

PHYSICS-INFORMED LEARNING-BASED RECOVERY TRAJECTORY PREDICTION FOR AUTOMATIC GROUND COLLISION AVOIDANCE

Jiho Han[1], Sungjoong Kim[1], Minhyuk Jung[1], Jinho Woo[1], Hyun-seok Kim[2] & Hyochoong Bang[1]

[1]Korea Advanced Institute of Science and Technology
[2]APPIA Engineering Co., Ltd.

Abstract

Automatic Ground Collision Avoidance Systems (AGCAS) require reliable recovery trajectory prediction to assess terrain collision risk and determine whether automatic recovery should be initiated. Although high-fidelity dynamics-based propagation can provide accurate recovery predictions, repeatedly performing such propagation online may be computationally demanding for real-time AGCAS applications. In contrast, simplified kinematics-based prediction methods are computationally efficient, but they may not sufficiently capture aerodynamic effects, aircraft response delays, transient pull-up dynamics, and velocity variations during aggressive recovery maneuvers. This paper proposes a physics-informed supervised learning framework for recovery trajectory prediction in AGCAS applications. Recovery maneuver data are generated offline using an F-16 six-degree-of-freedom (6-DOF) simulator with an aerodynamic database, and the altitude-change response during recovery is represented by a compact polynomial model. A global learning model is trained to map the recovery-initiation state, consisting of flight speed, flight-path angle, and roll angle, to the polynomial coefficients of the recovery trajectory. To improve prediction accuracy and consistency, two physics-informed constraints are incorporated into the learning model: an initial altitude-rate constraint and a roll-symmetry constraint. Simulation results show that the proposed physics-informed learning model reduces the median trajectory prediction root mean square error (RMSE) by approximately 89% compared with a simplified kinematics-based model and by approximately 46% compared with an unconstrained baseline learning model. The learned recovery trajectory model is further applied to an AGCAS scenario in which delayed recovery trajectories are compared with a buffered terrain profile for terrain collision assessment and automatic recovery initiation, demonstrating its applicability to ground collision avoidance.

Keywords: Automatic Ground Collision Avoidance System, Recovery Trajectory Prediction, Supervised Learning, Physics-Informed Learning, Data-Driven Modeling

1. Introduction

Automatic Ground Collision Avoidance Systems (AGCAS) have been developed to prevent Controlled Flight Into Terrain (CFIT) accidents by autonomously initiating recovery maneuvers when an impending terrain collision is detected and pilot response is insufficient [1, 2]. In a typical AGCAS architecture, the system scans the terrain along the predicted flight path, predicts the aircraft recovery trajectory, and assesses whether the predicted trajectory intersects the terrain profile. The F-16 Auto GCAS, for example, employs a pre-planned recovery maneuver consisting of a roll to wings level followed by a high-load-factor pull-up. Accordingly, recovery trajectory prediction plays a central role in AGCAS because collision assessment directly depends on the expected aircraft motion after recovery initiation.

A reliable recovery trajectory prediction model is particularly important during aggressive recovery maneuvers. Prior GCAS nuisance-criteria studies have shown that recovery prediction errors, terrain database uncertainty, and navigation errors all contribute to the activation margin and must be

managed to avoid both ground impact and nuisance recoveries [3]. An optimistic prediction may delay automatic intervention and increase the risk of terrain impact, whereas an overly conservative prediction may trigger unnecessary automatic recovery and degrade operational acceptance. This activation margin can be expressed in terms of Time Available (TA), or equivalently as the allowable recovery delay for which the predicted recovery trajectory still clears the terrain. In this study, delayed recovery trajectories are evaluated to assess the remaining recovery margin, and the AGCAS application uses a prescribed 1.5 s activation delay as a representative simulation setting motivated by prior GCAS nuisance-criteria studies [3]. Therefore, an accurate representation of recovery behavior is essential for robust AGCAS operation.

High-fidelity dynamics-based propagation can provide accurate recovery trajectory predictions by accounting for detailed aircraft dynamics, aerodynamic characteristics, and control responses. However, repeated online propagation may be computationally demanding for real-time AGCAS applications, especially when multiple candidate initiation times or terrain conditions must be evaluated. Simplified kinematics-based methods reduce this computational burden, but they may not sufficiently capture aerodynamic effects, response delays, transient pull-up dynamics, and velocity variations during aggressive recovery maneuvers. Optimization-based approaches, including optimal-control and model-predictive-control methods, have also been investigated to generate or predict terrain-avoidance trajectories under aircraft performance and terrain constraints [4, 5]. Although these methods can provide adaptive trajectory generation, they may require repeated online optimization during flight. These considerations motivate a surrogate prediction approach that retains the recovery characteristics obtained from high-fidelity simulation data while enabling efficient online trajectory prediction.

To address this need, this paper proposes a supervised learning-based recovery trajectory prediction approach for AGCAS applications. The proposed method learns recovery trajectory characteristics from offline recovery maneuver data generated using an F-16 6-DOF simulator with an aerodynamic database [6]. For each simulated recovery maneuver, the recovery-initiation state, including the flight speed, flight-path angle, and roll angle, is stored together with the altitude-change profile over the recovery horizon. The altitude-change profile is approximated using a polynomial representation, and the learning problem is formulated as a regression from the recovery-initiation state to the polynomial coefficients of the recovery altitude profile.

The main contribution of this study is the development of a physics-informed supervised learning framework for AGCAS recovery trajectory prediction. The proposed method compresses the recovery altitude response into a compact polynomial form and trains a global regression model using the initial speed, flight-path angle, and roll angle. Two physical properties are then incorporated into the model structure: the initial altitude-rate condition, imposed by fixing the first polynomial coefficient using the initial vertical velocity, and the roll-symmetry property, reflected by using the roll-angle magnitude in the regression model. These constraints reduce unnecessary learning freedom and improve prediction accuracy.

The learned recovery trajectory model is evaluated in terms of both prediction accuracy and AGCAS applicability. Prediction accuracy is assessed by comparing a simplified kinematics-based model, an unconstrained baseline learning model, and the proposed physics-informed learning model. In the AGCAS simulation, delayed recovery trajectories are compared with the buffered forward terrain profile, and automatic recovery is initiated when the trajectory predicted with a prescribed 1.5 s delay no longer remains clear of the terrain buffer. These results demonstrate that the proposed learning-based model can support both trajectory prediction and terrain collision assessment in an AGCAS framework.

The remainder of this paper is organized as follows. Section 2. presents the AGCAS problem formulation and the recovery trajectory prediction problem. Section 3. describes simulation-based dataset generation and the learning-based prediction framework. Section 4. introduces the physics-informed constraints based on the initial altitude rate and roll symmetry. Section 5. analyzes the recovery trajectory prediction accuracy and presents the AGCAS application results including automatic recovery simulation. Finally, Section 6. concludes the paper.

2. Problem Formulation

2.1 Overview of AGCAS

An AGCAS is designed to prevent terrain impact by monitoring the aircraft state and the forward terrain profile, assessing potential collision risk, and initiating an automatic recovery maneuver when necessary. A typical AGCAS architecture consists of four main functions: terrain scanning, recovery trajectory prediction, collision assessment, and automatic recovery command generation. The terrain scanning constructs the forward terrain profile along the predicted flight path using navigation and terrain elevation information, with safety margins included when necessary to account for system uncertainties. The recovery trajectory prediction function estimates the aircraft motion that would result if a recovery maneuver were initiated from the current or a delayed time. The collision assessment function compares recovery trajectories predicted under assumed initiation delays with the terrain profile to determine how long recovery can be delayed while maintaining terrain clearance. Automatic recovery is initiated when the allowable delay margin is exhausted. Figure 1 illustrates the four main AGCAS functions and their information flow.

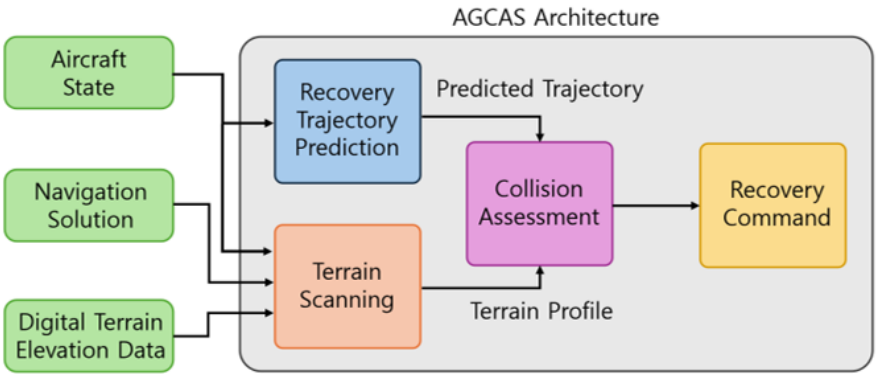


Figure 1 – Conceptual overview of AGCAS architecture

The recovery scenario considered in this study follows a standard AGCAS-type sequence. In the trajectory prediction process, the system evaluates the recovery trajectory that would result if automatic recovery were initiated after a prescribed delay from the current time. By considering different assumed delays, AGCAS can assess how long recovery initiation can be delayed while still maintaining terrain clearance. The predicted recovery scenario is composed of the following stages:

1. **Recovery delay:** The aircraft continues along its current flight condition for a prescribed delay, representing the assumed time margin before automatic recovery initiation. This stage is used to evaluate whether terrain clearance can still be maintained for the assumed recovery delay.

2. **Roll recovery:** After recovery initiation, the aircraft rolls toward a wings-level attitude so that the subsequent pull-up maneuver can generate lift primarily in the upward direction.

3. **Load-factor build-up:** The commanded normal load factor is gradually increased toward the prescribed maximum value, representing the transient response to the pull-up command.

4. **Maximum-load-factor pull-up:** Once the maximum load factor is reached, the aircraft maintains the maximum load factor to increase the flight-path angle and recover altitude. In this study, the maximum commanded load factor is set to $5\,G$.

5. **Flight-path-angle hold:** When the flight-path angle reaches its prescribed maximum, the recovery command switches to a flight-path-angle hold mode. The maneuver terminates or transitions to another control mode upon satisfying completion conditions, such as sufficient terrain clearance.

In this formulation, the recovery delay represents the assumed time margin before automatic recovery initiation. Therefore, determining the recovery activation timing is equivalent to determining how

long this delay can be allowed while still maintaining terrain clearance. If the allowable delay is overestimated, automatic recovery may be initiated too late, leaving insufficient terrain clearance and increasing the risk of ground impact. Conversely, if the allowable delay is underestimated, recovery may be initiated unnecessarily early, which can interfere with the mission and reduce pilot trust. Therefore, reliable recovery trajectory prediction is required to evaluate the allowable recovery delay and to support collision assessment in AGCAS.

2.2 Recovery Trajectory Prediction Problem

This study focuses on the recovery trajectory prediction function in the AGCAS framework. Although the aircraft recovery motion is three-dimensional, the terrain collision assessment considered in this study is mainly based on the vertical separation between the predicted altitude profile and the forward terrain profile. Therefore, the prediction problem is formulated in terms of the altitude-change response during recovery.

Let $h(t)$ denote the aircraft altitude, and let $t = 0$ represent the recovery initiation time. The altitude change during recovery is defined as

$$\Delta h(t) = h(t) - h(0). \tag{1}$$

Given the predicted altitude-change response, the corresponding aircraft altitude profile can be reconstructed as

$$\hat{h}(t) = h(0) + \Delta\hat{h}(t), \tag{2}$$

where $\Delta\hat{h}(t)$ denotes the predicted altitude change.

The recovery altitude response depends strongly on the flight condition at the moment of recovery initiation. In particular, the flight speed affects both the initial descent distance and the subsequent pull-up response, the flight-path angle determines the initial vertical motion, and the roll angle affects the duration and effectiveness of the roll recovery phase. Accordingly, the prediction model uses the flight speed V_0, flight-path angle γ_0, and roll angle ϕ_0 at the recovery initiation point as input variables. Given V_0, γ_0, and ϕ_0, the objective of the prediction model is to estimate the altitude-change profile over a finite recovery horizon $0 \le t \le T$. This prediction problem is expressed as

$$(V_0, \gamma_0, \phi_0) \mapsto \Delta\hat{h}(t), \qquad 0 \le t \le T. \tag{3}$$

The predicted altitude profile is then compared with the forward terrain profile in the collision assessment process. In this paper, the altitude-change prediction model is constructed using a supervised learning approach trained on recovery maneuver data generated from an F-16 6-DOF simulation model.

3. Learning-Based Recovery Trajectory Prediction

3.1 Simulation Model for Dataset Generation

Recovery maneuver data are generated using a nonlinear 6-DOF rigid-body simulation model of the F-16 fighter aircraft. The simulation model is implemented in MATLAB/Simulink based on the mathematical model and aerodynamic data reported by Nguyen et al. [6]. The purpose of the simulation model in this study is not to develop a new aircraft model, but to generate physically consistent recovery maneuver trajectories over a wide range of initial flight conditions.

Aircraft dynamics. The aircraft equations of motion are formulated in a body-fixed axis frame. Let u, v, and w denote the body-axis velocity components, and let p, q, and r denote the roll, pitch, and yaw rates, respectively. The translational equations of motion are given by

$$\begin{aligned}
\dot{u} &= rv - qw - g\sin\theta + \frac{\bar{q}S}{m}C_{X,t} + \frac{T}{m}, \\
\dot{v} &= pw - ru + g\cos\theta\sin\phi + \frac{\bar{q}S}{m}C_{Y,t}, \\
\dot{w} &= qu - pv + g\cos\theta\cos\phi + \frac{\bar{q}S}{m}C_{Z,t},
\end{aligned} \tag{4}$$

where m is the aircraft mass, g is the gravitational acceleration, ϕ and θ are the roll and pitch Euler angles, $\bar{q}$ is the dynamic pressure, S is the wing reference area, T is the engine thrust, and $C_{X,t}$, $C_{Y,t}$, and $C_{Z,t}$ are the total aerodynamic force coefficients along the body axes.
For the laterally symmetric aircraft model, $I_{xy} = I_{yz} = 0$, while I_{xz} is retained. The rotational equations of motion are written as

$$\begin{aligned}
\dot{p} &= \frac{I_Y - I_Z}{I_X} qr + \frac{I_{XZ}}{I_X}(\dot{r} + pq) + \frac{\bar{q}Sb}{I_X} C_{l,t}, \\
\dot{q} &= \frac{I_Z - I_X}{I_Y} pr + \frac{I_{XZ}}{I_Y}\left(r^2 - p^2\right) + \frac{\bar{q}S\bar{c}}{I_Y} C_{m,t} - H_e\, r, \\
\dot{r} &= \frac{I_X - I_Y}{I_Z} pq + \frac{I_{XZ}}{I_Z}(\dot{p} - qr) + \frac{\bar{q}Sb}{I_Z} C_{n,t} + H_e\, q,
\end{aligned} \tag{5}$$

where I_X, I_Y, and I_Z are the moments of inertia about the body axes, I_{XZ} is the product of inertia, b is the wing span, $\bar{c}$ is the mean aerodynamic chord, $C_{l,t}$, $C_{m,t}$, and $C_{n,t}$ are the total aerodynamic moment coefficients, and H_e is the engine angular momentum.

Aircraft parameters. The mass, inertia, and geometric parameters used in the simulation are summarized in Table 1. The values are listed in SI units for consistency with the rest of this paper.

Table 1 – F-16 mass, inertia, and geometric parameters used in the simulation.

Parameter	Value
Mass, m	9.295×10^3 kg
Moment of inertia, I_X	1.2875×10^4 kg m^2
Moment of inertia, I_Y	7.5674×10^4 kg m^2
Moment of inertia, I_Z	8.5552×10^4 kg m^2
Product of inertia, I_{XZ}	1.331×10^3 kg m^2
Engine angular momentum, H_e	216.9 kg m^2/s
Wing span, b	9.144 m
Wing reference area, S	27.87 m^2
Mean aerodynamic chord, $\bar{c}$	3.45 m
Reference c.g. location	$0.35\bar{c}$

The aerodynamic force and moment coefficients are constructed from tabulated wind-tunnel data and coefficient build-up relations, including static aerodynamic terms, control-surface increments, rate-dependent damping derivatives, and center-of-gravity corrections. The detailed coefficient equations are not repeated here because the aerodynamic model is used only for dataset generation.
The engine is modeled as an afterburning turbofan engine with a first-order thrust response to throttle command. The thrust force is applied along the body X-axis, and the engine gyroscopic effect is included through the angular momentum term H_e in (5).

Closed-loop controller. The recovery maneuver data are generated in closed-loop simulation rather than by open-loop propagation of the aircraft dynamics. Figure 2 shows the closed-loop simulation architecture used for data generation. In the simulator, a cascaded nonlinear dynamic inversion (NDI)-based flight-control structure is used to track the commanded recovery maneuver. The outer loop generates attitude commands, while the inner loop regulates angular-rate responses and allocates the resulting commands to the control surfaces. A separate velocity-hold loop commands the engine throttle to regulate the flight speed. The actuator, engine, sensor, and rigid-body aircraft dynamics are included in the closed-loop simulation.
This closed-loop implementation allows the generated recovery trajectories to include not only the nonlinear aircraft dynamics and aerodynamic characteristics, but also the effects of control response, actuator dynamics, thrust response, and feedback behavior during the recovery maneuver. Since the focus of this paper is recovery trajectory prediction rather than flight-control design, the detailed NDI controller formulation is not presented here.

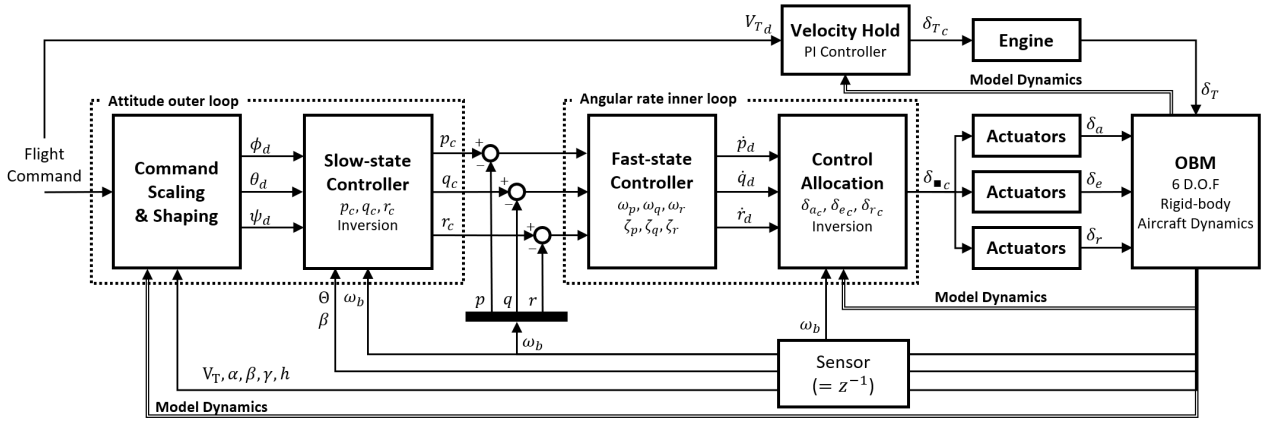


Figure 2 – Closed-loop simulation architecture for recovery maneuver data generation.

3.2 Training Dataset Generation

The training dataset is constructed by simulating the AGCAS recovery maneuver over a wide range of recovery-initiation conditions. During data collection, reference commands are sampled from uniform distributions over prescribed ranges of roll angle, pitch angle, and flight speed. An attitude-hold mode is then used to bring the aircraft to the commanded flight condition prior to recovery initiation. The commanded roll angle ϕ_{cmd}, pitch angle θ_{cmd}, and flight speed V_{cmd} are sampled over the following ranges:

$$-89^\circ \le \phi_{\text{cmd}} \le 89^\circ, \quad -89^\circ \le \theta_{\text{cmd}} \le 10^\circ, \quad 150\ \text{m/s} \le V_{\text{cmd}} \le 350\ \text{m/s}. \tag{6}$$

For each commanded condition, the recovery maneuver is simulated, and the actual recovery-initiation state variables, V_0, γ_0, and ϕ_0, are stored together with the resulting altitude-change profile $\Delta h(t)$. Although the data are collected by commanding specific attitude and speed setpoints, the learning model uses the actual flight state at the recovery initiation point rather than the commanded state. The flight-path angle γ_0 is used in place of the pitch angle because it directly characterizes the initial vertical motion of the aircraft during recovery initiation.

The training dataset consists of 2,000 recovery trajectories, and a separate held-out test set of 157 trajectories is reserved for performance evaluation. For each trajectory, the altitude-change profile $\Delta h(t)$ is recorded over a horizon of $T = 20$ s with a sampling interval of $\Delta t = 0.04$ s, yielding $N_t = 501$ samples per trajectory. Figure 3 shows the distribution of the reference commands and the actual recovery-initiation states used for dataset generation.

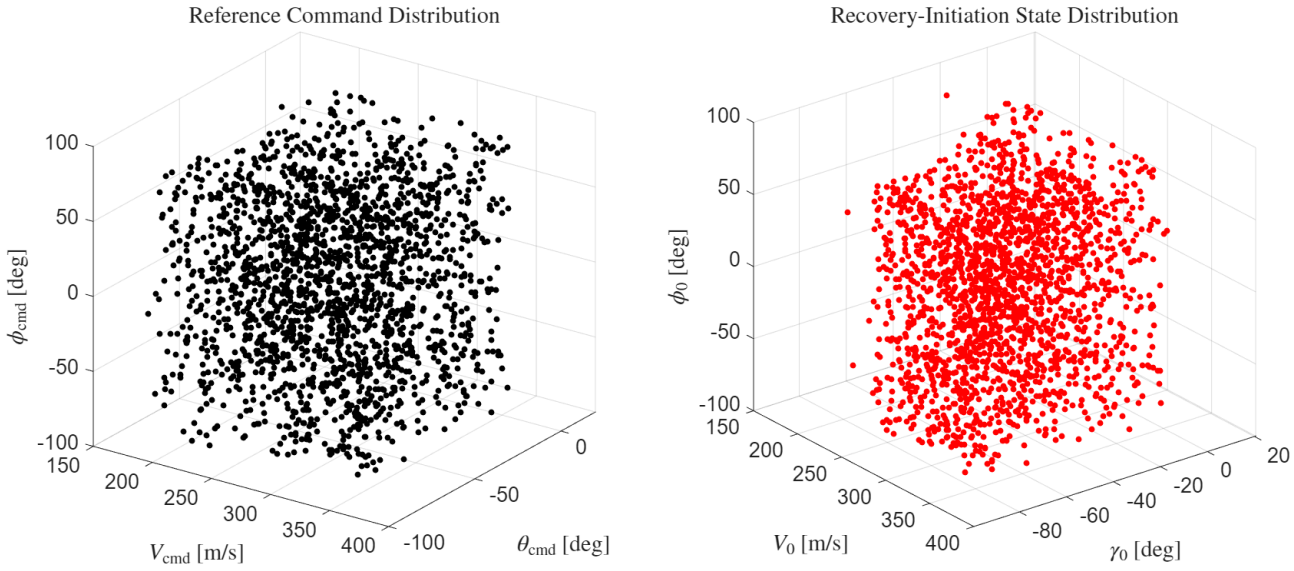


Figure 3 – Distribution of reference commands and actual recovery-initiation states.

3.3 Polynomial Representation of Recovery Trajectories

Each simulated recovery trajectory consists of $N_t = 501$ altitude-change samples. Directly predicting these time-series profiles would require a high-dimensional regression target. Therefore, to obtain a compact and differentiable representation, the altitude-change profile is approximated by an M-th order polynomial in time:

$$\Delta h(t) \approx \Delta h_{\text{fit}}(t;\mathbf{c}) = \sum_{j=1}^{M} c_j t^j = c_1 t + c_2 t^2 + \cdots + c_M t^M, \qquad 0 \le t \le T. \tag{7}$$

Here, $\Delta h_{\text{fit}}(t;\mathbf{c})$ denotes the polynomial approximation of the simulated altitude-change profile for a given coefficient vector $\mathbf{c}$. The constant term is excluded so that $\Delta h_{\text{fit}}(0) = 0$, consistent with $\Delta h(0) = 0$ by definition. Through this representation, the learning problem is transformed from high-dimensional time-series regression into the prediction of M scalar polynomial coefficients, which can be performed efficiently at inference time.

For each simulated trajectory, the coefficient vector $\mathbf{c} = [c_1, c_2, \ldots, c_M]^T$ is obtained by solving the least-squares fitting problem:

$$\mathbf{c} = \arg\min_{\alpha} \sum_{i=1}^{N_t} \left(\Delta h(t_i) - \Delta h_{\text{fit}}(t_i;\alpha)\right)^2 = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\Delta\mathbf{h}, \tag{8}$$

where $\alpha = [\alpha_1, \alpha_2, \ldots, \alpha_M]^T$ denotes the optimization variable for the polynomial coefficients, $\mathbf{X} \in \mathbb{R}^{N_t \times M}$ is the Vandermonde matrix with elements $X_{ij} = t_i^j$, and $\Delta\mathbf{h} = [\Delta h(t_1), \ldots, \Delta h(t_{N_t})]^T$.

The polynomial order M determines the trade-off between representation accuracy and model complexity. A low-order polynomial may fail to capture the recovery altitude-change profile, whereas an unnecessarily high order increases the number of regression targets with limited accuracy improvement. To select an appropriate order, the training trajectories are fitted using polynomial orders from $M = 1$ to $M = 8$, and the resulting fitting RMSE is evaluated.

As shown in Fig. 4, both the mean and median fitting RMSE decrease rapidly as the polynomial order increases. At $M = 6$, the mean and median RMSE are reduced to 2.02 m and 1.85 m, respectively. Increasing the polynomial order beyond $M = 6$ further reduces the fitting error, but the improvement is relatively small compared with the increase in the number of regression targets. Therefore, $M = 6$ is selected in this study to provide an accurate yet compact representation of the recovery altitude-change profiles.

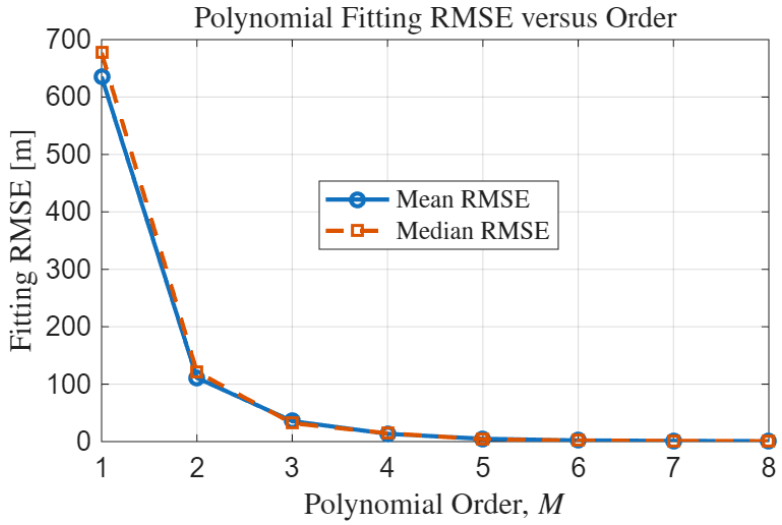


Figure 4 – Polynomial fitting RMSE for different polynomial orders.

Figure 5 shows representative altitude-change samples and their sixth-order polynomial approximations for different recovery-initiation states (V_0, γ_0, ϕ_0). These examples show that the sixth-order polynomial provides a compact representation for diverse altitude-change profiles.

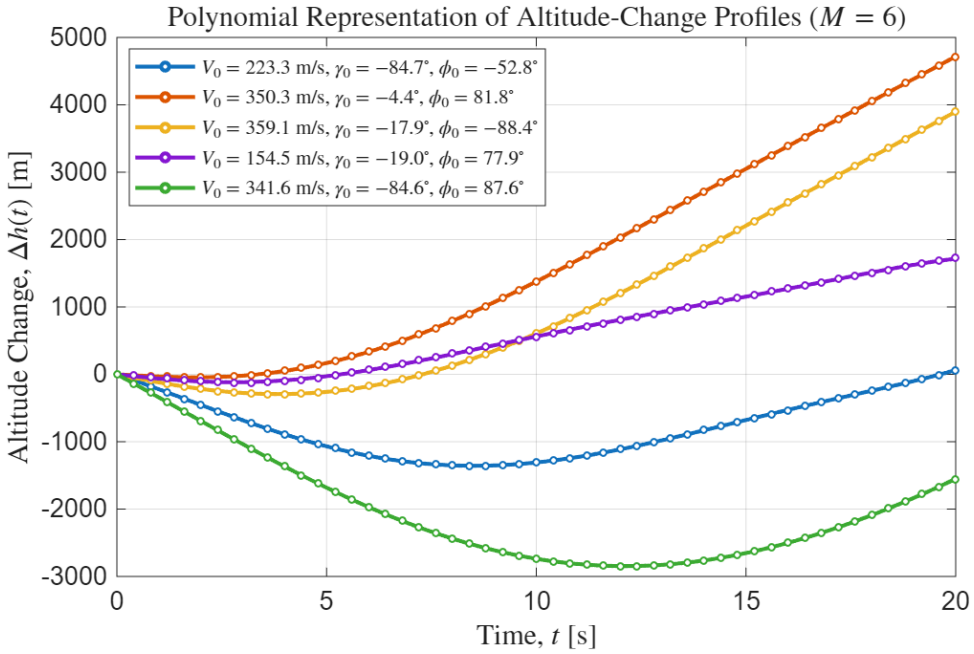


Figure 5 – Representative altitude-change profiles. Markers represent simulated altitude-change samples $\Delta h(t_i)$, while solid lines represent the fitted polynomial approximations $\Delta h_{\text{fit}}(t;\mathbf{c})$.

3.4 Baseline Coefficient Learning Model

Using the polynomial representation introduced in Section 3.3, the trajectory prediction problem in (3) is reformulated as a coefficient prediction problem. Instead of directly predicting the full altitude-change profile, the model learns the mapping from the recovery-initiation variables to the polynomial coefficient vector:

$$(V_0, \gamma_0, \phi_0) \mapsto (c_1, c_2, \ldots, c_M). \tag{9}$$

A global ridge regression model is constructed for this mapping.

1) Input normalization. The input variables are first normalized using the mean and standard deviation computed from the training dataset:

$$\tilde{V}_0 = \frac{V_0 - \mu_V}{\sigma_V}, \qquad \tilde{\gamma}_0 = \frac{\gamma_0 - \mu_\gamma}{\sigma_\gamma}, \qquad \tilde{\phi}_0 = \frac{\phi_0 - \mu_\phi}{\sigma_\phi}, \tag{10}$$

where μ and σ denote the mean and standard deviation of each input variable in the training dataset. The tilde notation is used to indicate normalized variables. Normalization accounts for differences in physical units and numerical scales among the input variables and improves the conditioning of the subsequent polynomial regression.

2) Polynomial feature expansion. To capture nonlinear interactions between the input variables, the normalized variables are augmented with monomial terms up to total degree D. The resulting feature vector is written as

$$\mathbf{f}\left(\tilde{V}_0, \tilde{\gamma}_0, \tilde{\phi}_0\right) = \left[1, \tilde{V}_0, \tilde{\gamma}_0, \tilde{\phi}_0, \tilde{V}_0^2, \tilde{V}_0\tilde{\gamma}_0, \tilde{V}_0\tilde{\phi}_0, \tilde{\gamma}_0^2, \tilde{\gamma}_0\tilde{\phi}_0, \tilde{\phi}_0^2, \ldots\right], \tag{11}$$

where the ellipsis denotes the remaining monomial terms up to total degree D. Here, the total degree refers to the sum of the exponents in each monomial term. In this study, the maximum total degree is set to $D = 6$. This feature expansion allows the regression model to approximate nonlinear mappings from the recovery-initiation variables to the polynomial coefficients without requiring a more complex model architecture.

3) Ridge regression. Let $\mathbf{F} \in \mathbb{R}^{N_{\text{train}} \times N_f}$ be the feature matrix, where each row corresponds to the feature vector of one training sample, and let $\mathbf{C} \in \mathbb{R}^{N_{\text{train}} \times M}$ be the target coefficient matrix. The regression parameter matrix $\Gamma \in \mathbb{R}^{N_f \times M}$ is identified from the training data using ridge regression:

$$\Gamma = \left(\mathbf{F}^T\mathbf{F} + \lambda \mathbf{I}\right)^{-1} \mathbf{F}^T \mathbf{C}, \tag{12}$$

where $\lambda > 0$ is the regularization parameter, $\mathbf{I}$ is the identity matrix, and N_f is the dimension of the feature vector. The ℓ_2 regularization term improves numerical conditioning and helps prevent overfitting when the feature dimension is large relative to the number of training samples.

4) Inference. For a new recovery-initiation condition $(V_0^*, \gamma_0^*, \phi_0^*)$, the input variables are normalized using the training-data statistics, yielding $(\tilde{V}_0^*, \tilde{\gamma}_0^*, \tilde{\phi}_0^*)$. The predicted polynomial coefficient vector is then computed as

$$\hat{\mathbf{c}}^T = \mathbf{f}\left(\tilde{V}_0^*, \tilde{\gamma}_0^*, \tilde{\phi}_0^*\right) \Gamma. \tag{13}$$

The predicted recovery altitude-change profile is reconstructed as

$$\Delta\hat{h}\left(t; V_0^*, \gamma_0^*, \phi_0^*\right) = \sum_{j=1}^{M} \hat{c}_j\left(V_0^*, \gamma_0^*, \phi_0^*\right) t^j. \tag{14}$$

This expression can be evaluated at any desired time point within the recovery horizon with negligible computational cost, enabling efficient online recovery trajectory evaluation without requiring real-time dynamic propagation.

3.5 Inspection of the Baseline Coefficient Mapping

After constructing the baseline polynomial regression model, the learned coefficient mapping is inspected to examine how each polynomial coefficient depends on the recovery-initiation variables. Since the baseline model treats all coefficients as free learning targets, the learned surfaces provide qualitative insight into the dependencies identified from the training data.

Figure 6 shows the learned coefficient surfaces of $c_1, \ldots, c_6$ with respect to the initial flight-path angle γ_0 and roll angle ϕ_0 at a fixed flight speed of $V_0 = 300\ \text{m/s}$. The first coefficient c_1 varies predominantly along the flight-path-angle direction, while remaining nearly invariant with respect to the roll angle. Since c_1 is the coefficient of the linear time term in $\Delta h_{\text{fit}}(t; \mathbf{c})$, it determines the initial slope of the fitted altitude-change profile. This observation suggests that the learned c_1 is more closely associated with the initial vertical-motion tendency than with roll-dependent transient recovery dynamics. This interpretation is also consistent with the role of the roll angle in the recovery sequence: the initial roll angle mainly affects the subsequent roll-recovery duration and the development of the later pull-up response, rather than the instantaneous altitude-change rate at recovery initiation. The kinematic relation used to constrain c_1 is derived in the next section.

In contrast, the remaining coefficients $c_2, \ldots, c_6$, which shape the subsequent transient and later portions of the recovery altitude-change response, show more visible variations along the roll-angle direction. However, their overall trends are approximately symmetric about $\phi_0 = 0$. In other words, positive and negative roll angles with similar magnitudes tend to produce similar coefficient values. This trend is physically reasonable because the recovery maneuver first rolls the aircraft toward a wings-level attitude before the pull-up maneuver. Under a laterally symmetric aircraft model and a symmetric recovery control law, the altitude-change response is expected to be similar for left and right initial roll angles of the same magnitude.

These observations reveal two useful structures in the learned coefficient mapping. The first coefficient c_1 is nearly invariant with respect to the roll angle, while the higher-order coefficients exhibit approximately symmetric trends with respect to the sign of the initial roll angle. Although the baseline model can learn these trends from data, incorporating such physically motivated structures into the model can reduce unnecessary learning degrees of freedom. The following section introduces two physics-informed constraints based on these observations.

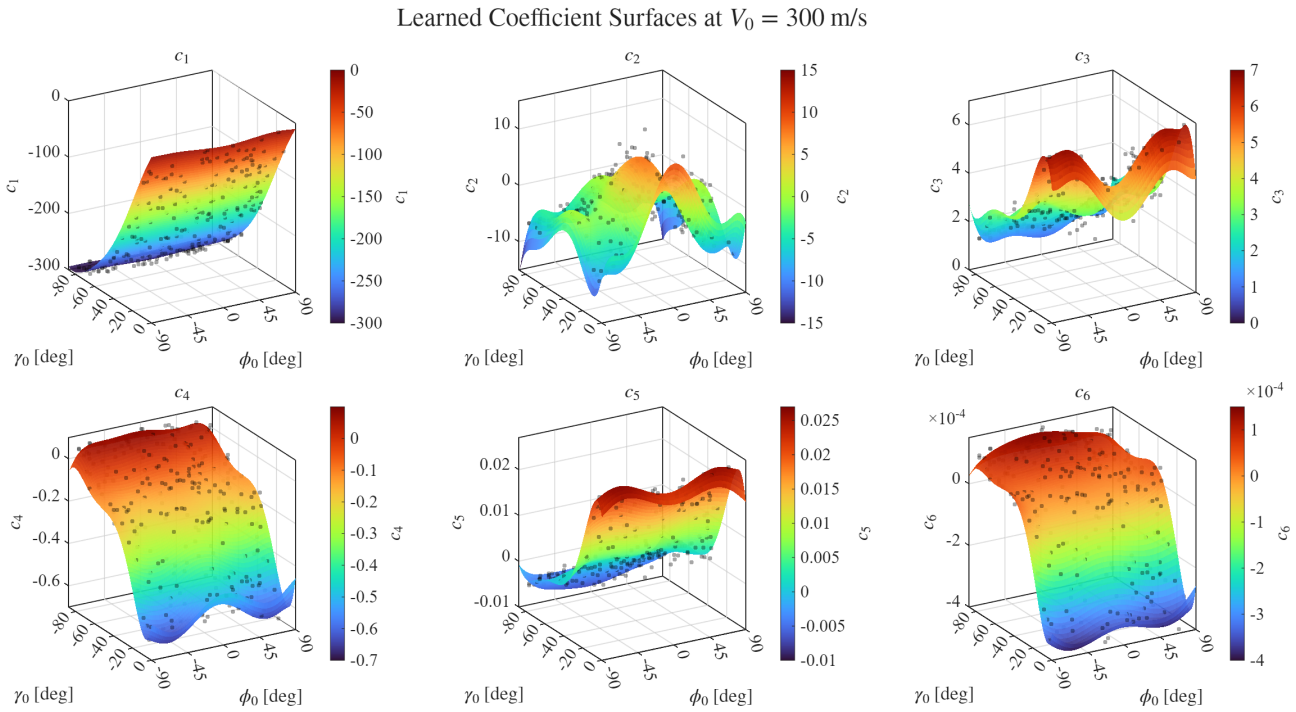


Figure 6 – Learned coefficient surfaces of the unconstrained baseline regression model at $V_0 = 300$ m/s. The surfaces are plotted with respect to the initial flight-path angle γ_0 and roll angle ϕ_0.

4. Physics-Informed Learning

The baseline coefficient mapping inspected in Section 3.5 reveals physically interpretable structures in the learned coefficients. In particular, the first polynomial coefficient is related to the initial slope of the fitted altitude-change profile, while the higher-order coefficients exhibit approximate symmetry with respect to the sign of the initial roll angle. This section embeds these structures into the model through two physics-informed constraints: an initial altitude-rate constraint and a roll-symmetry constraint.

4.1 Initial Altitude-Rate Constraint

From flight kinematics, the initial altitude rate of an aircraft flying at speed V_0 with flight-path angle γ_0 is given by

$$\dot{h}(0) = V_0 \sin \gamma_0. \tag{15}$$

Since $\Delta h(t) = h(t) - h(0)$, the initial slope of the altitude-change profile is identical to the initial altitude rate. In the polynomial representation, this initial slope corresponds to the first polynomial coefficient. For an individual fitted trajectory, differentiating the polynomial approximation $\Delta h_{\text{fit}}(t;\mathbf{c})$ in (7) with respect to time and evaluating it at $t = 0$ gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\Delta h_{\text{fit}}(t;\mathbf{c})\bigg|_{t=0} = c_1. \tag{16}$$

Therefore, by matching the initial slope of the fitted altitude-change profile to the kinematic altitude rate, the physically constrained first coefficient is given by

$$c_1 = V_0 \sin \gamma_0. \tag{17}$$

The same relation is imposed at the prediction stage. In the physics-informed model, the predicted first coefficient is not learned by the regression model, but is assigned analytically as

$$\hat{c}_1(V_0, \gamma_0) = V_0 \sin \gamma_0. \tag{18}$$

As a result, the regression model predicts only the remaining coefficients $\hat{c}_2, \ldots, \hat{c}_M$. The altitude-rate-constrained predicted trajectory is then given by

$$\Delta\hat{h}(t; V_0, \gamma_0, \phi_0) = V_0 \sin\gamma_0 t + \sum_{j=2}^{M} \hat{c}_j(V_0, \gamma_0, \phi_0) t^j. \tag{19}$$

This constraint enforces the correct initial altitude-change rate by construction and prevents the model from learning an analytically known kinematic relationship.

4.2 Roll-Symmetry Constraint

During the AGCAS recovery maneuver, the aircraft first rolls toward a wings-level attitude before performing the pull-up maneuver. For a laterally symmetric aircraft under a symmetric recovery control law, changing the sign of the initial roll angle reverses the roll-recovery direction, but the vertical altitude-change profile is expected to remain approximately unchanged for the same roll-angle magnitude. This approximate roll-symmetry property can be expressed as

$$\Delta h(t; V_0, \gamma_0, \phi_0) \approx \Delta h(t; V_0, \gamma_0, -\phi_0). \tag{20}$$

To embed this property into the regression model, the signed roll angle ϕ_0 is replaced by its magnitude $|\phi_0|$ in the polynomial feature expansion. The regression model is then trained to predict the remaining coefficients $c_2, \ldots, c_M$ using V_0, γ_0, and $|\phi_0|$ as input variables.
Combining the initial altitude-rate and roll-symmetry constraints, the complete physics-informed recovery trajectory model is written as

$$\Delta\hat{h}(t; V_0, \gamma_0, \phi_0) = V_0 \sin\gamma_0 t + \sum_{j=2}^{M} \hat{c}_j(V_0, \gamma_0, |\phi_0|) t^j. \tag{21}$$

This formulation enforces the correct initial altitude-change rate and incorporates roll-sign symmetry by construction, thereby improving the consistency of the learned recovery trajectory model.

5. Simulation Results and AGCAS Application

The proposed method is evaluated in two parts. First, recovery trajectory prediction accuracy is assessed on the held-out test dataset by comparing a simplified kinematics-based model, the unconstrained baseline learning model, and the proposed physics-informed learning model. Second, the learned trajectory model is applied to an AGCAS recovery simulation, in which delayed recovery trajectories are compared with a buffered forward terrain profile to determine automatic recovery initiation. Together, these evaluations assess both the prediction accuracy of the proposed model and its applicability to AGCAS recovery activation.

5.1 Recovery Trajectory Prediction Accuracy

The recovery trajectory prediction accuracy is evaluated on the held-out test dataset consisting of 157 trajectories. Three prediction models are compared: a simplified kinematics-based model, the unconstrained baseline learning model, and the proposed physics-informed learning model. The simulated 6-DOF recovery trajectories are used as the reference data for evaluation.
The kinematics-based model serves as a simplified reference model. It predicts the recovery altitude-change profile using point-mass kinematics, assuming constant flight speed and ideal tracking of the commanded recovery sequence described in Section 2.1. This model captures the nominal recovery trend, but does not explicitly account for closed-loop response delays, transient aircraft dynamics, aerodynamic effects, or velocity variations.
For the learning-based models, the polynomial order is set to $M = 6$, and the maximum total degree of the polynomial feature basis is set to $D = 6$. The baseline learning model predicts all six polynomial coefficients, whereas the physics-informed learning model fixes $\hat{c}_1$ using the prediction-stage altitude-rate constraint in (18) and predicts only $\hat{c}_2, \ldots, \hat{c}_6$ using $|\phi_0|$ as the roll-related input.

The quantitative prediction accuracy is evaluated using the RMSE between the simulated altitude-change profile and the predicted profile:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N_t}\sum_{i=1}^{N_t}\left(\Delta h(t_i) - \Delta\hat{h}(t_i)\right)^2}. \tag{22}$$

Figure 7 shows a representative trajectory prediction result for a held-out test condition. The left panel compares the simulated altitude-change profile with the predictions obtained from the kinematics-based, baseline learning, and physics-informed learning models. The right panel shows the corresponding pointwise prediction errors. The kinematics-based model captures the overall recovery trend, but exhibits noticeable deviation because of the simplified assumptions on aircraft response and velocity variation. The baseline learning model reduces this discrepancy by learning recovery characteristics from the 6-DOF simulation data.

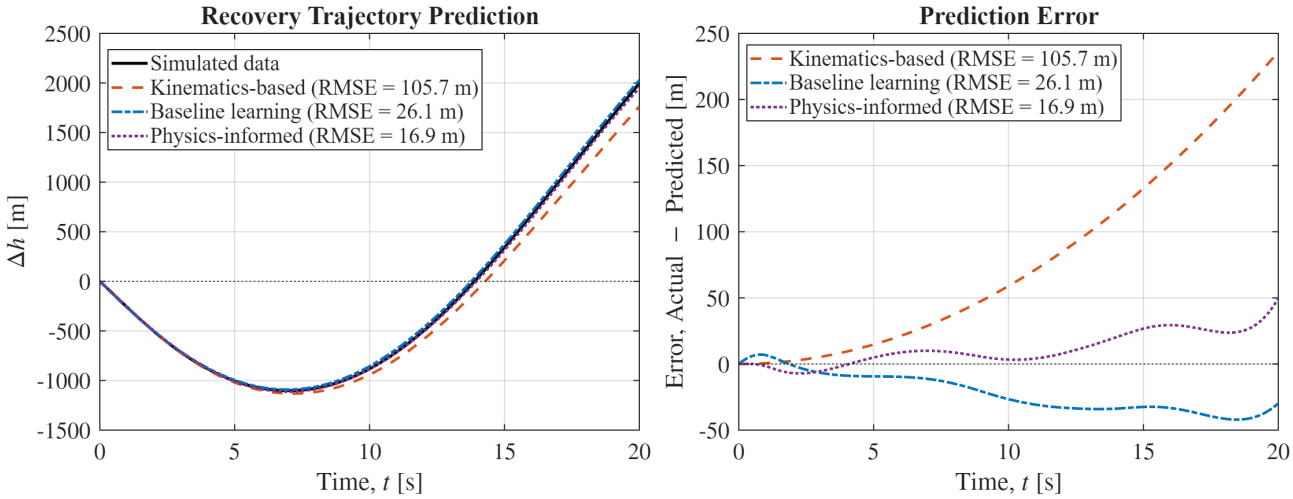


Figure 7 – Example recovery trajectory prediction on the held-out test dataset with $V_0 = 360$ m/s, $\gamma_0 = -45.5°$, and $\phi_0 = 33.2°$ at recovery initiation.

A notable difference between the baseline and physics-informed models appears in the early part of the prediction horizon. Since the baseline learning model is trained only to minimize the overall trajectory prediction error, the initial slope of the predicted altitude-change profile is not explicitly constrained. Consequently, its prediction error grows immediately after recovery initiation when the learned initial slope differs from the kinematic altitude rate. In contrast, the physics-informed learning model enforces the initial altitude-rate condition by construction, so the prediction error remains close to zero in the initial segment of the trajectory. This behavior shows that the initial-slope constraint improves the consistency of the predicted recovery trajectory near recovery initiation.

Table 2 summarizes the prediction performance of the three models on the held-out test dataset. The kinematics-based model yields a median RMSE of 105.70 m and a mean RMSE of 118.43 m, reflecting the limited accuracy of the simplified point-mass prediction. The baseline learning model reduces the median and mean RMSE to 21.63 m and 26.39 m, respectively, by learning recovery characteristics from the 6-DOF simulation data. With the physics-informed constraints, the median RMSE further decreases to 11.58 m, and the mean RMSE decreases to 15.85 m. This corresponds to an approximately 89.0% reduction in the median RMSE compared with the kinematics-based model and an approximately 46.4% reduction compared with the unconstrained baseline learning model.

Table 2 – Recovery trajectory prediction RMSE comparison on the held-out test dataset ($N = 157$).

Trajectory prediction model	Median RMSE [m]	Mean RMSE [m]
Kinematics-based model	105.701	118.432
Baseline learning model	21.625	26.387
Physics-informed learning model	11.582	15.845

Overall, the learning-based models provide substantially more accurate recovery trajectory predictions than the simplified kinematics-based model by capturing recovery characteristics contained in the 6-DOF simulation data. The physics-informed learning model further improves the prediction accuracy and consistency by reducing unnecessary learning freedom. In particular, the initial altitude-rate constraint prevents the model from learning a kinematic relationship that is already known, while the roll-symmetry constraint avoids learning separate dependencies for positive and negative initial roll angles.

5.2 AGCAS Recovery Maneuver Simulation

To demonstrate the applicability of the learned recovery trajectory model in an AGCAS framework, an automatic recovery simulation is conducted. At each simulation time step, the learned model predicts the recovery altitude profile that would result if automatic recovery were initiated after a prescribed delay. During this delay, the aircraft is assumed to continue along its current flight condition before the recovery command is applied. The predicted recovery trajectory is then compared with the forward terrain buffer, which represents the terrain profile augmented by a safety margin. If the delayed recovery trajectory intersects the terrain buffer, the corresponding delay is considered unsafe. This procedure allows the learned trajectory model to be used directly in the collision assessment process without performing online 6-DOF trajectory propagation.

In this simulation, the prescribed activation delay is set to 1.5 s. At each simulation time step, the model evaluates whether the aircraft would maintain clearance from the terrain buffer if recovery were delayed by 1.5 s. Automatic recovery is initiated when the recovery trajectory predicted with this delay no longer remains clear of the terrain buffer. The delay threshold is selected as a representative value for the present simulation and is not optimized in this study.

Figure 8 shows the predicted recovery trajectories generated at a representative prediction time. The trajectories correspond to different assumed recovery delays and are overlaid with the forward terrain buffer. These trajectories illustrate how the available terrain clearance changes as the assumed recovery delay is varied. As the assumed recovery delay increases, the predicted recovery trajectory shifts downward relative to the terrain buffer, indicating reduced terrain clearance. In the tested scenario, automatic recovery is triggered when the trajectory associated with the prescribed delay threshold intersects the terrain buffer.

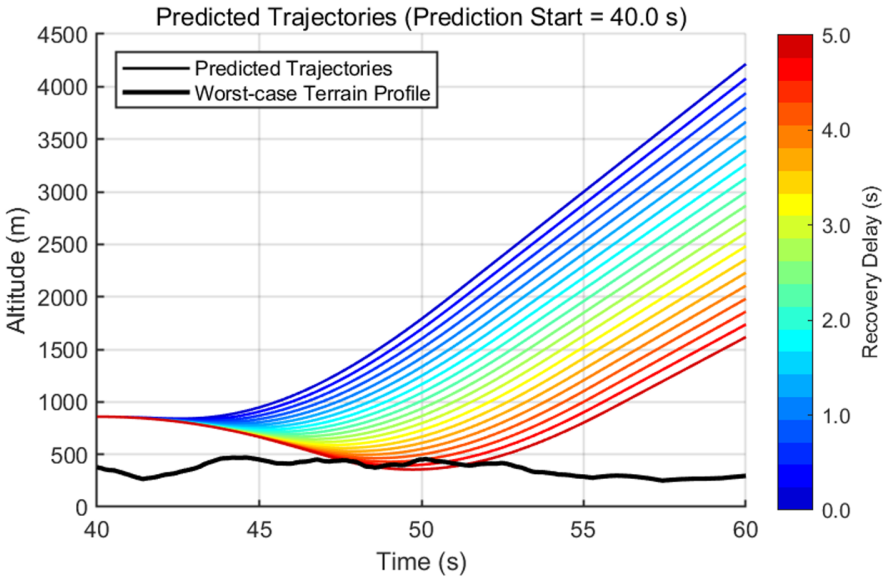


Figure 8 – Predicted recovery trajectories at a representative prediction time for different assumed recovery delays.

Figure 9 presents the time histories of key aircraft state variables during the AGCAS recovery maneuver. The responses follow the recovery sequence described in Section 2.1. Before recovery activation, the aircraft continues along its current flight condition, and the altitude decreases as the

aircraft approaches the terrain. After activation, the roll angle rapidly converges toward a wings-level attitude, indicating the roll-recovery phase. The pitch angle then increases during the pull-up maneuver, while the vertical velocity and acceleration responses indicate the build-up of the commanded load factor. As the pull-up develops, the altitude stops decreasing and subsequently increases, showing that the aircraft successfully transitions from descent to recovery. These responses are consistent with the prescribed recovery sequence consisting of recovery delay, roll recovery, load-factor build-up, maximum-load-factor pull-up, and flight-path-angle hold.

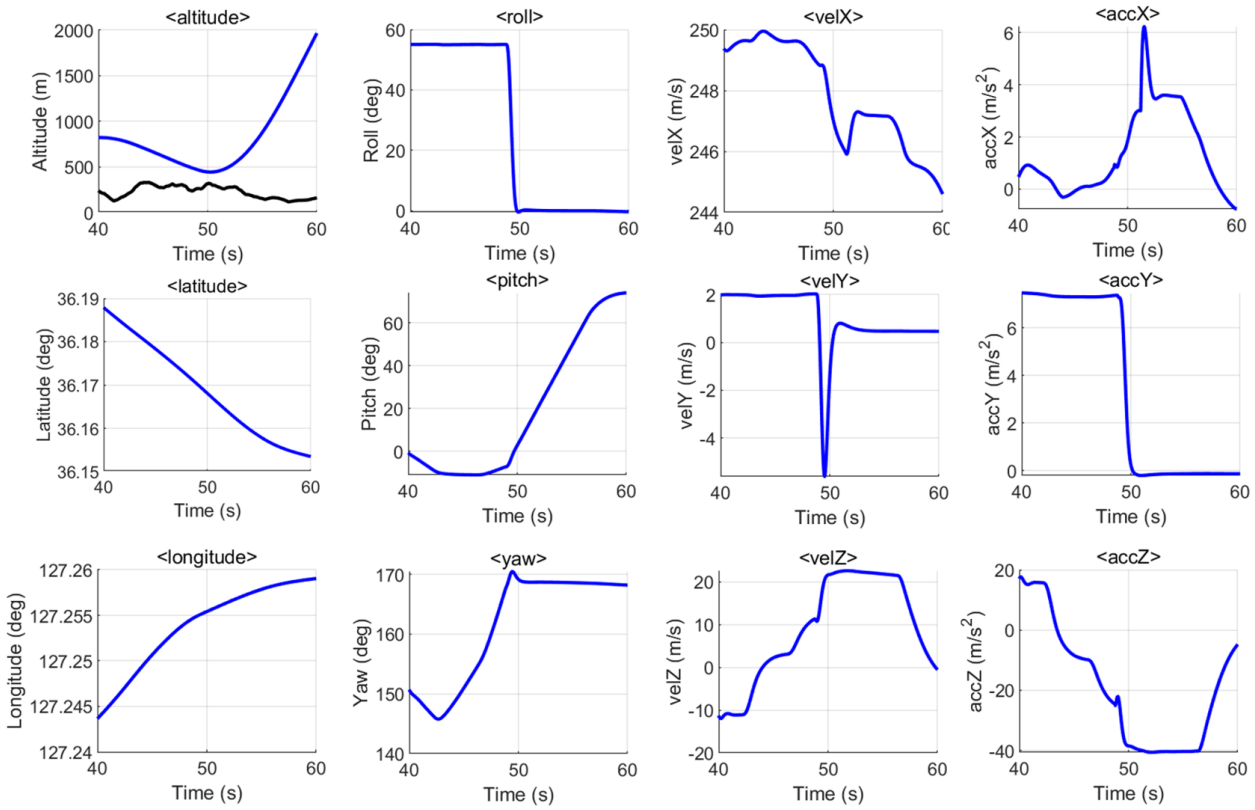


Figure 9 – Aircraft state responses during the AGCAS recovery maneuver with a prescribed activation delay of 1.5 s.

These results demonstrate that the proposed learning-based recovery trajectory model can support terrain-buffer-based collision assessment and automatic recovery activation in the tested AGCAS scenario. The simulation results, covering delayed recovery trajectory prediction, terrain-buffer-based collision assessment, and automatic recovery execution, show the applicability of the proposed framework to ground collision avoidance scenarios.

6. Conclusions

This paper proposed a physics-informed supervised learning framework for recovery trajectory prediction in Automatic Ground Collision Avoidance Systems. Recovery maneuver data were generated offline using an F-16 6-DOF flight simulator with an aerodynamic database over a wide range of recovery-initiation conditions. The recovery altitude-change profiles were represented using compact sixth-order time-polynomial models, and a global ridge regression model with a sixth-degree polynomial feature basis was trained to map the recovery-initiation variables, consisting of flight speed, flight-path angle, and roll angle, to the corresponding polynomial coefficients.

Two physics-informed constraints were incorporated into the learning model. First, the initial altitude-rate constraint analytically fixed the first polynomial coefficient using the flight-kinematic relation $\hat{c}_1 = V_0 \sin\gamma_0$, eliminating the need to learn a kinematically determined relationship from data. Second, the roll-symmetry constraint used the roll-angle magnitude $|\phi_0|$ as the roll-related input, reflecting the approximate symmetry of the recovery altitude response with respect to the sign of the initial roll

angle. These constraints reduce unnecessary learning freedom and improve the consistency of the learned recovery trajectory model.

The proposed model was evaluated on a held-out test dataset by comparing three prediction models: a simplified kinematics-based model, an unconstrained baseline learning model, and the physics-informed learning model. The kinematics-based model yielded a median RMSE of $105.70\ \mathrm{m}$, whereas the baseline learning model reduced the median RMSE to $21.63\ \mathrm{m}$ by learning recovery characteristics from the 6-DOF simulation data. With the physics-informed constraints, the median RMSE was further reduced to $11.58\ \mathrm{m}$. This corresponds to an approximately 89.0% reduction in median RMSE compared with the kinematics-based model and an approximately 46.4% reduction compared with the unconstrained baseline learning model.

The learned recovery trajectory model was further applied to an AGCAS automatic recovery simulation. In the simulation, delayed recovery trajectories predicted by the learned model were compared with a forward terrain buffer to assess terrain clearance and determine automatic recovery initiation. A prescribed activation delay of $1.5\ \mathrm{s}$ was used as a representative simulation setting, and automatic recovery was initiated when the recovery trajectory predicted with this delay no longer remained clear of the terrain buffer. The resulting aircraft responses followed the intended recovery sequence, including recovery delay, roll recovery, load-factor build-up, maximum-load-factor pull-up, and flight-path-angle hold. These results demonstrate that the proposed physics-informed learning framework can provide an accurate and compact recovery trajectory representation and can support terrain-buffer-based collision assessment and automatic recovery activation in AGCAS applications.

7. Acknowledgement

This research was supported by APPIA Engineering Co., Ltd.

8. Contact Author Email Address

If you have any questions, please send an email to: rkdxk7085@kaist.ac.kr or hcbang@kaist.ac.kr.

9. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] D. E. Swihart, A. F. Barfield, E. M. Griffin, R. C. Lehmann, S. C. Whitcomb, B. Flynn, M. A. Skoog, and K. E. Prosser. Automatic ground collision avoidance system design, integration, and flight test. *IEEE Aerospace and Electronic Systems Magazine*, 26(5):4–17, 2011.

[2] D. E. Swihart and E. M. Griffin. Automatic ground collision avoidance technology for fighter aircraft. In *Proceedings of the International Conference on Integrated Navigation Systems*, pages 1–12, Saint Petersburg, Russia, 2010.

[3] R. Huffman, M. Skoog, L. Piñeiro, and M. Fern. Application of ground collision avoidance system nuisance criteria. In *Proceedings of the 21st Congress of the International Council of the Aeronautical Sciences*, pages 1–9, Melbourne, Australia, 1998. ICAS-98-1.8.3.

[4] Angela W. Suplisson. *Optimal Recovery Trajectories for Automatic Ground Collision Avoidance Systems (Auto GCAS)*. Ph.d. dissertation, Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, USA, 2015.

[5] Shiyi Yuan, Qifu Li, Bei Lu, Xingjie Niu, Yishu Liu, and Wei Gao. Trajectory prediction for fighter aircraft ground collision avoidance based on the model predictive control technique. *Aerospace Systems*, 8:61–70, 2025.

[6] Luat T. Nguyen, Marilyn E. Ogburn, William P. Gilbert, Kemper S. Kibler, Phillip W. Brown, and Perry L. Deal. Simulator study of stall/post-stall characteristics of a fighter airplane with relaxed longitudinal static stability. Technical Paper NASA-TP-1538, NASA Langley Research Center, December 1979.