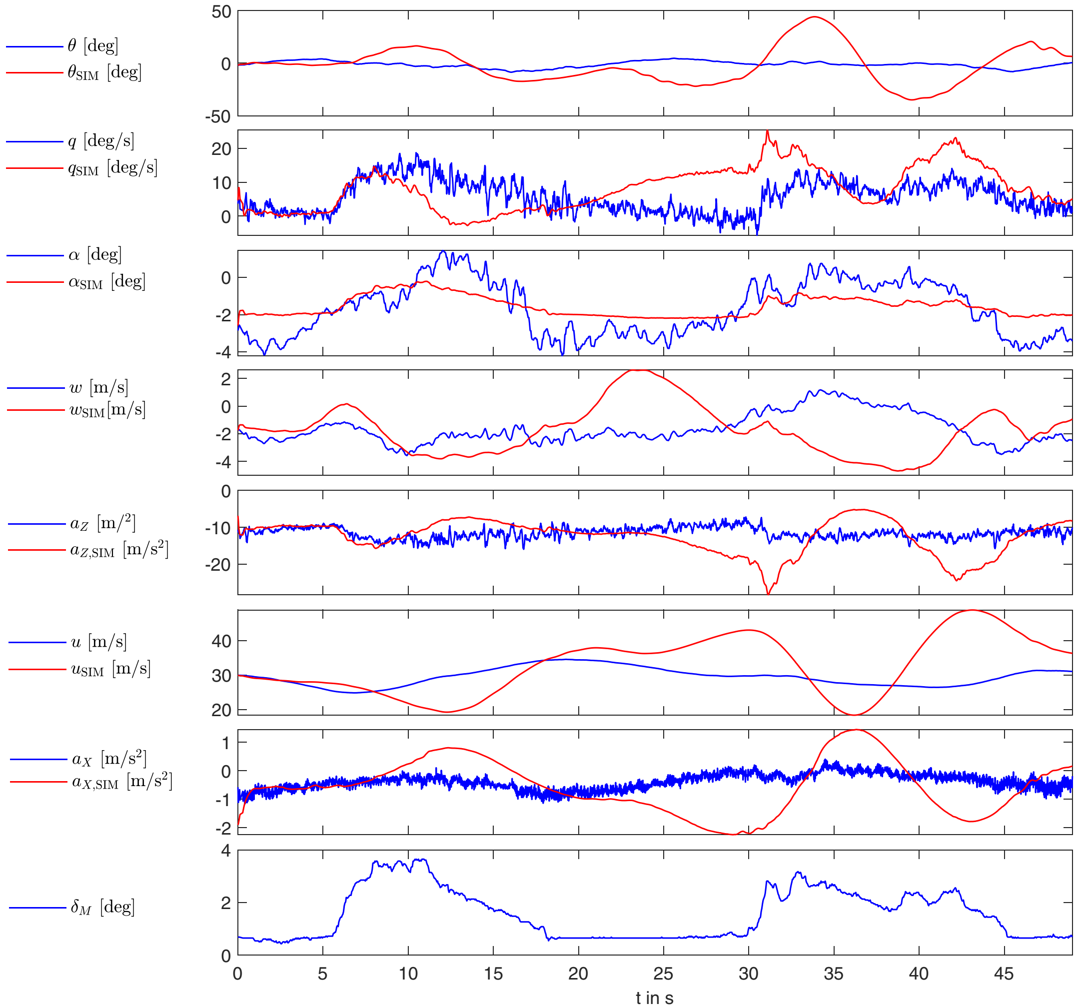

Figure 8. *Model fit on validation data: longitudinal dynamics.*

4.1.1 Synthesis framework

The tuning of the controller gains has been carried out via a mixed-sensitivity $\mathcal{H}_\infty$ formulation [30]. The $\mathcal{H}_\infty$ synthesis problem consists in finding a stabilising controller that minimises the influence of exogenous inputs (reference signal and disturbances) on the signals of interest (plant outputs or actuator inputs). A multi-objective synthesis framework [31] has been implemented using the weighting filters pictured in red in Fig. 12. It allows for explicitly defining the control performance requirements and the actuator constraints on a channel-by-channel basis.

For each output channel $y \in \{\beta, \phi, \theta\}$, the performance requirement is expressed as a disturbance rejection constraint using the output sensitivity function $\mathrm{S} = T_{y_d \to y}$ multiplied by the weighting filter W_S to form the performance channel $T_{y_d \to z_\mathrm{DR}}$. Equations (1) and (2) introduce two types of frequency templates, which define the inverse of the weighting filter W_S. Both $\mathrm{W_{S,1}}^{-1}$ and $\mathrm{W_{S,2}}^{-1}$ are shaped like high-pass filters in order to provide good reference tracking and disturbance rejection performance. $\mathrm{W_{S,2}}^{-1}$ mainly differs from $\mathrm{W_{S,1}}^{-1}$ in that it features a $+40\,\mathrm{dB}$ per decade slope instead of $+20\,\mathrm{dB}$ per decade but it shares the same low-frequency gain $k_{S,LF}$, high-frequency gain $k_{S,HF}$ and $-3\,\mathrm{dB}$ bandwidth ω_S. The template from Equation (1) is used for the sideslip channel as its sensitivity shows a first-order behaviour in the frequency region of interest. Conversely, the pitch and bank channels exhibit second-order behaviour, which warrants the use of the template given in Equation (2).

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

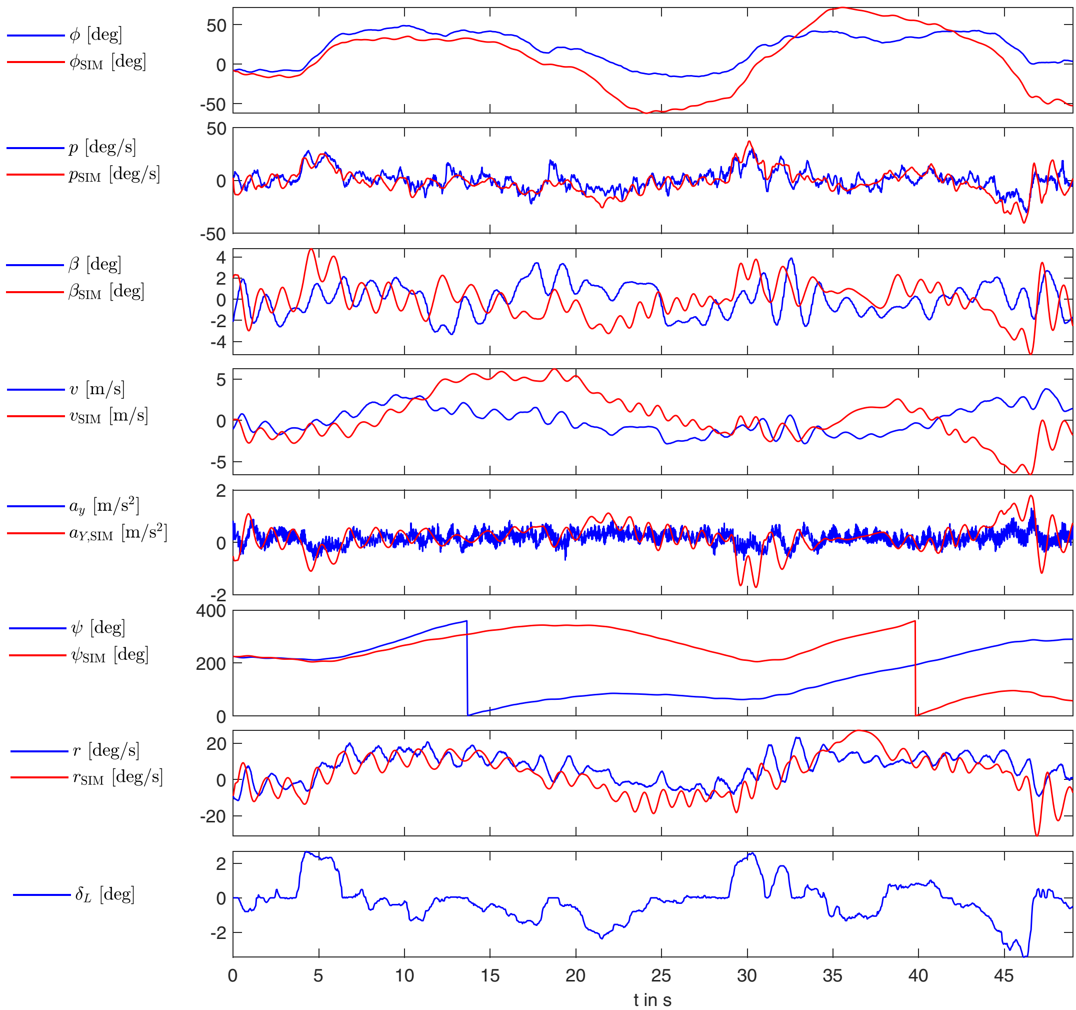

Figure 9. *Model fit on validation data: lateral-directional dynamics.*

$$\mathrm{W_{S,1}}^{-1} = \frac{k_{S,HF}\, s + k_{S,LF}\, \omega_S}{s + \omega_S} \tag{1}$$

$$\mathrm{W_{S,2}}^{-1} = k_{S,HF} \frac{\left(s + \sqrt{\frac{k_{S,LF}}{k_{S,HF}}}\, \omega_S \sqrt{\sqrt{2} - 1}\right)^2}{\left(s + \omega_S \sqrt{\sqrt{2} - 1}\right)^2} \tag{2}$$

Table 4 summarises the numerical values chosen for the template parameters and stability margins. For all channels, the low-frequency bound on S is chosen to be very small in order to suppress steady-state error, while the high-frequency asymptote is set as low as possible to maximise closed-loop stability under the limitation imposed by the 'waterbed' effect [32, chap. 5]. The bandwidth with respect to the template asymptote is set to 8 rad/s to promote fast reference tracking and adequate disturbance rejection properties. Because of the different template shapes, this last requirement translates to -3 dB bandwidths of 4.3 rad/s on β and 4.8 rad/s on ϕ and θ.

The actuator limitations are enforced by the closed-loop transfer function from the reference to the actuator input $\mathrm{KS} = T_{y_r \to \delta_{act}}$ multiplied by the filter W_{KS} to form the control effort channel $T_{y_r \to z_{act}}$. The

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

Table 4. *Parameters of the synthesis constraints*

Constraint Type	Parameter	Value	Unit
$W_S\left(T_{y_r\to y_e}\right)$	$k_{S,LF}$	0.01	- -
	$k_{S,HF}$	1.5	- -
	ω_S	5	rad/s
$W_{KS}\left(T_{y_r\to \delta_{act}}\right)$	$k_{KS,LF}$	2	- -
	$k_{KS,HF}$	0.001	- -
	ω_{KS}	15	rad/s
$W_{coord}\left(T_{\phi_r\to \delta_n}\right)$	$k_{coord,LF}$	0.316	- -
	$k_{coord,HF}$	0.001	- -
	ω_{coord}	15	rad/s
Out. margins	M_G	6	dB
	M_ϕ	45	deg

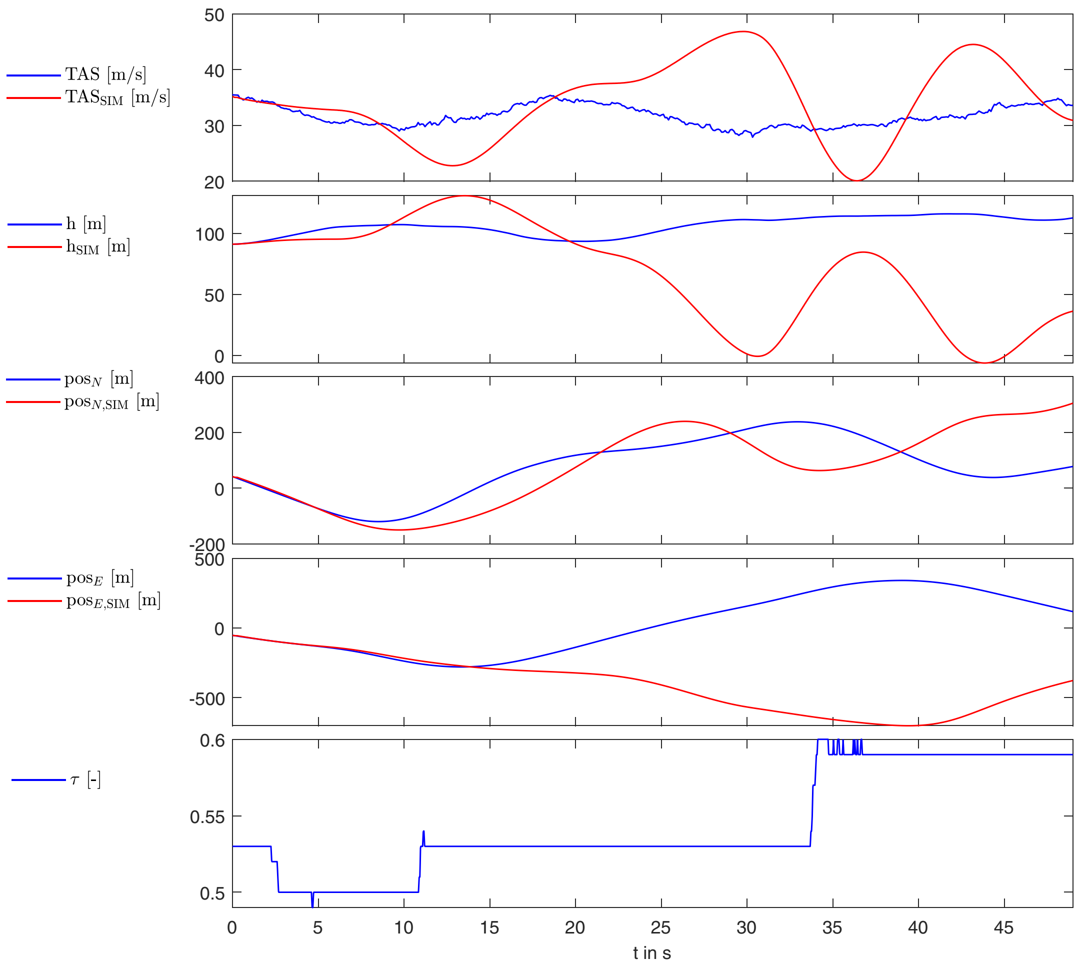

Figure 10. *Model fit on validation data: airspeed and position.*

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

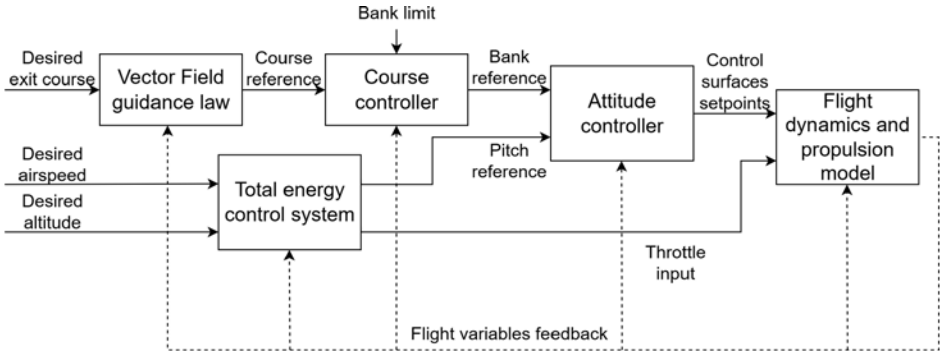

Figure 11. *Path following autopilot structure.*

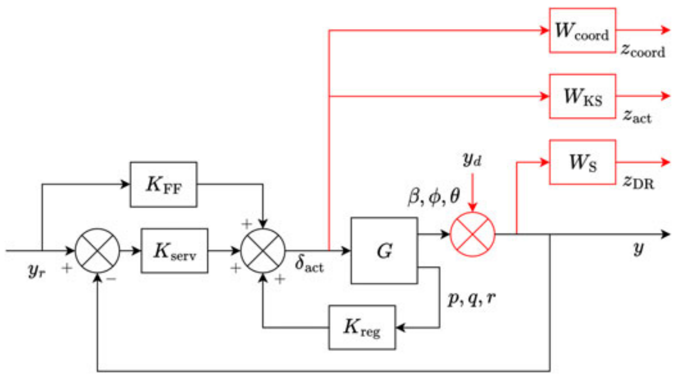

Figure 12. *Layout of the attitude controller. The components represented in red are described in Section 4.1.1.*

template $W_{KS}{}^{-1}$ follows the layout of Equation (1) but features a low-pass behaviour to attenuate the effect of measurement noises, preserve the actuators from high-frequency content and limit the control energy. The maximum bandwidth requirement is set to 15 rad/s, which corresponds to one third of the bandwidth achievable by comparable-size actuators [28]. The high-frequency constraint on KS is set to a very small nonzero value to cut-off the high-frequency content of the control signal. Finally, the bound on the low-frequency gain is set to avoid amplitude saturation in the usual range of reference inputs, i.e. $|\beta_r| \leq 5^\circ$, $|\phi_r| \leq 30^\circ$ and $|\theta_r| \leq 10^\circ$. Although the attitude of the UAV may commonly exceed these ranges, the upper-level controllers (course controller for ϕ and TECS for θ) rarely inject large amplitude steps into the attitude controller. Hence, moderate amplitudes have been chosen for synthesis purposes. In addition, trial-and-error has shown the necessity to constrain the bank angle reference to the rudder transfer $T_{\phi_r \rightarrow \delta_n}$ because of the large control action required to counter adverse yaw and maintain a coordinated turn. Thus, the weighting filter W_{coord} has been introduced to form the turn coordinator channel $T_{\phi_r \rightarrow z_{\text{coord}}}$. Its template W_{coord}^{-1} shares the same HF gain and bandwidth with W_{KS}^{-1}. However, its low frequency gain has been set to -10 dB in order to avoid rudder saturation for bank angle changes larger than 60°.

The margins at each channel output $y \in \{\beta, \phi, \theta\}$ were constrained to improve the robustness and stability properties of the control system. The values of 6 dB gain and 45° phase were chosen according to the recommendations of SAE Aerospace Standard 94900.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

***Table 5.** Gains of the tuned attitude controller*

Channel	FF Gain	P Gain	I Gain
β	8.307	-8.297	-11.70
ϕ	-1.667	1.663	2.077
θ	-2.003	2.001	12.81
p	- -	$-9.234e^{-2}$	- -
q	- -	$-1.686e^{-1}$	- -
r	- -	-1.501	- -

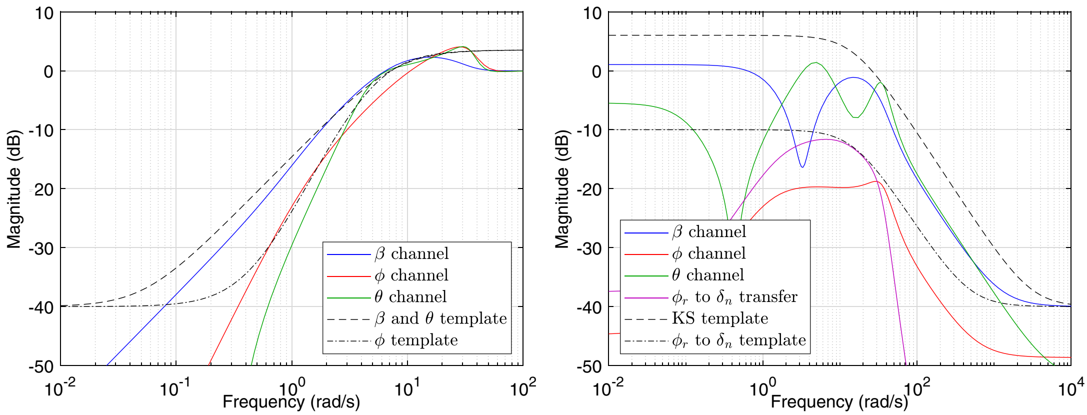

***Figure 13.** Output sensitivities (left) and reference-to-actuator transfers (right) of the attitude controller channels.*

4.1.2 Optimisation results

In the $\mathcal{H}_\infty$ optimisation problem, the norm of the performance transfer functions represents the cost functions and shall be minimised under the condition that the actuator constraints and stability requirements are met. Thus, the disturbance rejection constraint has been declared as ‘soft’ in the synthesis framework. Conversely, the actuator usage template and the stability margins prevent unsatisfactory behaviour of the closed-loop system and often conflict with the performance criterion (for instance, the stability margins limit the peak of S). That is why these requirements have been marked as ‘hard’ constraints to avoid reducing system performance by minimising them beyond their maximum admissible levels. The resulting controller gains are given in Table 5. Note that the feedforward action is the opposite sign of the proportional action. This is due to the roll-off requirement on the KS transfer, which implies that the total contribution proportional to y_r (see Fig. 12) must be close to zero as y_r may contain high-frequency content.

4.1.3 Frequency-domain analysis

The Bode plots of the output sensitivity S and the reference-to-actuator transfer KS of Fig, 13 give insight into the respective disturbance rejection and actuator usage of the attitude controller. The first aspect is summed up by the -3 dB bandwidth of the output sensitivities: The sideslip, bank and pitch channels, respectively, come up at 4.02 rad/s, 6.84 rad/s and 4.50 rad/s, very close to the synthesis requirements for β and θ, and actually much faster for ϕ. The second aspect is represented by the shape of the reference-to-actuator transfers: All of the responses fall within the prescribed template. It should be noted that the constraint on rudder deflection for turn coordination is active, demonstrating the relevance of this specific consideration. As for the stability considerations, the channel-by-channel gain margins reach 11.8 dB, 8.41 dB and 8.43 dB for β, ϕ and θ while the corresponding phase margins come up to 61.2°, 48.4° and 48.5° implying that the synthesis requirements are met.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

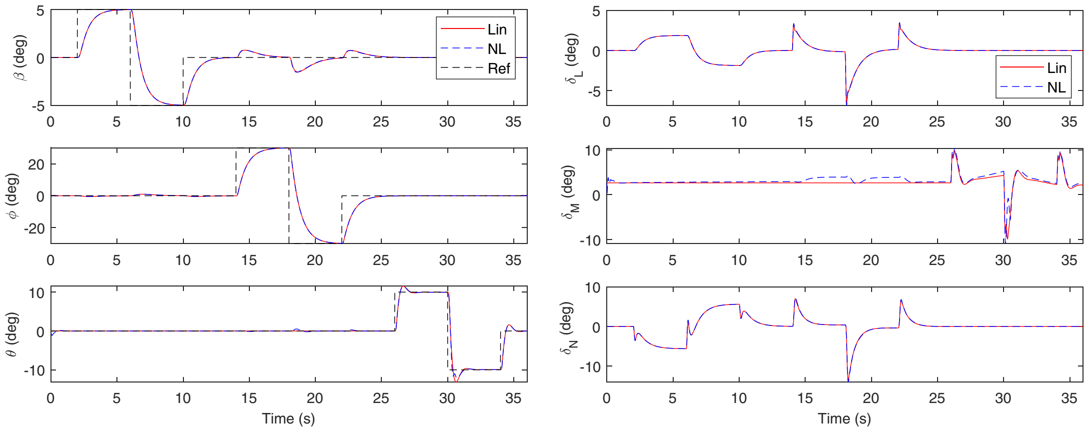

Figure 14. *Reference tracking (left) and control usage (right) of the attitude controller.*

4.1.4 Time-domain analysis

In order to complement the frequency-domain assessment of the control loop and provide more insight about potential cross-channel couplings, the closed-loop response to a multistep reference signal has been simulated and analysed. The reference amplitudes are chosen to be consistent with the assumptions made in Section 4.1.1. Figure 14 depicts the attitude response and the control surface deflections of the autopilot during the reference tracking task. On those plots, the response of the full non-linear model is compared to the output of the linear form used as a synthesis model. The left plot attests that the controller tracking performance is satisfactory, with no overshoots and a 95% response time of 2.11 s for β, 2.16 s for ϕ and 0.98 s for θ. It should be noted that the linear model does not capture the longitudinal-lateral couplings shown by the non-linear model. However, in most cases, given the very small magnitude of the control action required to counter those effects, one can safely assume that they can be neglected at the synthesis stage. The right plot shows that the actuator usage fits within the control deflection limits ($+17^\circ$ / -12°) for the elevator, $\pm 15^\circ$ for ailerons and rudders). There is, however, little margin left for both the elevator and rudder travels. It is also interesting to note that countering adverse yaw during a rolling manoeuvre appears to be more stringent for the rudders than tracking the sideslip angle. However, flight data shows that no significant sideslip excursion occurs when ailerons are applied without coordinated rudder inputs. This discrepancy could be attributed to the fact that, on the actual airframe, the ailerons are rigged to deflect more in the upwards direction than when traveling trailing-edge-down. Unfortunately, this mechanism cannot be modelled by the state-space representation of the plant chosen for this study.

4.2 Total Energy Control System (TECS)

The principle of the TECS is based on controlling both the total energy of the platform and its allocation into potential (height-dependent) and kinetic (airspeed-dependent) energies. Hence, it allows to indirectly control airspeed and altitude by taking advantage of the fact that the throttle influences the total energy of the UAV while the elevator mainly affects the energy balance of the platform. As shown by Spark [33], the non-dimensional rate of total energy change $\dot{E}_S$ can be approximated by the sum of the flight path angle γ and the time derivative of the airspeed TAS divided by the gravitational constant g. Conversely, the rate of change $\dot{L}_S$ of the Lagrangian of the system is given by the difference between those two quantities.

Hence, the selected TECS layout pictured in Fig. 15 features the proportional integral (PI) controllers $K_{\dot{E}_S}$ and $K_{\dot{L}_S}$ acting on their respective quantities. To avoid differentiation of noisy airspeed measurements, the true air speed (TAS) derivatives in $\dot{E}_S$ and $\dot{L}_S$ are replaced by the difference between the actual TAS

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

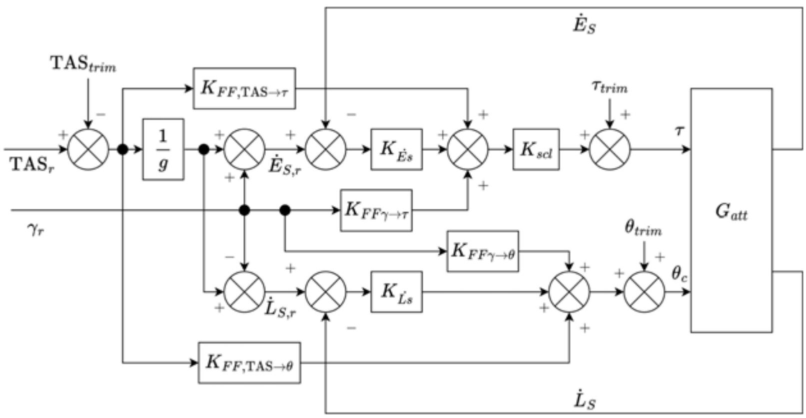

Figure 15. *Layout of the TECS.*

and its trim value. Feedforward action is added to both channels, with inputs from both γ and TAS references. In order to remove the high-frequency content from the reference signal before injecting it into the actuators, the four transfers $K_{FF,\mathrm{TAS}\to\theta}$, $K_{FF,\mathrm{TAS}\to\tau}$, $K_{FF,\gamma\to\theta}$ and $K_{FF,\gamma\to\tau}$ are chosen as second-order low-pass filters. The output thrust increment is scaled by K_{scl}, which is equal to $mg\mathrm{TAS}^2_{trim}$ divided by the square of the maximum engine power obtained in bench tests. Finally, the total pitch command θ_c is fed to the attitude control loop while the throttle fraction τ is directly passed onto the plant G. Note that the $\dot{E}_S$ channel is protected against integrator windup due to throttle command saturation. For readability purposes, the pre-conditioning of the input γ_r is omitted from Fig. 15: γ_r is taken as the altitude error $h_r - h$ divided by the trim TAS.

4.2.1 Synthesis framework

The gains of the TECS structure are tuned according to the same mixed-sensitivity principles applied to the attitude controller. However, a few differences should be mentioned. First, due to the different dynamics of the plant (UAV with its attitude controller), the previously used stability margins no longer provide a relevant tuning criterion. Hence, to limit the sensitivity peak on S and improve time-response damping, additional second-order step response matching requirements are added to TAS and h channels. Additionally, the low-pass filters on KS leading to a throttle action are replaced with static gain limits. This is because the propulsion model includes its own low-pass dynamics and is not susceptible to command-rate saturation. Finally, the placement of the feedforward controller poles is constrained to regions with damping ratios in excess of 0.6. This is to promote robustness to modelling uncertainties by preventing controllers with narrow resonance peaks. Table 6 lists the numerical values of the synthesis constraints, while Table 7 provides the numerical expressions of the tuned TECS transfers.

4.2.2 Frequency-domain analysis

Figure 16 displays the output sensitivities and reference-to-actuator transfers of the TECS. The left plot shows that the synthesis routine has more trouble meeting the sensitivity constraints compared to the attitude control case. Minor deviations from the templates can be observed, either at low frequencies for the TAS channel or slightly above the bandwidth for the h channel. The realised -3 dB bandwidths are slightly greater than the objective on the h channel (0.431 rad/s versus the prescribed 0.365 rad/s) but a bit narrower on the TAS channel (0.233 rad/s versus the prescribed 0.292 rad/s). The right plot reveals that the KS constraints are only active for both the TAS-to-elevator and h-to-throttle transfers. All constraints fit with prescribed templates.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

Table 6. *Parameters of the synthesis constraints*

Constraint Type	Parameter	Value	Unit
W_S (TAS channel)	$k_{S,LF}$	0.01	- -
	$k_{S,HF}$	1.2	- -
	ω_S	0.4	rad/s
W_S (h channel)	$k_{S,LF}$	0.01	- -
	$k_{S,HF}$	1.2	- -
	ω_S	0.5	rad/s
W_{mm} (TAS channel)	ω_0	0.8	rad/s
	Overshoot	10	%
W_{mm} (h channel)	ω_0	0.8	rad/s
	overshoot	10	%
$W_{KS}\left(T_{TAS_r \to \tau}\right)$	$k_{KS,\max}$	0.1	- -
$W_{KS}\left(T_{TAS_r \to \delta_M}\right)$	$k_{KS,LF}$	0.03	- -
	$k_{KS,HF}$	$1e^{-6}$	- -
	ω_{KS}	2.4	rad/s
$W_{KS}\left(T_{h_r \to \tau}\right)$	$k_{KS,\max}$	0.02	- -
$W_{KS}\left(T_{h_r \to \delta_M}\right)$	$k_{KS,LF}$	0.05	- -
	$k_{KS,HF}$	$1e^{-6}$	- -
	ω_{KS}	2.4	rad/s

Table 7. *Transfers of the tuned TECS*

$K_{\dot{E}s}$	$9.163 + \frac{0.2690}{s}$
$K_{\dot{L}s}$	$-4.164e^{-3} - \frac{0.05829}{s}$
$K_{FF,\text{TAS}\to\tau}$	$\frac{7.893e^{-4}}{s^2+7.790e^{-2}s+4.147e^{-3}}$
$K_{FF,\text{TAS}\to\theta}$	$\frac{-0.1745}{s^2+3.425s+8.139}$
$K_{FF,\gamma\to\tau}$	$\frac{1846}{s^2+450.9s+1730}$
$K_{FF,\gamma\to\theta}$	$\frac{499.3}{s^2+719.6s+982.2}$

4.2.3 Time-domain analysis

The tracking performance in response to a multistep input signal on TAS and altitude channels is displayed in Fig. 17. On those plots, the response of the full non-linear model (in dashed blue) is overlaid with the output of the linear model (in solid red). The TECS tracking performance is globally satisfactory, save for a slow convergence of the TAS, which is consistent with the sensitivity template violation identified in the previous section. The non-linear model (in blue) shows that throttle saturation degrades controller performance when a large step in airspeed is required. The linear model does not consider saturations, hence, the discrepancies observed between the two models. However, the anti-windup allows a prompt recovery when the reference airspeed is increased again. Note that the usage of the elevator for short-term energy exchange provides a quick response at the price of some cross-coupling between the altitude and airspeed channels. The airspeed undershoot in response to a climb request is moderate and does not jeopardise flight safety. One can also notice that the minimum airspeed transient happens as the pitch attitude is already being reduced and thus is unlikely to induce an excessive angle-of-attack leading to a stall.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

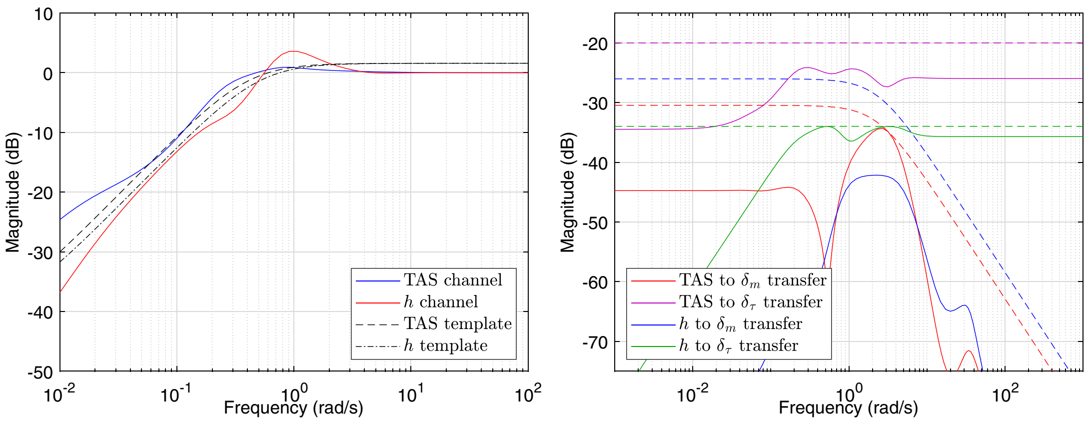

Figure 16. *Output sensitivities (left) and reference-to-actuator transfers (right) of the TECS channels.*

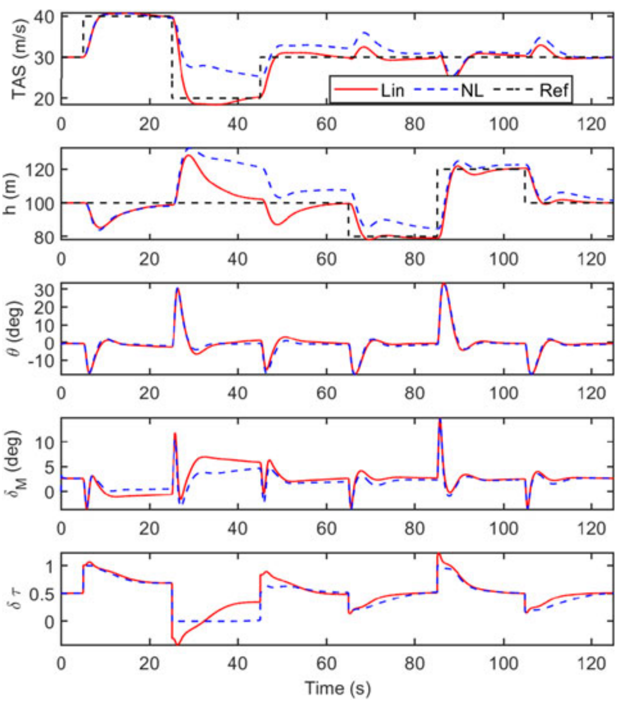

Figure 17. *Reference tracking of the TECS.*

4.3 Course controller and vector field horizontal guidance law

The course controller features a simple proportional action fed by the course error into the bank angle reference of the attitude controller. A bank limiter set at $\pm$ 60° is implemented to avoid inducing excessive disturbances on the longitudinal dynamics through lift vector inclination. The tuning procedure is conducted in a similar fashion to the two previous controllers and poses no particular challenge. A value of 1.259 is found for the proportional gain.

$$\chi_d = \chi_q - \chi_\infty \frac{2}{\pi} \tan^{-1}(k_{sl} e_{py}) \tag{3}$$

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

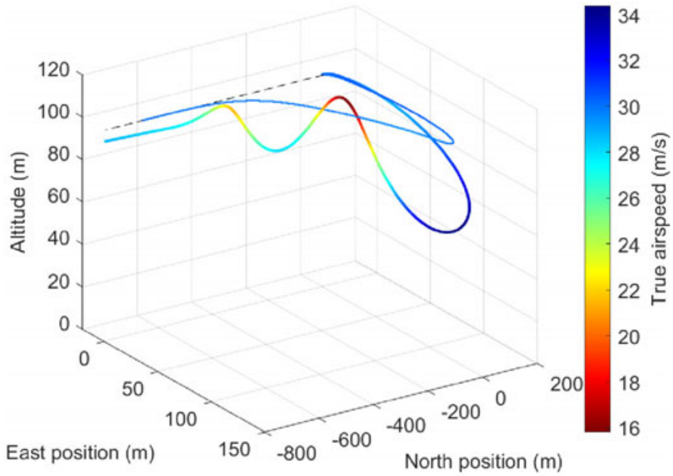

Figure 18. *Effect of icing on a level 180° turn manoeuvre: clean trajectory is shown in light blue while iced trajectory is multicoloured.*

The vector field horizontal guidance law is based on the works of Nelson [34]. It is selected due to its smooth course commands and satisfactory convergence properties in the case of wind disturbances. It has been implemented according to Equation (3) given by Fari [35], which is well-suited to the straight path following task. In this equation, e_{py} represents the lateral deviation of the UAV from the path while χ_q is the desired course and both χ_∞ and k_{sl} are user-defined parameters. χ_∞ represents the angle at which the path should be approached when the lateral separation is infinite and is set to 60°. k_{sl} is a tuning constant and is set to 0.008 to provide quick convergence while avoiding overshoots.

4.4 Simulation of an egress manoeuvre with clean and iced airframes

A simple sample case is introduced to check the operation of the baseline path-following autopilot and assess the impact of icing on the autopilot trajectory. The platform is commanded to perform a level 180° turn and level off to a southwards heading. The airspeed hold is set to 30 m/s, and the altitude to be maintained is 100 m. Figure 18 depicts the simulated trajectory in both clean (no ice) and iced conditions (as described in Section 3.2). Note that, for readability purposes, the axis scales differ between the North and East directions. Figures 19 and 20 provide more insight into the evolution of the flight variables and the control inputs during the manoeuvre for both cases.

The simulation results confirm that the path following autopilot is able to adequately steer the clean aircraft on the prescribed exit route. Altitude and airspeed excursions are minimal, and application of controls is smooth and moderate in magnitude. The UAV responds swiftly and requires less than 20 s to converge within five meters of the egress path. In the iced case, the trajectory projection into the North-East plane is not significantly affected, with a turn radius only 5% greater, but the time required to complete the manoeuvre increases to 30 s. Moreover, significant altitude and airspeed excursions are predicted, with a maximum altitude loss of 41 m and a minimum airspeed of only 16 m/s. The causes of this poor energy management are underlined in Figs 19 and 20. It is reasonable to assume that the iced airframe is unable to produce enough normal force to compensate for lift vector inclination, leading to a lift deficit and loss of altitude. In response to the undesired sink rate, the TECS commands a pitch and power increase. This causes the angle-of-attack to exceed the wing stall limit (dashed red line on Fig. 20), aggravating the lift deficit and altitude loss. The stall condition persists even though the bank angle is reduced, as the TECS struggles to converge to the target altitude. This manifests as poorly

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

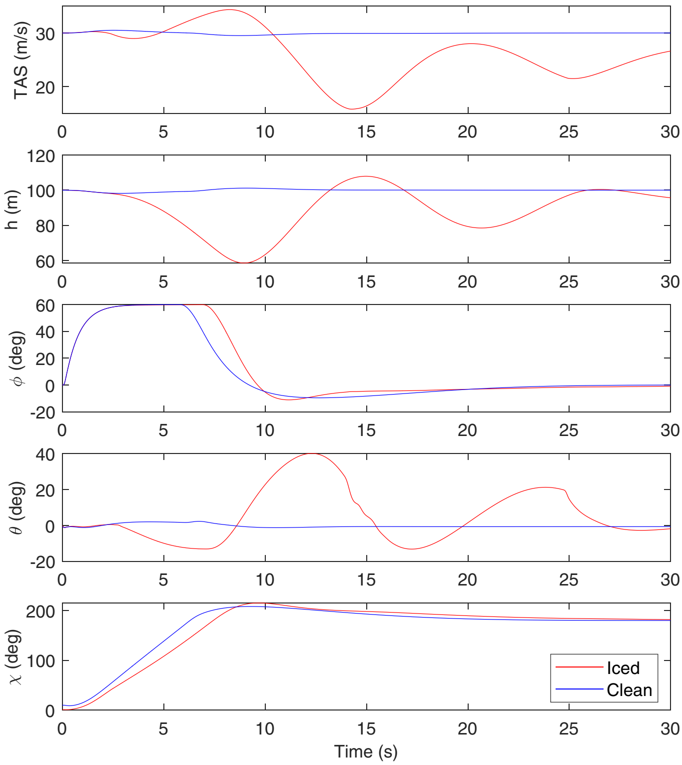

Figure 19. *Flight states during 180° turn manoeuvre.*

damped altitude and airspeed oscillations. As the ice accretion also reduces the lift of the horizontal tail and the control authority of the elevator, a saturation occurs, preventing the controller from increasing the angle-of-attack beyond the domain of validity of the longitudinal model.

One could argue that the outcome of the iced platform simulation is overly optimistic. In fact, it is likely that the wing stall would have affected the lateral dynamics, especially the aileron's effectiveness. This unmodelled effect could have delayed or even prevented recovery. On top of the aforementioned modelling limitation, it should be kept in mind that the impact of icing on the platform's aerodynamics remains highly uncertain as the CFD solver used for predicting the shapes of the ice accretions on the aerofoil has not yet been fully validated for Reynolds numbers as low as those experienced by small UAVs [36]. Moreover, even though the icing case has been identified as the worst possible case in the frame of the CS-25 Appendix C scenarios, more severe conditions may be encountered at low altitudes where atmospheric conditions are less known. In any case, it can be concluded that the behaviour exhibited by the baseline autopilot in icing is inadequate and needs improvement.

5.0 Improvement of autopilot performance in icing conditions

5.1 Modifications to the autopilot structure

The simulation results obtained in the previous section provide some valuable insight for improving the autopilot resilience in icing conditions. The inadequate management of the platform energy appears to be caused by the inability of the baseline controller to account for the platform lift limit and regulate the bank angle accordingly. The first step to more resilient flight control is then to prevent the controller from commanding angles of attack beyond the stall limit. Even though this limit varies according to the icing level of the airframe, the upper part of the lift polar is not required for manoeuvreing at reasonable

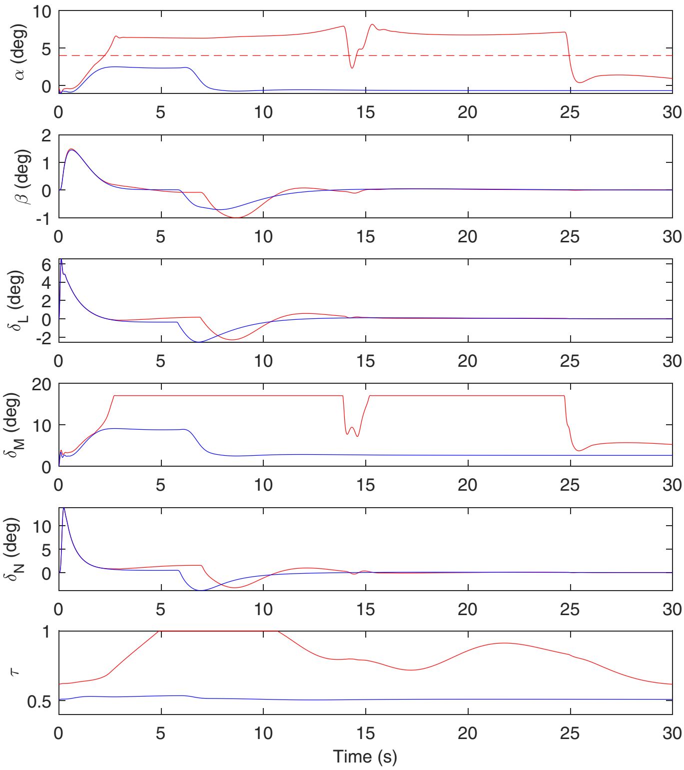

Figure 20. *Airflow angles and controls during 180° turn manoeuvre.*

bank angles and typical cruise airspeeds. Hence, the maximum elevator travel has been reduced to 10° nose-up instead of the original 17°. This modification, named elevator clamping, results in a reduction of the maximum angle-of-attack that can be induced by the attitude controller.

As the autopilot does not require such large deflections in normal flight, elevator saturation can be regarded as an indicator of degraded flight performance. An adaptive bank angle limiter can then be designed around this Boolean saturation signal. Its concept is straightforward: as long as the elevator saturation condition is true, an integrator continuously decrements the bank limit at a predetermined rate. When the condition disappears, another integrator re-increments the bank limit up to its maximum value of 60°. Increment and decrement rates can be tuned heuristically to compromise between mitigating the altitude excursion h_{min} and keeping the manoeuvre completion time t_{mvr} low. In this example, the rates were chosen to be equal, and the influence of this parameter on the previously mentioned trade-off was assessed by means of batch simulations. Figure 21 illustrates this compromise and shows that the adaptive bank limiter is able to bring the performance one step closer to the levels observed in the clean scenario. A new figure of merit was introduced to discriminate between the performance levels offered by discrete values of the limiter adaptation rate. It can be described as a non-dimensional distance between the performance point of interest and the clean reference point in the $(t_{\text{mvr}}, h_{\text{min}})$ plane. Its formula is defined according to Equation (4) and leveraged to compile Fig. 22. The bar chart shows a local minimum of the performance metric, which reveals the most interesting compromise. Given that the purpose of the adaptive bank limiter is to reduce the altitude loss while minimising the associated manoeuvre time penalty, a value of 3°/s, which corresponds to the second most optimal choice, is retained for the limiter adaptation rate. This value is used to obtain the trajectories labeled ‘clamping + adaptive’ in the remaining of this paper.

$$\mathrm{D}_{\text{perf}} = \sqrt{\frac{(t_{\text{mvr}} - t_{\text{mvr,clean}})^2}{t^2_{\text{mvr,clean}}} + \frac{(h_{\text{min}} - h_{\text{min,clean}})^2}{h^2_{\text{min,clean}}}} \tag{4}$$

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

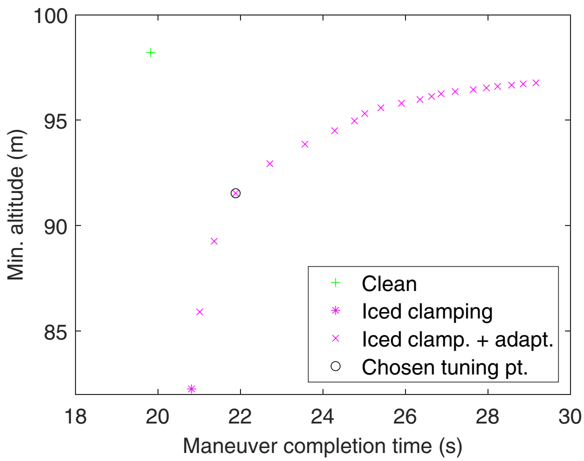

Figure 21. *Compromise between altitude loss h_{min} and manoeuvre duration t_{mvr}.*

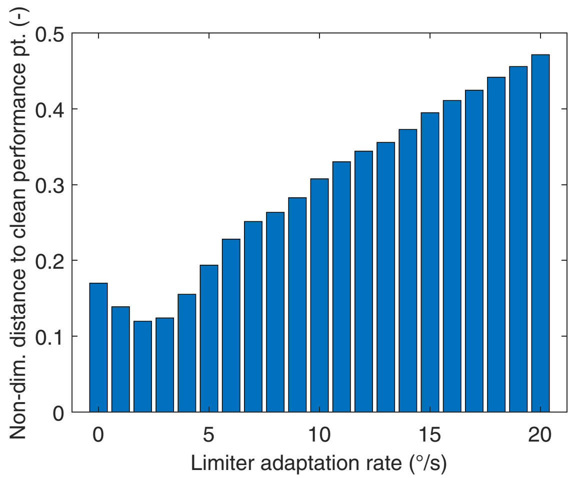

Figure 22. *Performance metric for selection of the bank limiter adaptation rate.*

It should be noted that the tuning of the adaptive bank limiter has been performed under the assumption of a given level of ice accretion, considered to be the worst-case scenario. If the actual icing level differs from the one used for tuning, the trade-off between minimum altitude and manoeuvre completion time may not be as relevant. For instance, if the icing level is less than expected when the manoeuvre is started, the penalty induced by the limiter on the completion time may be excessive with regard to the benefit of the reduction of altitude excursions. Hence, the icing level used for limiter tuning should be matched with the expected level at which the ice detection routine would trigger the manoeuvre. Regarding the elevator clamping threshold, it has been set to enable an adequate maximum lift to be generated in both clean and fully iced conditions. Hence, as the change of the provided lift polars with ice build-up is gradual, any intermediate icing condition is unlikely to result in a significant degradation of the trajectory due to insufficient lift.

5.2 Egress simulation in iced configuration

The egress manoeuvre presented in the previous section has been used as a benchmark to evaluate the effectiveness of both elevator clamping and the adaptive bank limiter. Figure 23 showcases the effect of

***Table 8.** Performance metrics for iced egress manoeuvre*

Case	Manoeuvre Time (s)	Min. Alt. (m)
Clean	19.8	98.2
Iced	29.7	58.6
Iced w. clamping	20.8	82.3
Iced w. clamp. + adapt.	21.9	91.5

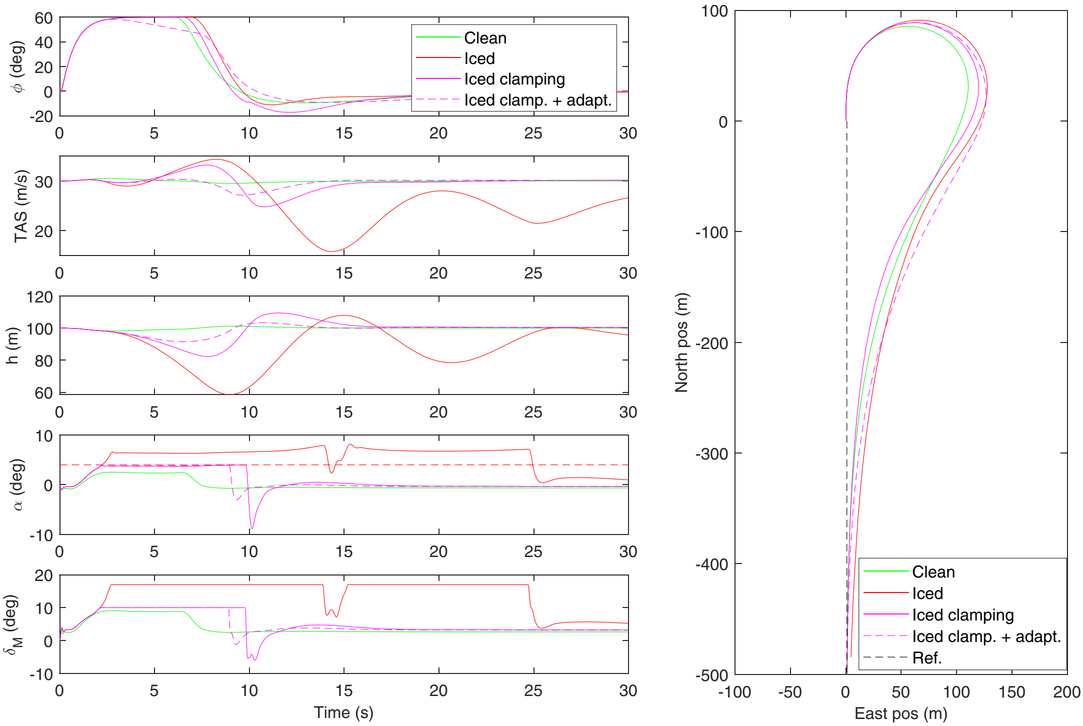

***Figure 23.** Longitudinal dynamics (left) and horizontal flight path (right) of improved autopilot in icing.*

both elevator clamping and adaptive bank limiter on the respective vertical and horizontal trajectories followed during the manoeuvre. One can notice from the first figure that the elevator clamping method is effective at keeping the angle-of-attack of the platform below its stall threshold of 4° (dashed red line). This allows the wing to operate at maximum lift coefficient, significantly reducing the lift deficit in the vertical plane and thus the altitude loss during the turn. As a consequence, the response from the autopilot is more gentle, and the magnitude of the resulting airspeed loss is decreased. The addition of the adaptive bank limiter leads to an early reduction of the platform bank angle, further contributing to the reduction of the lift deficit and the altitude excursion. The impact of the bank reduction on the turn radius of the UAV is minimal.

Table 8 quantifies the effect of the autopilot improvements on the altitude excursion and manoeuvre completion time. Note that the criterion for manoeuvre completion is fulfilled when the lateral deviation relative to the commanded exit path reaches less than five meters. The performance recovery offered by the combination of both solutions is significant, with the iced UAV losing less than 10 meters of attitude and needing only 2.1 s more to complete the manoeuvre compared to the clean platform. These values are to be compared to the more than 40 meters lost and about 10 seconds added to the manoeuvre when the baseline autopilot is operating in iced configuration. Under the light of those results, the authors conclude that the objective of designing an icing-resilient autopilot has been achieved.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

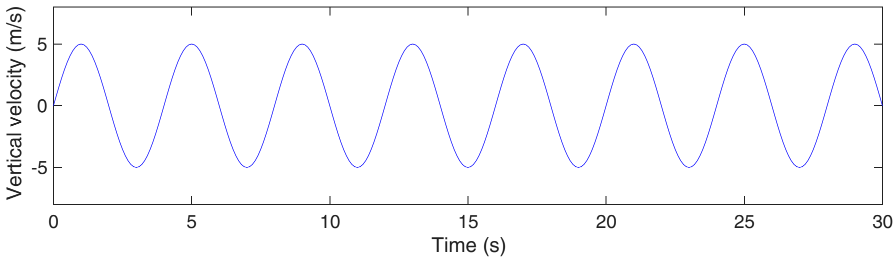

Figure 24. *Vertical velocity gust profile.*

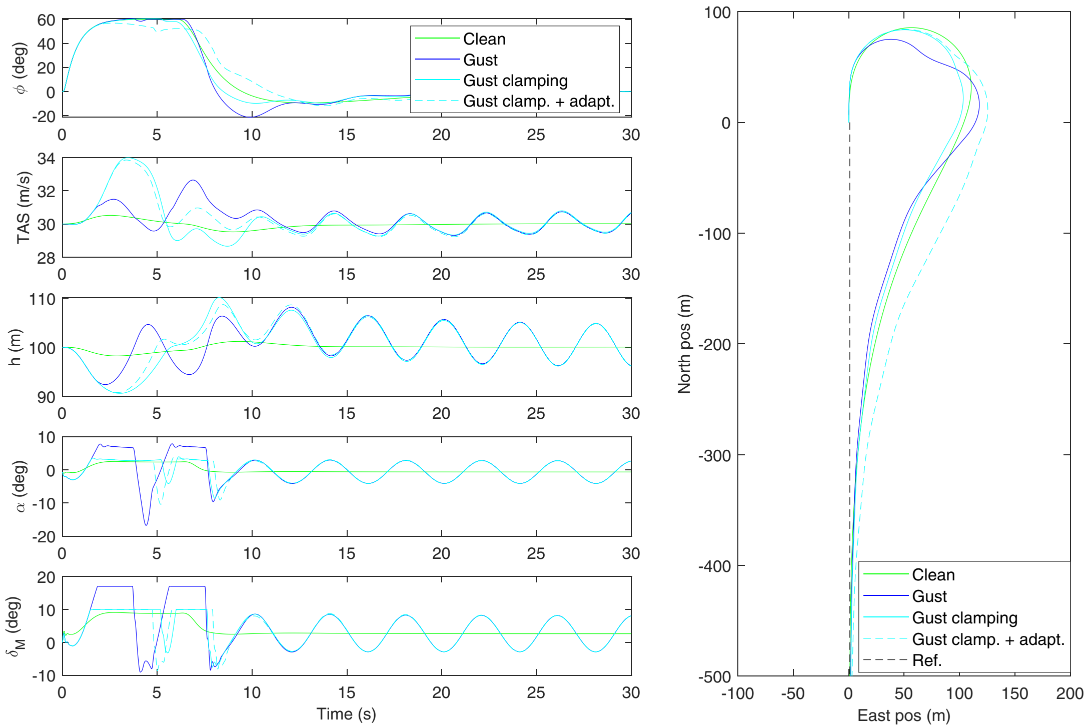

Figure 25. *Longitudinal dynamics (left) and horizontal flight path (right) of improved autopilot in sinusoidal gust.*

5.3 Gust rejection simulation in clean configuration

Analysis of manoeuvre simulation has shown that the threshold chosen for elevator clamping does not affect the behaviour of the platform in nominal conditions. However, one can wonder if the disturbance rejection properties of the controller may have been affected by elevator clamping. To investigate this question, a new simulation case was designed on the basis of the egress manoeuvre. This case is independent of the iced case as it features a clean airframe subjected to a wind disturbance. A sinusoidal gust pictured in Fig. 24 has been injected on the local vertical axis for the whole manoeuvre duration to evaluate the disturbance rejection properties of the baseline and modified autopilots. Figure 25 depicts the trajectory of the platform under those conditions and the action of the pitch control system. The green line corresponds to the clean airframe with its baseline autopilot performing the manoeuvre in the absence of wind. The dark blue line shows the effect of the gust on the baseline autopilot. The cyan lines show the effect of the gust on the modified autopilots.

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

Table 9. *Performance metrics for gusty egress manoeuvre*

Case	Manoeuvre Time (s)	Min. Alt. (m)
Clean	19.8	98.2
Gust	19.6	92.4
Gust w. clamping	20.1	90.6
Gust w. clamp. + adapt.	22.8	90.8

All autopilots are able to successfully complete the manoeuvre while attenuating the effect of the disturbance on the flight states. The difference between the different architectures is only apparent during the transients in the first 10 seconds of the manoeuvre. Because of the restriction in control power due to elevator clamping, the first couple of peaks of the response are slightly more pronounced, even though the amplitude of the deviations compared to steady state remains reasonable. Table 9 shows the same metrics as for the iced case, this time under the influence of the sinusoidal gust. As expected, the effect of the gust on the baseline autopilot is minimal, with a minimum altitude delta of less than 6 meters and an approximately equal manoeuvre completion time. The elevator clamping modification barely affects the performance metrics, while the addition of the adaptive bank limiter does slightly increase manoeuvre time. In the authors' opinion, this mild performance degradation in a fairly uncommon edge case does not question the relevance of the bank limiter modification. Overall, this test case shows no major regression in the disturbance rejection properties of the controller that could be attributed to the modification of the autopilot. This confirms the fitness for purpose of the improved solution.

6.0 Summary

This paper introduces the idea of sense-and-egress to mitigate the threat posed by icing conditions to small UAVs. The relevance and feasibility of the concept were investigated, and an architecture candidate was proposed. A non-linear model of the Falk UAV flight dynamics was developed and refined through system identification from flight test data. The model was implemented into a controller synthesis framework in order to tune the path following autopilot tasked with conducting the escape manoeuvre. The design and tuning process of both the attitude and total energy controllers was detailed, including frequency- and time-domain analyses. A sample case based on a 180° level turn showed that the baseline autopilot is adequate for steering the clean airframe but exhibits significant shortcomings when aerodynamic penalties due to ice accretion are introduced. Two simple solutions were developed to improve the autopilot performance in icing: elevator clamping and adaptive bank limiting. The combination of both modifications allows to speed up the egress manoeuvre and limits the altitude excursion during its execution, while preserving the gust rejection capabilities of the autopilot. Compared to the baseline controller, the manoeuvre time penalty induced by icing has been cut down from +50% to +10.6%, while the relative amplitude of the altitude excursion has been reduced from 41% to only 8.5%. The impact of icing on both of those metrics has been mitigated by a factor of almost five, confirming the effectiveness of the controller modifications. Overall, the proof of concept of an icing-resilient autopilot able to conduct simulated egress manoeuvres has been successfully established.

Several additional perspectives are left for future work: the suitability of the improved autopilot could be assessed on more realistic egress trajectories involving altitude changes. For instance, a diving turn would allow building up vertical separation with respect to a hypothetical icing layer while conserving the platform's kinetic energy. Additionally, the impact of icing on the trajectory may be analysed for different accretion times corresponding to the reaction time of the ice detection subsystem. This could be complemented by a study on the sensitivity of various ice detection methods and the choice of an adequate threshold for the ice detector. In order to further improve the performance of the autopilot in icing, the authors suggest additional research to be conducted on the topics of flight condition monitoring and flight envelope prediction. More complex control methods, such as model predictive control, should also be investigated for their ability to explicitly account for time-dependent constraints on the flight

https://doi.org/10.1017/aer.2026.10199 Published online by Cambridge University Press

envelope bounds. Finally, to validate the conclusions of this study, the simulated egress manoeuvre should be test flown on the actual UAV fitted with artificial ice shapes replicating the aerodynamic penalties induced by ice accretion.

Acknowledgements. This work has been conducted under a DLR-DAAD fellowship programme. The authors wish to thank the German Academic Exchange Service (DAAD) for their financial support.

References

[1] Sørensen, K.L. et al. *UAV Atmospheric Icing Limitations, Climate Report for Norway and Surrounding Regions*. Tech. rep. UBIQ Aerospace, 2021.
[2] Hann, R. and Johansen, T.A. UAV icing: the influence of airspeed and chord length on performance degradation, *AEAT*, 2021, **93**, (5), p 832841.
[3] Hann, R. et al. *Experimental investigations of an icing protection system for UAVs*. Tech. rep. SAE Technical Paper, 2019.
[4] Deiler, C. and Fezans, N. Performance-based ice detection methodology, *J. Aircr.*, 2020, **57**, (2), pp 209–223.
[5] Sachs, F., Schwarz, C.W. and Deiler, C. Flight testing the indirect ice detection system in the Horizon 2020 project SENS4ICE, *AIAA Aivation Forum and Ascend 2024*. Las Vegas: American Institute of Aeronautics and Astronautics, 2024.
[6] Deiler, C. Aerodynamic modeling, system identification, and analysis of iced aircraft configurations, *J. Aircr.*, 2018, **55**, (1), pp 145–161.
[7] Lindner, M., Wallisch, J. and Hann, R. *UAV icing: Numerical simulation of icing effects on wing and empennage*. Tech. rep. SAE Technical Paper, 2023.
[8] Skogan, E.I. Modelling and Safety-based Path Planning of Fixed-Wing Unmanned Aerial Vehicle in Icing Conditions. MA thesis. NTNU, 2023.
[9] Løw-Hansen, B. et al. *Identification of an Electric UAV Propulsion System in Icing Conditions*. Tech. rep. SAE Technical Paper, 2023.
[10] Hann, R. *UAV icing: ice accretion experiments and validation*. Tech. rep. SAE Technical Paper, 2019.
[11] Hann, R. et al. *UAV Icing: Experimental Validation Data for Predicting ice Shapes at Low Reynolds Numbers*. Tech. rep. SAE Technical Paper, 2023.
[12] Deiler, C. Testing of an indirect ice detection methodology in the Horizon 2020 project SENS4ICE, *CEAS Aeronaut. J.*, 2024, **16**, (1), pp 213–229.
[13] Beard, R.W. and McLain, T.W. *Small Unmanned Aircraft: Theory and Practice*, Princeton University Press, Princeton, 2012.
[14] Zogopoulos-Papaliakos, G., Karras, G.C. and Kyriakopoulos, K.J. A fault-tolerant control scheme for fixed-wing UAVs with flight envelope awareness, *J. Intell. Robot. Syst.*, 2021, **102**, (2), p 46.
[15] Fischenberg, D. and Jategaonkar, R. Identification of Aircraft Stall Behavior from Flight Test Data, System Identification for Integrated Aircraft Development and Flight Testing. *RTO Meeting Proceedings*. Madrid, Spain: RTO/NATO, 1998.
[16] Saderla, S., Kim, Y. and Ghosh, A.K. Online system identification of mini cropped delta UAVs using flight test methods, *Aerosp. Sci. Technol.* 2018, **80**, pp 337–353.
[17] de Aguiar, R.S.S., Apkarian, P. and Noll, D. Structured robust control against mixed uncertainty, *IEEE Trans. Control Syst. Technol.*, 2018, **26**, (5), pp 1771–1781. Conference Name: *IEEE Transactions on Control Systems Technology*.
[18] Apkarian, P. Tuning controllers against multiple design requirements, *American Control Conference*, 2013, pp 3888–3893.
[19] Gavilan, F., Vazquez, R. and Camacho, E.F. A High-level model predictive control guidance law for unmanned aerial vehicles, *European Control Conference (ECC)*, 2015, pp 1362–1369.
[20] Eren, U. et al. Model predictive control in aerospace systems: current state and opportunities, *J. Guid. Control Dyn.* 2017, **40**, (7), pp 1541–1566. https://doi.org/10.2514/1.G002507.
[21] Hann, R. and Trolle Borup, K. Technical requirements for RPAS operations in cold climates, *International Workshop on Atmospheric Icing of Structures (IWAIS)*, 2024.
[22] Ratvasky, T., Zante, J.V. and Riley, J. NASA/FAA tailplane icing program overview, *Aerospace Sciences Meeting and Exhibit. American Institute of Aeronautics and Astronautics*, Reno, NV, USA, 1999.
[23] Bromfield, M.A. et al. Loss of control in flight accident case study: icing-related tailplane stall, *Aeronaut. J.* 2023, **127**, (1315), pp 1554–1573.
[24] Zipfel, P.H. *Modeling and Simulation of Aerospace Vehicle Dynamics*, American Institute of Aeronautics and Astronautics, Reston, 2000.
[25] de Sousa, R.S.C. and da Silva, R.G.A. NASA-CRM lateral-directional stability derivatives comparison between semi-empirical, vortex lattice and panel method, *32th ICAS Congress 2020/21*. Shanghai, China, 2021.
[26] Cook, M.V. *Flight Dynamics Principles: A Linear Systems Approach to Aircraft Stability and Control*, Butterworth-Heinemann, Oxford, 2012.
[27] Mracek, C. and Ridgely, D. Missile longitudinal autopilots: comparison of multiple three loop topologies, *AIAA Guidance, Navigation, and Control Conference and Exhibit*. American Institute of Aeronautics and Astronautics, San Francisco, California, 2005.
[28] Strub, G. et al. Pitch axis control for a guided projectile in a wind tunnel hardware-in-the-loop setup, *J. Spacecr. Rockets* 2015, **52**, (6), pp 1614–1626.

[29] Johansen, T.A. et al. On estimation of wind velocity, angle-of-attack and sideslip angle of small UAVs using standard sensors, *International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2015, pp 510–519.
[30] Kwakernaak, H. Robust control and H-infinity-optimization—tutorial paper, *Automatica* 1993, **29**, (2), pp 255–273.
[31] Doyle, J. et al. State-space solutions to standard H2 and H-infinity control problems, *American Control Conference*, 1988, pp 1691–1696.
[32] Skogestad, S. and Postlethwaite, I. *Multivariable Feedback Control: Analysis and design*, John Wiley & Sons, Hoboken, 2005.
[33] Spark, H. et al. *An Assessment of Aircraft Control via SISO Control Loops and Total Energy Control*, Deutsche Gesellschaft für Luft-und Raumfahrt-Lilienthal-Oberth eV, Bremen, Germany, 2021.
[34] Nelson, D.R. et al. Vector field path following for miniature air vehicles, *IEEE Trans. Rob.* 2007, **23**, (3), pp 519–529.
[35] Fari, S. Guidance and control for a fixed-wing UAV. *MA thesis*. Politecnico di Milano, 2016.
[36] Hann, R. UAV Icing: challenges for computational fluid dynamic (CFD) tools, *International Conference on Computational Fluid Dynamics (ICCFD11)*, 2022.

Cite this article: Riss Z. and Deiler C. Development of an icing-resilient autopilot for automated egress of small UAVs from dangerous icing conditions. *The Aeronautical Journal*, https://doi.org/10.1017/aer.2026.10199