DOT/FAA/TC-24/4

Federal Aviation Administration
William J. Hughes Technical Center
Aviation Research Division
Atlantic City International Airport
New Jersey 08405

Analysis of the Ground Collision Avoidance System Within NASA's Expandable Vehicle Autonomy Architecture

March 2025

Final report

U.S. Department of Transportation
Federal Aviation Administration

NOTICE

This document is disseminated under the sponsorship of the U.S. Department of Transportation in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof. The U.S. Government does not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to the objective of this report. The findings and conclusions in this report are those of the author(s) and do not necessarily represent the views of the funding agency. This document does not constitute FAA policy. Consult the FAA sponsoring organization listed on the Technical Documentation page as to its use.

This report is available at the Federal Aviation Administration William J. Hughes Technical Center's Full-Text Technical Reports page: actlibrary.tc.faa.gov in Adobe Acrobat portable document format (PDF).

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
DOT/FAA/TC-24/4		

4. Title and Subtitle	5. Report Date
Analysis of the Ground Collision Avoidance System Within NASA's Expandable Vehicle Autonomy Architecture	March 2024
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Loyd R. Hook, Jude M. Urban, Alan M. Hubbard, Patrick Maley	

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
The University of Tulsa 800 S. Tucker Drive Tulsa, OK 74104-9700	11. Contract or Grant No.

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
FAA Central Regional Office 901 Locust St Kansas City, MO 64106-2325.	14. Sponsoring Agency Code

15. Supplementary Notes

The FAA William J. Hughes Technical Center Aviation Research Division COR was Robert McGuire. FAA sponsors are David Sizoo and Ross Schaller of the Policy and Standards Section, AIR-714, in the Compliance and Airworthiness Division.

16. Abstract

This document provides comprehensive methodology and results in the evaluation of a Ground Collision Avoidance System (GCAS) tuned for a general aviation airplane. The GCAS under investigation was designed by the Armstrong flight research center at the National Aeronautics and Space Administration (NASA). This research focused on two characteristics of the GCAS: protection and nuisance. 'Protection' refers to the GCAS' ability to detect and prevent an airplane from a ground collision. 'Nuisance' refers to any inappropriate intervention or inhibition from the GCAS inflicted on the pilot.

To evaluate the system protection, a recovery autopilot was built and coupled alongside the GCAS under investigation. The system-wide analysis employed Monte Carlo techniques, simulating over 60,000 unique conditions. Our findings showed that GCAS' protection capabilities are exceptional, with an overall protection rate of 98.5%. Nuisance evaluation utilized over 3,000 hours of ADS-B flight data from more than 8,000 flights to ascertain whether a GCAS maneuver would be commanded. Our findings show that the GCAS version under test was quite nuisance-prone, especially around terminal areas with over 30% of flights containing unrequired GCAS activations. Overall, GCAS has very strong protection abilities but further tuning to alleviate prominent nuisance of this version is required.

Throughout the duration of this project, the authors wrote four different academic conference papers pertaining to these two categories. This document aggregates these four papers into chapters and provides a high-level summarization of the project.

17. Key Words	18. Distribution Statement
General Aviation -- Ground Collision Avoidance System (GCAS) General Aviation -- Controlled Flight into Terrain (CFIT)	This document is available to the U.S. public through the National Technical Information Service (NTIS), Springfield, Virginia 22161. This document is also available from the Federal Aviation Administration William J. Hughes Technical Center at actlibrary.tc.faa.gov.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	138	

Acknowledgements

The authors extend their gratitude to David Sizoo and Robert McGuire from the FAA for their exemplary leadership and unwavering commitment to general aviation safety. Additionally, the valuable insights and constructive feedback from Richard Adler, Scott Fohrman, and Dan Dellmyer at the FAA were instrumental in enriching this evaluation. We also appreciate the expertise provided by Mark Skoog and Ethan Williams from NASA on Ground Collision Avoidance Systems (GCAS) and Trajectory Prediction Algorithm (TPA) tuning. Furthermore, we are thankful to Raven LeClaire and Rachael Winiecki from Blackbird Aviation for their engaging technical discussions and essential data contributions, which greatly enhanced the quality of this report.

Contents

Figures

Tables

Acronyms

Acronym	Definition
ADS-B	Automatic Dependent Surveillance-Broadcast
CDTM	Compressed Digital Terrain Map
CFIT	Controlled Flight Into Terrain
EVAA	Expandable Variable Autonomy Architecture
EGPWS	Enhanced Ground Proximity Warning Systems
FAA	Federal Aviation Administration
FDM	Flight Dynamics Model
GA	General Aviation
GCAS	Ground Collision Avoidance System
GPS	Global Positioning System
ICD	Initial Condition Domain
KCAS	Knots Calibrated Airspeed
KEDW	Edwards Air Force Base
KIAS	Knots Indicated Airspeed
KML	Keyhole Markup Language
KVNY	Van Nuys Airport
LAO	Low Altitude Operations
LOC-I	Loss Of Control - Inflight
MAT	Minimum Approach to Terrain
NASA	National Aeronautics and Space Administration
NTSB	National Transportation Safety Board
PCC	Pearson Correlation Coefficients
PID	Proportional-Integral- Derivative
RTA	Run-Time Assurance
SCF-NP	System Component Failure Not Related to the Powerplant
TPA	Trajectory Prediction Algorithm
UIMC	Unintended Flight into Instrument Meteorological Conditions
USGS	United States Geological Survey
UTM	Universal Transverse Mercator

Executive summary

This document provides comprehensive methodology and results in the evaluation of a Ground Collision Avoidance System (GCAS) tuned for a general aviation airplane. The GCAS under investigation was designed by the Armstrong Flight Research Center at the National Aeronautics and Space Administration (NASA). More specifically, the system is a multi-monitor run-time assurance system designed to enable aircraft autonomy, called Expandable Variable Autonomy Architecture (EVAA). EVAA, which was provided by the FAA, contains a GCAS module. The GCAS module within EVAA version 3.3 R1 was the subject of this research.

This research focused on two characteristics of the GCAS — nuisance and protection. Nuisance refers to any inappropriate intervention or inhibition from the GCAS inflicted on the pilot. Protection refers to the ability of the GCAS to detect and prevent an airplane from a ground collision. Throughout this project, the authors wrote four different academic conference papers pertaining to these two categories. This document aggregates these four papers and provides a high-level summary of the project as a whole.

The findings show that EVAA GCAS is quite nuisance-prone particularly around terminal areas, showing marginally acceptable results in the category of nuisance. In all Automatic Dependent Surveillance-Broadcast (ADS-B) flights processed, 34.05% had one or more unrequired GCAS activations with the vast majority of those occurring within the terminal area. A GCAS whose activation rate is approximately thirty percent warrants further investigation in the category of nuisance. Using ADS-B data from real-life flights to test nuisance was a very effective approach to understand undesirable interference of the GCAS (see Paper 1: ADS-B Nuisance Evaluation).

To evaluate the system protection, a recovery autopilot was built and coupled alongside the GCAS under investigation. In all protection-centric analyses, the recovery autopilot was used as a ‘real-world’ airplane which the GCAS attempted to protect. Following the creation of the autopilot, the trajectory prediction algorithm (TPA) was tuned to match the simulated airplane’s flight path (see Paper 2: Creating the Recovery). This research explored two different avenues of protection assessment — TPA analysis and system-wide analysis.

The TPA analysis compared the true simulated flight path with the anticipated flight path of the TPA. This metric is significant because terrain not residing within the anticipated flight path is not assessed for potential ground collision risk. If the anticipated flight path is inaccurate, the decision-making capabilities of GCAS are corrupted. This research found that the TPA exhibits exceptional performance in most flight envelopes. However, the TPA veers away from the

airplane's true flight path (see Paper 3: TPA Evaluation) specifically in the context of left turning maneuvers.

The system-wide analysis involved simulating many random and unique scenarios to reveal a big-picture representation of system performance. This was achieved using Monte Carlo analysis techniques, which aim to reveal a high-level understanding of complex, multi-variable systems. After simulating over 60,000 unique conditions, findings showed that EVAA GCAS protection capabilities are exceptional, with an overall protection rate of 98.45% (see Paper 4: System-Wide Monte Carlo Evaluation).

The results from the system-wide evaluation corroborated the findings from both the nuisance and TPA evaluations. The ADS-B nuisance study showed an activation rate which was higher than desirable, primarily due to activations in the terminal area. The system-wide assessment showed that most activations occurred with a very comfortable 'gap' between the aircraft and the terrain. In addition, the TPA evaluation found that the left turning recovery was potentially less effective. The system-wide study showed a left-turning climb protection rate of 96.61% as opposed to 97.65% and 98.78% for straight climb and right-turning climb, respectively. Of the small number of failed protection scenarios, many were at very low speed (near the stall speed) and in highly mountainous terrain. These areas are very dangerous and difficult to protect against. Overall, GCAS has very strong protection abilities, but further tuning to alleviate prominent nuisance may prove beneficial to the system. All tasks outlined as a part of the University of Tulsa's role in the evaluation of GCAS, were successfully completed. Using real-world ADS-B to evaluate nuisance of GCAS proved extremely effective as did Monte Carlo analysis of the protection capabilities. Future work to expand these technologies to additional safety systems should be performed.

1 Introduction

General aviation, or GA, the segment of aviation typically associated with small personal aircraft, is much more dangerous, given almost any metric, compared to commercial aviation. In the period between 2000 and 2019, GA accounted for 62% of all hours flown in civil aviation but over 95% of all civil aviation fatalities. The most recent finalized statistics on GA fatalities give figures through 2019. In that year, GA accidents took the lives of 414 people. Research shows that up to 25% of these fatal accidents might be prevented with the inclusion of a properly utilized system to avoid terrain collision. This 25% comes from accidents with defining events of controlled flight into terrain (CFIT), low altitude operations (LALT), unintended flight into instrument meteorological conditions (UIMC), and loss of control – inflight (LOC-I), amongst others. Fortunately, systems to avoid ground collision are relatively mature and have been in operational use for many decades. Even automatic variants are beginning to be included in operational military aircraft, and commercial aircraft are exploring their use. For general aviation however, current terrain warning and avoidance systems are often unsatisfactory when installed and are only installed on a small number of general aviation aircraft. One issue with these systems is that there is no well-established method of evaluation for the GA community, therefore, current methods are expensive and oftentimes unsatisfactory. Researchers at the University of Tulsa have 15 years of experience in designing and evaluating Ground Collision Avoidance Systems (GCAS) and will bring that experience to evaluate the effectiveness of GA GCAS. The outcomes of this system evaluation will then be disseminated to the larger community so that regulators and industry can benefit from the experience. The goal is the proliferation of GA GCAS to the general aviation fleet to potentially save nearly 100 lives a year.

1.1 GA GCAS evaluation approach

The Armstrong Flight Research Center at the National Aeronautics and Space Administration (NASA) developed a multi-monitor run-time assurance system (Hook, et al., 2018) to enable aircraft autonomy (Skoog, Hook, & Ryan, 2020). The system, called the Expandable Variable Autonomy Architecture (EVAA) contains a GCAS and is the subject of this research. For this project, EVAA was provided by the FAA for the sole purpose of a comprehensive system evaluation. This assessment will occur in two high level parts.

1. The nuisance evaluation determined how well GA GCAS allowed normal operations without any unnecessary warnings.

2. The protection evaluation determined how well GA GCAS protected an airplane and pilot in the event of an actual and imminent ground collision.

Details for these evaluations are included in the following sections, as are details for work tasks centered on tool development and software preparation including the conversion of EVAA to a software in the loop version.

1.2 Nuisance evaluation overview

To perform the nuisance evaluation, actual real-life airplane operations were used. In this way, we assessed the actual nuisance potential of GCAS, and not any unrealistic scenarios. Data to perform the nuisance evaluation was derived from FAA Automatic Dependent Surveillance-Broadcast (ADS-B) datasets for Cessna 172 aircraft in the vicinity of Van Nuys airport (KVNY) in LA County, California. This data contains latitude, longitude, altitude, and ground speed. Additional information required of the GCAS algorithm, such as bank angle, pitch angle, throttle position, et cetera, which are input into the GA GCAS algorithm, were developed with the use of a custom model created as a part of the project. Because of the size of the ADS-B dataset, the team began with hand-picked flights, performing the requisite modeling and state output. Then, as algorithms became more mature, larger scale evaluations were used, which included a single airframe over the course of many months. After data extraction and modeling, the resultant data was input into the GA GCAS algorithm to see if it would have caused a nuisance warning. This data was then combined into a statistical analysis of the results.

In general, the goal of this research project was to develop and produce tools that provide a functional analysis of a GCAS (in this case the EVAA GA GCAS). Throughout the duration of the project, four different academic conference papers were written, each signifying a new milestone of the project. Table 1 provides an overview of each paper topic.

Note: Each one of the following sections is an independent paper submitted to one of three conferences. Therefore, each has an independent introduction, background, results, and conclusions. This report summarizes all four papers into an overall conclusion.

Table 1. Papers that comprise this report

Paper Topics	**Category**	**Description**
1. ADS-B Nuisance Evaluation	Nuisance	Assess the EVAA GCAS nuisance by feeding real-world ADS-B data to the GCAS and looking at inappropriate activations.

Paper Topics	Category	Description
2. Recovery Autopilot	Protection	Details the development of the custom C172 GCAS autopilot that supported the protection and nuisance evaluations.
3. TPA Evaluation	Protection	Evaluates the EVAA GCAS trajectory prediction algorithm (TPA) module and its implications on full-system protection.
4. Monte Carlo Analysis	Protection	Explores a system-wide approach to evaluate the EVAA GCAS protection abilities by generating and simulating a set of over 60,000 randomly derived initial flight conditions.

1.3 Protection evaluation overview

The protection evaluation utilizes Monte Carlo analysis methods which vary several aircraft state and environmental parameters and compare resultant simulated trajectories with underlying terrain heights to determine whether a collision would have occurred. Terrain heights were pulled from publicly available terrain datasets given the latitude and longitude of the aircraft state during the maneuver.

The software structure used to generate the Monte Carlo analysis was broken into three technical parts and a user interface. The user interface allows the user to select which datasets and which parameters to vary. Other options will include the number of tests to be performed, the model setup for the runs, and any configuration data necessary for the simulation and evaluation. After the parameters are inputted, a Cessna 172 model is used to generate a flight path and flight parameters required by the EVAA GA GCAS algorithm. When an imminent simulated ground collision was detected by the EVAA algorithm, a commanded recovery maneuver was generated and flown with the intent of clearing the local terrain. After the simulated avoidance maneuver was complete, the post flyup trajectory was compared with terrain and obstacle data heights to determine if the flyup would have been successful, or whether terrain impact had occurred. After all results were obtained, they were stored and categorized so that statistical analysis could be performed. If a flyup was unsuccessful, these runs were plotted and graphed to understand commonalities of all failed GA GCAS activations.

2 Paper 1: ADS-B Nuisance Evaluation (presented at the 2023 IEEE Digital Avionics Systems Conference, Barcelona, Spain, 2023 (Hubbard, 2023))

2.1 Using ADS-B data to analyze the potential for nuisance in a Ground Collision Avoidance System

Travel via air carrier is undoubtedly one of the safest forms of long-range transportation available. In contrast, every year an unacceptable number of fatalities still occur in GA, aviation's most utilized segment. Research has shown that safety technologies that exist for military aviation have the potential to reduce this rate by nearly 71% (Hook, Ryan, Skoog, & Fuller, 2023), little progress has been made to adapt them to GA. In response, the FAA has begun funding research into both the development and verification of these systems to hasten their availability. As part of this research, this work has developed a methodology for evaluating the occurrence of false alarms in a GCAS using a large number of actual ADS-B flight logs. These methods were then applied to the evaluation of NASA's EVAA GCAS, which is an example of a mature GCAS solution. The first results statistically show that EVAA's current suppression methods reduce nuisance activations considerably, but improvements can be made to further reduce activations during take-off and landing procedures. Furthermore, it was discovered that most cases of nuisance activations occur during pattern work or training maneuvers. In this scenario, a false activation could be hazardous to the crew of the aircraft. It will be shown that the utilization of ADS-B flight data has proven highly effective in detecting the potential for nuisance activations within a GCAS. This suggests that the methodology employed in this study can be applied to similar safety systems, thereby enhancing safety measures in the realm of GA across the board.

2.2 Introduction

Several studies have shown unacceptable numbers of yearly US fatalities in Federal Acquisition Regulation (FAR) Part 91, GA (Hook, Sizoo, & Fuller, 2022; Fuller & Hook, 2023). Fortunately, new digital safety systems may hold the key to significantly reducing fatalities in this mode. In fact, Hook, Sizoo, & Fuller (2022) performed data mining on National Transportation Safety Board (NTSB) general aviation accident reports (National Transportation Safety Board, 2021) to determine the potential of existing safety technologies to otherwise prevent or save passenger lives. Analysis showed that of all NTSB reports investigated, 71.2% (Hook, Sizoo, & Fuller, 2022) could potentially have been saved by the incorporation of one or more automatic digital safety systems.

Of those systems, the most promising, in terms of preventing fatal accidents, may be GCAS. GCAS was found to be an effective solution to prevent up to 26.2% of all fatal accidents in GA (Hook, Sizoo, & Fuller, 2022), or approximately 110 preventable fatalities per year on average, as shown in Figure 1 (Hook, Ryan, Skoog, & Fuller, 2023). This is because GCAS can prevent accidents that are classified in several accident categories. The most obvious of these is CFIT; however, other categories such as LALT and UIMC are also prime candidates for GCAS saves. Even many LOC-I accidents, which represents the largest fatal accident category by a wide margin, are candidates for GCAS saves.

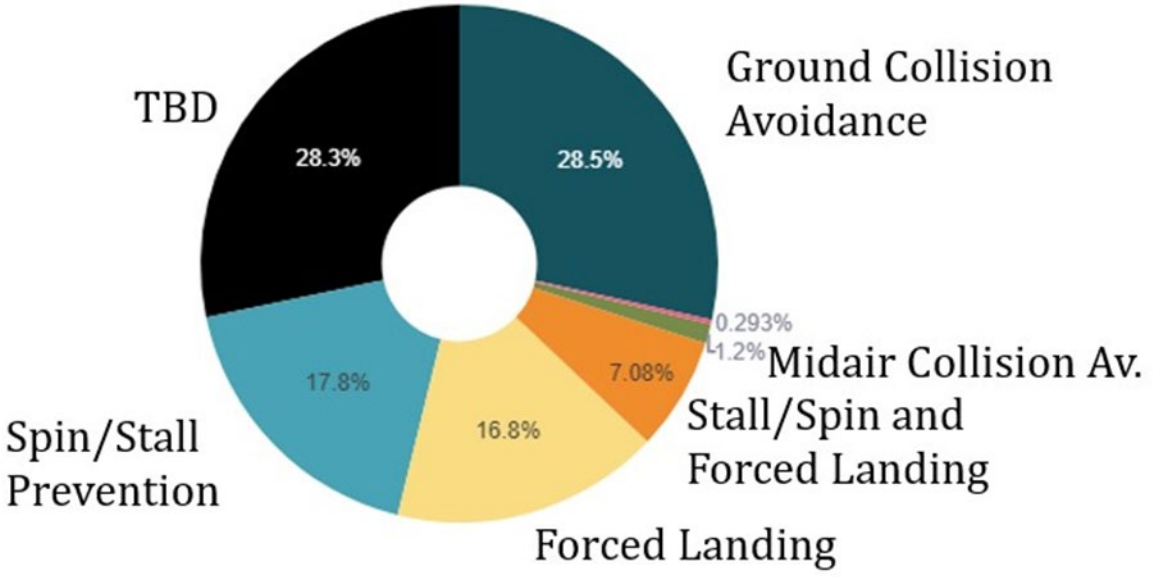

Figure 1. Total preventable accidents by safety system

In response to this reality, work to design and incorporate GCAS into GA airplanes is underway. As a part of this process, techniques to verify the acceptable operations of GCAS are also being developed. This research is a part of this larger effort and is specifically interested in developing methods to determine how GCAS may interfere with normal, safe GA flight.

In general terms, digital safety systems should not interfere with regular flights if there are no imminent and genuine safety risks. This is especially true when performing GCAS-based flight maneuvers, as an unneeded GCAS avoidance activation could lead to many undesirable effects. Due to this, there is a need to determine the potential for unwanted or nuisance-based activations that can occur in the system (Swihart, et al., 2011).

The goal of this paper is to provide techniques to analyze GCAS activations, developing a system to assess the potential for nuisance during routine flight operations. To accomplish this,

ADS-B data for a large number of Cessna 172 aircraft was utilized to develop a flight model which was fed into a GCAS. For this study NASA's EVAA GCAS system was used as it is a mature GCAS solution and was developed with small aircraft in mind. The ADS-B data was interpolated to ensure that a minimum of 1 Hz update rate was well exceeded for GCAS operation and analysis. Additionally, this work implemented a tuned Cessna-172 trajectory model developed by the Vehicle Autonomy and Intelligence Lab (VAIL) at the University of Tulsa for GCAS specific maneuvers (Maley, Hubbard, Urban, & Hook, 2023). This trajectory model was used by the GCAS module to generate the TPA predictions for the specified airframe.

2.3 Background

2.3.1 General aviation

The majority of aviation can be broadly deconstructed into two separate categories of operation, as defined by the International Civil Aviation Organization (ICAO), Commercial Air Transport and GA (Hook, Ryan, Skoog, & Fuller, 2023). Commercial air transport covers civilian aerial operations related to passenger transport, cargo, or mail services for remuneration. GA however is defined as civil operations that are neither air transport nor other air services for hire.

As a result, GA operations offer a wide range of mission objectives and crew proficiency. Instructive flights, private business flights, recreational flights, and personal transport are some examples. Additionally, professional commercial pilots are also able to fly under the GA classification so long as they are not under commercial obligations, further diversifying the proficiency of the potential crew.

Differences between the aircraft categories can become even more complex when examining the FAA's classification of aviation operations. Part 91, GA, must follow a separate set of regulations than both Part 121, commercial, and Part 135, air taxi. This has led to commercial charters being more closely regulated from an operational perspective than to GA operations. Beyond that there remains a significant disparity in research funds invested in GA-based safety systems compared to equivalent systems on commercial aircraft.

2.3.2 Ground collision avoidance systems

Ground avoidance systems have been used in aviation for over 50 years. Traditional systems, such as Terrain Avoidance and Warning System (TAWS), involve manual warnings to pilots. This is usually an auditory "Altitude!" alert or similar visual or auditory warning. These systems, however, are reaching the limit at which they can prevent fatalities. Swihart et al. (2011) found that pilots quickly 'tuned out' manual warnings or would disable the system entirely due to the

frequency of false activations. As such, automatic systems that activate only when the pilot is unaware or unable to perform the recovery maneuver themselves are the next natural step to increase flight safety.

Traditional GCAS methods of collision detection involves the use of a trajectory prediction algorithm (TPA) to predict potential future states of an aircraft, given that it begins to perform a preset avoidance maneuver (Suplisson, 2015). This trajectory is then compared to a virtual elevation database; any overlap or touching of the TPA against a buffer region is considered a collision with terrain (Swihart, et al., 2011). The buffer region, commonly referred to as the terrain clearance buffer (TCB), is tuned specifically to the type of airframe that is in operation. Figure 2 (Maley, Hubbard, Urban, & Hook, 2023) outlines the detection process of a typical GCAS activation, where the TPA is projected to enter within the TCB region. The red circle indicates where a collision detection occurs. Modern systems generate several TPA paths for different avoidance maneuvers to maximize the potential for safe recovery. For these systems only the last viable TPA is selected as the avoidance maneuver. Once the need to avoid has been established, the high-level decider switches from the nominal controller, in this case the pilot, to the GCAS controller. This method of high-level detection and switched control functions inside a so called run time assurance (RTA) control architecture. The role of an RTA system is to ensure proper and timely handling between different safety modules. Figure 3 (Maley, Hubbard, Urban, & Hook, 2023) shows an overview of a basic RTA system with only a GCAS monitor integrated into the control loop. The introduction of safety monitors allows for minimal interference during routine operation. This is because the monitor is designed to deactivate and return authority to the nominal controller once avoidance has been achieved and no threat is detected. The monitor is also unable to directly assume control, as that decision is made by the decider, thereby preventing an overlap or struggle for control from several safety monitors at the same time. An advantage of this is that the structure provides a natural hierarchy of control between separate monitors, allowing one to take priority over another based on the changing threat assessment values the monitors produce.

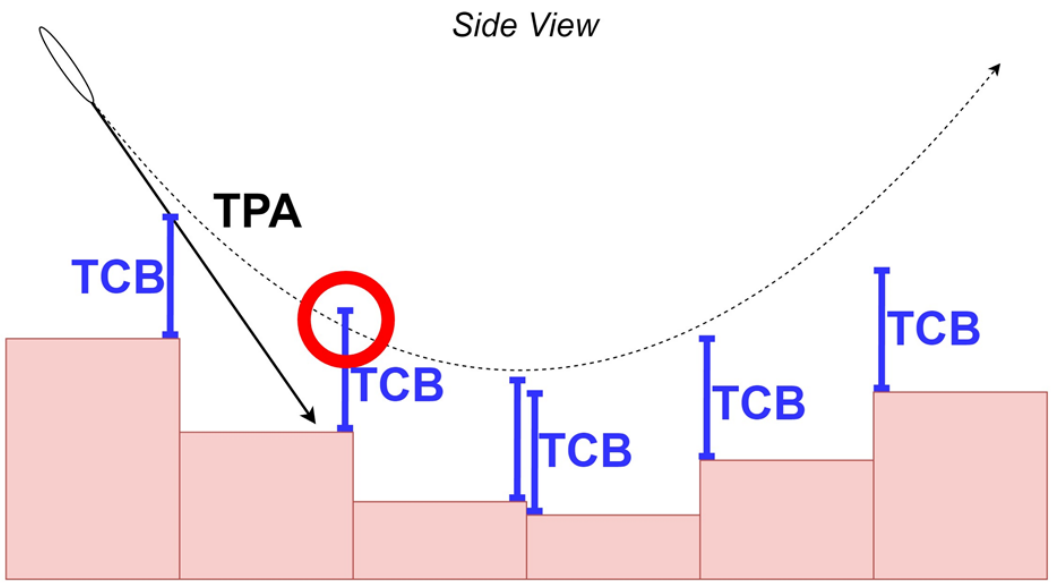

Figure 2. GCAS Terrain clearance and TPA prediction

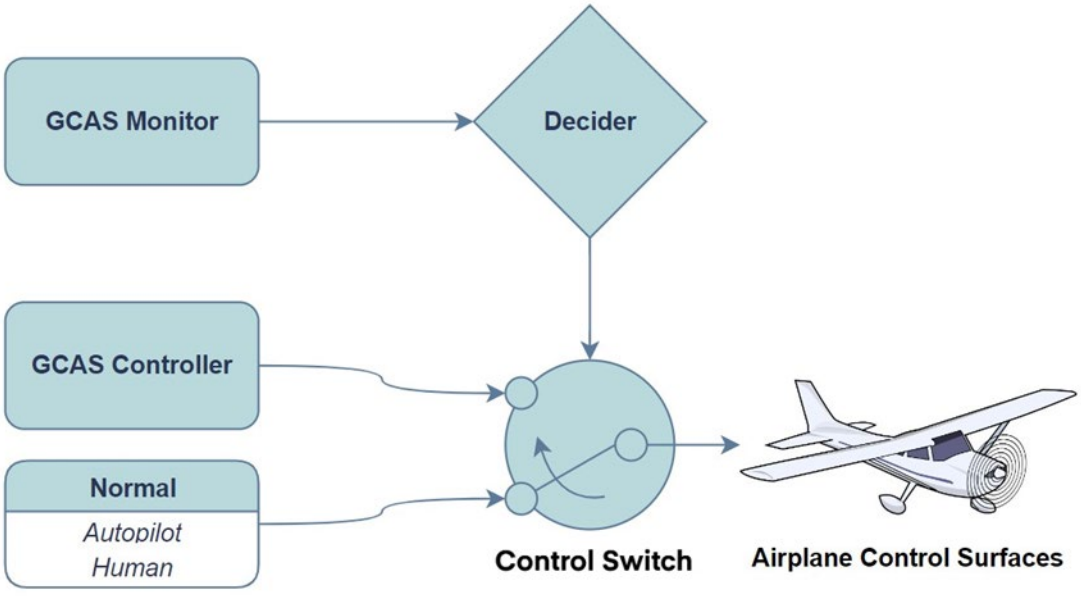

Figure 3. GCAS RTA decision structure

2.3.3 Expandable variable autonomy architecture

The National Aeronautics and Space Administration's (NASA)'s new EVAA, produced under the Resilient Autonomy Project at Armstrong Flight Research Center, offers a centralized control platform for the integration of multiple high-level, autonomous safety modules (Connor, 2020). From a high-level, EVAA operates as a control structure, monitoring the state of the aircraft at all times. It then determines which safety modules to activate based on current mission-safety priorities. The architecture is designed to be self-deterministic and adhere to the following priority structure:

1. Human safety
2. Damage to property (both aircraft and civilian infrastructure)
3. Mission completion

The rigid set of behaviors coded into the platform's high-level monitor management allows EVAA to have greater control over mission intent and aircraft limits, just as a pilot would.

EVAA, while primarily being designed for unmanned vehicles, can be configured to operate on manned aircraft. To do so additional control is implemented giving the pilot the ability to disengage any emergency override executed by EVAA. Moreover, embedded systems in EVAA follow strict rules to only interfere in nominal flight as absolutely necessary. This means that safety modules are only executed seconds before it is determined 'point of no return' is reached.

2.3.4 Automatic Dependent Surveillance-Broadcast

Today, nearly all commercial and GA aircraft have some form of onboard ADS-B technology, allowing for remote monitoring and recording of flight data. The data collected by ADS-B includes relevant information describing the aircraft's position, velocity, and identifying information. This assists Air Traffic Control (ATC) in managing air traffic and the local airspace. Positional information is collected by navigational satellites and then broadcast out to other aircraft and ground tracking stations from the aircraft (Sun, Ellerbroek, & Hoekstra, 2016). The information broadcast using this system is unencrypted and can be freely retrieved from any party with an appropriate ground collection station.

ADS-B can perform data collection through several different methods, each of which has their own limitations. The primary source of ADS-B data is collected through on-board GPS readings that are relayed by a transponder back to remote ground stations. However, in the case of GPS failure, ADS-B data can be supported by supplemental navigation systems, onboard sensors, and using older radar tracking systems when possible (Amin, Clark, Offutt, & Serenko, 2014). This

in-turn leads to a wide range of error thresholds within the ADS-B data itself, as each aircraft can be flying with a varied array of commercial sensors. To assist in identifying the source, each set of ADS-B data includes tags inside the data streams to identify the source sensor.

2.4 Data extraction and interpolation

2.4.1 ADS-B pre-processing

Generally, when working with ADS-B data, there are two different approaches to evaluating the data: real-time processing and post-processing (Ostroumov, Kuzmenko, & Kyzymchuk, 2022). Real-time processing is relevant to monitoring systems that need to use ADS-B data sequentially, with reference to previous state information. Alternatively, post-processing is not restricted by sequential ordering, allowing any discrete point of the flight to be processed without reference to a previous state. GCAS systems make predictions based on the current state of the aircraft, without reference to previous states. However, all parameters used to generate corresponding trajectories are not provided by the initial ADS-B data stream and must be calculated, of which some values reacquire reference to future states. As a result, a mix of sequential and post-processing is used to first generate missing parameters and then scrub the results for GCAS activations.

Sun, Ellerbroek, & Hoekstra (2016) undergo large-scale flight phase identification using machine learning techniques on ADS-B data. From this, they provide recommendations and suggestions when handling and preparing unrefined ADS-B data. Collected data is accumulated and stored in an unordered manner, as ground stations record any signals received regardless of completion, origin, or time of retrieval. Because of this, pre-sorting and handling of data is required for large-scale processing. Additionally, Sun et.al (2016) found unique data entries can surpass the hundreds of millions each day, providing an intimidating amount of data to sift through. Partial signals are also very common during collection, as the transponders are unable to transmit outside line-of-sight.

Steps were taken to perform initial pre-sorting and filtering of the ADS-B data to resolve many of these issues before interpolation and GCAS testing were conducted. Unrefined ADS-B data from 2021 was sourced from a region in the United States and utilized for this analysis. Because the GCAS controller had been tuned to match the flight dynamics of a Cessna 172, only Cessna 172 data from those respective flights is of interest. A list of all Cessna 172 aircraft operating within the region of interest was obtained to allow isolation of desired aircraft. This list was then compared to the unrefined ADS-B data to identify all flights with the airframe of interest. Flights that do not correspond to the matching airframe were dropped from the source data file.

After data has been refined to include only aircraft of interest, sorting methods were applied to order data sequentially, set entries with missing fields were removed (partial messages), and any duplicate messages were dropped. Individual aircraft were then isolated by their tail numbers and examined to identify significant gaps within the data stream. A threshold of 30 seconds of data loss was found to be acceptable, any data entries that occur with a greater delay were cut and considered a separate flight segment. After pre-sorting was completed and each individual flight's data isolated, interpolation and GCAS detection was then initiated.

2.4.2 Parameter calculations

As mentioned previously, ADS-B data only updates at a rate of 1Hz under ideal conditions. It is, however, more desirable to test for activations at a higher data resolution, as the aircraft may make unexpected movements during update periods. Additionally, ADS-B data typically only includes GPS-based positional and velocity information, which does not contain information concerning the aircraft's attitude. To resolve this issue, a model was implemented and refined to calculate approximate attitude information, given available parameters within ADS-B data. Equations used to derive the model are based on Beard & McLain (2012). The finalized equations used for generic GA aircraft GCAS parameter calculations are outlined below:

Airspeed, 'V_a':

ADS-B data provides north, east, down coordinate frame velocities based on GPS sensors; these are converted into a magnitude airspeed for GCAS detection and operation. Here V_N, V_E, and V_D are the velocity of the aircraft in the north, east, and down directions, respectively. Unfortunately, the scope of this project did not allow for localized winds and weather from the flight period to be incorporated into the dynamics. However, the modified model developed can be configured to incorporate wind data into parameter calculations for future analysis.

$$V_a = \sqrt{V_N^2 + V_E^2 + V_D^2} \tag{1}$$

Track, 'χ':

This formula uses the following equation to calculate the respective χ values for each of the newly interpolated time steps. For calculation purposes latitudinal and longitudinal information is converted into Universal Transverse Mercator (UTM) format so that unit meters may be used.

$$\chi = \tan^{-1}((E_n - E_o)/(N_n - N_o)) \tag{2}$$

Here E_o and N_o are the previous state's easting and northing values. E_n and N_n represent the airstate's easting and northing values two-time steps from the current. This was done to better approximate the heading angle of the aircraft, as direct point to point approach would fail to approximate the true heading. It is important note that for data collection, the last data point is considered to be forfeit. This is due to the breakdown of the equation and model when there is no future airstate information. This single point of data loss is carried over in all subsequent calculations.

Gamma, 'γ':

This value is needed to properly calculate TPA maneuvers for prediction. It is the angle between the aircraft vertical component of the velocity vector and the local horizon, commonly referred to as the flight path angle. Here it is determined as a function of the lateral and vertical velocities.

$$\gamma = \tan^{-1}\left(-\frac{V_D}{\sqrt{V_N^2 + V_E^2}}\right) \tag{3}$$

Normal Acceleration, 'N_zG':

This value is the acceleration experienced in the vertical direction of the aircraft body. This will be used by the GCAS controller and used to determine the approximate roll angle, ϕ. Here N_zG is determined by calculating the horizontal and vertical components and then using equation 4 to determine the resulting normal acceleration.

$$N_zG = \frac{\sqrt{((V_a \cdot \dot{\chi})^2 + (V_a \cdot \dot{\gamma})^2)}}{G} \tag{4}$$

$\dot{\chi}$ represents the approximate rate of change of χ and G represents the value of gravity.

Roll Attitude, 'ϕ':

This value is the last parameter to be calculated as it relies on both NzG and γ. Because this parameter is determined through 2nd level calculations, it is further subject to noise from the sensors.

$$\phi = \tan^{-1} \frac{(V_a \cdot \cos\gamma \cdot \dot{\chi})}{N_z G \cdot G} \tag{5}$$

2.5 Model development

2.5.1 Stage I: Initial model development

The first stage of development involved generating a generic model to approximate the dynamics of GA aircraft. A basic aircraft simulator, constructed based upon control models outlined in Beard & McLain (2012), was used to generate mock flight data that would resemble what would be received from ADS-B sources. The simulation ran with a time step of 10Hz, for ten seconds, generating two sets of data. The first set has all calculated data points and is labeled as the 'Full-Sample'. The second set is a 'Sub-Sample' of the first set, in which only data points with integer time values are recorded. See Figure 4 for an example of mock ADS-B Altitude data generated.

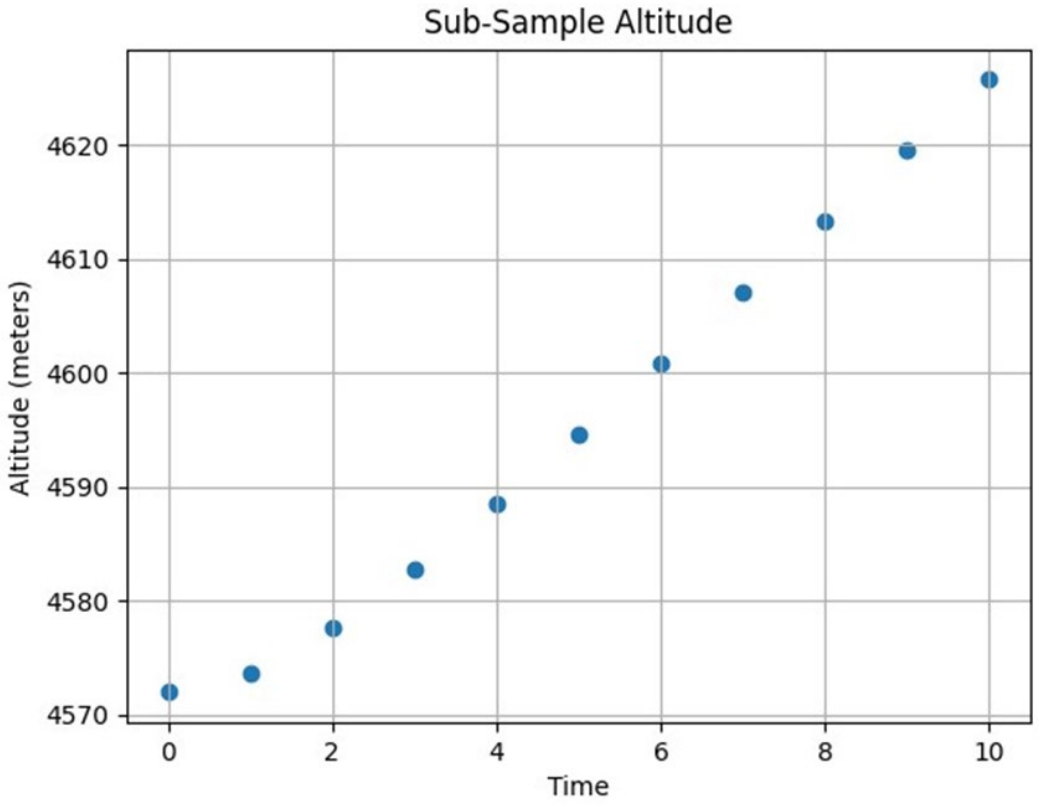

Figure 4. Pre-interpolated ADS-B data

The Sub-Sample's time data is then interpolated to increase the resolution by a factor of five. From here a cubic spline function is then applied across the airstate parameters to generate

values corresponding to the new time indexes. Figure 5 shows the resulting interpolation compared to both the Full-Sample and Sub-Sample sets.

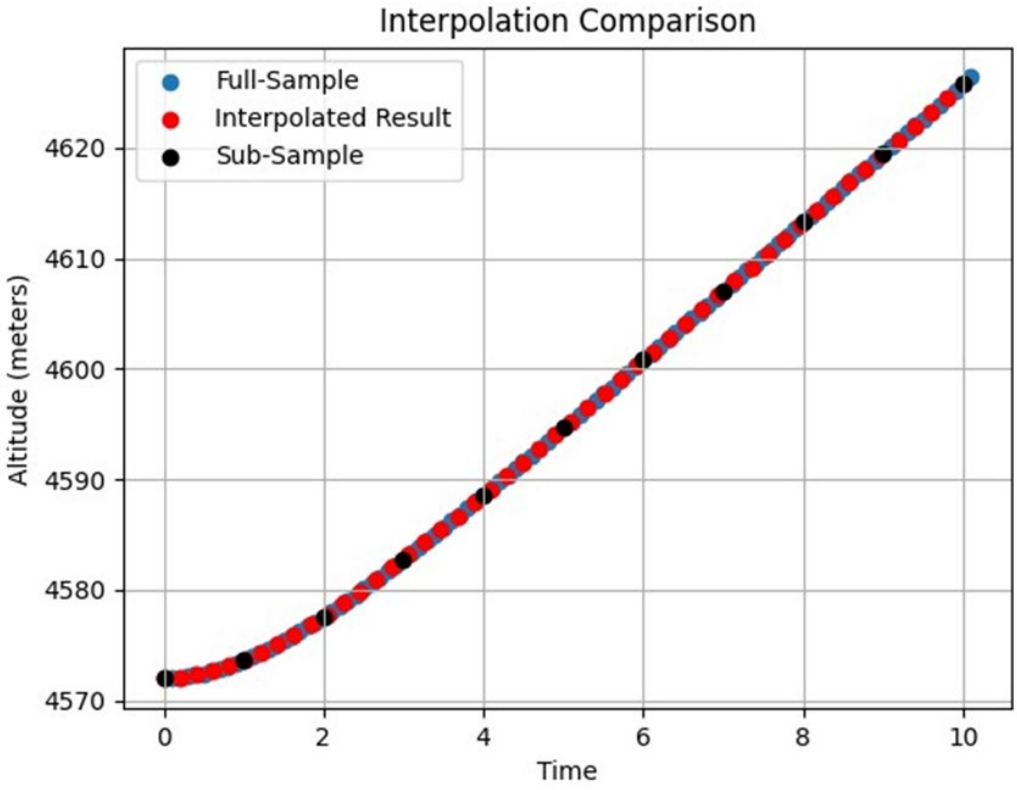

Figure 5. Post-interpolated ADS-B data compared to Full-Sample and Sub-Sample

2.5.2 Stage II: Simulation testing

The second stage of development moved from initial model creation to simulation testing. For this, JSBSim, an open-source flight dynamics simulator, was selected to generate test flight data. JSBSim is highly configurable, allowing for explicit testing conditions to be set and for control surfaces to be monitored directly.

Similar to the first stage, mock flight and ADS-B data were generated in two sets. However, the Full-Sample is taken in this case at the maximum collection rate, 120 Hz, instead of the previous 10Hz. Several test flights were generated with a variety of maneuvers to examine the performance of the parameter interpolation and calculations. These flights were flown manually and a variety of maneuvers were performed to verify the interpolation algorithm. Six simulation flights, each with increasingly more extreme maneuvers, were created. Figure 6 and Figure 7

show one of the flight's resulting γ and χ calculation and interpolation compared to the real, recorded Full-Sample values.

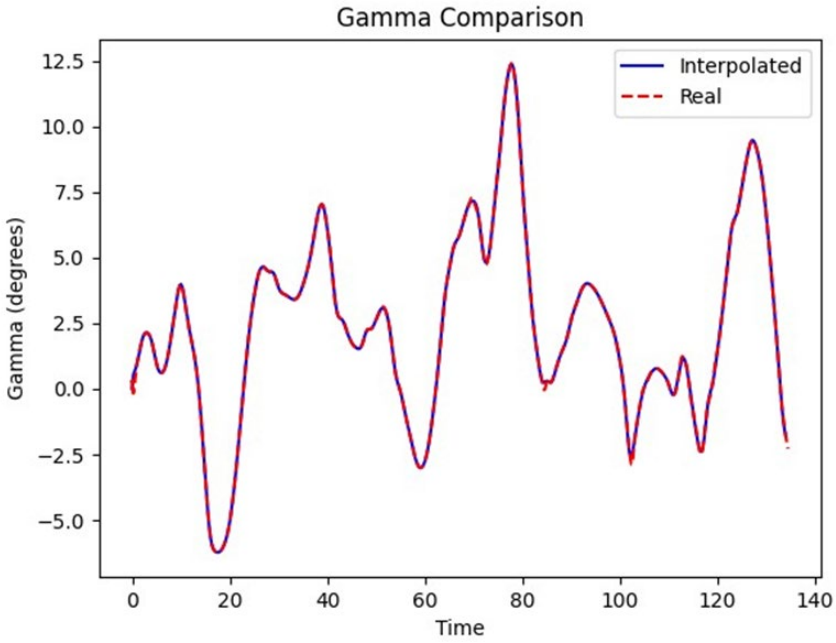

Figure 6. γ Calculation and Interpolation from simulated ADS-B data

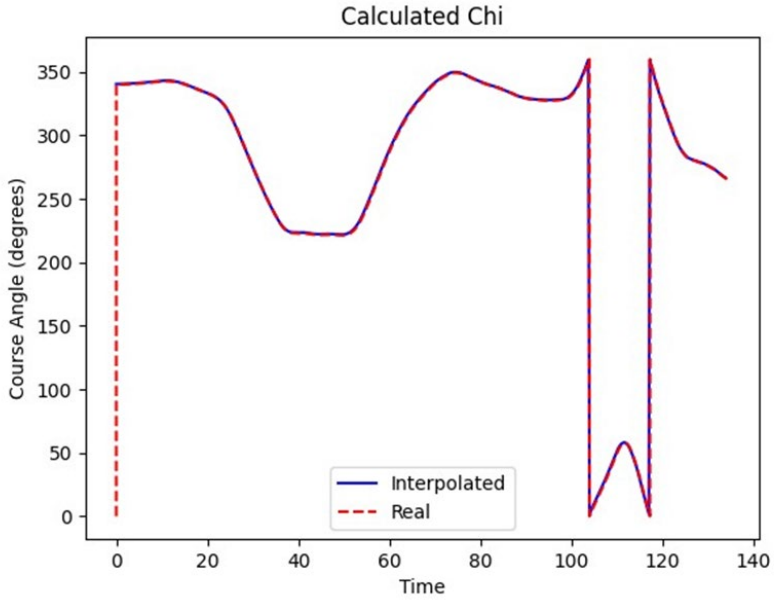

Figure 7. χ Calculation and interpolation from simulated ADS-B data

Based on the results, it was found that flights with the most deviation from the 'real' values are those that perform tight, continuous spiral maneuvers. Additionally, nearly every flight had brief deviations during the initial few points. When investigating this, it was determined that the loss of accuracy is unavoidable due to the nature of the developed model when it begins.

2.5.3 Stage III: ADS-B implementation

The last stage of development applies the interpolation and parameter calculations to the sorted ADS-B data. Here, further model tuning is needed, as real ADS-B data contains both sensor noise and occasional false points. To mitigate this, an averaging filter was applied to ϕ, as it displayed the most noise interference of all parameters. Additionally, GCAS detection was integrated into the testing process. After interpolation has been completed, the GCAS monitor is called to examine each time step of the processed flight segment and detect if an activation is required. If activation is detected, the activation period and current airstate is recorded for later analysis.

Figures 8 through 10 display processed parameters from the same training flight. Figure 8 shows the resulting altitude interpolation performed on ADS-B data over time in seconds. The interpolated result was then compared to the original ADS-B data to identify any extreme deviation from the source data. Throughout this process, further tuning was conducted to the interpolation process. Specifically, the interpolation count was reduced from five points to four. This was done in part to reduce the noise present within the data itself.

Figure 9 displays the flightpath of the aircraft in UTM projection. Detected GCAS activations are recorded as markers on the altitude and flightpath charts. Red upper arrow markers indicate that a GCAS activation has been detected and the system has entered recovery mode. Green 'x' markers represent the point when the GCAS system has determined the aircraft has been removed from danger. Figure 10 shows χ and ϕ attitudes throughout the duration of the flight. Due to the nature of the designed data displays, headings that shift between ±0° cause a jump between 0° and 360°.

The process was then scaled up and automated to pre-sort a month's worth of ADS-B data, which then analyzed accepted flight segments. This process is highly memory-intensive, as the unrefined ADS-B files for this period can be over 1GB per day. Depending on how many flight segments are accepted and how many activations are detected, each month of data takes approximately 22-26 hours to process.

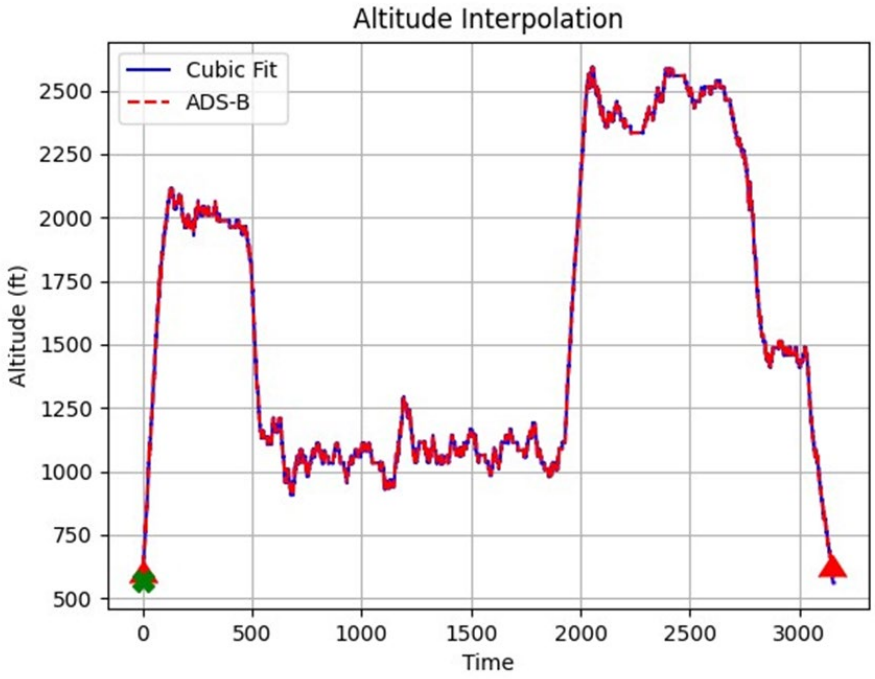

Figure 8. Altitude interpolation from ADS-B data

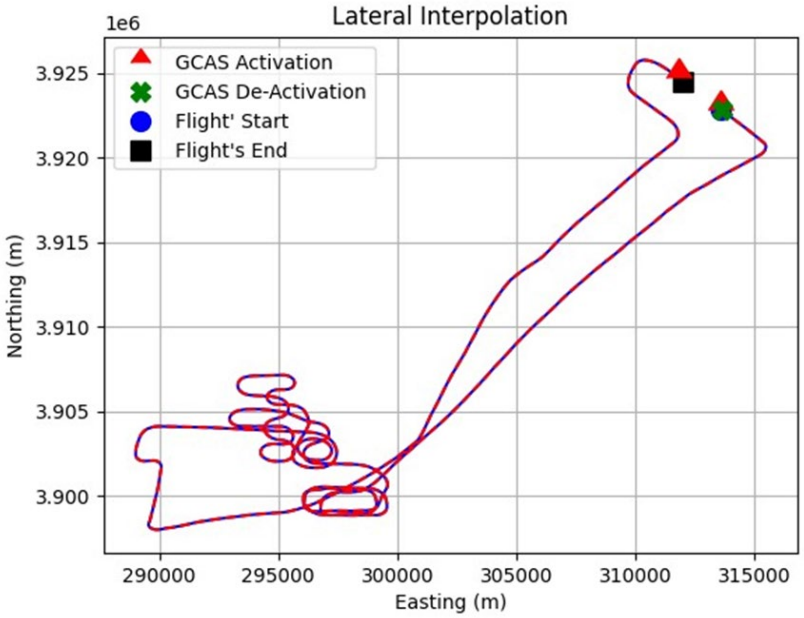

Figure 9. Lateral position interpolation from ADS-B data

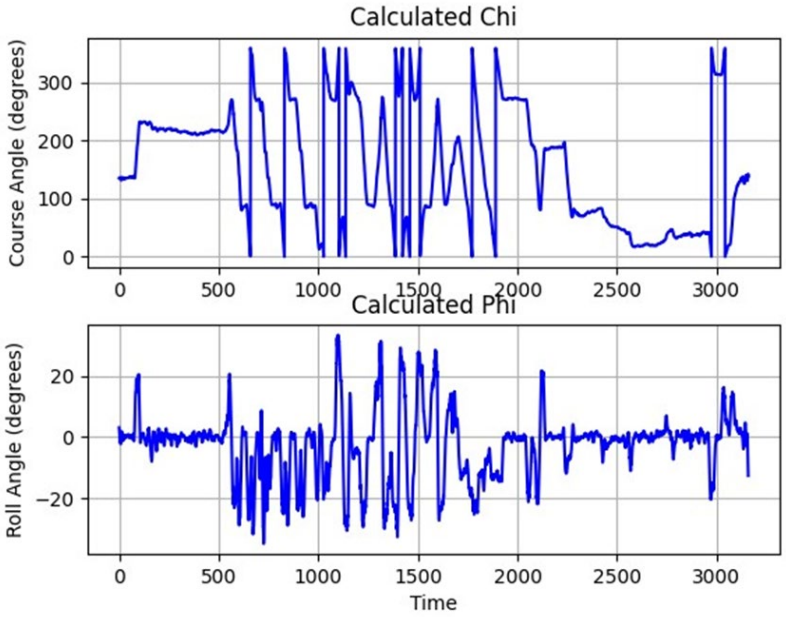

Figure 10. χ and ϕ Calculation and interpolation from ADS-B data

During this batch process, each stage of processing is saved in case of unexpected errors occurring before it can be complete. Additionally, statistics regarding the number of activations and their frequency are recorded and later discussed.

2.6 Nuisance analysis

2.6.1 Runway suppression

Traditional GCAS systems for military aircraft are designed to be nuisance resistant while flying at low altitude. A large majority of GA aircraft, however, are unlikely to be flying intentionally at low altitude with the exception of take-offs, landings, and mission specific objectives such as crop dusting. As a consequence of this, the locations with the highest probability of false activations are around airports and airstrips. A basic runway suppression system was implemented inside of EVAA's GCAS monitor to attempt to counter this. The suppression system records rectangular cutouts of runways across the United States. Within these zones activations are suppressed, preventing an avoidance warning from triggering inside the GCAS monitor. See Figure 11 for examples of suppression zones in blue, flight path outlined in yellow, and GCAS activation points marked by pins.

The effectiveness of this suppression method was tested using the interpolated ADS-B data for the March 2021. This month included 2,378 distinct flight segments. Results from suppression effectiveness are summarized in Table 2.

Figure 11. Example flight segments and runway suppression zones

Table 2. Runway suppression vs. no suppression

	Suppression ON	**Suppression OFF**
Accepted Flight Segments	2,318	2,318
Segments w/o Activations	1,536	932
Activations Recorded	1,819	3,482
Activations in Flight	**Total Percentage of Flights**	
No Activations	66.26	40.20
Single Activation	19.37	32.44
2 Activations	6.60	12.60

3 Activations	2.46	5.44
4 Activations	1.25	2.16
5 or more Activations	4.06	7.16

While testing with runway suppression activated and de-activated, it was found that 2,318 flight segments met conditions for interpolation with nearly 930 flight hours of data. Of these flight segments 1,536 had no GCAS activations with suppression activated, while no suppression had only 932 flights without activations. A resulting breakdown of the number of activations while using suppression shows that roughly 66% of flights complete without an activation. For flight segments with activations, nearly 20% contain at least one avoidance command, while subsequent activation counts decrease further. It should be noted that five or more activations are grouped together in this statistic as a majority of examined flights are pilots performing pattern work. As such, repeated activations occur at the same point causing the activation count to skew, an example of which can be seen in Figure 12..

Figure 12. Repeated activations due to failed suppression

According to the evaluation performed, the current suppression method is a significant improvement to nuisance reduction but can be refined for improved take-off and landing protection. In the remaining activations that were not suppressed, an overwhelming amount occurs near runways, but just outside of the suppression zone. Figure 13 and Figure 15in Section 2.6.2 below show this occurring. There are consistent activations right before the aircraft enters the runway while performing ‘touch-and-go’ maneuvers.

A possible alternative to strict runway segmentation is a lateral distance approach. In this, an acceptable distance from the closest airstrip is set and any activations that occur within the range are suppressed. This would expand the suppression zone from the airstrip itself, preventing activations when aircraft drift outside of the hard-set zone.

2.6.2 Large-scale GCAS statistics

ADS-B data from 2021 for the months of January-April, was run through this process then examined by EVAA's GCAS monitor. Statistics on the activation frequency recorded during this period are shown in Table 3.

Table 3. January-April 2021 potential GCAS activation during flight segments

ADS-B Flights: January-April		
Accepted Flight Segments	8,241	
Segments w/o Activations	5,435	
Activations Recorded	6,757	
Total Flight Hrs. Processed	3,377	
Activations in Flight	**Total Percentage of Flights**	**Percentage of Activations**
No Activations	65.95%	—
Single Activation	19.14%	56.20%
2 Activations	6.25%	18.35%
3 Activations	2.86%	8.41%
4 Activations	1.53%	4.49%
5 or more Activations	4.27%	12.54%

In total 8,241 flight segments were found to be acceptable, totaling approximately 3,377 flight hours processed. During this period over 6,757 activations were recorded, with the largest portion of activations only occurring once during a given flight segment, roughly 56%. While examining the frequency of activations, it was observed that the overall activation count during offending flight segments decreases as there are more activations throughout the flight. However, this trend stops holding true once the count surpasses five or more activations, instead beginning to increase in frequency. When examining an array of offending flights, there are specific findings and observations..

Figure 13 represents the middle of the envelope, with only one-two activations, for all flights processed to date. These flights are typically pilots switching between different nearby airports, performing training patterns, or undertaking personal transit. Activations related to relocating airstrips and personal transit make up the minority of the total activations observed. Flights that

indicate pilot training consistently contain the highest number of activations. The most commonly seen cause of activations in these cases is during touch-and-go maneuvers as the aircraft gets low and attempts to line up with the airstrip. Figure 14 shows an Altitude plot of such maneuvers along with corresponding GCAS activation periods.

Figure 13. Touch-and-go false activation

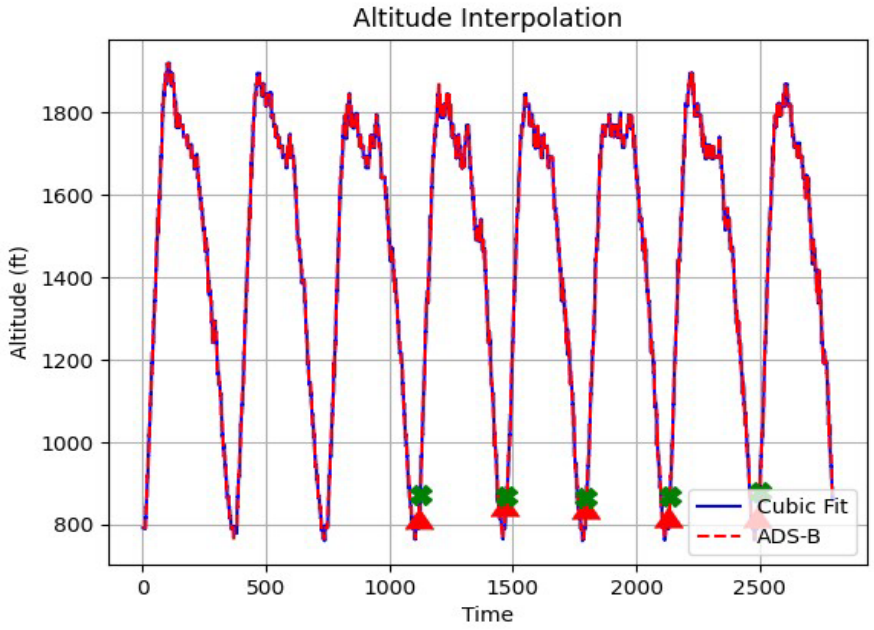

Figure 14. Touch-and-go altitude example

During these touch-and-go maneuvers, the aircraft either makes direct contact with the runway or maintains a low altitude over it before climbing back to cruising altitude. This process repeats several times as the aircraft circles around the runway. Activations during these maneuvers are shown to be inconsistent, as the varied approaches do not always trigger the GCAS monitor. Due to this, excessive activations would occur if the system was implemented onboard aircraft performing similar maneuvers. Furthermore, as these activations are occurring during take-off/landing the initiation of auto-GCAS maneuvers could potentially place the aircraft into harm.

While performing analysis of several of these flight paths, using a produced Keyhole Markup Language (KML) of relevant data inside of Google Earth, investigators found that there was significant noise during several of the touch-and-go maneuvers. This was found to be caused by the GPS data indicating the aircraft had left the runway while taking-off or landing. It was concluded that these activations should be considered nuisance, as the GCAS module would have falsely read the GPS position. The main causes of nuisance occurrence during these maneuvers are the sensor noise, combined with the strict runway suppression method.

Figure 15 shows the flight path and activation points of a training flight, in which several activations can be detected. During investigation, it was determined that several of these activations were in fact warranted and not a nuisance activation. Throughout this flight the pilot practices emergency landing approaches, emergency landings, and touch-and-go maneuvers. This was determined by several factors, the respective aircraft's identification, a short drop in ADS-B data collection during e-landing, altitude on approach, and landing location. Figure 16 shows the initiation of the aircraft taking off after the emergency landing, and the resulting activations while at low altitude. Here the activations begin immediately upon take-off as the aircraft is not on a recorded runway.

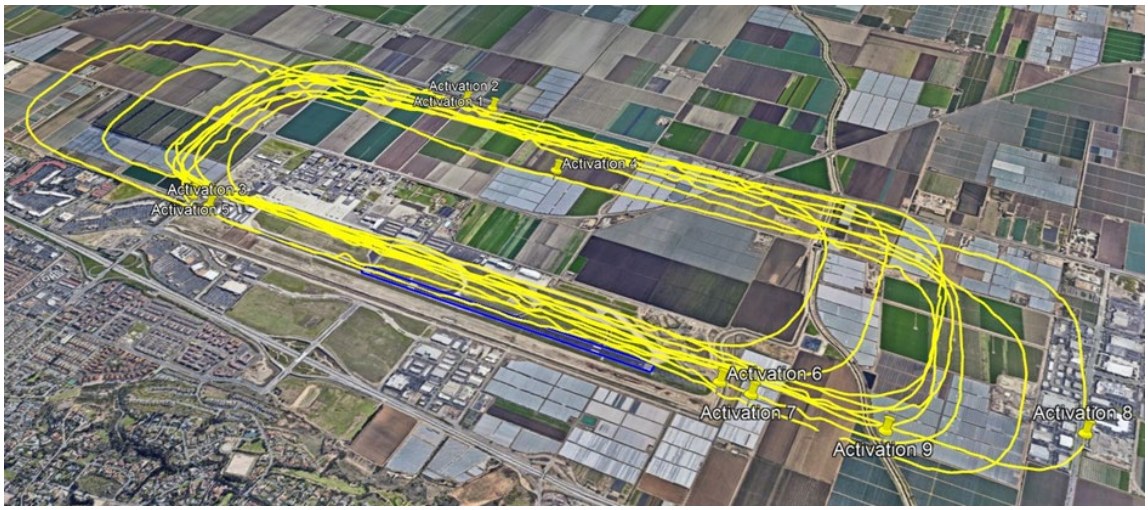

Figure 15. Training pilot flight path

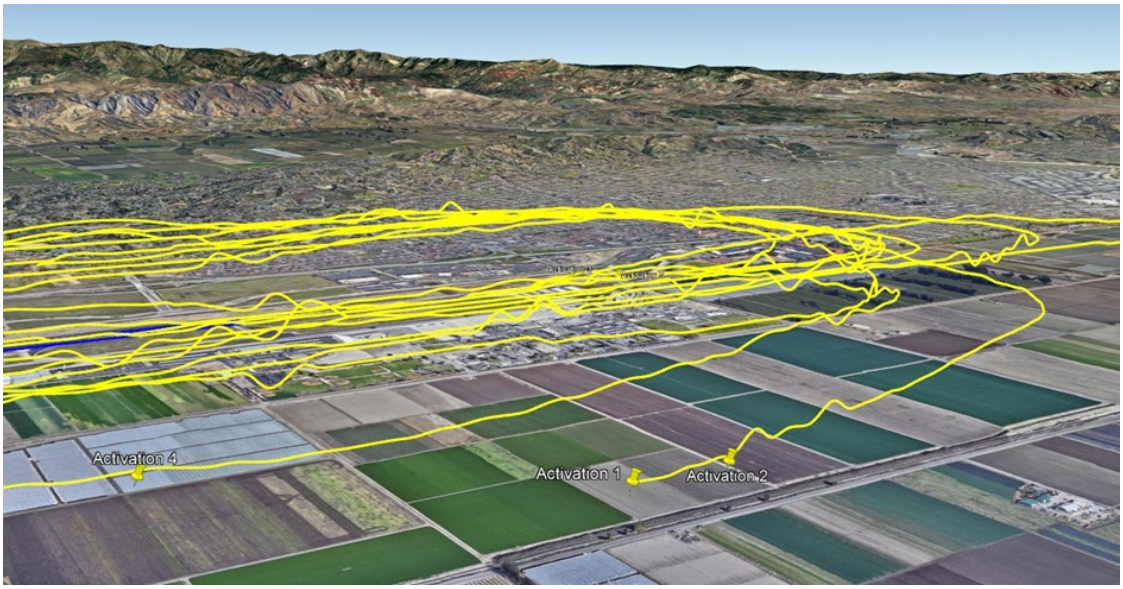

Figure 16. Activation during emergency landing training

Figure 17 shows the altitude graph of another selected flight segment with two separate GCAS activations. In it, the altitude suddenly drops to nearly 0 feet within mere seconds on two separate occasions.

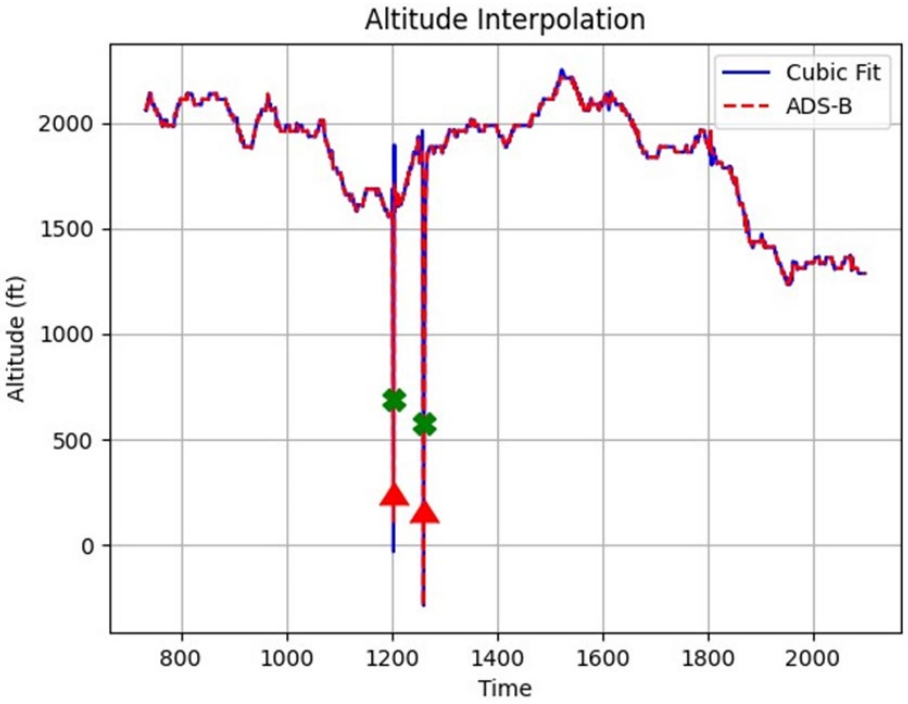

Figure 17. False activation due to ADS-B elevation

It was determined that this was not due to an error with parameter interpolation, but was in fact incorrectly recorded parameters within the ADS-B data itself. Though this occurrence is rare within the data, only accounting for ∼< 1% of all processes runs, it is still prevalent enough to be a cause for concern as it can falsely increase the nuisance estimate on the system. Due to this, additional steps should be made during either the pre-sorting processes or by the GCAS monitor to detect and remove offending data points or discredit the respective activation.

2.7 Conclusions

This paper presented results on the development of techniques to analyze potential nuisance in general aviation GCAS. Testing results applicable to NASA's EVAA GCAS indicated a remarkable drop in activation count when implementing runway suppression, but there remains a need for improvement to better prevent nuisance occurrences around runways. Improved methods should consider GA pilots who perform training maneuvers in and around airstrips as they contain the highest number of false activations. Furthermore, suppression methods that do not require aircraft to be at recorded runways should be implemented, as a non-insignificant portion of the GA fleet operates outside usual airports.

We have found that utilizing ADS-B data to simulate real flight maneuvers for safety systems offers an effective approach to isolating and analyzing the potential to cause nuisance activations. The use of ADS-B data to retroactively examine flights for danger provides a true representation of GA pilot behavior while in the air. Additionally, it removes the need for test data to be manually generated, allowing for an incomparable amount of flight time testing. Lastly, this approach to nuisance detection can be extended to other comparable automatic safety systems for aircraft in general.

3 Paper 2: Creating the Recovery Autopilot (Presented at the 2023 IEEE Aerospace Conference, Big Sky, MT, USA (Maley P. D., 2023))

3.1 Recovery autopilot analysis for a General Aviation Ground Collision Avoidance System

General aviation, the mode of air travel typified by small personal aircraft, is responsible for roughly 19 of every 20 fatalities every year in the US. However, recent technological solutions are becoming available which may bring the total number of fatalities in GA down considerably. Potentially the most effective of these solutions is the GCAS. GCAS avoids ground collision in a large number of cases: pilot error, disorientation, and temporary incapacitation. However, GCAS does not yet exist for general aviation, despite the wide benefit from its implementation. In response to this reality, GCAS development for general aviation has begun. This paper describes the design and verification of a GCAS controller on a simulated Cessna 172 aircraft. The GCAS controller provides the ability for the aircraft to automatically avoid terrain and is an important step in the initial phases of GCAS design. Considerations for the design of the controller's lateral and longitudinal axes are provided along with discussions on the overall controller structure. Controller modes and limiters have been designed and described to ensure safe operation of the controller, along with considerations for transient switching effects. Analysis of controller response as well as verification of mode and limiter operation are provided. Also included is a comparison between the trajectory generated by the GCAS controller and one predicted by the GCAS system. This paper provides important considerations on the initial stages of GCAS design for general aviation, and the beginning of safety assurance for GA.

3.2 Introduction

A dichotomy exists in aviation with regards to safety. While aviation classified as commercial transport is one of the safest forms of long-distance mass transit across the world, "GA is one of

the most dangerous. In fact, from the years 2010 to 2020 commercial transport operations accounted for roughly 51% of the flying hours, but only 6% of the aviation fatalities. On the other hand, general aviation accounted for 49% of the flying hours but 94% of the fatalities (National Transportation Safety Board, 2021). What contributes to this large difference in safety outcomes between aviation modes? Certainly, the aircraft themselves have a part to play, as GA aircraft commonly have only one engine and some may lack maintenance scrutiny. However, the majority of the difference in safety can be directly attributed to the pilot.

3.2.1 General aviation

According to an International Civil Aviation Organization (ICAO) working paper (2009), general aviation is defined in relation to commercial air transport as follows:

- Commercial Air Transport- the transport of passengers, cargo, or mail for remuneration or hire.
- GA – all civil operations other than scheduled air services and non-scheduled air transport operations for remuneration or hire.

In this sense, GA can have a huge range of missions and crew proficiency. For example, instructional flying, business flying, pleasure flying, and personal transportation can all be considered general aviation. Both non-professional and professional pilots may fly GA as long as it is not for commercial opportunity. In addition, there is a wide spectrum of aircraft and aircraft equipage in general aviation. As a result of this wide spectrum, there is a huge challenge to target enhancements for GA operations.

However, the differences between GA and commercial aviation go further than these technical definitions. The FAA has different rules for GA operations (14 CFR Part 91) as opposed to commercial air carrier operations (14 CFR Part 121), commercial air charter and air taxi operations (14 CFR Part 135). Commercial aviation is more highly regulated from an operational perspective, especially in regards to the numbers and experience level required of the pilots and the maintenance of the airplanes. However, do these factors fully account for the marked difference in safety?

3.2.2 General aviation accident categories

Looking deeper into GA accidents, one finds that of the more than 30 possible categories that investigators at the NTSB can use to classify fatal accidents, 75% are classified with one of five

categories. Figure 18 shows that over 45% of the accidents are classified as LOC-I. LOC-I is defined by the CAST/ICAO Common Taxonomy Team (CICTT) (International Civil Aviation Organization, 2009) as the “Loss of aircraft control while, or deviation from intended flight-path, in-flight." Thus, the pilot intention (intended flight-path) is a major factor in the determination of this category.

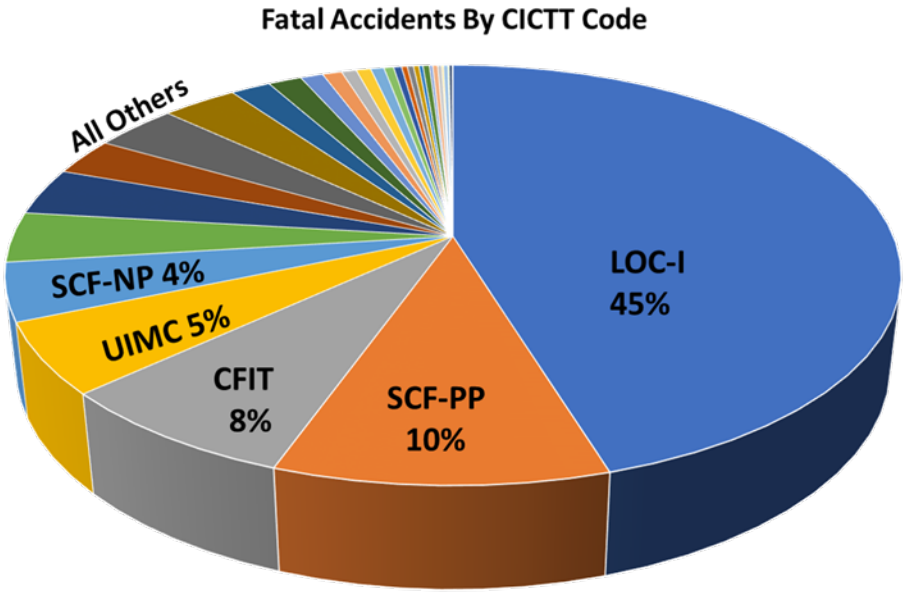

Figure 18. Classification of fatal accidents by CICTT code (CAST/ICAO Common Taxonomy Team (CICTT), 2011)

The next most frequently applied accident category is system component failure of the power plant (SCF-PP). This category refers to “Failure or malfunction of an aircraft system or component related to the power plant” (International Civil Aviation Organization, 2009). Thus, the fact that GA airplanes commonly have only a single power plant, with more relaxed maintenance requirements, may contribute to more accidents. The next two categories are CFIT, and UIMC. These accident categories relate directly to mistakes made by the pilot, and generally result in fatal accidents into terrain. The fifth most common is system component failure not related to the powerplant (SCF-NP). Here again is a GA risk factor relating to system redundancy and maintenance. Of these top five categories, three are directly related to pilot performance, and two are related to aircraft failure. However, the three top pilot related accident categories account for 58% of the total aircraft fatalities, and if one included all pilot related causes the percentage would be much higher.

3.3 Background

3.3.1 Ground collision avoidance system

So, what is to be done about the high number of pilot related fatalities in general aviation? Recently released research has shown that a GCAS may be an effective technological solution for as much as 26.2% of the fatalities in GA (Hook, Sizoo, & Fuller, 2022; Fuller & Hook, 2020). This works out to as many as 110 lives saved a year if every general aviation aircraft had a GCAS installed. Interestingly, these potential saved lives come from several different accident categories. The research predicts that over 95% of CFIT accidents can be saved with a GCAS, 16% of LOC-I, and 54% of UIMC. Other categories, such as LALT, are candidates for large percentage (95%) of GCAS saves as well.

Currently, GCAS are only available on larger aircraft with complex avionics packages. Automatic versions are only available in military aircraft such as the F-16 (Figure 19). Examples of warning systems include TAWS (Anderson, Jones, & Beamon, 2011), flown on many commercial aircraft, and Enhanced Ground Proximity Warning Systems (EGPWS) (Breen, 1999) flown on aircraft such as US Navy fighters. Going beyond warnings and including automatic control, automatic GCAS is operational on the US Air Force F-16 and F-35 aircraft and has been an overwhelming success (Swihart, et al., 2011). Since its installation, the system has saved the lives of eleven fighter pilots (Wetsig, 2021). In addition, versions of this technology have passed into flight testing for military transport category aircraft (Suplisson, 2015) and small unmanned aircraft with NASA's iGCAS project (Sorokowski, Skoog, Burrows, & Thomas, 2015).

Even with the success of military automatic systems, there are unfortunately no automatic systems that are currently available for use on GA and only a limited number of GA aircraft with TAWS, even though accident rates in GA are much higher. This reality has led the FAA to fund three separate but complimentary projects seeking to fast-track the development of GCAS for GA. This paper will explore the motivation and viability of implementing GCAS for use by general aviation (GA-GCAS).

Figure 19. US Air Force F-16 shown with an automatic GCAS system installed

3.3.2 Run time assurance

To understand design philosophies used for GCAS, it is important to know that GCAS can be generally described as an RTA system (Skoog, Hook, & Ryan, 2020; Nagarajan, Kannan, Torens, Vukas, & Wilber, 2021). The role of an RTA system is to ensure proper, timely handling between different control systems including ones for safety assurance such as an automatic GCAS. Before understanding the specific implementation of GA-GCAS, an understanding of a generic RTA system is required. Figure 20 (ASTM International, 2017) contains a diagram of a general high level RTA functional diagram. The Complex Function (human or auto-pilot) will be in control of the aircraft at all times unless the Safety Monitor deems an unfolding situation to be imminently dangerous. If danger is imminent, the system's switching functionality will give control to the Recovery Control Functions. These functions are generally autopilot functions that have a pre-determined envelope of operation or other mechanisms whose sole purpose is to return the aircraft to a safe state.

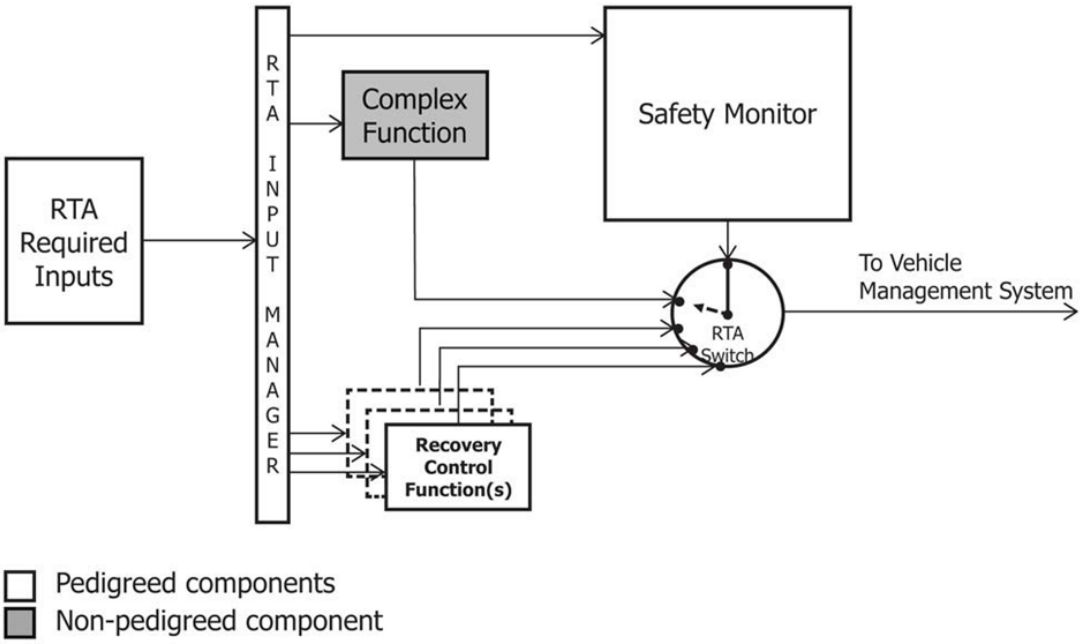

Figure 20. Real time assurance architecture from ASTM 3269

Applying the RTA concept to GCAS, the safety monitor becomes the GCAS monitor, which is always evaluating the future trajectory of the aircraft to determine if an imminent ground collision is likely. If so, the system gives control to the Recovery Control Function(s), which is the GCAS controller in this case. Figure 21 reflects this specific nomenclature for the GA-GCAS implementation. The monitor and controller are the pivotal components in GCAS.

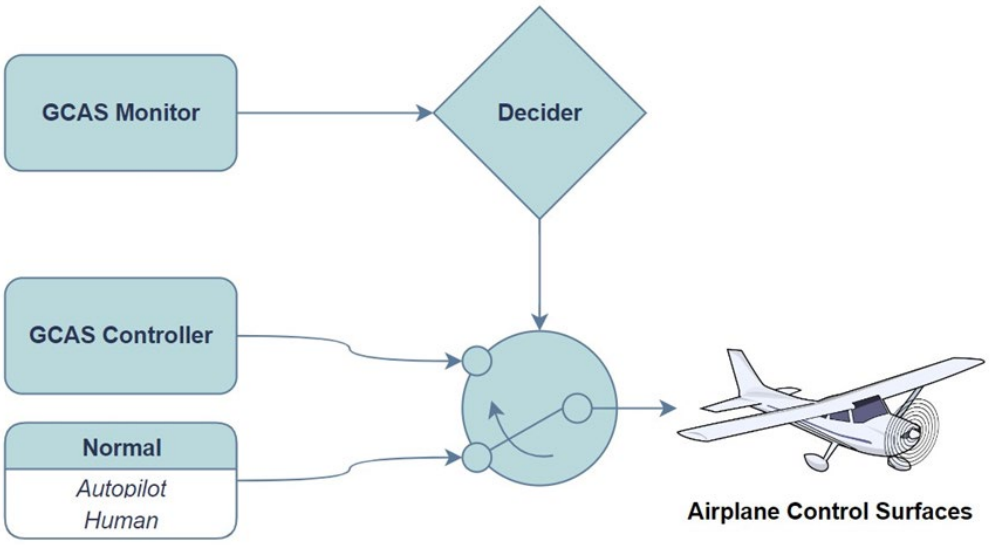

Figure 21. RTA decision structure for GA-GCAS

3.3.3 GCAS monitor

The GCAS monitor is generally composed of three subsystems. First, for all but the most rudimentary GCAS, a prediction of future trajectory is required. The prediction of the future trajectory requires aircraft-specific dynamic considerations to be incorporated. At the most basic level the trajectory prediction projects the current future state of the aircraft. This could include vertical velocity or the dynamics of a measured terrain separation sensor (such as a radar altimeter). More advanced predictions include control mechanisms, such as the initiation of a GCAS controller. In these cases, the system does not extend the current state, but instead assumes a new dynamic state will be brought on with the initiation of the new controller.

In addition to knowing the trajectory, information related to the aircraft's surroundings is also required. In particular, the proximity of terrain relative to the aircraft needs to be known. Generally, the local terrain model collects this information through active sensors, such as forward-looking radar or, for more advanced implementation, through a digital terrain database stored on the aircraft. By coupling the TPA and the Map Manager together, the third subsystem, the comparator, can recognize when a ground collision is a threat. If the difference between the TPA output and the local terrain model falls below a minimum distance threshold, the GCAS monitor will initiate a takeover from the complex function and give control to the GCAS controller.

3.3.4 GCAS maneuver selection

Previous implementations of GCAS, like the F-16's Auto GCAS, used only one possible recovery maneuver (a wings level climb). On the other hand, GA aircraft can consider three different recovery maneuvers due to its slower nature. Each of these maneuvers, listed below, will be continuously evaluated by the TPA to detect potentially fatal scenarios.

1. Left climbing turn
2. Wings level climb
3. Right climbing turn

If an imminent collision is anticipated after comparing the TPA with a section of terrain data, the TPA will flag the boolean need to avoid (NTA) value as true for the particular maneuver being evaluated. Each maneuver has an NTA associated with it and will be flagged "true" independently of each other.

3.3.4.1 Last man standing

Keeping in mind that three potential maneuvers are being evaluated at all times, the concept of last man standing now becomes a pivotal element in the system. Last man standing is a pseudo-algorithm that will be used to choose one of the three recovery maneuvers under consideration. Simply put, the algorithm states that the answer to a system of choices is the final choice to enter the discussion.

By default, each maneuver's NTA value is false, but can be flagged true by a determine need to avoid (DNA) function. As a fatal scenario begins to unfold, each controller's NTA will independently flag as true if the defined altitude threshold is surpassed. The final maneuver's NTA to be flagged as true will be selected and used for the collision avoidance maneuver.

There are two primary benefits that arise from using last man standing to select a recovery maneuver. Firstly, it maximizes the amount of time that the primary controller has before a recovery maneuver is required. Secondly, it increases overall safety by increasing the takeover threshold. This prevents any unnecessary seizure of control from the primary controller. It also decreases nuisance by avoiding a hyper-sensitive takeover threshold.

3.3.4.2 Defining the takeover threshold

Figure 22 depicts a single-maneuver TPA encountering a potentially fatal scenario. In this scenario, the GCAS monitor will flag the NTA as true, allowing the GCAS controller to take over the aircraft. There are a few critical variables that influence this decision.

- Closest point of approach (CPA): shortest distance between the TPA and the terrain data.
- Terrain clearance buffer (TCB): predefined, minimum altitude threshold between the airplane and the terrain.

The trajectory prediction produced by the TPA is shown as the dotted line in the image. The CPA is also shown, which is any distance between the dotted line and the red terrain data. The TCB can be seen in blue. Notice that the trajectory clears all but one TCB instance. In this example, the maneuver being evaluated will flag NTA as true because it surpasses the minimum altitude threshold. Because NTA is evaluated as true, the GCAS monitor will direct control from the primary controller to the GCAS controller to execute a recovery maneuver.

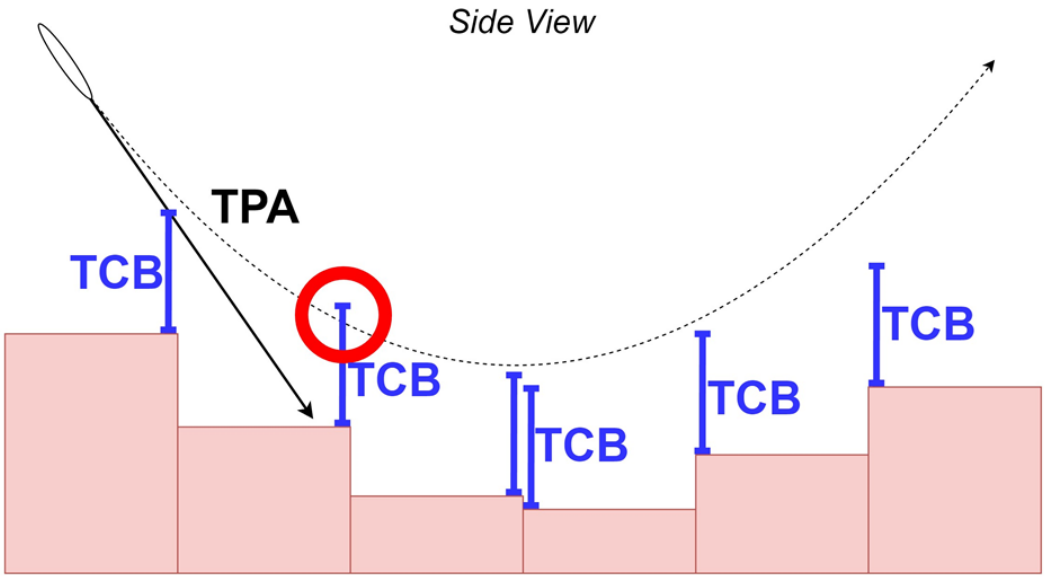

Figure 22. Decider in the event of a fatal scenario

3.3.5 The GCAS controller

After the GCAS monitor decides that a takeover of the aircraft is required, the GCAS controller may have authority to control the aircraft and execute the recovery maneuver. For the 3-maneuver GA-GCAS being proposed (Figure 23), all maneuvers have a pull-up component and a bank angle component. To perform these recovery maneuvers, control in both the lateral and longitudinal axes is required. The GCAS controller can be separated into these axes, called the lateral controller and the longitudinal controller. The lateral axis is concerned with the bank angle and sideslip of the aircraft. For the scope of GA-GCAS, the controller will command a zero bank angle for wings-level or a non-zero bank angle for bank left or right. It also commands

zero sideslip angle. The longitudinal controller is responsible for getting the airplane away from the ground. The first phase of the maneuver is the initial pull-up, or recovery pull. In this phase, it is important to not "over G" the aircraft. The second phase, called the zoom climb, trades airspeed for altitude. After reaching a desirable state, the recovery enters phase three, executing a steady sustainable climb to increase the distance from the ground.

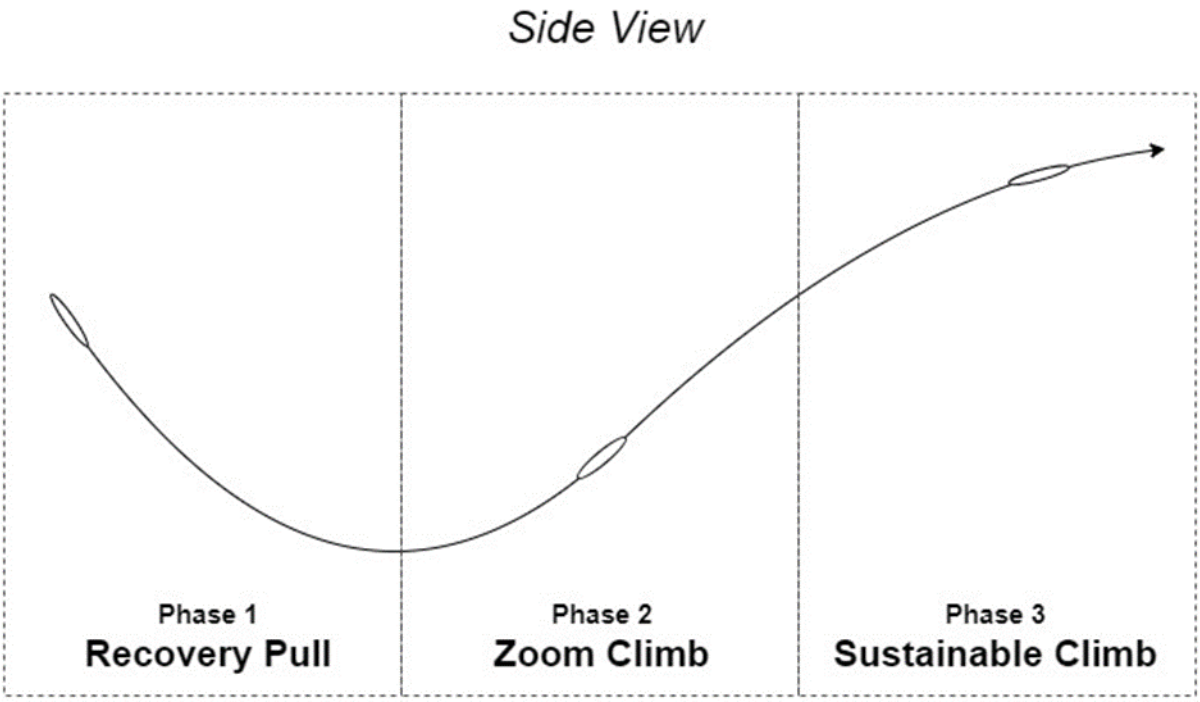

Figure 23. Three maneuver phases

It should be noted that due to structural concerns related to performing a loaded roll, GA-GCAS will initiate a roll first, then a pull-up second. Both will not be commanded simultaneously. This functionality is not in the current implementation of the system, but future revisions will reflect this behavior.

3.4 General aviation GCAS controller dynamic considerations

The dynamics of the vehicle must first be well understood to develop an effective control system for a vehicle. Without this, the dynamic response when being conducted by the controller cannot be correctly interpreted, making the controller tuning process extremely difficult. Furthermore, accurately considering the vehicle's dynamics will allow the developer to know which controllers can be more aggressive without jeopardizing the safety of the vehicle. To develop an in-depth understanding of aircraft dynamics, Beard & McLain (2012) and Stevens, et al. (2015) can be studied in detail. This material can be supplemented with material from the *Pilot's*

Handbook of Aeronautical Knowledge (Federal Aviation Administration, 2009) for factors such as asymmetric thrust and the gyroscopic effect of the propeller.

3.4.1 Lateral considerations

Lateral dynamics affect the orientation of the aircraft about the X and Z axes. For a ground collision avoidance scenario, sideslip angle, β, and bank angle, ϕ, are what we want to control. Course angle, χ, control is also normally handled by the lateral controllers, but it is not needed in a ground collision avoidance scenario.

Since the aircraft must be able to change its heading for some maneuvers, roll control is needed. By banking the aircraft to the left or to the right from center, the lateral component of the lift produced by the wings will change. This will cause the flight path vector of the aircraft to accelerate in the direction of the bank. This then results in a circular motion in the lateral plane from the perspective of the inertial frame. However, for a properly coordinated turn, the sideslip induced from rolling the aircraft must also be considered.

Nonzero sideslip causes the nose of the aircraft to point in a slightly different direction than the vehicle's direction of travel. Having an excessive sideslip angle is detrimental for the quality of the control of the aircraft and puts the aircraft in danger of a spin or other dangerous condition. Furthermore, a larger sideslip angle will induce more drag due to a larger surface area of the aircraft being incident to the relative wind.

The amount of induced sideslip on an aircraft depends on several factors, with the main ones being thrust setting, propeller characteristics, angle of attack, aileron deflection angles, and environmental conditions.

Internal factors relating to the characteristics of the aircraft are mainly responsible for inducing sideslip. These factors most commonly manifest when significant rates, angles, or thrust settings are commanded. When a positive pitch moment is applied to a clockwise (relative to the pilot) spinning propeller, a leftward yawing moment is induced on the aircraft. This moment is caused by the propeller being forced out of its plane of rotation, causing a gyroscopic action to apply a torque 90° from the direction of rotation. Similarly, when an aircraft with a clockwise-spinning propeller has a positive angle of attack, a leftward yawing moment is also applied due to asymmetric thrust. Asymmetric thrust is caused by the downward motion of the propeller being incident to a larger air mass than the upward motion. This causes the downward motion to create more lift, producing more thrust to the right of the aircraft.

Rolling moments and excessive sideslip angles can also excite a Dutch Roll dynamic mode that will cause the aircraft to bank to the left with positive sideslip or induce a negative sideslip when the aircraft banks left, and vice-versa. A left rolling moment is caused with positive sideslip due to the right wing leading into the relative wind, meaning the trailing left wing is creating less lift. Conversely, when the ailerons deflect to cause a rolling moment to the left, the right wing is creating more lift than the left wing and, therefore, more drag. This induced drag differential then causes a yawing moment that pushes the right wing back, inducing negative sideslip.

Due to this dynamic mode, bank and sideslip control are directly linked. If a new bank angle is commanded, then the beta controller must work to keep the sideslip zeroed. Furthermore, sideslip induced by any change in throttle setting, angle of attack, or pitch moment must also be accounted for. Because of the many ways in which sideslip can be induced on an aircraft and its effects on the quality of control, sideslip control is very important.

3.4.2 Longitudinal considerations

Longitudinal dynamics concern the orientation of the aircraft about its y-axis. For a GCAS controller, maneuvering the aircraft to its maximum published longitudinal performance limits is necessary. However, due to imperfections with controller tuning and the possibility of unpredictable atmospheric conditions, a buffer from these limits must be considered. This is a delicate balance, as a longitudinal controller being given an excessive buffer from the performance limits will result in the aircraft losing more altitude than necessary.

Given that the GCAS controller must be able to push the aircraft close to its published performance limits, the authors of this paper decided to track the following for the longitudinal controllers:

- Angle of Attack (α)
- Calibrated Airspeed (V_c)
- G-loading
- Vertical Flight-Path Angle (γ)

The first three of these variables were chosen to be tracked because there are published maximums for these numbers for the C172. If these maximums are surpassed, either stall conditions or airframe damage may occur. The vertical velocity vector of the aircraft is tracked for near-stall airspeed scenarios.

Angle of attack must be tracked because it is directly related to the amount of lift that is created by the wings. If the angle of attack surpasses a critical value, then the amount of lift begins to decrease as angle of attack continues to increase. This is commonly referred to as an "aerodynamic stall". By tracking angle of attack and tuning our controller appropriately, we can ensure that the aircraft does not enter an aerodynamic stall state given that the aircraft remains above the C172's normal operating range of 48 KIAS.

Airspeed is tracked to ensure that the aircraft does not rise above the C172's never exceed speed of 163 KIAS and stays above the cruise-configuration stall speed of 48 KIAS. . Tracking airspeed also allows the controller to capture the best rate of climb and the best angle of climb airspeeds. This allows the aircraft to either gain altitude as quickly as possible or gain as much altitude while covering as little horizontal distance as possible.

Should the aircraft's airspeed begin to fall below or close to the stall speed while it is in a cruise configuration, the aircraft would need to quit climbing. If the aircraft is at full throttle and airspeed is approaching a stall state, the common instinct from a pilot may be to point the nose down to more quickly regain airspeed. However, any altitude loss should be avoided in a ground collision avoidance scenario. Therefore, the vertical velocity vector is tracked for scenarios where the aircraft is close to stall speed.

G-loading was tracked to ensure that the aircraft was not overloaded during pull-up maneuvers. G-loading is the amount of acceleration being applied to the aircraft in the Z (down) direction in the body frame. This is measured in units of G (gravitational acceleration) where 1G is approximately 9.81 $m/s^{2.}$

When an aircraft is flying straight and level, it is experiencing 1G of G-loading due to the natural gravitational pull of Earth. When pulling a high angle of attack, more lift is created, which then causes more G-loading. This is due to the force of lift overcoming the Z components of the weight and drag of the aircraft, causing an acceleration in the direction of the lift vector. As the lift vector continues to increase past the drag and weight threshold, a higher loading is put on the airframe. Too high of a loading, and the airframe will be overly stressed and either become damaged or structurally fail.

Overloading the airframe is only an issue whenever the aircraft's airspeed exceeds the maneuvering airspeed (VA). Below VA, a full elevator deflection angle can be commanded without overloading the airframe. This is because the critical angle of attack will be reached before the maximum load is put on the airframe when at or below maneuvering airspeed. Maneuvering speed is a function of the aircraft's weight. A lighter aircraft will have a lower

maneuvering speed and vice-versa. At the same airspeed, a lighter aircraft will pull less angle of attack to overcome the Z components of weight and drag when compared to a heavier aircraft. A lighter aircraft flying at a heavier aircraft's maneuvering speed will overload its airframe if a full elevator deflection angle is commanded.

3.4.3 Auto-pilot design

With a thorough understanding of the dynamics of an aircraft, an effective control system can be designed to control the orientation of the aircraft. Proportional -Integral-Derivative (PID) controllers comprise the GCAS controller. An in-depth understanding of PID control can be gained by studying Franklin et al. (2002).

3.4.4 Lateral auto-pilot

In the lateral axis, the two state variables that are controlled are sideslip angle (β) and bank angle (ϕ). The β controller is required to keep the sideslip angle at zero. To do this, it must react quickly to turbulent wind or other internal factors that would induce sideslip. β is controlled by commanding the deflection angle of the rudder. The ϕ controller is required to command bank angles somewhat aggressively, as well as achieving and maintaining an accurate steady-state value. ϕ is controlled by commanding the deflection angle of the ailerons.

Unlike with the system of controllers that control the longitudinal axis orientation, these two controllers are both the highest-order about their respective axes. Together, their main functionality is to perform coordinated turns. The ϕ controller rolls the aircraft to the desired angle while the β controller counteracts the adverse yaw. The β controller is also used to account for sideslip induced from asymmetric thrust, gyroscopic effects, and turbulent wind.

3.4.5 Longitudinal auto-pilot

Having both good longitudinal control and lateral control is important for ground collision avoidance scenarios. However, a greater amount of time was focused on the development of the longitudinal controllers. This is due to heightened time constraints associated with attempting to achieve a positive climb-rate attitude in a ground collision avoidance scenario.

To maximize longitudinal performance and ensure that the aircraft remains out of a stall state, an angle of attack controller was chosen as the lowest level longitudinal controller. With the α controller, the maximum angle of attack can be set, therefore ensuring that an α past the critical angle of attack cannot be commanded. Combined with a properly tuned controller, we can be confident that the critical angle of attack will not be surpassed in standard atmospheric conditions.

While the α controller ensures that a stall angle of attack cannot be commanded, the controller itself cannot directly control the airspeed of the aircraft. To prevent the airspeed from entering a dangerous state, and to capture optimal climb rates, a calibrated airspeed (*Vc*) controller was introduced. This controller keeps the commanded airspeed within safe margins and facilitates the capture of best-rate and best-angle climbs.

However, having the airspeed controller as the sole top-level controller in the longitudinal axis presented a problem. The airspeed controller commands an α, which then commands an elevator deflection angle, meaning the elevator deflection angle is based on the error in the aircraft's calibrated airspeed. Therefore, if the aircraft's airspeed is below the commanded airspeed, the nose will be pointed downwards so that altitude can be exchanged for airspeed. This behavior, caused by the Vc controller, is undesirable for a ground collision avoidance scenario. To ensure that the aircraft would not enter a nose-down attitude when the actual airspeed was less than the commanded airspeed, a vertical flightpath controller was introduced. This controller was placed in parallel with the Vc controller so it would override any nose-down commands. However, placing two high-level controllers in parallel presented another issue. A system of controllers was needed to facilitate switching between two controllers in real time.

The hierarchy of the system of controllers can be seen in Figure 24. In this system, both the Vc and γ controllers return a commanded α, but only one is sent to the α controller. If Vc is greater than or equal to the best-rate airspeed, the commanded α calculated by the Vc controller is used. Else, the commanded α calculated by the γ controller is used. The resulting behavior is, if the aircraft is below the best-rate airspeed when the ground collision avoidance scenario begins, the aircraft will start the scenario by holding a 0° γ. Once Vc is greater than or equal to the best-rate airspeed, the Vc controller then supplies the commanded α calculations. This prevents the Vc controller from commanding a nose-down attitude to gain airspeed.

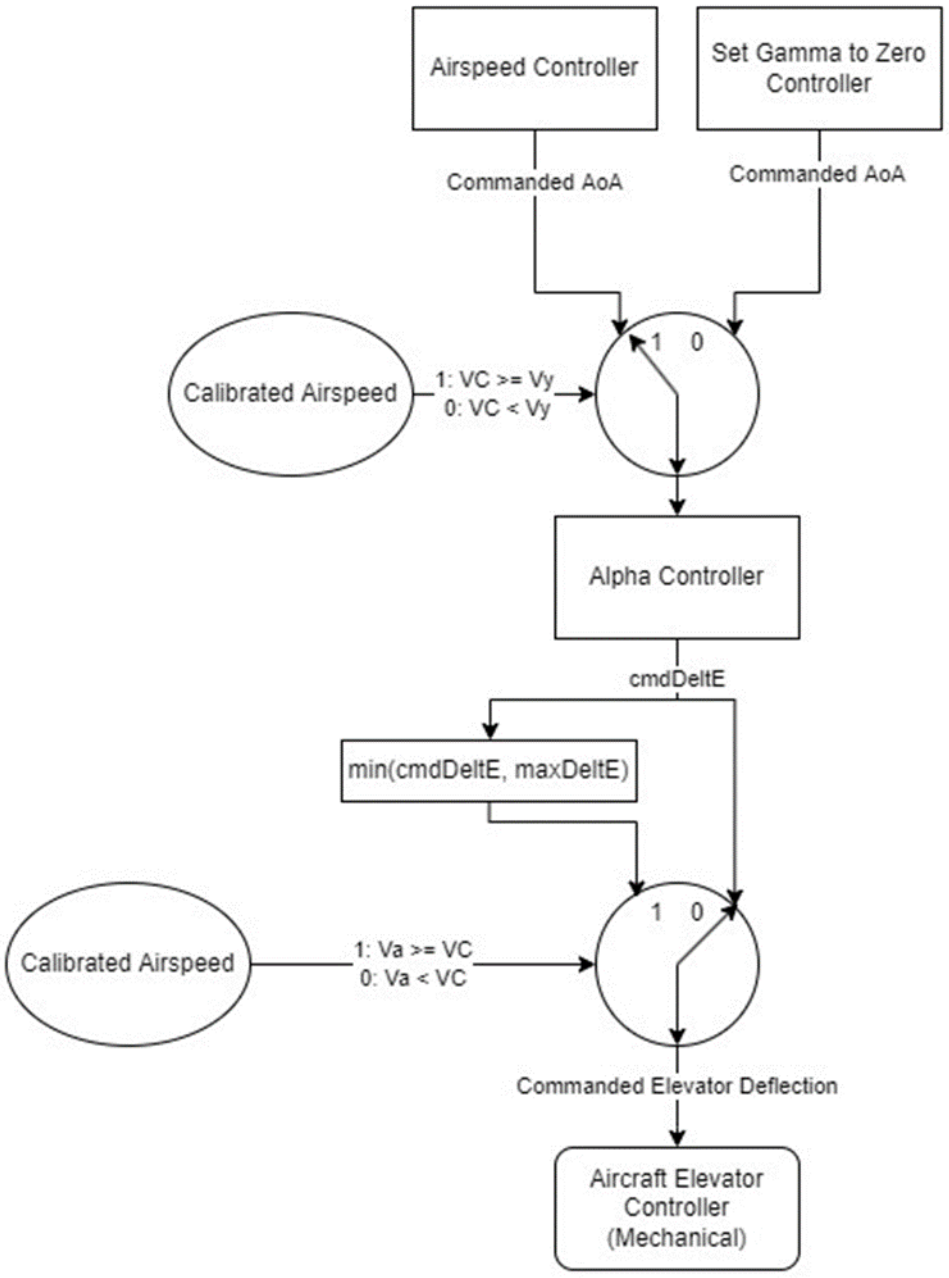

Figure 24. Structure of the longitudinal controller system

In the actual implementation of this system, only one controller is running at any time. This prevents the inactive controller's integral value from spooling up while it is not in use. However, an issue occurred when switching between the controllers where the commanded α would suddenly decrease at the switch. This was due to the inactive controller's integral value being zero at switching time. This would result in a brief nose-down attitude and would cause a noticeable loss in altitude. To remedy this, an integral bias was added when the controllers switch, allowing the aircraft to maintain its altitude.

After the commanded α is computed from either the Vc or γ controller, a commanded deflection angle for the elevator (δe) is calculated by the α controller. Before the commanded elevator deflection angle is sent to the aircraft, it is first run through a G-limiter to prevent stressing the airframe beyond its published limits. The maximum allowable load in the -Z direction is 3.8G for a C172 (Cessna 172N pilot's operating handbook, 1977). A buffer of 10% was used for the G-limiter, giving a load limit of 3.42G.

To prevent over-G, an accurate relationship between the amount of Gs pulled, airspeed past maneuvering speed, and elevator deflection angle needed to be known. To find this relationship, data generated from JSBSim was used, since no such relationship or data can be found in the Cessna C172 *Pilot Operating Handbook* (POH) (1977). The data used to interpolate the equation was found by initializing the aircraft at airspeeds varying from 100 to 160 knots calibrated airspeed (KCAS) in 10 knot increments. For each starting airspeed, the amount of load applied to the airframe for different elevator deflection angles was sampled. With this data, the maximum amount of elevator deflection was found as a function of Vc and VA and can be found in Equation 6.

$$\delta_{e,\max} = -8.2 \cdot 10^{-5}(V_c - V_A)^2 + 1.07 \cdot 10^{-2}(V_c - V_A) - 0.35 \quad (6)$$

3.5 Auto-pilot performance

To test the response of the GCAS controller, a series of tests were created that would place the aircraft in various starting orientations, then initiate a recovery maneuver. While the controller was designed to work within a $\gamma \in \pm 20°$, $V_c \in [50,120]$ KCAS, and $\phi \in \pm 30°$ envelope, some tests measure the response of the aircraft when outside of this envelope. In all of the tests, the response of the aircraft as conducted by the GCAS controller was favorable, even in most of the outside-the-envelope tests. The only major shortcoming was with the G-limiter. Past 140 KCAS, the limiter would still allow the aircraft to overload by about 0.2G. Despite this, the fundamental

behavior of the aircraft under instruction by the GCAS controller was favorable, though some tuning is required.

3.5.1 Lateral

3.5.1.1 Sideslip controller analysis

To test the performance of the β controller (Figure 25), step wind inputs were applied to the aircraft while it held a steady and level flight with a constant course angle. The magnitudes of the winds applied, which varies by test, are 15 knots, 25 knots, and 35 knots. The winds are applied at 10 and 15 seconds. The 10 second wind comes from the east, while the 15 second wind, equal in magnitude, comes from the west.

The eastern wind caused the aircraft to experience instantaneous 10° to 25° of sideslip, varying with the magnitude of the wind. The aerodynamics of the aircraft alone would counteract this sideslip but would cause oscillation about the z-axis. In all cases, the sideslip controller dampens this response in about three seconds. The western wind, applied at 15 seconds, has a more profound effect on the initial magnitude of the sideslip angle of the aircraft. This is due to the β controller having already reached a steady-state from the eastern wind and having to account for an equal and opposite change. The initial sideslip angle on the western wind application ranges from about -22° to -36°and settles to 0° in about four seconds.

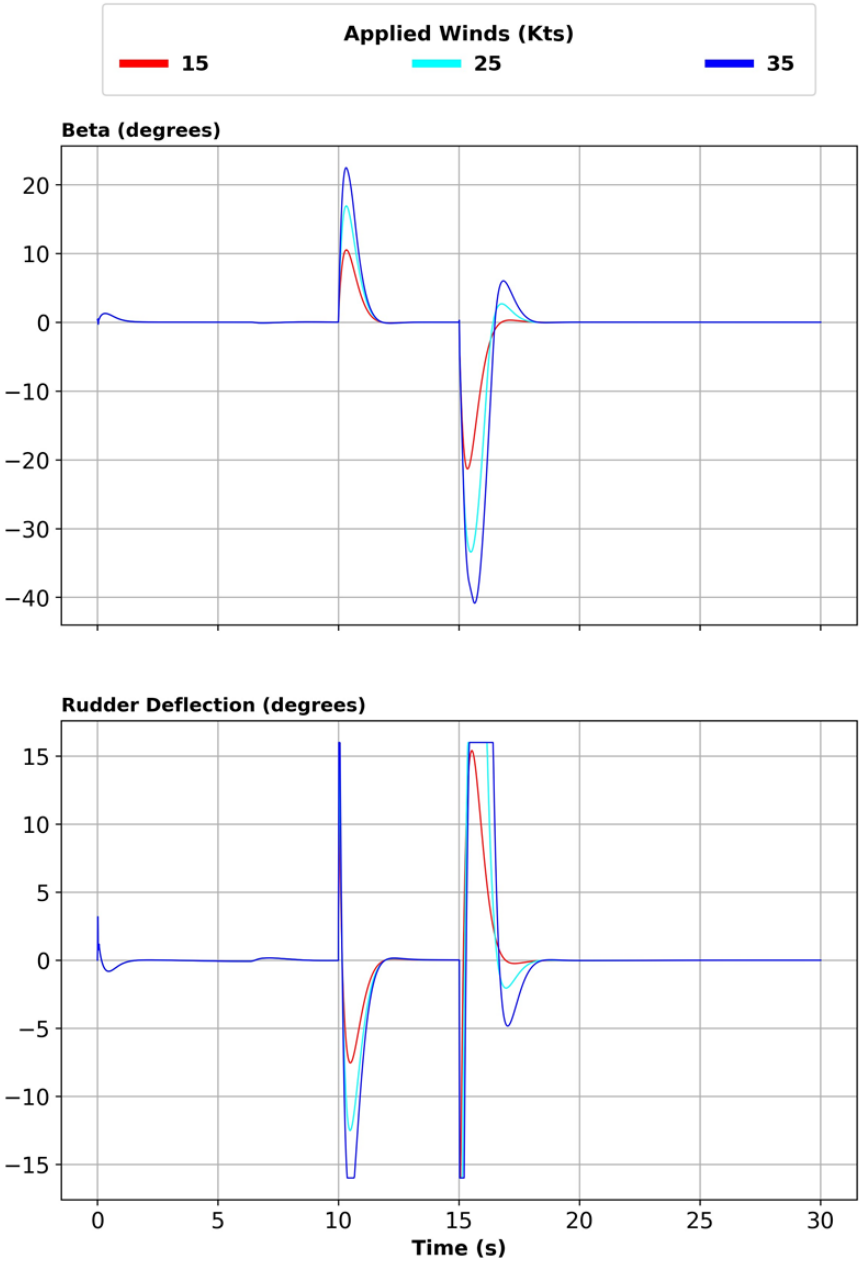

Figure 25. β controller response to opposing disturbances

3.5.1.2 Roll controller analysis

The behavior of the ϕ controller was observed by commanding the aircraft to different bank angles while the aircraft maintained a constant altitude. The aircraft starts the test with wings-level, and several roll doublets are commanded with maximum bank angles of ±20°, ±35°, and ±60°. The response from these tests is seen in Figure 26.

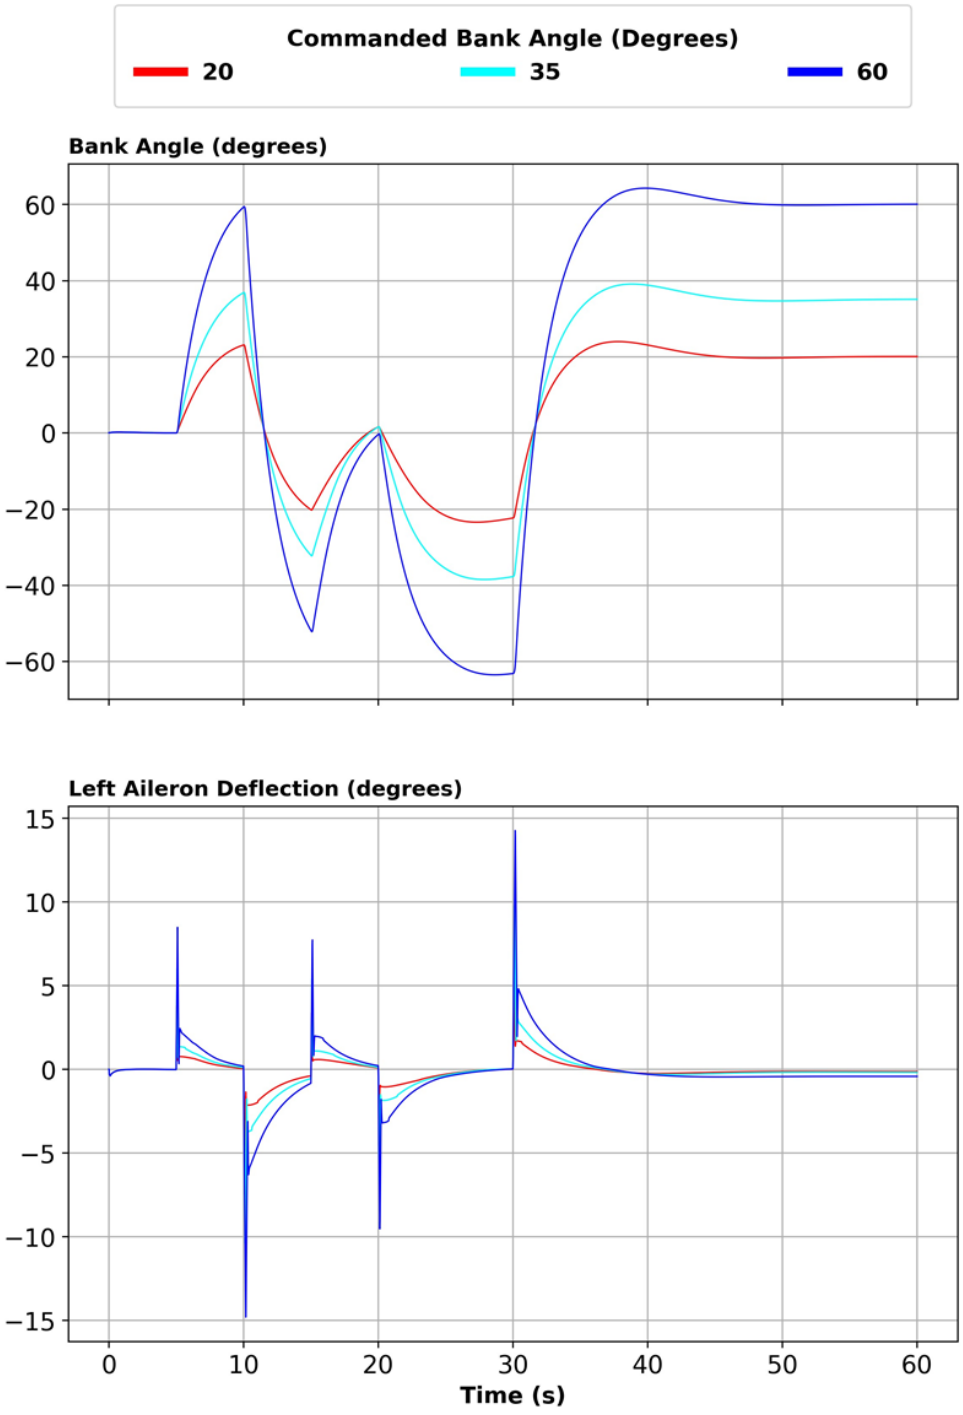

Figure 26. ϕ controller response to roll doublets

3.5.2 Longitudinal test setup

To test the longitudinal performance of the GCAS controller, tests were performed that vary initial airspeed, initial vertical flight path angle, initial altitude, and the type of maneuver being commanded.

Initial airspeeds for the tests varied from 50-150 KCAS in 10 knot increments. Initial vertical flight path angles were +20°, -20°, 0°, and -80°. Maneuvers tested include a level pull-up, a +35° bank pull-up, and a -35° bank pull-up. The ultimate objective of the longitudinal controller system in these tests is to acquire and maintain the best-rate airspeed for the aircraft while throttle is at 100%. Since the responses from each of the maneuvers are similar in their longitudinal response, only the 0° bank pull-up results will be shown for the sake of brevity.

3.5.3 Longitudinal analysis

The first tests shown are level pull-up maneuvers performed with initial airspeeds ranging from 50 to 90 KCAS, all with initial vertical flight path angles of -20°. The response observed in this test can be seen in Figure 27. In all cases, recovery occurs in less than four seconds, resulting in a maximum altitude loss for all cases of around 80 feet. The attitude-hold functionality of the longitudinal controller can be seen in this figure when the starting airspeed is below 75 KCAS.

The tests that start with airspeeds below the best-rate airspeed of 75 KCAS start with the γ controller being the active highest-order longitudinal controller. Once the airspeed is at or above the best-rate airspeed, the Vc controller becomes the active highest-order longitudinal controller. This behavior can be seen best in Figure 27 in the 50 KCAS start case. In this start case, the controller first recovers from the nose down attitude, then maintains a level vertical flight path while airspeed is regained. Once the airspeed reaches 75 KCAS, the Vc controller takes over and achieves, then maintains, the best-rate airspeed.

The longitudinal controller was designed to work within a $\pm 20°$ γ and 50-120 KCAS envelope. However, it was determined that if this envelope was pushed, some otherwise uncovered behavior of the GCAS controller could be discovered. To gain a more expansive understanding of how the GCAS controller behaved in ground collision scenarios, tests were performed to measure the response of the aircraft when operating outside this envelope under the instruction of the GCAS controller.

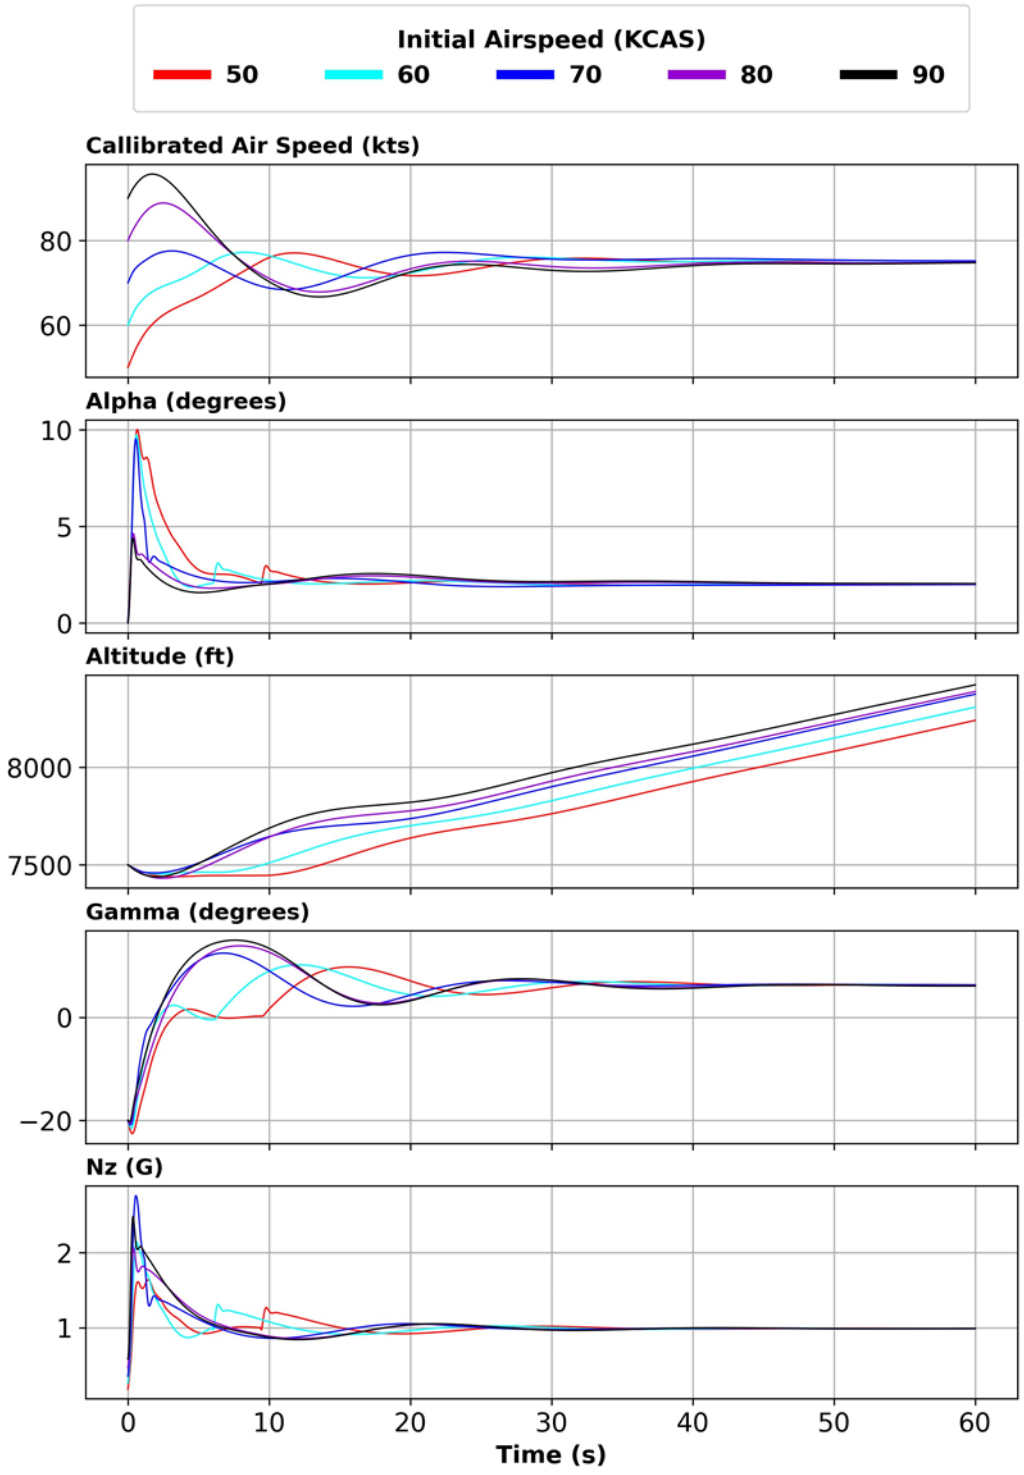

Figure 27. Level pull-up maneuver being performed for 50-90 KCAS initial airspeeds with an initial vertical flight path of -20°

Several tests were run that start the aircraft outside of the GCAS controller envelope. The first one has the aircraft start in a steep dive at -80° γ with airspeeds ranging from 50-110 KCAS. The response of this test is shown in Figure 28. In this scenario, the aircraft recovers from the nose-down attitude in about 4-5 seconds depending on the initial airspeed. In the 110 KCAS start case, the aircraft reaches a maximum airspeed of 140 KCAS, which is around 20 KCAS below the published never-exceed speed for the C172. Furthermore, the aircraft is G-limited to about 3.5G when pulling up at 140 KCAS. Altitude loss in this scenario ranges from around 300 feet to around 550 feet depending on the starting airspeed. At around the 12 second mark, the airspeed from the 110 KCAS start case drops below 55 KCAS, and the γ controller becomes the active highest-level longitudinal controller. The γ controller levels the vertical flight path angle of the aircraft until the best-rate airspeed is regained at about 25 seconds.

In the next set of outside-the-envelope tests, the aircraft was placed with an initial γ of 0° with initial airspeeds ranging from 100 - 150 KCAS. The results of these tests can be seen in Figure 29. Zoom climbing behavior can be seen in this figure, especially for the higher speed start cases. On the 150 KCAS start case, a G-loading of around 4G can be seen, which exceeds the C172's published G-limit by 0.2G. This is due to the G-limiter only being designed to handle up to 45 KCAS of overspeed, and the 150 KCAS start case starts with around 60 KCAS of overspeed.

The last extreme scenario test performed was a repetition of the -20° γ start test, but with higher start speeds. Seen in Figure 30, the aircraft started with a γ of -20° and airspeeds ranging from 100 - 150 KCAS. Like in the previous test that also had initial airspeeds ranging from 100 - 150 KCAS, the aircraft experienced over-G by around 0.2G for the 150 KCAS start case. Despite starting at such high speeds with a nose-down attitude, only around 100 feet of altitude is lost. This is due to the quick recovery from the nose-down attitude, which happens in less than four seconds for all start cases.

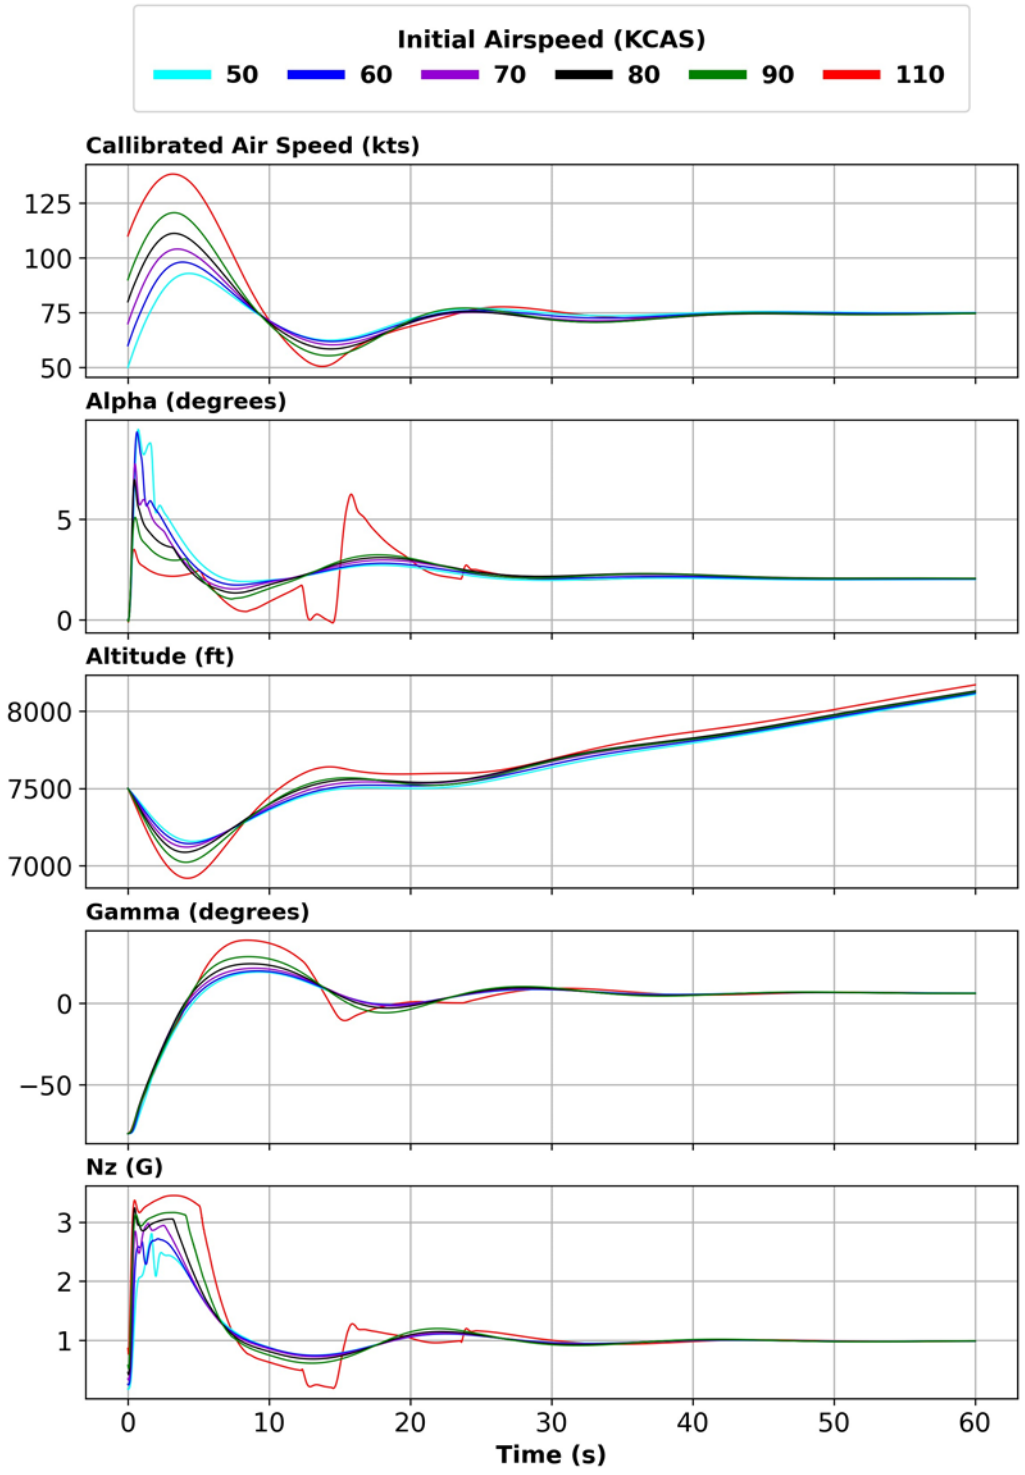

Figure 28. Level pull-up maneuver being performed for 50-110 KCAS initial airspeeds with an initial vertical flight path of -80°. range for a GCAS controller

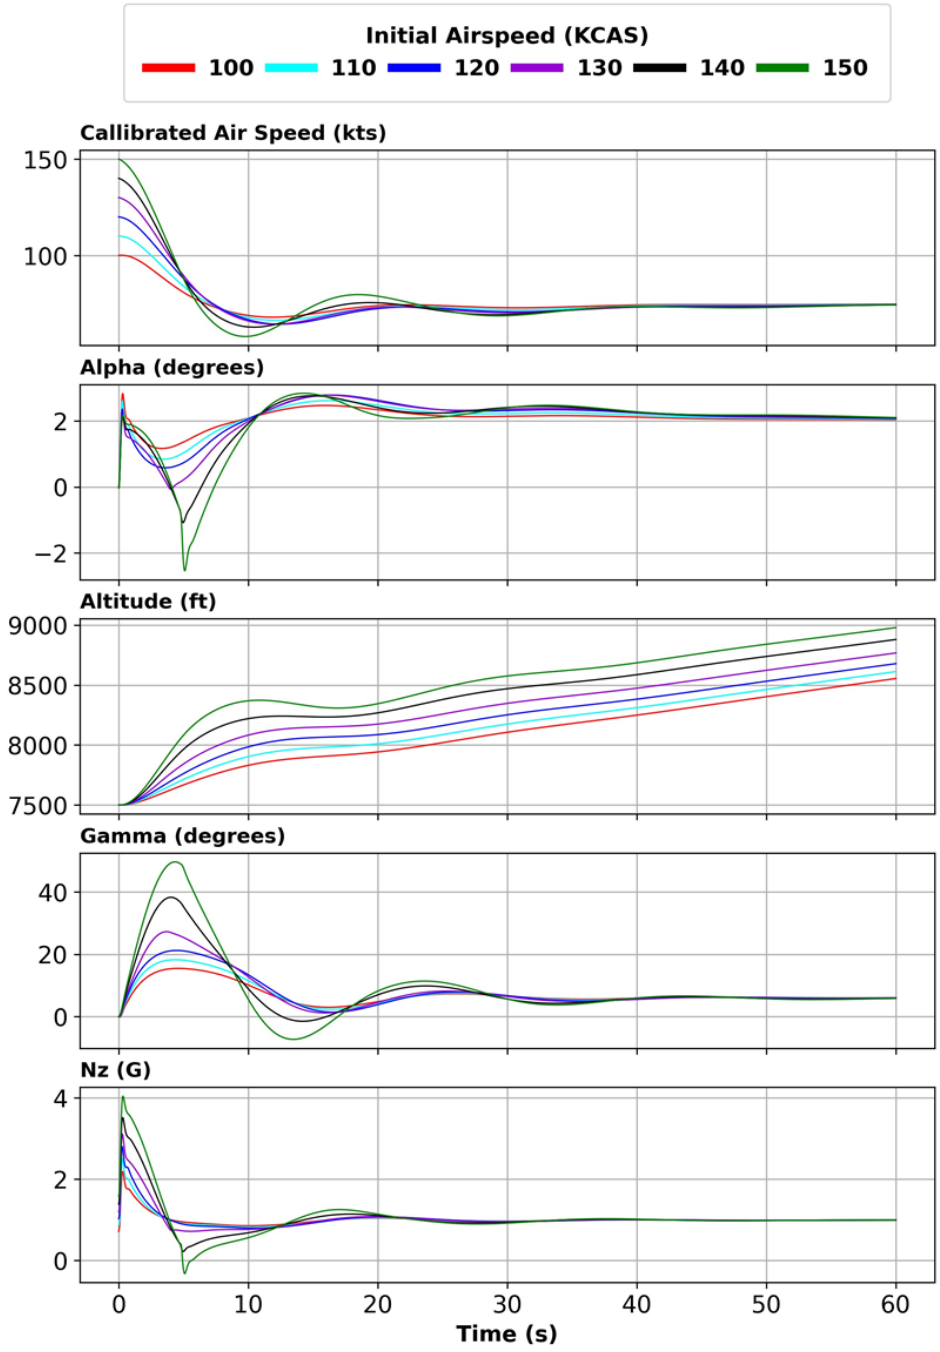

Figure 29. Level pull-up maneuver being performed for 100-150 KCAS initial airspeeds with an initial vertical flight path of 0°

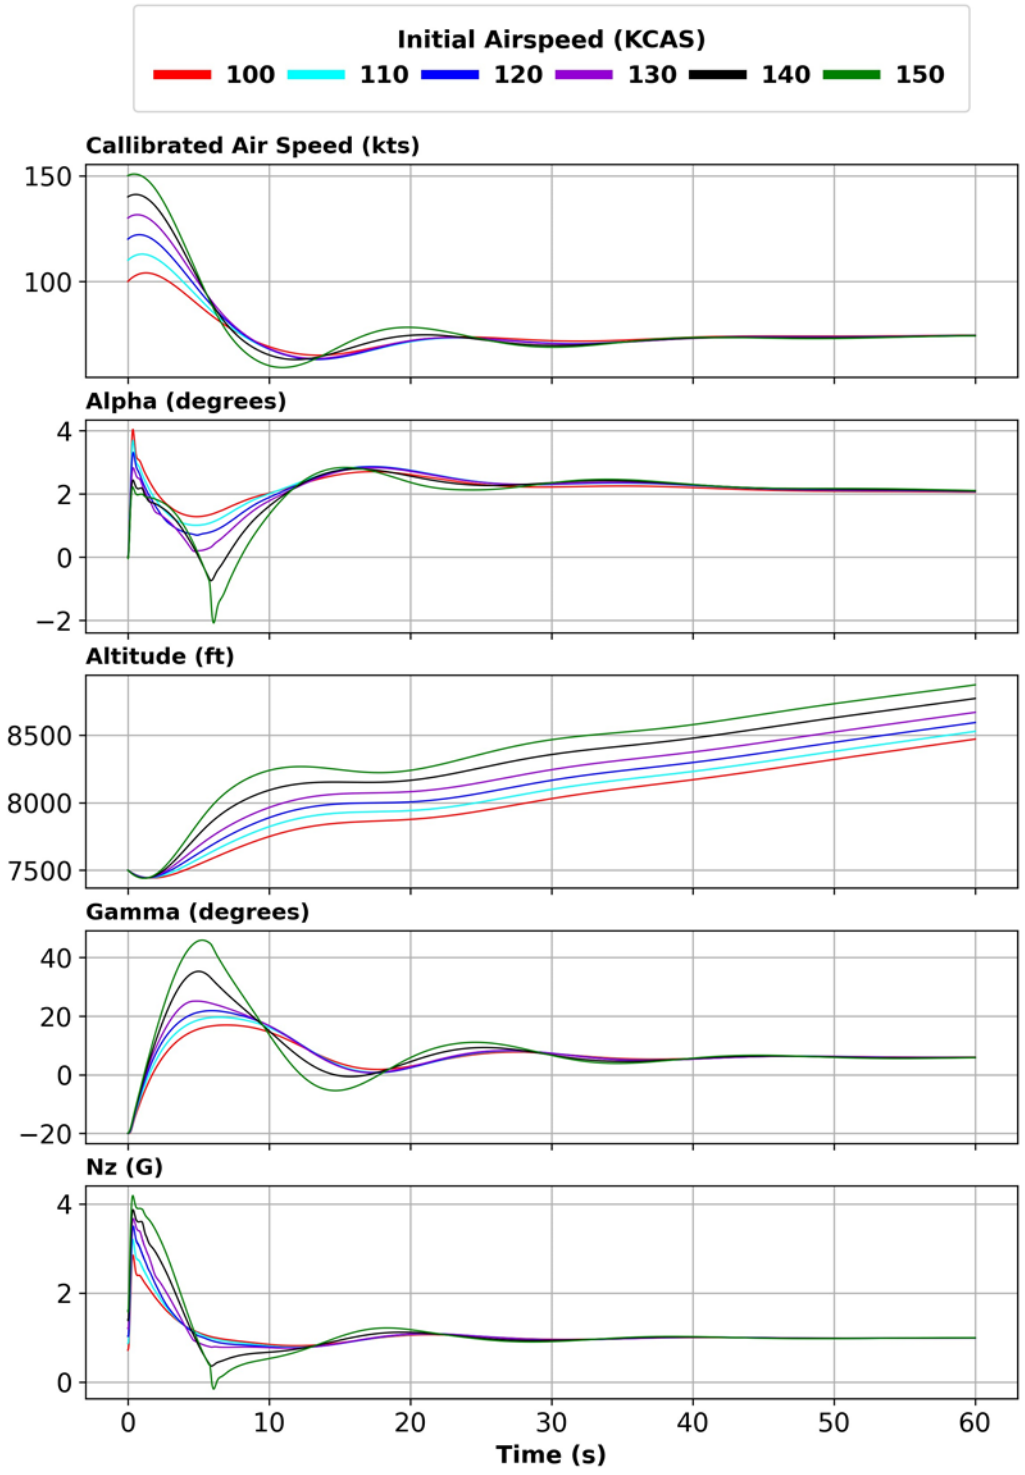

Figure 30. Level pull-up maneuver being performed for 100-150 KCAS initial airspeeds with an initial vertical flight path of -20°

3.5.4 Analyzing maneuvers with the TPA

To fill this void and create a suitable metric for measuring the effectiveness of the auto-pilot, we propose using the TPA to evaluate the maneuvers flown by the GCAS controller to avoid terrain. In order to do this, deviations from the TPA are measured in a statistical manner. In this case, the TPA predictions will be considered conservative estimates. Therefore, any maneuver that under performs the TPA by more than the terrain clearance buffer will be considered unacceptable. On the other hand, any maneuver that exceeds the TPA will be considered acceptable, with any maneuver in between considered marginal.

On the other end of the analysis spectrum, we propose a deviation metric for evaluating auto-pilot performance above the TPA. This metric requires a comparison of the TPA and auto-pilot trajectory at specific maneuver distances and the deviation will be allowed to increase at greater distances. For instance, if a positive 30-foot deviation from the TPA was allowed at 500 foot downrange from the initialization point, a larger value (maybe positive 100 foot for example) may be allowed at 2000 foot downrange. Similarly, deviation is expected for horizontal movement of the aircraft. Both horizontal and vertical deviation need to be monitored to ensure that maneuvers will neither under or over avoid, potentially causing another collision scenario.

3.5.5 TPA tuning

For the TPA to correctly predict the trajectory of a given aircraft, the TPA model must to be tuned to the specific aircraft and auto-pilot combination. In this particular case a test matrix was developed, using achievable control targets for a Cessna 172. The test matrix was then run using an open-source flight dynamics engine, JSBSim. Data collection of the flights was delayed until initial startup transients were minimalized. Once the transients had been minimized the GCAS maneuver was initiated, and data collection began. The maneuver continued until a steady state was reached (the aircraft rates reached at or around zero). After steady state was achieved, the test instance was terminated, and the next test condition began with a new state configuration. Once the full test matrix had completed, the full set of flight data was compiled through a tuning program. The results were examined, tested, and then re-tuned until the TPA closely matched the flight characteristics and maneuvers of the modeled aircraft.

3.5.6 Evaluation results

. Test cases to evaluate the autopilot performance were developed using the same parameters varied during the tuning process of the TPA. These included varying bank angles, angle of attack, initial altitudes, airframe weight, and indicated airspeed.

Following the evaluation method outlined, the aircraft's altitude over cumulative distance from the initiation point was examined. In all cases it was found that the altitude gain of the GCAS controller outperformed that of the TPA. Once the vertical performance of the GCAS controller was verified, the horizontal components were then examined. Similarly the GCAS controller performed tighter turns when conducting the maneuver.

After both the vertical and horizontal components were considered, the horizontal deviation from the TPA was determined. From the deviation measurements, a standard deviation from TPA over cumulative distance was found. Sub-sampling from the created chart at 100 feet after initiation of the standard deviation for a left emergency turn, with initial bank, was about 2.5. At 900 feet the standard deviation increased to nearly 6.5.

While the standard deviation at longer ranges increases greatly, the GCAS controller performed with much better avoidance margins than the TPA had predicted. This suggests that the TPA needs to be further tuned to the aircraft's maneuvers to provide a more accurate prediction.

3.6 Takeaways and conclusions

General aviation aircraft are currently one of the most dangerous methods of travel. Fortunately, there are many proven safety systems that could be implemented to significantly reduce the fatality rate in GA travel. Of these, GCAS has a proven record for saving several military aircraft from a fatal collision, and if implemented for GA, could reduce the fatality rate by over 26% alone. To help realize this reduction, the authors of this paper are developing a GCAS meant for GA aircraft. The design philosophy behind the two main components of the GCAS, the controller and the GCAS monitor, has been discussed. Promising results produced by the GCAS controller show that even dire initial orientations are recoverable.

The design target of the GCAS controller, the C172, is the most produced aircraft in history. This ensures that the largest portion of the GA community will be able to use this GCAS. However, in the system's current state, there are many other popular airframes that cannot be properly controlled. Therefore, a tuning process must be outlined for the GCAS controller so new tuning profiles can be easily created. Furthermore, more work needs to be done to improve the tuning process for the TPA. Currently, tuning the TPA is a lengthy and partially manual process, and has to be done for specific maneuvers. Automating the process would greatly shorten the development time for new controllers and for new development of other popular airframes.

4 Paper 3: TPA Evaluation (Presented at the 2023 IEEE Digital Avionics Systems Conference, Barcelona, Spain (Urban, Hubbard, Hook, Mark, & Sizoo, 2023))

4.1 Evaluation of a trajectory prediction algorithm within a ground collision avoidance system

While transport category aircraft have maintained an excellent safety record for over 20 years, small GA aircraft pose significantly higher risks. To address this, NASA and the FAA are funding research to adapt technologies originally developed for military fighter aircraft to enhance safety in small general aviation aircraft. One promising technology is the GCAS. This paper focuses on evaluating a critical component of GCAS, the TPA, for implementation in general aviation airplanes. The TPA approximates the aircraft's future trajectory, which is cross-referenced with terrain data to assess the need for ground avoidance maneuvers.

This paper presents a detailed description of a comprehensive simulation-based evaluation tool-set designed for the verification of the TPA. Specifically, it provides techniques and analysis of the NASA EVAA TPA, which has been tuned for a 1982 Cessna 172P aircraft. Results of over 600 individual test cases, considering varied initial states and environmental conditions, are presented. These tests reveal the strengths of NASA's TPA and areas where further tuning could be beneficial. The findings highlight the effectiveness of the evaluation tools in the development and verification of GCAS for small general aviation aircraft.

4.2 Introduction

Aviation safety is a critical concern across all aircraft types, encompassing commercial, military, and general aviation operations. One technology aiming to increase aviation safety is the GCAS. It has proven to be a game-changer in mitigating the risk of aircraft ground collisions. Since its 2014 installation on the fleet of U.S. Air Force F-16s, the Auto GCAS has saved the lives of eleven pilots (Wetsig, 2021).

In response to the major success that Auto GCAS has provided for military aircraft, recent strides have been made to provide a similar safety system for general aviation. Research suggests that if installed on the United States general aviation fleet, up to 110 deaths could be prevented every year. This would equate to a 26.2% yearly reduction in general aviation fatalities (Hook, Sizoo, & Fuller, 2022; Fuller & Hook, 2023). Despite this, a GCAS for general aviation does not exist for public use at this time.

NASA is actively developing the EVAA software system that accomplishes the previously mentioned task (Connor, 2020). Though it aims to prevent both air and ground collisions, EVAA's GCAS module will be the focus of this report.

In response to the creation of a GCAS like EVAA's GCAS, evaluation techniques will be required to assess its safety. The intent of this paper was to explore the evaluation techniques used to assess EVAA's GCAS. The following methodology involved rigorous testing and simulation techniques specifically tailored for general aviation aircraft, with a particular focus on evaluating EVAA's GCAS for a Cessna 172P. The assessment techniques presented in this research are not limited to EVAA, but can also be utilized to evaluate other safety systems. Thoroughly evaluating and refining these algorithms through simulation-based testing will give general aviation enhanced safety and reliability.

4.3 Ground collision avoidance system

A typical GCAS has three or four modules.

1. Map Manager: responsible for loading local terrain data into memory. The terrain data will be utilized when determining the need for a collision-avoidance maneuver.
2. Trajectory Prediction Algorithm(TPA): approximates the aircraft's future trajectory. It provides an approximation for the foreseeable future, around twenty seconds in the case of EVAA.
3. Decider: cross-references the trajectory from the TPA and the terrain from the Map Manager. It determines if the aircraft's future trajectory will come close to intersecting the local terrain. If so, a collision may be imminent, and a recovery maneuver must be performed.
4. Autopilot (Optional): executes the avoidance maneuver by taking control of the aircraft.

To understand how these modules work together, consider how each module processes the data to contribute to the outcome for the following fictional scenario shown in Figure 31.

1. Map Manager - The Map Manager has buffered local terrain data into memory. In this example, the aircraft is flying dangerously close to Mt. Connes in California.
2. TPA - The three blue traces show the output from the TPA. This is the approximated future trajectory for three different avoidance maneuvers - left climbing turn, wings level climb, and a right climbing turn.

3. Decider - In this scenario, the Decider chose the left maneuver. Left was chosen because the straight trajectory was very close to the mountainous terrain, and the right trajectory intersected the terrain.
4. Autopilot - When the Decider commanded a left avoidance maneuver, the Autopilot executed the maneuver by temporally taking control of the aircraft away from the pilot. The aircraft's "True" trajectory is shown in orange.

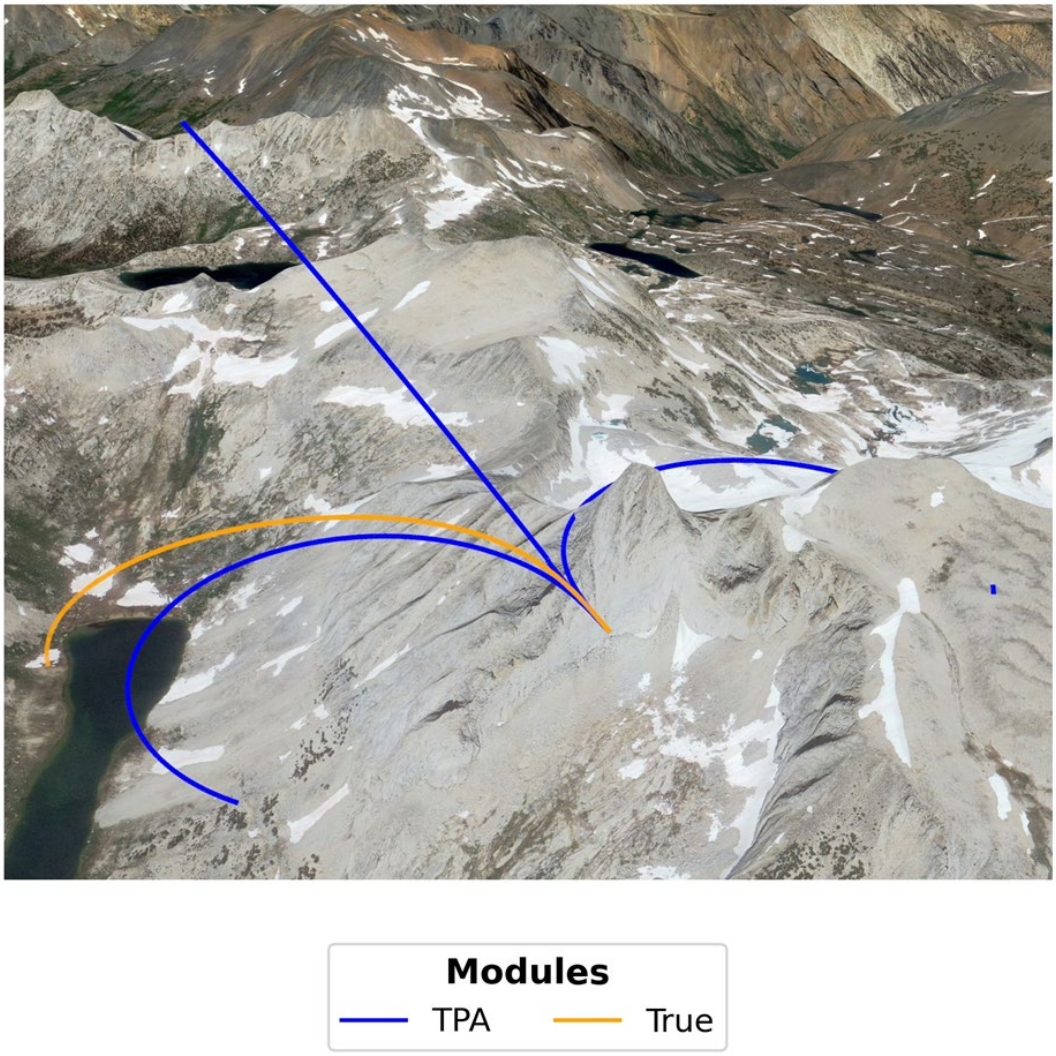

Figure 31. Fictional GCAS scenario over Mt. Connes in California

4.4 The trajectory prediction algorithm

The TPA is responsible for predicting the future trajectory of the aircraft based on its current state. The Decider compares the resultant trajectory with terrain data near the trajectory's proximity. From this comparison, the Decider makes a decision on whether or not a GCAS maneuver should be performed.

There is a noticeable dependency with this architecture. The Decider relies on an accurate trajectory to make a proper assessment of the situation. If the TPA is accurate, the Decider is accurate. Accordingly, if the TPA is inaccurate, the Decider is inaccurate. This dependency could be a life-or-death scenario for a pilot. This vital dependency necessitates giving the TPA extra attention during the development and safety evaluation stages of a GCAS.

In light of this necessity, the scope of this research only focuses on the TPA. By subjecting the TPA to comprehensive testing and analysis, this study aims to identify any outstanding anomalies or inaccuracies within the algorithm. This paves the way for its improvement and refinement. It is important to reiterate that the techniques presented in this research are not limited to the TPA under investigation but can be applied to other TPAs as well.

4.5 Bimodal simulations

To evaluate EVAA's TPA was run in parallel with a completely separate Flight Dynamics Model (FDM) provided by JSBSim (Berndt, 2004). This structure allowed the TPA to be compared against highly-dependable flight dynamics. It is important to keep in mind that in any given experiment, both EVAA's TPA and the FDM ran simultaneously. Figure 32 highlights the bimodal architecture.

The TPA and JSBSim needed to be given matching conditions of the initial aircraft state. Because the TPA is a nimble approximation to a full FDM, it does not have a complete collection of state variables. JSBSim has hundreds of variables that the TPA does not have.

An example of this is aircraft weight configuration and distribution. JSBSim models the individual weight, inertia, and location of aircraft empty weight, fuel weight, passenger weight, point masses, and more. The TPA only models gross weight. If one point mass is not properly approximated by the TPA, a potentially lethal scenario might arise. Discrepancies in weight configurations are discussed in Section 4.8.4. The common state aims to address this variable mismatch.

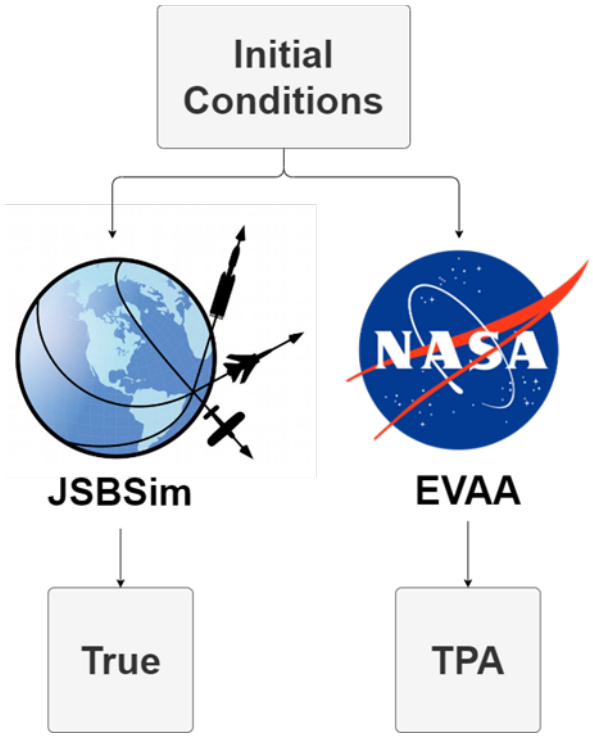

Figure 32. Bimodal architecture

4.5.1 The common state

The TPA and FDM are two distinct entities, each with their own respective sets of variables. The intersection of these two sets was identified and coined the “common state.” The common state is shown in Table 4, the TPA-only parameters in Table 5, and the JSBSim-only parameters in Table 6.

In all experiments that follow, the values of the common state were set identically between both modes.

Table 4. Common state

Variable	**Units**
Latitude	degrees
Longitude	degrees
Altitude	feet
Density Altitude	feet
True Airspeed	knots
Roll	degrees
Pitch	degrees
Yaw	degrees
Flight Path Angle	degrees
Course Angle	degrees
Winds Direction	degrees
Wind Speed	feet per second
Winds Down	feet per second
Vertical Velocity	feet per second
Roll Rate	degrees per second
Pitch Rate	degrees per second
Z-acceleration	g

Table 5. TPA-only parameters

Variable	**Units**
Indicated Airspeed	knots
Gross Weight	lbs

Table 6. JSBSim-only parameters

Variable	**Units**
Calibrated Airspeed	Knots
Complex Weight Configuration	N/A
Remaining Variables	N/A

4.5.2 Cessna 172P autopilot

This research utilizes an autopilot developed by Patrick Maley, et al. (2023) and the Vehicle Autonomy and Intelligence Laboratory (VAIL) at the University of Tulsa. The autopilot, designed for a 1982 Cessna 172P, uses JSBSim (Berndt, 2004) as its FDM. The output from the autopilot is considered the True trajectory. In all experiments, the TPA's trajectory was compared to the autopilot's trajectory.

4.5.3 Base conditions

Before conducting any simulations, it was essential to establish initial base conditions for all experiments. These are the default initial conditions for all simulations. The base parameter value for each variable is documented in Table 7.

Table 7. Initial base conditions

Variable	**Value**
Aircraft State	
Latitude	34.5°N
Longitude	-118°W
Altitude	5000 feet
Density Altitude	5000 feet
True Airspeed	90 knots
Roll	0°
Pitch	0°
Course Angle	0°
Flight Path Angle	-5°
Winds Direction	0°
Wind Speed	0
Winds Down	0
Weight Configuration	
Empty Aircraft	1454 lbs
Pilot	180 lbs
Co-pilot	0 lbs
Luggage	0 lbs
Left Wing Contents	60 lbs
Right Wing Contents	60 lbs
Gross Weight	1754 lbs

4.6 Required data processing

Before performing any analysis, pre-processing was required to format the data into a comparable medium.

4.6.1 Linear interpolation

JSBSim and EVAA produce data in different formats and granularities. It is not possible to perform element-wise comparisons between arrays of different lengths.

To combat this hurdle, linear interpolation was performed on the TPA's trajectories using SciPy's interp1d class. The higher-fidelity output produced by JSBSim was used as the reference when interpolating, with cumulative range as the independent variable. After interpolation, element-wise comparisons between the two trajectories could be performed. Figure 33 shows the TPA's trajectory after interpolation.

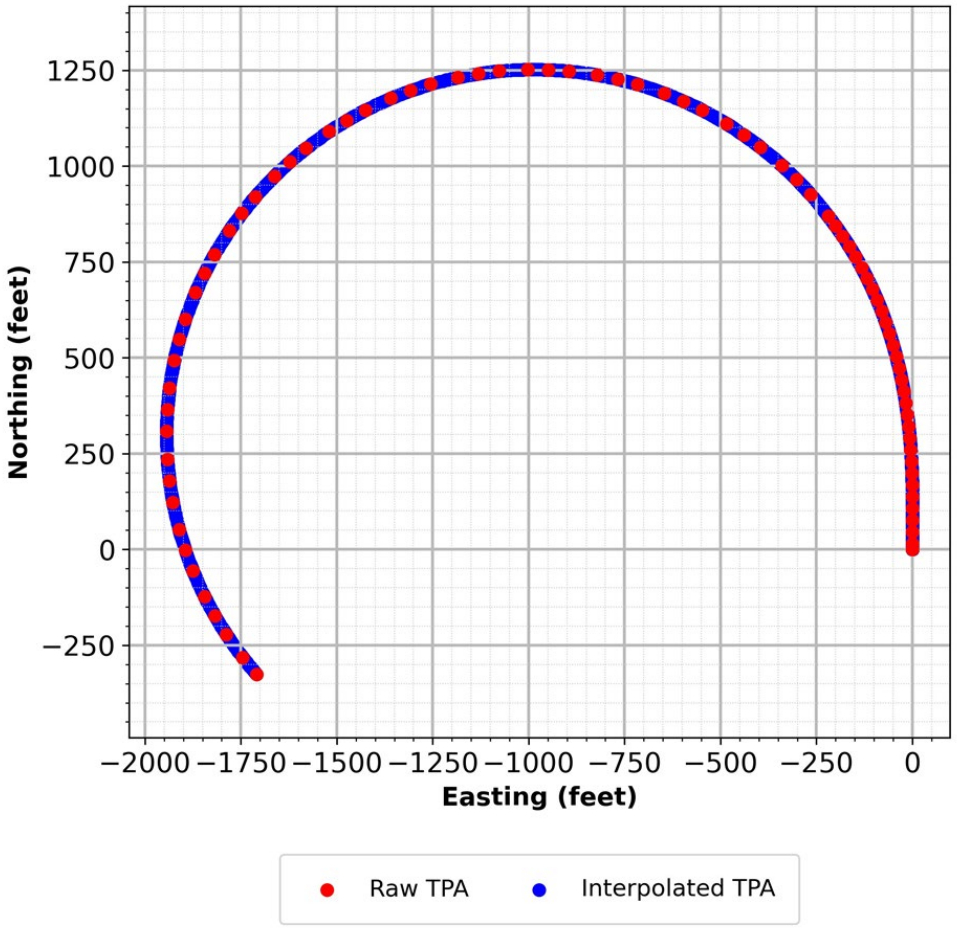

Figure 33. Linear interpolation of the TPA

4.6.2 Localization

In addition to interpolation, each data point's location was translated to a local coordinate frame, as depicted in Equation 7. While retaining location information is crucial when evaluating a complete GCAS, the starting location of the aircraft is not relevant for evaluation of the TPA itself. This research focuses solely on assessing the performance of the TPA within the GCAS, rather than the overall performance of the GCAS.

$$\begin{pmatrix} N_{idx} \\ E_{idx} \\ D_{idx} \end{pmatrix} = \begin{pmatrix} N_{idx} - N_o \\ E_{idx} - E_o \\ D_{idx} - D_o \end{pmatrix} \quad (7)$$

4.6.3 Defining the deviation metric

This research aims to study the accuracy of the TPA's trajectories. A single metric, deviation, was defined, comprising acceptable deviation and excess deviation, to assess the accuracy of the trajectory prediction. Before delving into the exact definition of the metric, it is essential to define its underlying components.

4.6.4 Metric components

Deviation is defined as the distance between two points along a single dimension. Although there are three spatial dimensions, they were simplified into just two - horizontal and vertical. In all trials, both trajectories experienced horizontal and vertical deviations.

The metric's second component is TPA dimensional uncertainty. In any future prediction, its accuracy tends to diminish as it extends further into the future. To address this, the TPA incorporates an uncertainty variable, which produces some interesting behavior. Rather than the TPA producing a linear, one-dimensional trajectory, the uncertainty variable expands it into a three-dimensional space, forming a cone-shaped trajectory volume. This volume increases in size as the prediction extends further into the future. The Decider operates on the assumption that any terrain encompassed within this volume may be encountered by the aircraft.

The uncertainty variable is of critical importance due to the limitations it imposes on the Decider. Any terrain information lying outside of the uncertainty volume remains invisible to the Decider. Consequently, if the True trajectory deviates beyond the boundaries of the uncertainty, it signifies a failure on the part of the TPA. Since the Decider lacks awareness of dangerous terrain beyond the uncertainty, it may fail to take preventive measures, leading to a possible collision. The TPA's uncertainties in the horizontal and vertical dimensions will be considered.

The TPA has a dedicated horizontal uncertainty variable. Its value indicates the horizontal uncertainty from the TPA's trajectory. Being a function of cumulative range, horizontal uncertainty grows in both directions along the aircraft's lateral axis. All terrain points that reside within the horizontal uncertainty will be evaluated by the Decider to determine the necessity of a

GCAS maneuver. The TPA's horizontal uncertainty can be observed in Figure 34. It is the green shaded region that mirrors the TPA's path.

The TPA does not have a dedicated vertical uncertainty variable. However, the Decider does utilize a parameter called the vertical clearance buffer (VCB). This variable represents the minimum allowable altitude difference between the TPA's trajectory and the underlying terrain before an avoidance maneuver is initiated. In EVAA, the VCB is constant and has a default setting of VCB = 50 ft. While the absence of a dedicated vertical uncertainty variable is acknowledged, the VCB serves as a substitute for this.

4.6.5 Acceptable deviation and excess deviation

When the TPA accurately predicts the aircraft's trajectory, it closely aligns with the True trajectory generated by the JSBSim autopilot. If the True trajectory falls within the uncertainty volume, it is considered a successful prediction by the TPA. This concept is referred to as acceptable deviation, indicating that the TPA effectively approximates the true flight dynamics since the True trajectory remains within the defined uncertainty bounds.

If the True trajectory extends beyond the horizontal uncertainty range defined by the TPA, it is referred to as excess deviation. Excess deviation indicates an unacceptable deviation and signifies a failure on the part of the TPA. This outcome indicates that the TPA does not accurately approximate the true flight dynamics. The presence of excess deviation provides the Decider with inaccurate information, which can potentially lead to a collision if not corrected.

Observe the top-down perspective shown in Figure 34. In this example, the unique initial aircraft state produced positive excess deviation. When the True trajectory (orange) stays within the bounds of the TPA's horizontal uncertainty (green), it is acceptable deviation. However, excess deviation is eventually attained when the True trajectory deviates outside the horizontal uncertainty.

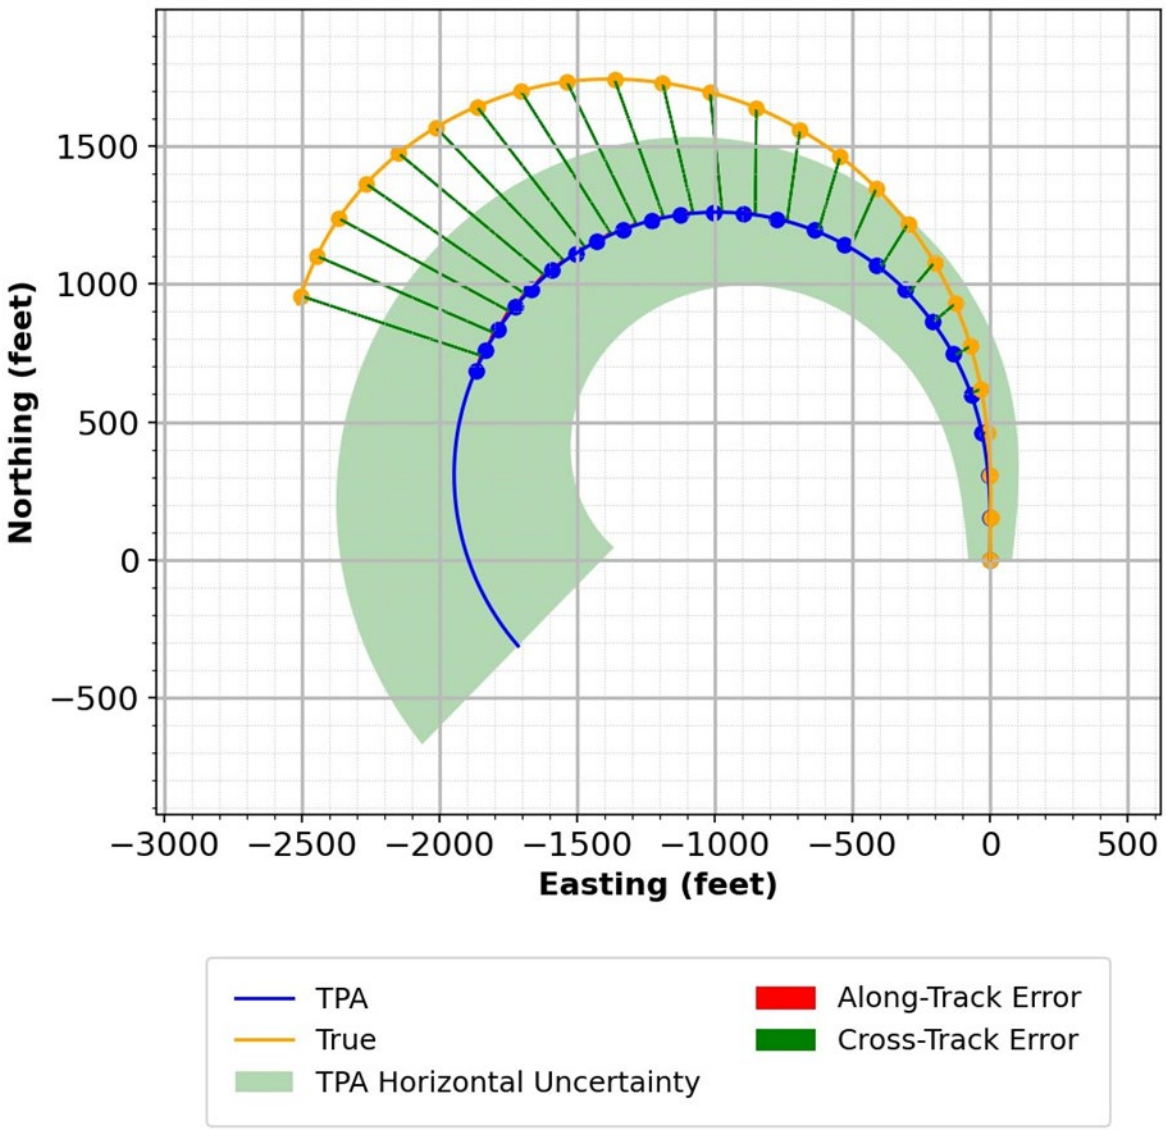

Figure 34. Top-down excess horizontal deviation for a left maneuver

While Figure 34 provides a top-down perspective of excess deviation, Figure 35 depicts the same scenario as a function of cumulative range. More importantly, it shows the relationship between the deviation and uncertainty.

When the deviation falls within the bounds of the uncertainty volume, it is considered acceptable deviation (shown in green). Conversely, if the deviation exceeds the limits of the uncertainty volume, it is classified as excess deviation (shown in red).

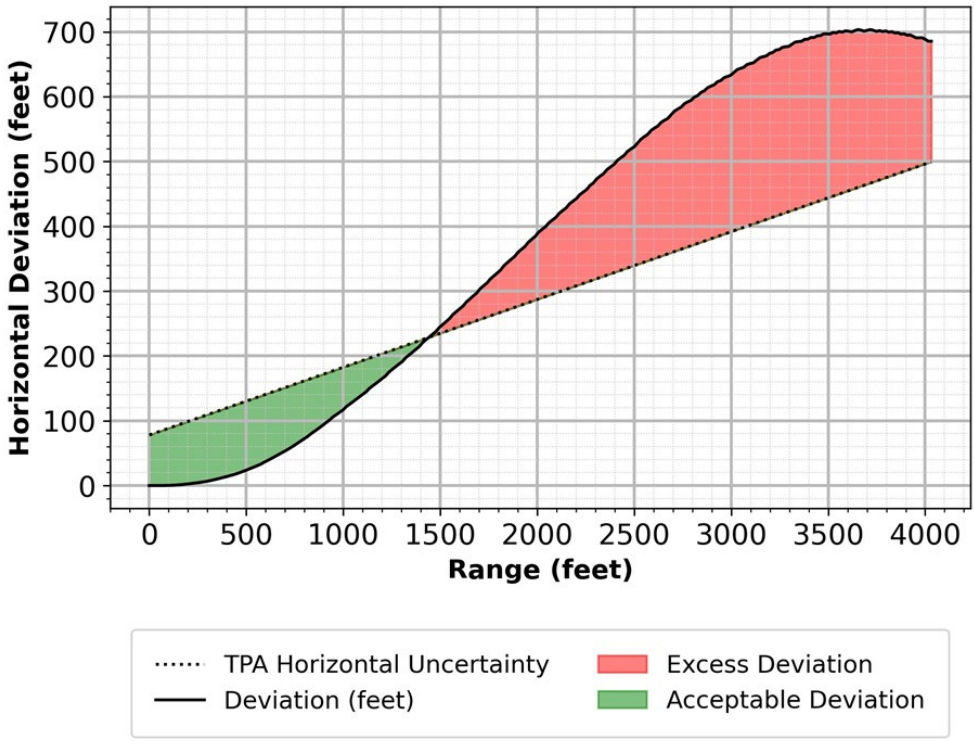

Figure 35. One dimensional excess horizontal deviation

Combining these two components into a single value is highly beneficial for analysis. By subtracting the TPA Uncertainty from the Horizontal Deviation, we obtain a simplified metric, as shown in Equation 8.

$$\text{excess\_deviation} = \text{deviation} - \text{tpa\_uncertainty} \tag{8}$$

Subtracting the TPA uncertainty from the deviation yields a result that can be used to determine whether the deviation falls within the acceptable range or exceeds it. A negative result indicates acceptable deviation, while a positive result indicates excess deviation. Figure 36 illustrates the result of this equation. The simplified subtraction model was adopted as the standard metric for evaluating excess deviation.

To facilitate the analysis of batch simulations, in certain cases, the maximum excess deviation was extracted from a single trial, depending on the experiment. This approach allowed for a straightforward comparison and identification of the most significant deviation observed in each case.

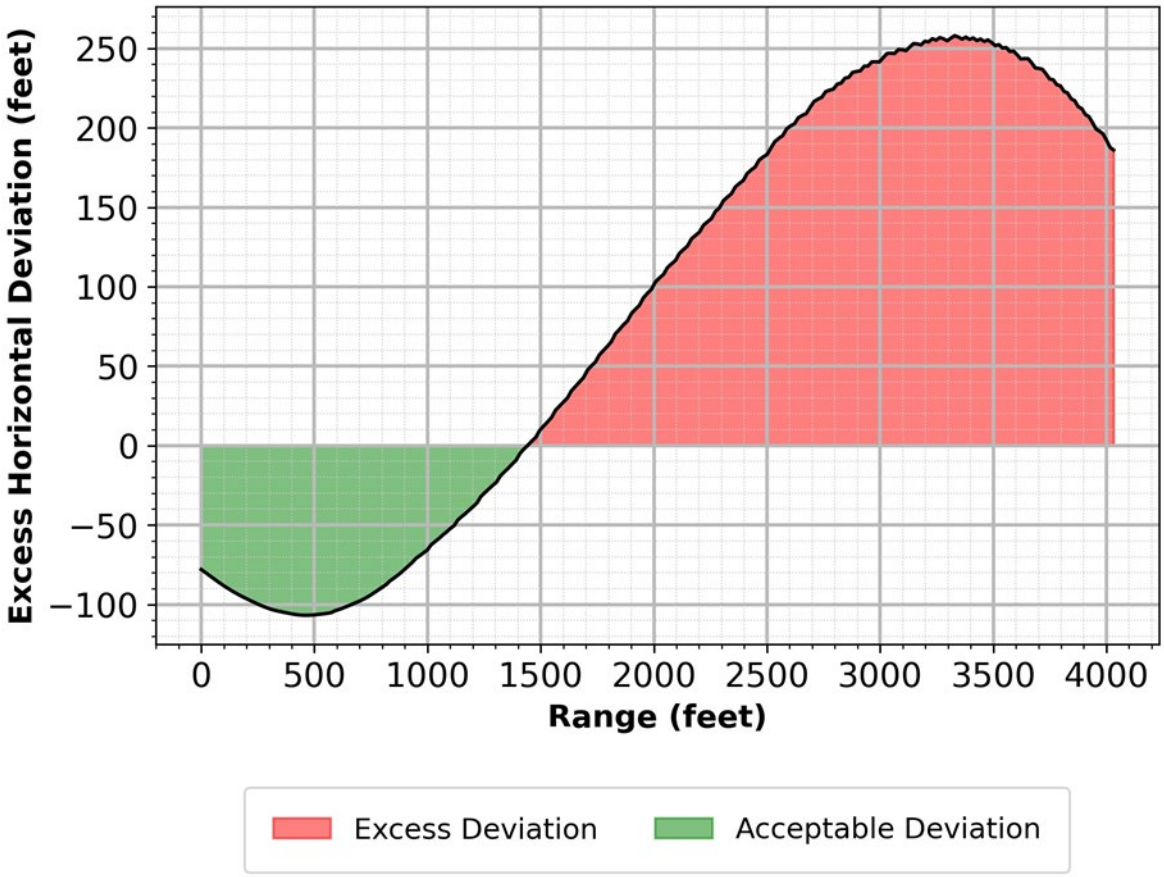

Figure 36. One dimensional excess horizontal deviation (simplified model)

4.7 Obtaining horizontal and vertical deviation

The objective of this section is to derive excess deviation along the horizontal and vertical dimensions. First, the standard terminology for dimensional error must be defined. The track errors, listed in Table 8 describe the deviation experienced by each dimension in the body-frame.

Table 8. Dimensional track errors

Name	Dimension
Along-track error	x-axis, or longitudinal error
Cross-track error	y-axis, or lateral error
Altitude-track error	z-axis, or vertical error

To facilitate the understanding of dimensional deviation, the following explanation will begin by focusing on just two points in 3D space. This idea can be extended to entire trajectories by applying the same process successively. In the example to follow, the first point is produced by the TPA, and the second point is produced by the JSBSim autopilot.

4.7.1 Assigning a transformation matrix

To determine the relationship between the two, both points are assigned a transformation matrix, which contains rotation and translation components. Because the excess deviation metric is only concerned with horizontal and vertical deviation, the transformation matrix will have one rotation about the z-axis only. ψ represents the horizontal rotation angle with respect to the north, east, down (NED) world-frame. For a purely translational matrix with no rotation, a $\psi = 0$ yields in an identity matrix with the translation components in the last column. This transformation matrix is shown in Equation 9.

$$T = \begin{pmatrix} \cos(\psi) & -\sin(\psi) & 0 & x \\ \sin(\psi) & \cos(\psi) & 0 & y \\ 0 & 0 & 1 & z \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{9}$$

The representation of the first point as a transformation matrix is shown in Equation 10. In this equation, ψ_{TPA} is the aircraft's heading within the TPA's trajectory.

$$T_{TPA}^{NED} = \begin{pmatrix} \cos(\psi_{TPA}) & -\sin(\psi_{TPA}) & 0 & N_{TPA} \\ \sin(\psi_{TPA}) & \cos(\psi_{TPA}) & 0 & E_{TPA} \\ 0 & 0 & 1 & D_{TPA} \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{10}$$

The representation of the second point as a transformation matrix is shown in Equation 11. In this equation, ψ_{True} is the aircraft's heading within the True trajectory.

$$T_{True}^{NED} = \begin{pmatrix} \cos(\psi_{True}) & -\sin(\psi_{True}) & 0 & N_{TPA} \\ \sin(\psi_{True}) & \cos(\psi_{True}) & 0 & E_{TPA} \\ 0 & 0 & 1 & D_{TPA} \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{11}$$

4.7.2 Computing the dimensional deviation between two points

The dimensional deviation between the two points can be obtained using simple linear algebra. Manipulating these matrices yields a new transformation matrix which represents the deviation from the first point's frame to the second point's frame. Equation 12 illustrates the matrix algebra required to obtain the transformation matrix from the TPA frame to the True frame.

$$T_{True}^{NED} = T_{TPA}^{NED^{-1}} \cdot T_{True}^{NED} = T_{NED}^{TPA} \cdot T_{True}^{NED} \quad (12)$$

Once complete, the track errors for all three dimensions can be extracted from the resulting transformation matrix. These track errors will be utilized when computing the horizontal and vertical deviations.

4.7.3 Transformations for trajectories

To extend the process from just two points to complete trajectories, an iterative approach was employed. The methodology described in the previous section was repeated for each data point within the trajectory. A transformation volume was generated for each trajectory by combining the transformation matrices for each point. Each transformation matrix within the volume corresponded to a point along the trajectory. The transformation volume encapsulated the collective set of transformations needed to map the world frame to the trajectory frame. The two transformation volumes were utilized when computing the horizontal and vertical deviations.

4.7.4 Collapsing track error into two dimensions

The deviation metric is only concerned with two dimensions - horizontal and vertical. To achieve this requirement, the along-track and cross-track dimensions were collapsed into one - horizontal track error. In doing so, only one horizontal dimension retains information about deviation from the TPA point to the True point. To attain this, an algorithmic approach was employed that utilized the transformation volumes for each respective trajectory. It searched for two points along the trajectory that yielded the smallest horizontal straight-line norm. When the two points were found, the horizontal error can be approximated with the cross-track error. The unoptimized $O(n^2)$ pseudo code for this algorithm is shown in Table 9.

Table 9. Computing horizontal deviation algorithm

Input: tpa_trajectory, true_trajectory
Output: horizontal_deviation

```
horizontal_deviation_arr = fill(n, inf)
for true_idx = 0 to n do
    true_point = true_trajectory[true_idx]
    for tpa_idx = 0 to n do
```

```
        tpa_point = tpa_trajectory[tpa_idx]
        horizontal_deviation = horizontal_norm(tpa_point, true_point)
        horizontal_deviation_arr[true_idx]
          = min(horizontal_deviation_arr[true_idx], horizontal_deviation)
      end for
end for
return horizontal_deviation_arr
```

When collapsing the along-track and cross-track dimensions into one, the altitude-track dimension remains alone by itself. The vertical deviation can be extracted using the altitude-track error component from each matrix in the transformation-volume.

4.8 Directed parameter variation

Having laid the foundation for all evaluation techniques, EVAA's TPA can now be analyzed byexecuting a thorough test plan to assess its viability as a TPA within a GCAS.

The objective of directed parameter evaluation is to isolate each variable and assess its unique impact on the accuracy of the TPA. In the following experiments, the initial conditions of the aircraft attributes were individually manipulated, while the values of all other variables remained at Algorithm 1 Computing Horizontal Deviation their default values as defined in the base conditions (Table 7). This approach ensured that only the variables under investigation were modified, while all other variables remained unaffected. Table 10 shows the variables that were tested, while Table 11 lists the weight variation experiments.

Table 10. Directed parameter variation experiments

Variable	**Minimum**	**Maximum**	**Increment**
KTA	55 knots	120 knots	5 knots
Bank Angle	-60°	60°	5°
Wind Speed	0 knots	30 knots	5 knots
Wind Direction	0°	360°	30°
Altitude	1000 ft	13000 ft	1000ft
V_z	-1000 ft/min	+500 ft/min	100 ft/min

Table 11. Weight directed parameter variation experiments

Weight Configuration	**Description**
Nominal	Eight different nominal weight configurations
Asymmetric	Laterally nominal, left-heavy, or right-heavy weight configurations
Extreme	159 different weight configurations, both safe and unsafe weight configurations

4.8.1 Forward velocity (KTAS)

The results from the forward velocity experiment are shown in Figure 37. Each maneuver was given 14 initial conditions with a KTAS between 55 and 120 knots. This allowed for a comprehensive analysis of the impact that KTAS has on the TPA.

The straight maneuver exhibited no excess horizontal deviation (EHD). In contrast, the left maneuver showed 10 instances of EHD, indicating frequent deviation from the desired path. Similarly, the right maneuver exhibited 5 instances of EHD, suggesting comparatively less deviation but still requiring attention. The largest EHD of 200 feet was encountered during a left maneuver at 55 knots.

One may notice an anomaly for banked maneuvers at a KTAS of 75. No EHD occurs at this particular airspeed. A KCAS of 75 also happens to be the best-rate for the Cessna 172P, which also serves as an input to the JSBSim autopilot's gamma controller (Maley, Hubbard, Urban, & Hook, 2023). In light of this, it should be mentioned that some outliers from these experiments may be attributed to the controller.

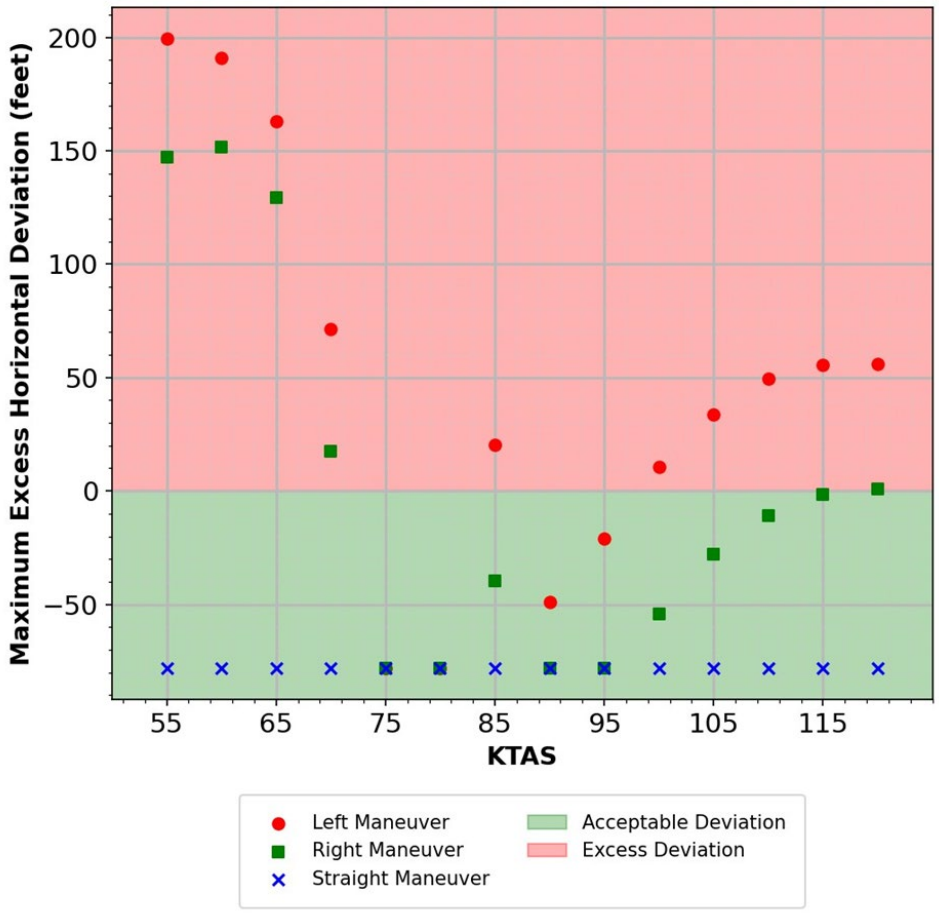

Figure 37. Initial forward velocity and its influence on the TPA's accuracy

4.8.2 Bank angle

The results from the bank angle experiment are shown in Figure 38. Each maneuver was simulated 25 times with an initial ϕ between −60° and 60°.

For nominal flight conditions, EVAA's TPA does a fine job at modeling trajectories with differing values of ϕ. All maneuvers between −35° and 35° have a nominal acceptable deviation. For more extreme initial bank angles, the True trajectory deviates significantly outside of the TPA's horizontal uncertainty. In the case of the straight and right maneuvers at $\phi = -60°$, the True trajectory deviates 600 ft away from the TPA's horizontal uncertainty. While its performance is exceptional for nominal bank angles, EVAA's TPA may benefit from further tuning for more extreme bank angles.

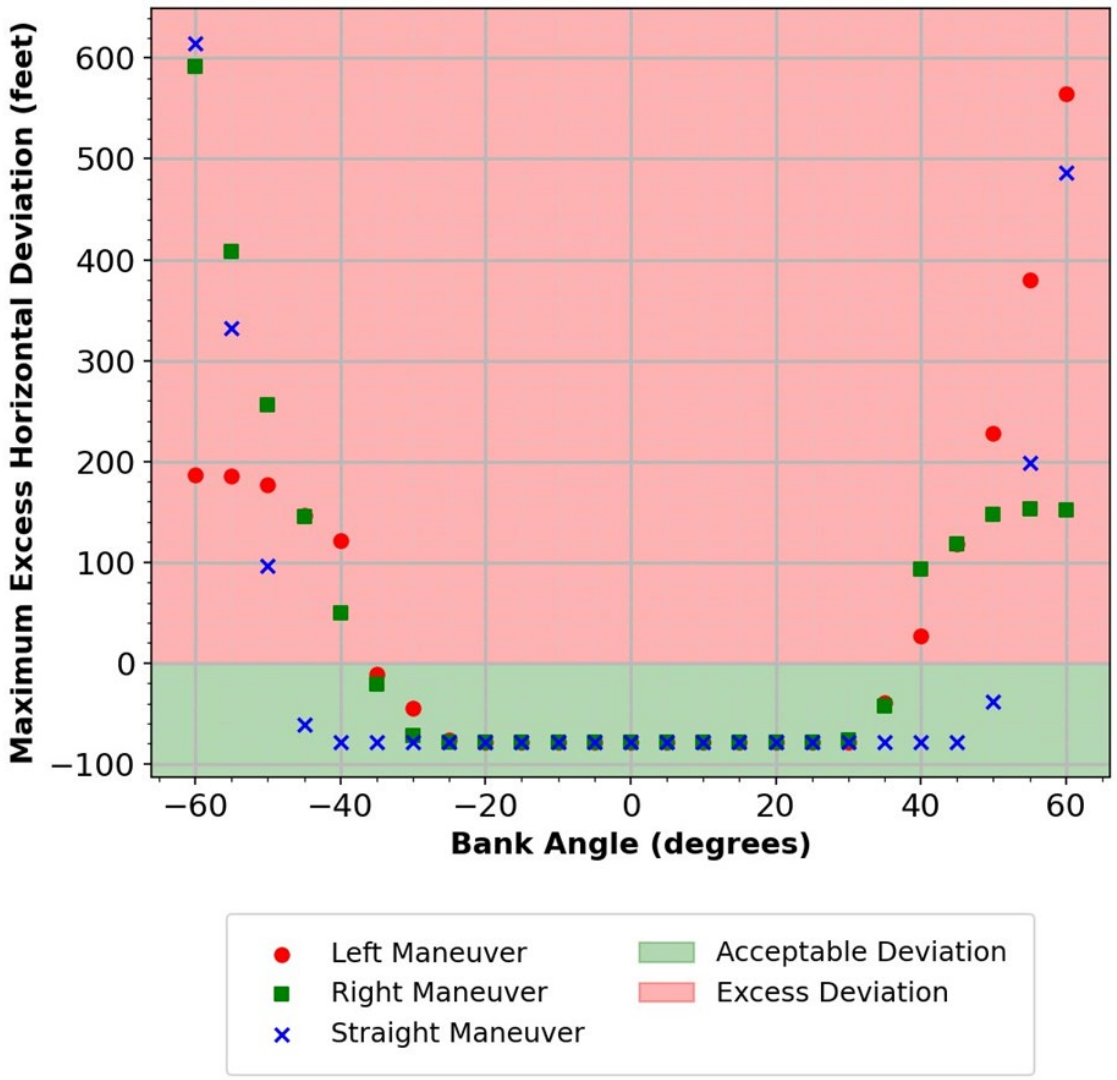

Figure 38. Initial bank angle and its influence on the TPA's accuracy

4.8.3 Vertical velocity

The results from the vertical velocity experiment are shown in Figure 39. Each maneuver was given 16 different initial vertical velocities that varied from −1000ft/min to 500ft/min.

Similar to the forward velocity experiment, the straight maneuvers exhibited no excess deviation, indicating nominal performance from the TPA. However, for the banked maneuvers, the results were less favorable. Only when vertical velocity (Vz) was set to −1000ft/min did the banked maneuvers show acceptable deviation. For all other Vz values, the banked maneuvers resulted in

excess deviation, indicating off-nominal performance from the TPA. It is worth noting that the largest excess deviation occurred during a banked maneuver with Vz = 500ft/min, reaching an excess vertical deviation (EVD) over 40 feet. The undesirable results from this experiment are likely attributed to the constant VCB = 50ft.

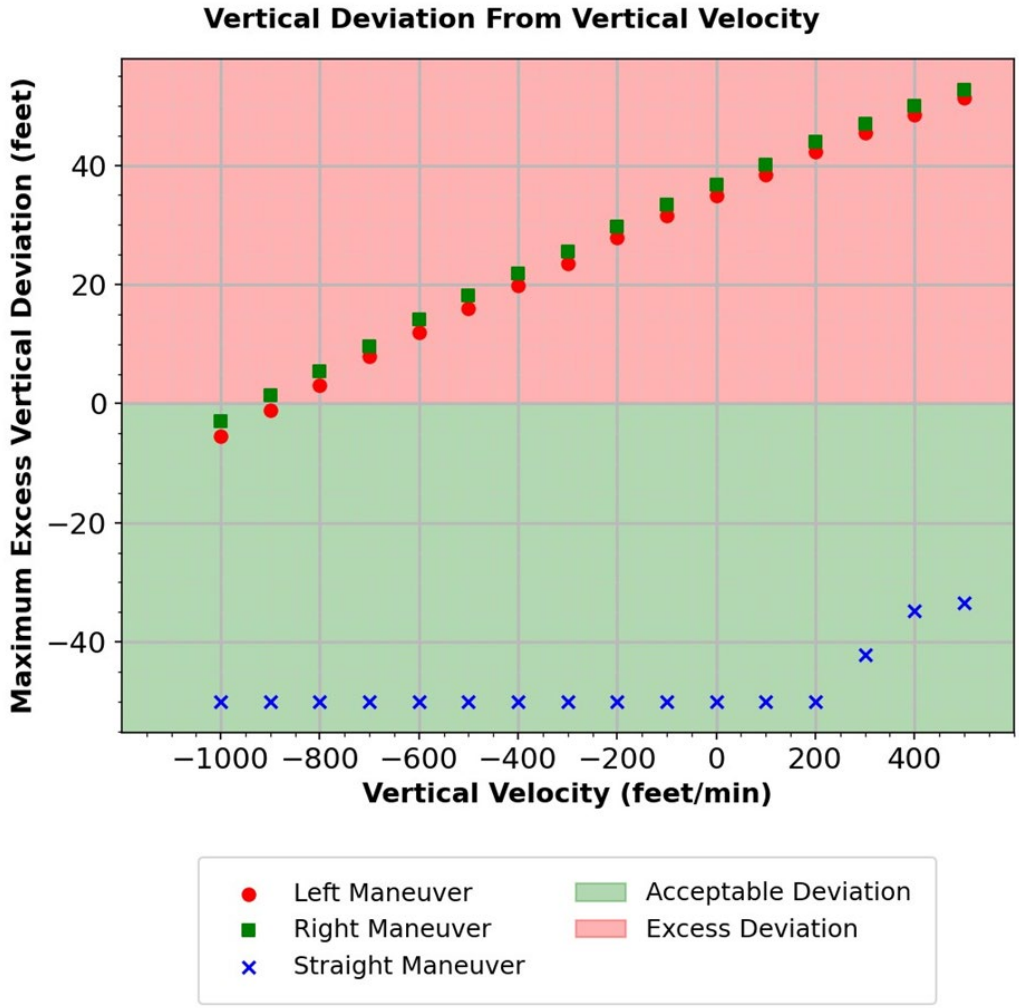

Figure 39. Initial vertical velocity and its influence on the TPA's accuracy

4.8.4 Weight

EVAA's TPA does not have complex point-mass modeling. Instead, it only models gross weight. In weight-directed testing, the goal was to determine if the gross-weight-only approximation served as an accurate model to true flight dynamics. In all simulations, the authors explicitly set all mass structures listed in Table 12.

Many complex weight distributions were simulated in hopes of unveiling any inaccuracies in the TPA's weight model. Weight experiments were compartmentalized into two categories - longitudinal and lateral.

Table 12. Mode weight structures

*JSBSim*Autopilot	TPA
Empty Aircraft Pilot Co-pilot Luggage Wing Fuel Contents	Gross Weight

4.8.5 Longitudinal weight experiments

The following experiments focus on the TPA's performance as a result changing the longitudinal weight distribution. Every C172 comes equipped with an operator's manual. In the manual, there is a section that describes proper weight and balance configuration along the longitudinal axis, similar to what is shown in Figure 40. In the figure, the x-axis is the CG Arm, measured in inches. This is the center of gravity along the longitudinal axis. The y-axis is the aircraft's gross weight, measured in pounds. Any configuration that intersects within the green is considered a proper loading configuration. Conversely, any configuration that intersects outside the green is considered improper, and is potentially hazardous.

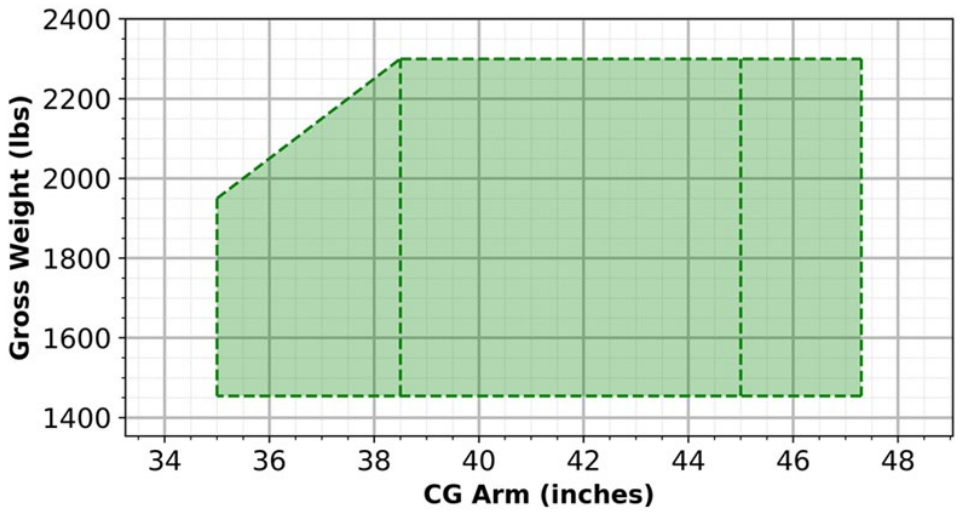

Figure 40. Standard C172 proper weight and balance chart

4.8.5.1 Nominal weight configuration

For the nominal weight configuration tests, eight different weight configurations were simulated. The combinations were all within standard Cessna 172P payload limits, shown in Table 13.

Table 13. Nominal weight configurations for this experiment

Variable	**Value(s) (lbs)**
Empty Aircraft	1454
Pilot	200
Co-pilot	{250, 0}
Luggage	{100, 0}
Wing Contents	{130, 60}

The results of this experiment are shown below in Figure 41 and Figure 42. Figure 41 shows eight different nominal weight configurations of the aircraft, where each color represents a unique payload. From each scatter point, the CG Arm and gross weight can be determined by its location on the graph. Figure 42 shows the individual simulations from those particular configurations. Note the unique configurations and their colors are consistent across both figures.

In all eight weight combinations, the straight and right maneuvers produced no EHD. On the other hand, half of the left maneuvers produced EHD, venturing outside of the TPA's Horizontal

Uncertainty. This experiment can conclude that the TPA accurately approximates a Cessna 172P's varying center of gravity along the longitudinal axis, but struggles to model trajectories with higher gross weights.

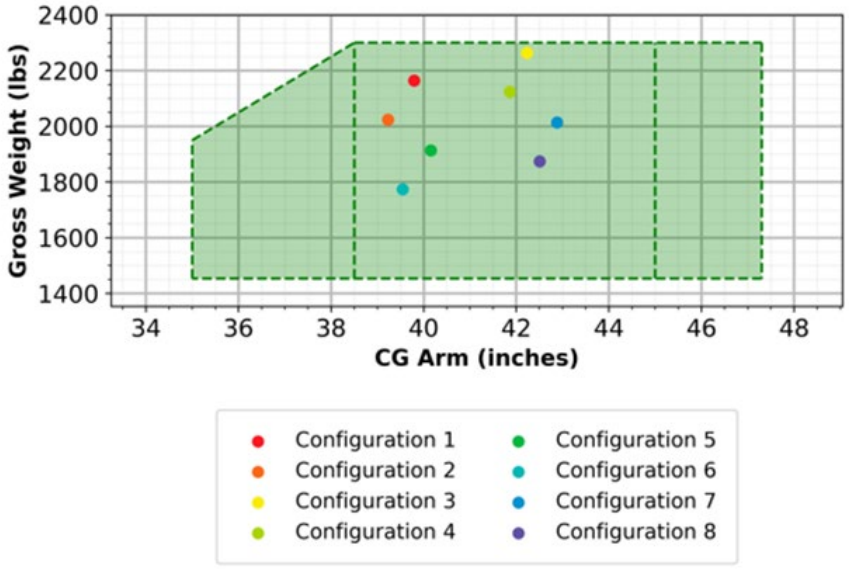

Figure 41. Weight and balance chart for eight nominal flights

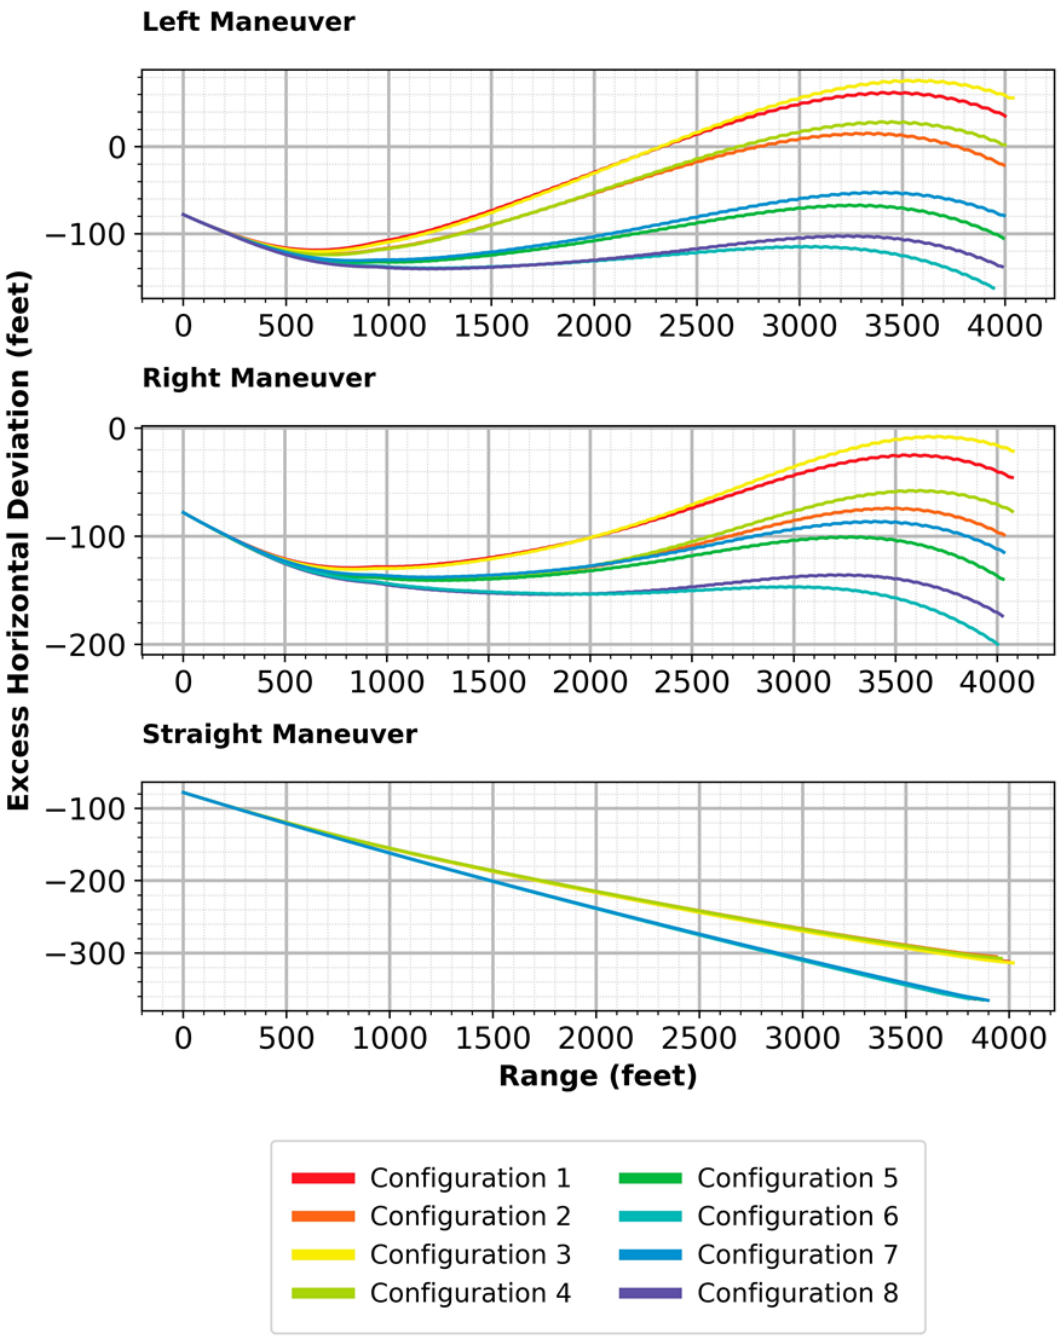

Figure 42. Excess horizontal deviation for eight nominal weight distributions

4.8.5.2 Extreme weight configuration

An extreme payload would have an impact on the TPA's accuracy. In total, 53 different weight configurations were simulated for each maneuver. This experiment assessed the TPA's ability to model the aircraft's longitudinal center of gravity, so only symmetric weight configurations were set. This means that the pilot/copilot had an identical weight, and each wing's fuel contents were identical. All combinations of weight configurations can be seen in Table 14. The results from the extreme weight configuration experiment can be observed in Figure 43.

Table 14. Extreme weight configuration

Variable	**Value(s) (lbs)**
Empty Aircraft	1454
Pilot/Co-pilot	0, 75, 150, 225, 300, 375
Luggage	0, 75, 150
Wing Contents	65, 130

As previously established, the location of the scatter point indicates the CG Arm and the gross weight of the aircraft. However, in this experiment, the color of the scatter point does not indicate a unique payload configuration. Instead, it represents the maximum level of EHD reached within that trial. The color scale ranges from violet, indicating the lowest amount of deviation, to red, indicating the highest amount of deviation.

During straight and right maneuvers, all nominal payload configurations produced acceptable horizontal. One particular off-nominal left maneuver produced the highest EHD at approximately 200 feet, indicating a significant level of deviation.

The color gradient in the Figure 43 progresses from bottom to top, rather than from left to right. This indicates that the trajectory changes were primarily influenced by gross weight rather than CG Arm. While the TPA demonstrates accurate modeling of changes in CG Arm, this experiment reiterates the limitation in its ability to accurately model gross weight. These findings further support the previous conclusion that further tuning of the TPA's modeling of gross weight is necessary to achieve more precise predictions in such scenarios.

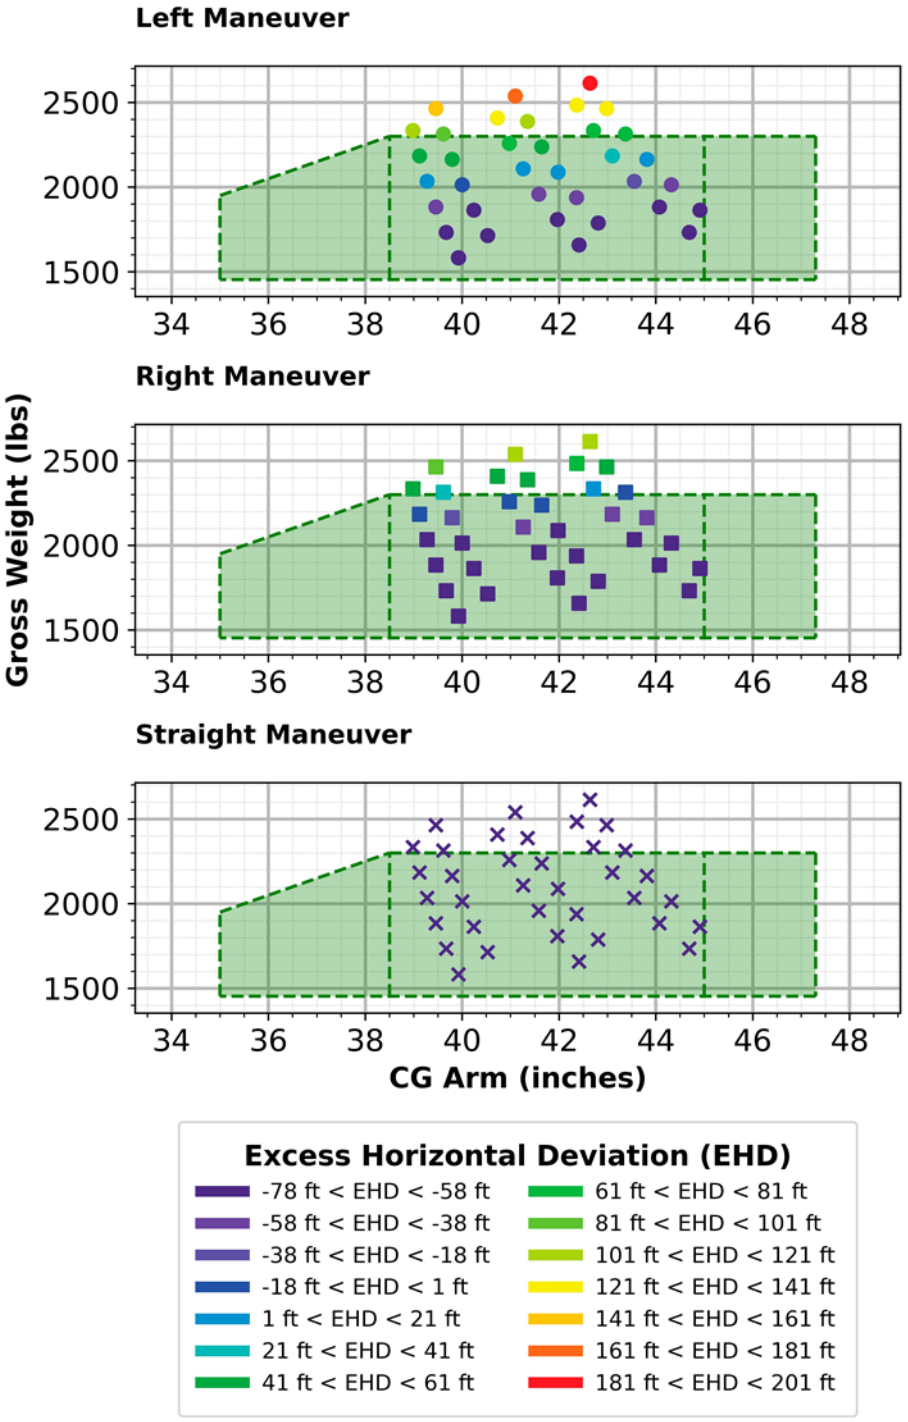

Figure 43. Excess horizontal deviation for extreme configurations

4.8.6 Lateral weight experiment

This experiment focused on the TPA's performance as a result changing the lateral weight distribution. The aircraft's weight lateral axis can be adjusted by modifying factors such as fuel tank contents and pilot/copilot status. Weight experiments will utilize this fact to assess the TPA's ability to model the aircraft's lateral axis.

An asymmetric weight configuration implies an imbalanced center of gravity along the lateral axis of the aircraft. To achieve this, pilot/copilot weight and right/left wing fuel contents were varied. Because CG Arm only describes the longitudinal axis, changes in the lateral axis will not influence the CG Arm. The weight configuration for the asymmetric experiment can be observed in Table 15. The results from the asymmetric weight experiment are shown in Figure 44.

All maneuvers exhibited instances of EHD. The left maneuver demonstrated EHD in all loading configurations except for left-heavy. The right maneuver only experienced EHD in the left-heavy configuration. Interestingly, the straight maneuver encountered EHD specifically in the right-heavy payload configuration. Both banked maneuvers resulted in EHD or near-EHD in the asymmetric nominal configuration. These findings once again highlight the TPA's difficulty in accurately modeling higher gross weight configurations. Furthermore, the TPA would benefit from improved support for accurately modeling asymmetric loads, as evident from the occurrence of EHD in asymmetrically imbalanced configurations for all maneuvers.

Table 15. Asymmetric weight configuration

	Value (lbs)
Lateral Base	
Empty aircraft	1454
Luggage	0
Asymmetric Nominal	
Pilot	200
Co-pilot	200
Left wing contents	130
Right wing contents	130
Heavy Left	
Pilot	200
Co-pilot	0
Left wing contents	130
Right wing contents	0
Heavy Right	
Pilot	0
Co-pilot	200
Left wing contents	0
Right wing contents	130

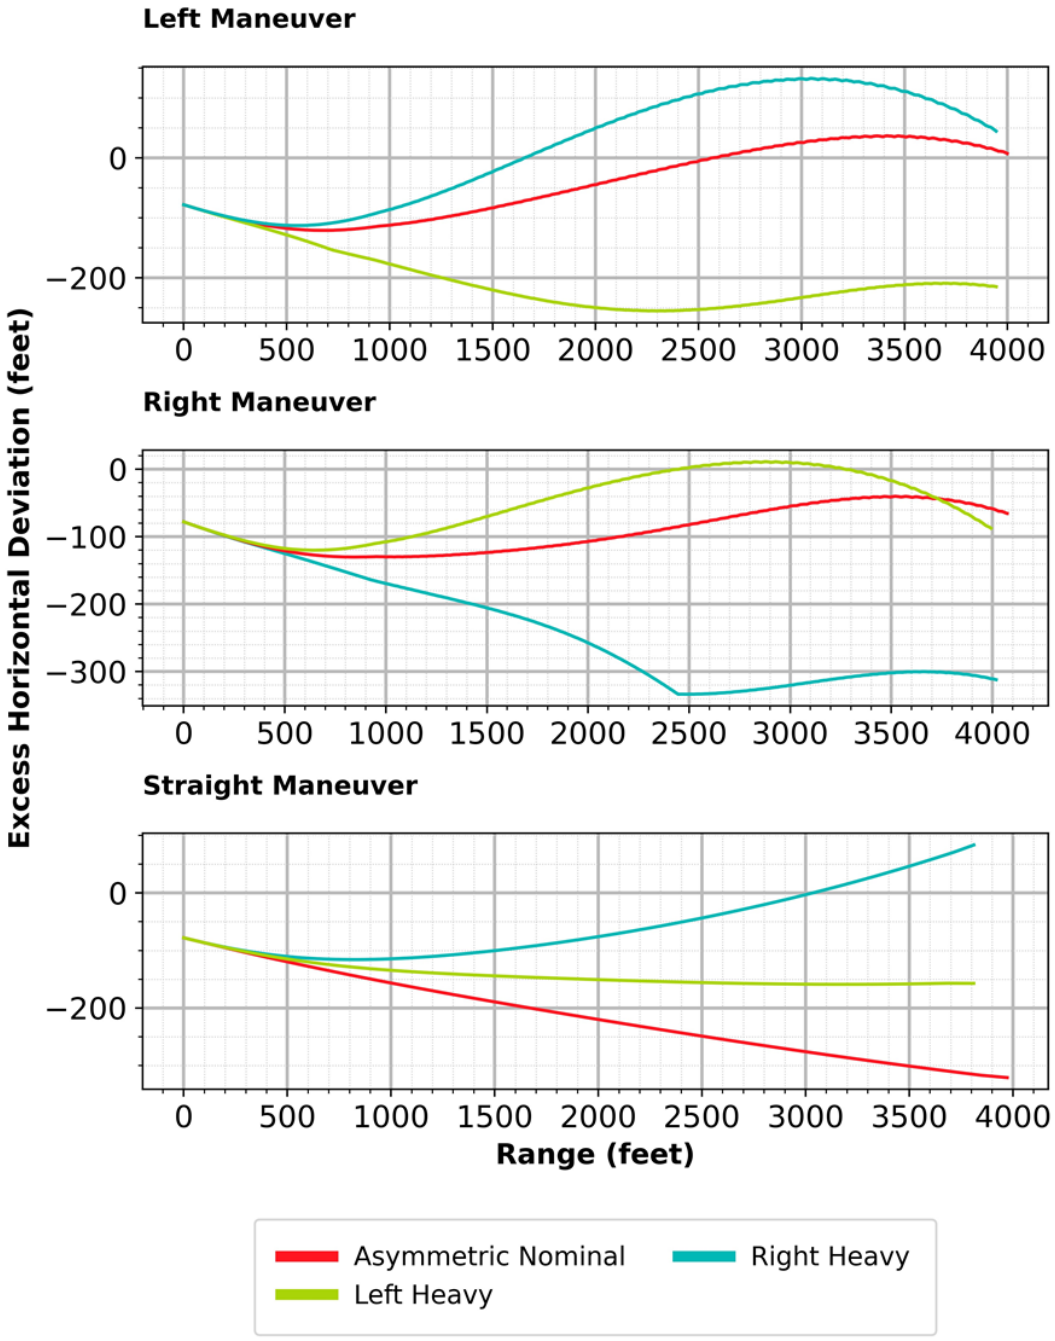

Figure 44. Excess horizontal deviation for asymmetric configurations

4.8.7 Sustained winds

The sustained winds experiment sought to reveal the TPA's ability to model trajectories in the midst of winds. Each maneuver was given 84 different wind conditions, totaling 252 trials overall for the experiment. The results of the winds experiment is shown in Figure 45. The color of the scatter point indicates the magnitude of the wind speed. The polar location of the scatter point indicates the direction the winds are coming from. The distance from the radial center to the scatter point indicates the maximum EHD experienced during the trial.

Not surprisingly, all straight maneuvers provided acceptable horizontal deviation. However, banked maneuvers exceeding 10 knots produced EHD. The largest recorded EHD of 200 ft occurred during a left maneuver with 30-knot winds blowing from the west. These results reveal that the TPA could aid in further tuning of its winds simulation during banked maneuvers.

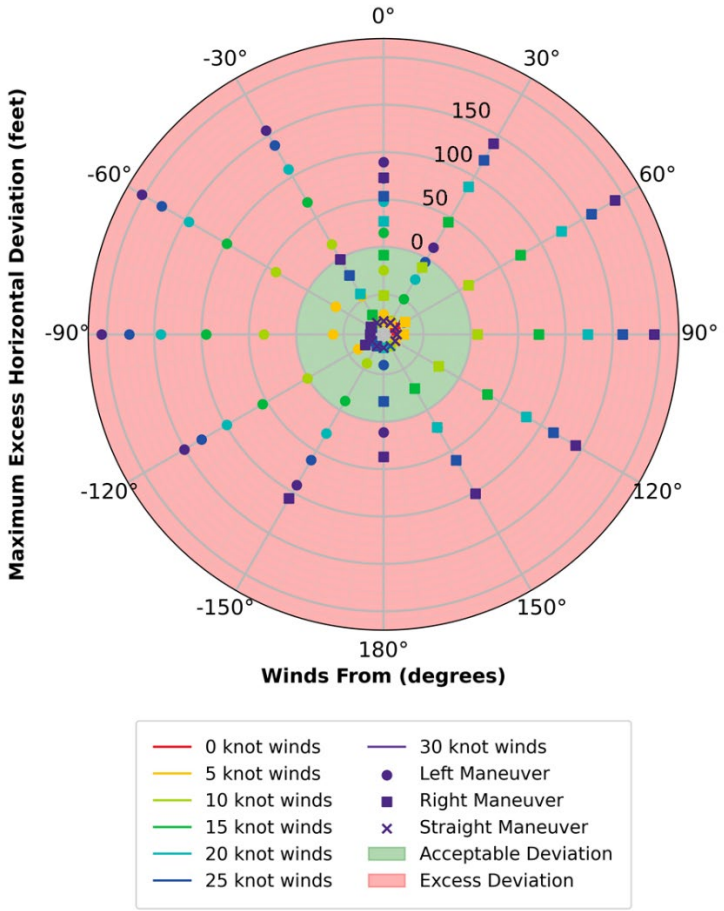

Figure 45. Sustained winds and its influence on the TPA's accuracy

4.8.8 Starting altitude

This experiment aims to evaluate the TPA's performance as a function of initial altitude. Each maneuver was given 12 different altitudes varying from 1,000 ft to 13,000 ft. The results of this experiment can be observed in Figure 46.

For the straight maneuver, no EHD was observed, regardless of the initial altitude. However, for the left maneuver, the TPA exhibited EHD at both lower and higher altitudes, with the lowest EHD occurring somewhere in the middle, around 6,000 ft. The straight maneuver began producing EHD for higher altitudes only, starting at 10,000 ft. These findings reinforce the TPA's strong performance in straight maneuvers while indicating poorer performance in left maneuvers at lower altitudes and banked maneuvers at higher altitudes. These findings highlight the need for further evaluation and refinement of the TPA, particularly for left maneuvers and higher altitudes.

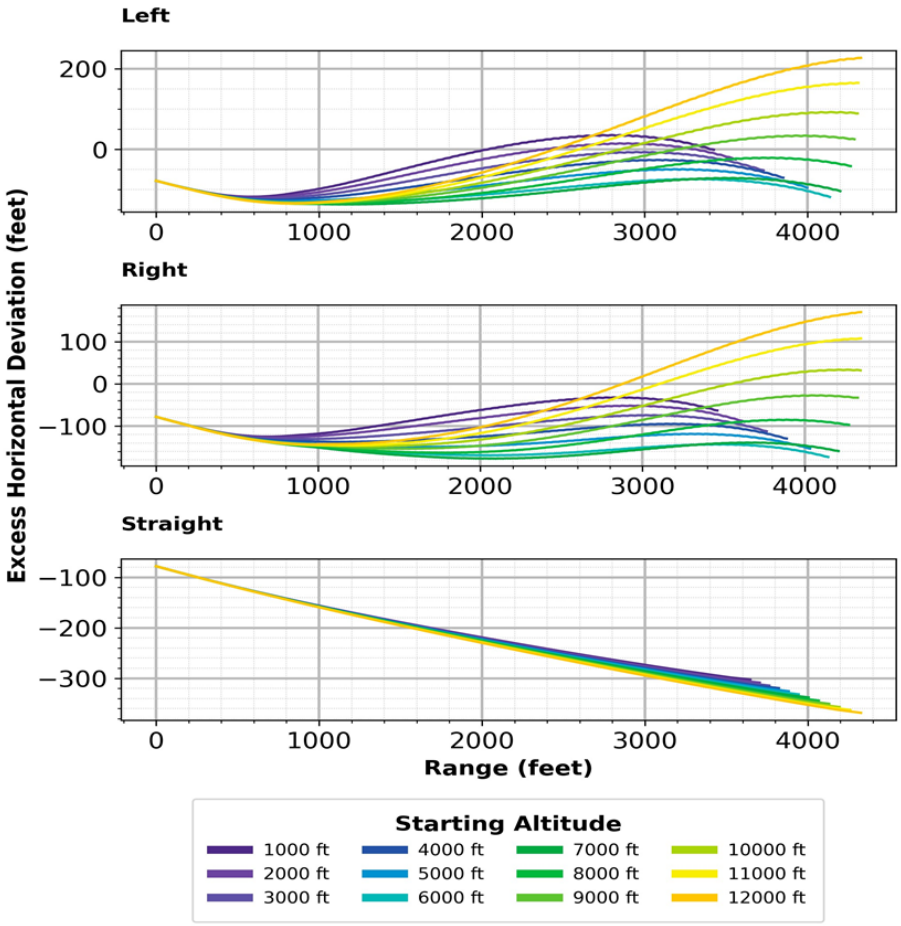

Figure 46. Excess horizontal deviation from varying starting altitudes

4.9 Future work

Aside from directed parameter testing, another technique to evaluate a TPA's performance may be a Monte Carlo-based approach. Monte Carlo analysis methods are computational techniques used to approximate or simulate the behavior of complex systems or processes involving random variables. By repeating the simulation many times with different random samples, Monte Carlo methods provide a probabilistic assessment of the system's performance or behavior.

Monte Carlo simulations could be employed to evaluate a TPA by generating random initial conditions for the aircraft. By sampling various parameters such as position, velocity, altitude, wind conditions, and other relevant variables, the Monte Carlo approach captures the full range of potential scenarios that the TPA may encounter in real-world flight situations. Running the TPA with these randomly generated initial conditions allows for a comprehensive assessment of

its performance under diverse and unpredictable conditions. By analyzing the results derived from the Monte Carlo simulations, insights can be gained into the TPA's robustness, accuracy, and potential areas of improvement.

4.10 Final remarks

Through the utilization of the evaluation techniques outlined in this study, we conducted a comprehensive analysis of over 600 trials, providing valuable insights into the capabilities and limitations of EVAA's TPA. Among the 645 trials examined, we observed exceptional performance of the TPA in accurately approximating the aircraft's trajectory. However, it is important to note that we also identified potential safety considerations, particularly in the context of left-maneuvers. The simulation-based evaluation tools presented in this research offer valuable insights into the performance of the TPA and provide a means for effective tuning. By leveraging these tools, the TPA can be refined to improve EVAA, and thereby contributing to the ongoing enhancement of aviation safety.

5 Paper 4: Evaluating a Ground Collision Avoidance System Using Simulation Based Techniques (Presented at the 2024 AIAA Aviation Conference, Las Vegas, NV, USA, July 30, 2024 (Urban, Hook, & Hubbard, Evaluation of a ground collision avoidance system using simulation based techniques, 2024)

5.1 The Monte Carlo approach

Small private aircraft transportation, commonly known as general aviation (GA), is governed in the US by Title 14 of the Code of Federal Regulations part 91, and by similar regulations in Europe. Statistical analysis shows that GA is considerably more dangerous than commercial aviation and 10 times more dangerous than automobile travel. Factors contributing to this include reduced pilot proficiency and aircraft maintenance standards. Attempts to enhance safety by targeting these factors have inadvertently increased the costs and complexity of flying, leading to a decline in the number of small aircraft pilots. However, technology presents an alternative means of enhancing safety in GA without imposing additional burdens on owners and pilots. One such technology developed by the National Aeronautics and Space Administration (NASA), called the Ground Collision Avoidance System (GCAS), functions by evaluating aircraft state-dependent ground risk and providing recovery trajectories to prevent ground collisions. Since these systems will be safety critical, new evaluation techniques will be required to verify their functionality. This paper defines techniques to verify the protection abilities of a GCAS using a simulation-based “Monte Carlo” approach. The Monte Carlo method seeks to understand the behavior of multi-variable complex systems by randomly visiting a large number of input states to reveal a big-picture representation of the system’s performance. The GCAS activation performance and protection rate is then assessed using multivariate statistical analysis to find the aircraft state variables most correlated with any failed protection events. In total, more than sixty thousand unique scenarios were simulated above a 2500 square mile region in southern California airspace. Our findings indicate that the GCAS protection abilities, coupled with a GCAS autopilot, are very effective, with a predicted successful protection rate of 98.45%. The simulation and statistical methods outlined in this paper could be used to evaluate and tune future autonomous systems, bringing increased safety to aviation.

5.2 Introduction

Using nearly any metric, aviation safety has consistently improved over the 120 years since the first airplane. This is particularly true in recent US commercial airline passenger transport, which has seen over 15 years of operations without a multiple fatality-producing crash. However, GA is still too dangerous to be used as an everyday travel mode for the general population. As evidence of this fact, GA experiences more than a 10 times higher fatality rate (per mile traveled) than automobiles (Hook, Clark, Sizoo, Skoog, & Brady, 2016; Fuller & Hook, 2023). Regulatory efforts to further increase safety in GA have led to the unintended side effect of reducing GA utilization over the last several decades (Fuller & Hook, 2020). This has prompted aviation researchers to propose and implement new technologies that have the potential to improve safety without adding to the regulatory burden of aircraft owners and pilots while also making flying easier.

One popular example is the GCAS. GCAS is designed to prevent ground collision from several underlying causes, including loss of control in-flight, controlled flight into terrain (CFIT), low-altitude operations, and others (Hook, Ryan, Skoog, & Fuller, 2023). GCAS works by detecting potential ground collision trajectories and either alerting the pilot to the danger or by engaging an autopilot to perform a preprogrammed avoidance maneuver. One notable example of an autopilot enabled system is the US Air Force (USAF) Automatic Ground Collision Avoidance System (Auto GCAS) installed on USAF F-16, F-35, and F-22 airplanes. Auto GCAS has saved the lives of 11 fighter pilots to date and has not allowed a single CFIT accident in airplanes on which it was installed (Wetsig, 2021). Further emphasizing the potential of GCAS, recent research suggests that if implemented across the United States general aviation fleet, a GCAS could save approximately 110 lives annually (Hook, Sizoo, & Fuller, 2022). Considering the significant success of Auto GCAS on military aircraft and its potential to reduce fatalities in general aviation, members of the University of Tulsa, the FAA, and NASA have begun efforts to bring this life-saving technology into the public sphere.

This paper details the implementation and results of applying a Monte Carlo-based testing framework to a NASA-derived GA GCAS system. The GCAS under test is a part of a multi-monitor run-time assurance system (Hook, et al., 2018) called the Expandable Vehicle Autonomy Architecture (EVAA). Originally designed for medium- to large-scale UAVs, EVAA has been modified to be implemented on GA aircraft. For this project, the GA aircraft selected was the very popular Cessna 172. All simulations in this research leveraged JSBSim (Berndt, 2004), a headless flight dynamics engine which contains a C172 model.

5.3 Background

5.3.1 Monte Carlo simulation

Monte Carlo simulation is a statistical technique used to understand the behavior of complex systems by generating a large number of random samples from a defined input space and analyzing the resulting outcomes. The essence of Monte Carlo lies in its ability to model the probability of different outcomes that are inherently uncertain due to the involvement of many random variables. By performing repeated random sampling and experimentation, Monte Carlo simulations provide a comprehensive picture of the possible outcomes and their associated probabilities, offering valuable insights into the system's performance under a myriad of conditions.

Monte Carlo simulation has been utilized in aviation safety research to evaluate and mitigate risks associated with various flight operations. The technique's ability to handle multi-variable complexities makes it particularly suitable for assessing the safety of critical aviation systems. For instance, the National Aerospace Laboratory (NLR) has employed Monte Carlo simulations to perform safety risk assessments of complex safety-critical operations. Their research demonstrated how this technique could quantify safety risks and support decision-making processes to enhance operational safety in traffic collision scenarios (Blom, Stroeve, & de Jong, 2006).

In our study, we leverage Monte Carlo simulations to evaluate the effectiveness of a GCAS. By simulating over 60,000 unique scenarios, we can assess the GCAS's performance in diverse conditions, identifying key factors that influence its protection capabilities.

5.3.2 Ground Collision Avoidance Systems

A typical GCAS system is shown in Figure 47. The diagram describes the functional relationship between the different sub-modules within GCAS. This research assesses the ability of GCAS to provide an apt and timely recovery maneuver based on the aircraft state. Let us consider each sub-module within GCAS.

- **Map Manager** takes location information and buffers local terrain data to make the system aware of its nearby surroundings. It has an interface to the "decider" module, which provides the highest point within a given location (Suplisson, 2015).
- **Trajectory Prediction Algorithm (TPA)** takes onboard sensor information and makes an approximation of the aircraft flight path into the near-term future. In other words, it predicts the future trajectory of the aircraft (Suplisson, 2015). The TPA considers three different recovery maneuvers: left climbing turn, wings level climb, and right climbing turn. Previous research performed an isolated evaluation of the TPA by itself.

- **Decider** takes information from the TPA's predicted flight path and the Map Manager's terrain and decides on the necessity of a recovery maneuver. If the future trajectory gets too close to the terrain, a recovery maneuver is initiated (Suplisson, 2015). More specifically, a recovery maneuver is initiated if the trajectory penetrates the Vertical Clearance Buffer (VCB), which is the minimum distance allowance between the underlying terrain and the TPA flight path. In the case of EVAA, a VCB = 50 ft was used. A depiction of the VCB is shown in Figure 51.

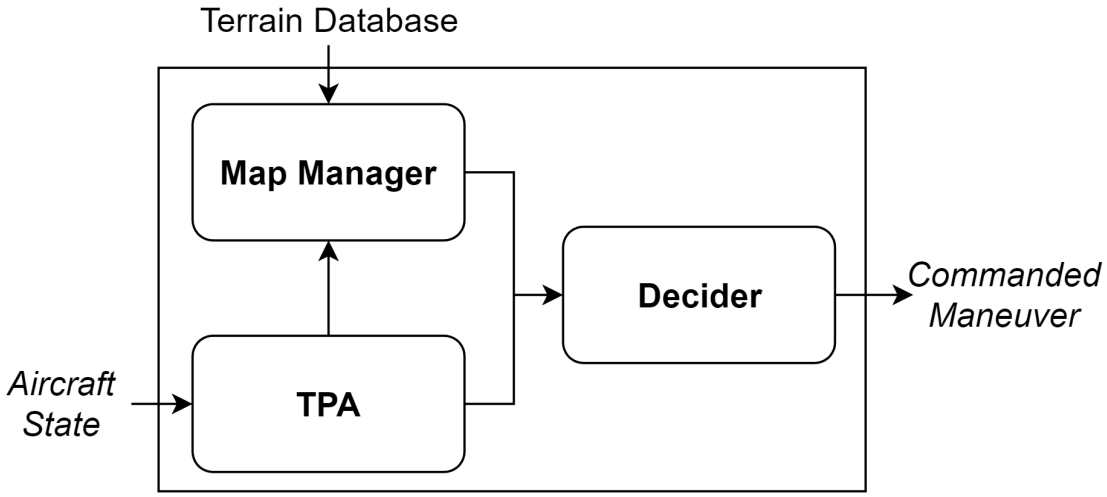

Figure 47. GCAS architecture

This research involved simulating an extensive and comprehensive set of unique scenarios to assess the GCAS ability to make an appropriate recovery decision. Post-simulation processing can determine if the simulated flight path penetrated the terrain, indicating a crash. From the crash status, we can infer if the GCAS would have prevented a collision if a real airplane found itself in that scenario.

5.4 Methodology

This section outlines the integration of the GCAS and Monte Carlo simulation techniques. By combining GCAS's real-time decision-making capabilities with Monte Carlo, we can evaluate the system's performance across a wide range of flight envelopes. This approach ensures a comprehensive assessment of GCAS's effectiveness in preventing ground collisions.

On a real airplane, a GCAS would come in one of two configurations.

- Open-loop control: GCAS detects a potential collision and suggests a recovery maneuver for a pilot to execute.
- Closed-loop control: GCAS detects a potential collision and commands a recovery maneuver for an autopilot to execute.

Because a human pilot's control and decision-making capabilities are not modeled in simulation, a recovery autopilot must be coupled alongside the GCAS. This research employs an "autopilot suite" which features two separate autopilot controllers. Each autopilot commands deflection angles on a Cessna 172P simulated within JSBSim.

- Combined Attitude Controller (CAC) autopilot: The CAC maintains a commanded attitude by controlling the following attributes of the aircraft state: flight path angle γ, roll angle ϕ, angle of attack α, side-slip angle β, and throttle command δ_t.
- Recovery autopilot: The recovery autopilot, developed by researchers at The University of Tulsa (Maley P. D., 2023), features full recovery functionality using turning-climb and straight-climb recoveries. It controls the following attributes of the aircraft state: best rate V_y, course angle χ, flight path angle γ, roll angle ϕ, angle of attack α, side-slip angle β, and throttle command δ_t.

To evaluate apt ground avoidance activations, the airplane needs to be close to the ground. The CAC autopilot accomplishes this by deliberately flying the airplane toward the ground. If the CAC autopilot commands a $\gamma<0$ attitude, it will eventually nose-dive the airplane into the underlying terrain. The GCAS is responsible for predicting a collision and commanding a recovery maneuver, switching to the recovery autopilot.

The combination of these two autopilots provides the system with a closed-loop circuit, where GCAS commands a recovery maneuver, and the autopilot suite produces the airplane's flight path. Figure 48 shows the relationship between the commanded maneuver and the aircraft's flight path.

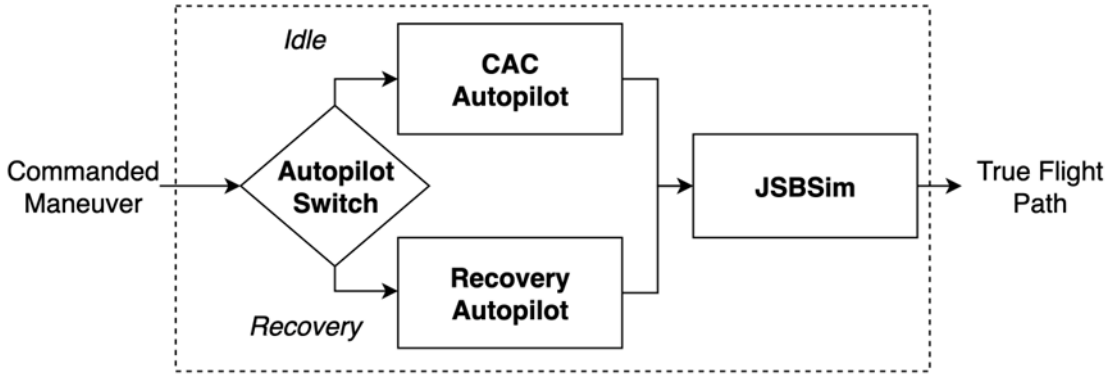

Figure 48. Autopilot suite

The steps to generate, simulate, and evaluate one trial within the context of a Monte Carlo process will be explored.

5.4.1 Pre-simulation: scenario generation

To generate a unique scenario, two conditions pertaining to the aircraft state must be met: validity and appropriateness. The criterion that classifies *valid* and *appropriate* will be defined.

Firstly, the aircraft state must be valid. In this case, *valid* means the generated aircraft state resides within a pre-defined valid state space. These were arbitrarily and pragmatically defined by the authors, as shown in Table 16. These domains describe the entire landscape of possibilities for the initial state of the aircraft, called the initial condition domain (ICD). With over 2500 square miles of starting locations, continuous distributions for attitude and other environmental factors, many different initial conditions can be spawned and simulated. For the Cessna 172P weight distribution, all simulations had a single 180-pound pilot and 60 pounds of fuel in each wing, resulting in a gross weight of 1754 pounds.

Table 16. Initial condition domain (ICD)

Variable	**Minimum**	**Maximum**
Latitude	34.1° N	34.9° N
Longitude	117.1° W	117.9° W
True Airspeed	55 knots	120 knots
AGL Alt.	$t - 30s$ to impact	$500ft$
Roll (ϕ)	$-60°$	$60°$
Course Angle (χ)	$0°$	$360°$

Variable	**Minimum**	**Maximum**
Vertical Velocity	$-1000 ft/min$	$500 ft/min$
Winds Direction	$0°$	$360°$
Wind Speed	0 knots	30 knots

Secondly, the aircraft must be appropriate. In this context, *appropriate* refers to the status of the GCAS activation at simulation start ($t = 0$). A GCAS activation at $t = 0$ is unfair to the safety system, as the situation is already deemed to be unsafe. If the GCAS is not actively commanding a recovery maneuver at $t = 0$, the scenario is valid. On the other hand, if the GCAS is actively commanding a recovery maneuver at $t = 0$, the scenario is invalid.

An aircraft state adhering to these two rules is a healthy initial condition. The two rules must be met before a generated condition is simulated.

5.4.2 Simulation: scenario execution

To simulate the trial, the two modules were coupled into the architecture shown in Figure 49. EVAA's GCAS is designed to operate at 1 Hz and was configured as such. The autopilot suite was configured to run at a 60 Hz resolution. The composite of these modules can simulate an entire GCAS scenario. A default flight simulation timeout of 120 seconds without activation was imposed on all simulations.

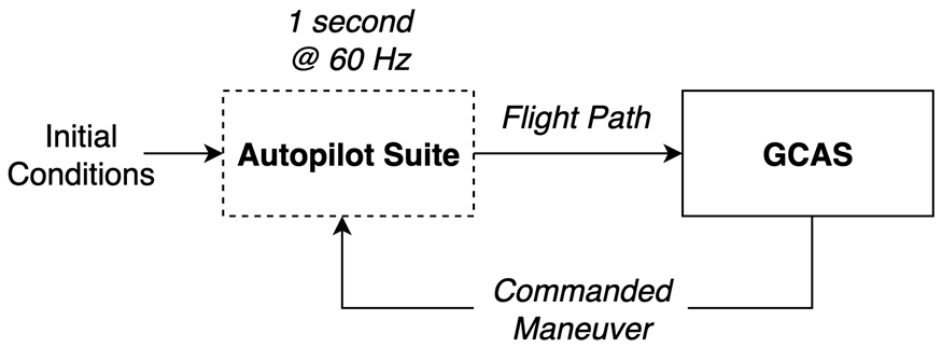

Figure 49. Monte Carlo control system and chronological flow diagram

5.4.3 Post-simulation: scenario assessment

After simulation, two Boolean variables can meaningly describe the outcome: activation status and crash status. These two "bits" can classify every simulation into four different categories, described in Table 17.

Table 17. Classification metrics

Metric Name	**Activation**	**Crash**	**Description**
Unnamed Metric	No	No	Indicates that the system did not activate, and the airplane did not crash. This metric, which was not given a name, is not useful for our research, as we are concerned with the protection abilities of the system. It will not be investigated.
Miss Activation Rate	No	Yes	Quantifies the instances in which the GCAS system should have activated to prevent a crash but failed to do

			so. This highlights situations where the system's failure to detect a collision risk resulted in a crash.
Protection Rate	Yes	No	Represents the percentage of cases where the GCAS system accurately detected an impending collision, successfully activated, and enabled the aircraft to recover without crashing. Trials designated within this category indicate nominal GCAS performance.
Protection Failure Rate	Yes	Yes	Measures the proportion of instances where the GCAS correctly activated but failed to prevent a crash, indicating when the system's activation did not lead to a successful aircraft recovery.

To detect any crashes, altitude/terrain comparison was employed. The "true" terrain database comparison was provided by United States Geological Survey (USGS) over the $1° \times 1°$ region specified in the ICD. After simulation, the aircraft's flight path was compared with the underlying terrain elevation at each time step. The equations below highlight the important relationship between aircraft altitude and the minimum approach to terrain (MAT). The MAT, analogous to the closest point of approach (CPA), is a meaningful metric that allows insight into the severity of the incident. It represents the smallest vertical distance between the aircraft and the underlying terrain. Because altitude/terrain comparison is computed post-simulation, the aircraft may fly underneath the terrain. Consequently, a $\text{MAT} \leq 0$ is perfectly valid in this research. If the $\text{MAT} \leq 0$, a crash occurred. Conversely, if the $\text{MAT} > 0$, a crash did not occur. Figure 50 depicts this important metric.

$$\text{agl}[] = \text{aircraft\_alt}[] - \text{terrain\_elv}[] \quad (13)$$

$$\text{MAT} = \min(\text{agl}) \quad (14)$$

$$\text{MAT} > 0 \rightarrow \text{Recovery} \quad (15)$$

$$\text{MAT} \leq 0 \rightarrow \text{Crash} \quad (16)$$

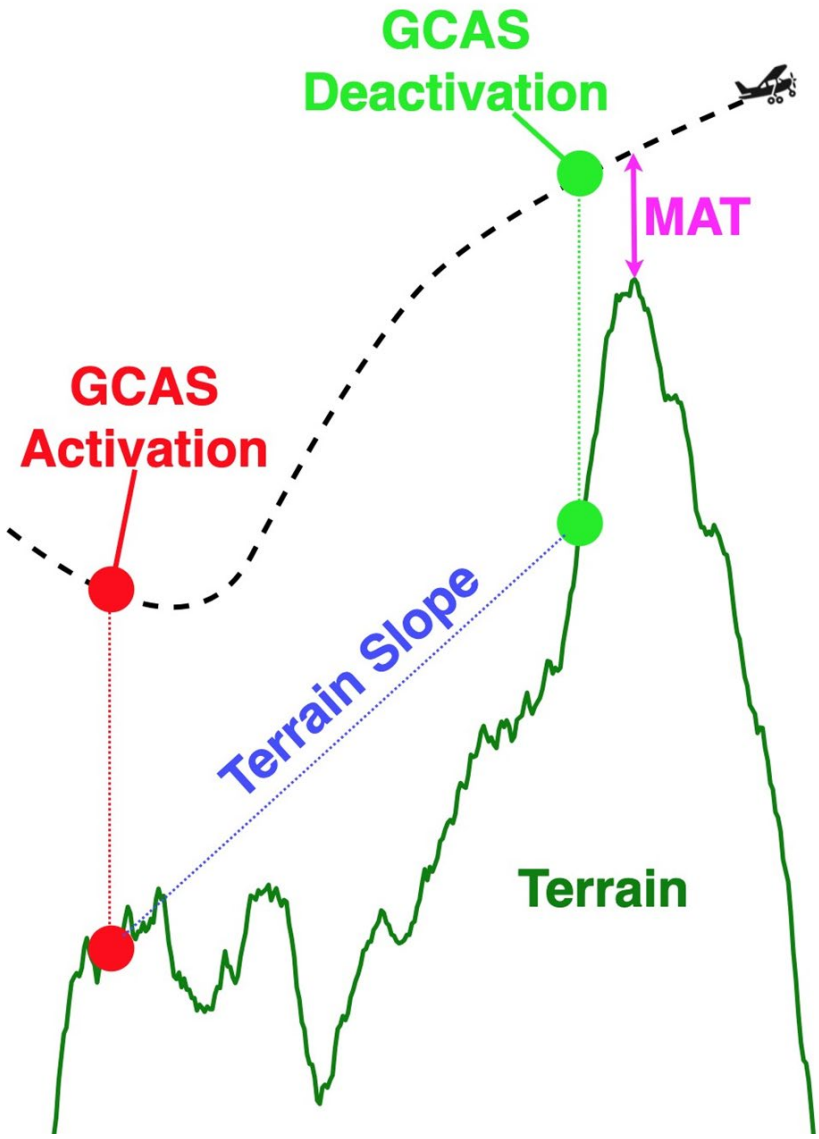

Figure 50. Scenario slope logic and the minimum approach to terrain (MAT)

5.5 Results

In total, 64493 unique scenarios were simulated above southern California airspace. Batch analysis showed insights into the high-level system performance of the GCAS under investigation. Table 17 reflects the high-level statistics obtained from all simulations. There was one trial where GCAS should have activated but did not — a *protection miss*. After investigation, it was determined that the activation condition was not appropriate because the aircraft spawned very close to terrain. Consequently, this trial was discarded from all analysis.

Table 18 shows the protection and failure rates for each maneuver. All maneuvers, though not evenly distributed, have comparable protection rates of approximately 97%. The maneuver with the lowest protection rate was the left-turning climb. This finding corroborates previous research which found poorer performance in left-turning recoveries. This is likely attributed to the TPA tuning, which should match the flight dynamics of a Cessna 172 simulated inside JSBSim.

Table 18. High-level statistics

Metric	**Value**
Trial Count	64493
Activation Count	49804
Collision Count	771
Miss Activation Rate	0.00%
Protection Rate	98.45%
Protection Failure Rate	1.55%

Table 19. Recovery maneuver statistics

Recovery Maneuver	**Count**	**Protection Rate**	**Protection Failure Rate**
Left Climb	7337	96.61%	3.38%
Straight Climb	598	97.65%	2.34%
Right Climb	41869	98.78%	1.21%

Figure 51 shows the 49804 activations from all 64493 trials. You can observe a densely populated square of data points, some red and some green. Green circles indicate a successful recovery, while a red x signifies a ground collision.

The color gradient across the map correlates to varying elevations, with the scale on the right indicating altitude in feet. The northern half of the region is characterized by predominantly flat terrain, evident by the uniform green color. In contrast, the southern half is a rugged, mountainous landscape, visually represented by the transition from green to brown and white, reflecting higher altitudes and steep slopes. This region was chosen by the authors to test the system’s safety performance amidst diverse terrain.

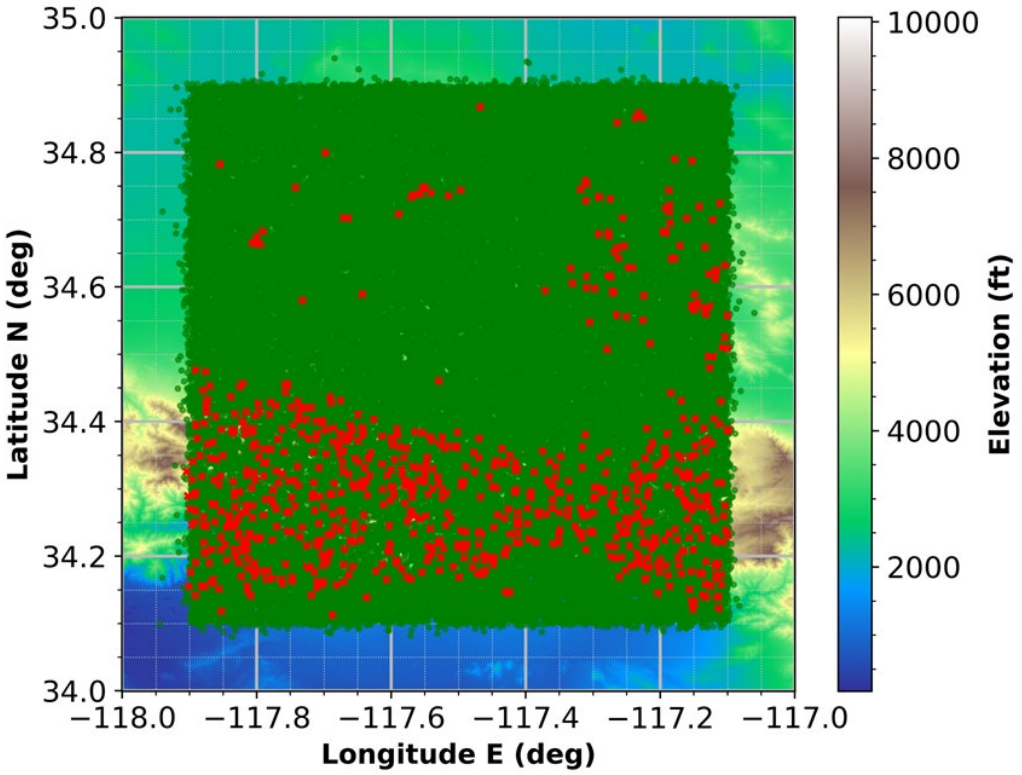

Figure 51. Activation map

The green circles (recoveries) are so numerous that they begin to obscure the terrain. To focus on the system's protection failures, Figure 52 hides all recoveries, showing only the ground collisions, each signified by a red x. The ground impacts are noticeably denser in the mountainous southern half. This pattern likely highlights the influence of strenuous terrain on aircraft performance. The suboptimal performance of the GCAS in the mountains can be attributed to the aircraft's reduced maneuverability and engine performance at the higher altitudes and steep inclines, which are typical characteristics of such challenging environments. These results are expected as a natural consequence of inhibited aircraft performance in strenuous environments.

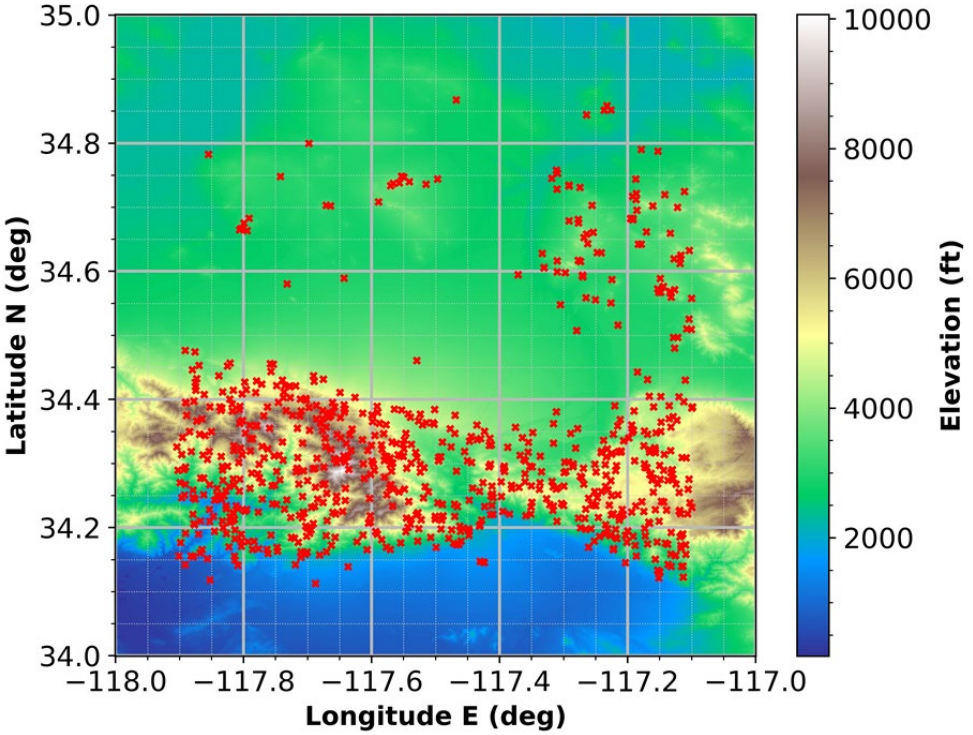

Figure 52. Crash map

5.6 Analysis

5.6.1 Aircraft state influence on minimum approach to terrain

This analysis aimed to find the relationship between the aircraft state variables and the resulting severity of the incident. In other words, which variables influence the GCAS's ability to protect, and which variables do not? This methodology does not account for interdependencies between variables. The correlation matrix for this research was computed using *pandas'* DataFrame.corr() method. Table 20 shows each variable's correlation with the MAT and its corresponding designation. For our research, equations 17, 18, and 19 define positive, non-influential, and negative correlations, correspondingly.

$$\text{PCC} > 0.05 \rightarrow \text{Positive Correlation} \quad (17)$$

$$|\text{PCC}| < 0.05 \rightarrow \text{Non-influential} \quad (18)$$

$$\text{PCC} < -0.05 \rightarrow \text{Negative Correlation} \quad (19)$$

Table 20. Pearson Correlation Coefficients (PCC) with the Minimum Approach to Terrain (MAT)

Metric Name	**PCC**
Positively correlated variables	
Minimum Approach to Terrain	1.000
Terrain Slope	0.448
Callibrated Airspeed	0.328
Density Altitude	0.320
Flight path angle γ	0.173
Terrain Database Discrepancy	0.126
Vertical velocity V_z	0.089
Non-influential variables	
Winds direction	0.001
Wind speed	-0.002
Pitch θ	-0.014
Z acceleration n_z	-0.023
Yaw ψ	-0.018
Ground track angle χ	-0.018
Roll ϕ	-0.046
Roll rate $\dot{\phi}$	-0.049
Negatively correlated variables	
Pitch rate $\dot{\theta}$	-0.086
Angle of attack α	-0.311

In the following figures, each trial is signified by a single scatter point. Each point provides valuable information regarding the aircraft's state at activation time and the trial's resulting outcome. The x-axis represents the aircraft state at the moment of GCAS activation, while the y-axis illustrates the resulting MAT for that specific trial.

Across all 64493 simulations, many aircraft states coincided with each other. This led to data points overlapping, which obscured valuable information about point distribution density. To mitigate this problem, a relative-density color map was incorporated into the plots. In this color map, dark purple represents low point density while bright orange signifies high point density. It

is important to note that the density within each figure is relative and not consistent across graphs.

What are we looking for? These figures represent big-picture, universal performance of the GCAS across many scenarios. In each scatterplot, we are aiming to identify any anomalies or patterns. For some variables, like ψ or χ, we expect a flat, uniform distribution of points. This is because heading should theoretically have no influence on the GCAS's ability to protect, assuming random and consistent terrain. On the other hand, variables like calibrated airspeed are expected to have a prominent influence on the GCAS's protection capabilities, as low airspeed will drastically reduce aircraft maneuverability and recoverability capabilities.

The dotted horizontal line at $y = 50ft$ represents the vertical clearance buffer (VCB) of the GCAS. The VCB defines the shortest vertical distance between the predicted flight path and underlying terrain before the GCAS triggers a recovery. In other words, if a predicted flight path comes within 50 ft of the underlying terrain, a recovery maneuver is initiated. The VCB, sometimes referred to as the terrain clearance buffer (TCB), is shown in Figure 53.

Most activations experienced a MAT above the VCB. This is evident by the dense cloud of activations shown in bright orange directly above the dotted horizontal line. This fact indicates that the system may be more prone to nuisance.

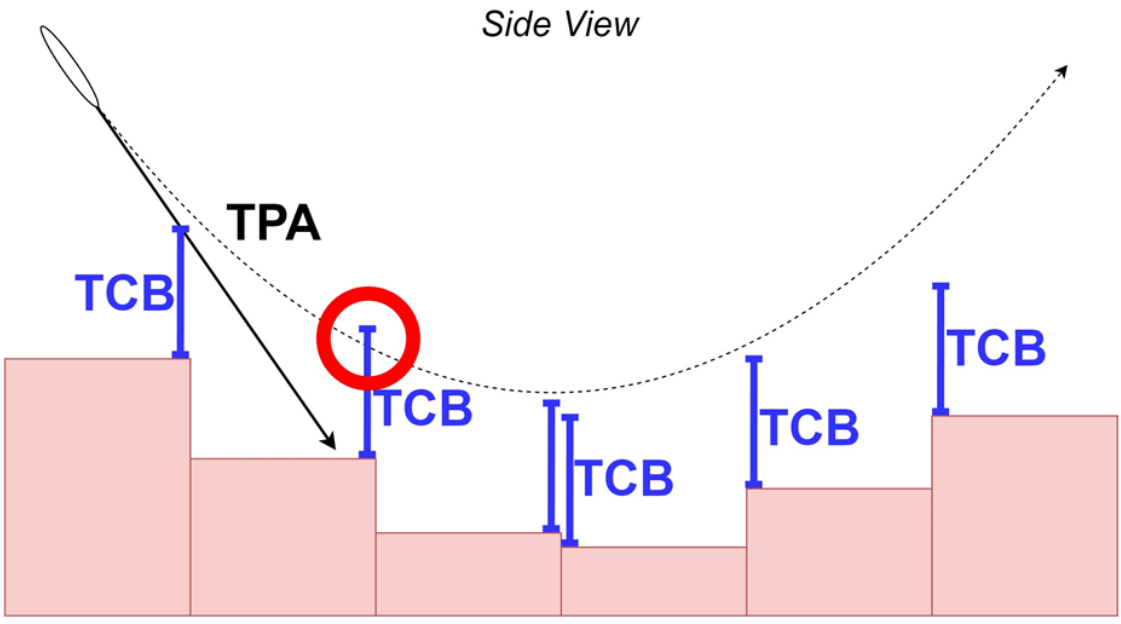

Figure 53. Vertical clearance buffer (VCB) or terrain clearance buffer (TCB)

Having laid the foundation to understand the following figures, consider the variables that have a positive correlation with the MAT. Note, to increase paging space on this document, the legend key, which is consistent all figures, is shown Figure 54.

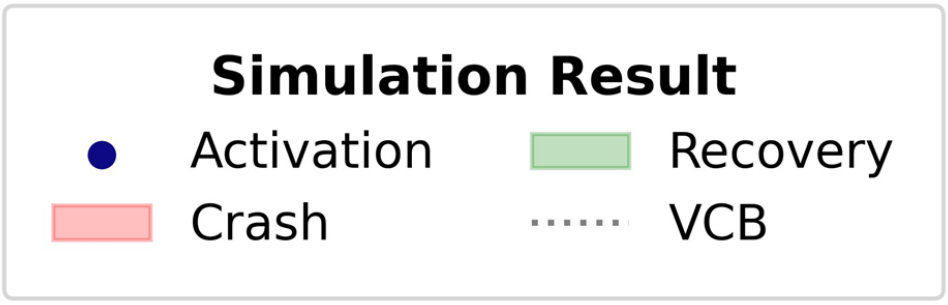

Figure 54. Legend key for MAT correlation analysis figure

5.6.2 Positively correlated variables

This section explores the aircraft state variables which have a positive relationship with the MAT.

5.6.2.1 Terrain slope

Figure 55 displays the MAT as a function of scenario slope. The scenario slope calculation is shown in Equation 20. It is a simple linear approximation of the terrain slope throughout the duration of the recovery, shown in Figure 55. In this research, a positive correlation between terrain slope and MAT was not expected. One hypothesis is that EVAA has an activation bias with steeper terrain. Another reasonable hypothesis is that the TPA undershoots the aircraft's true flight path. Previous experiments of the TPA found that all trajectories resided beneath the true flight path (Urban, Hubbard, Hook, Mark, & Sizoo, 2023). If a low TPA trajectory were to penetrate a terrain, it would likely trigger a recovery, leading to higher slope regions to cause more activations.

$$\text{slope} = \frac{\text{elev[deactivation]} - \text{elev[activation]}}{\text{cumulative distance traveled}} \tag{20}$$

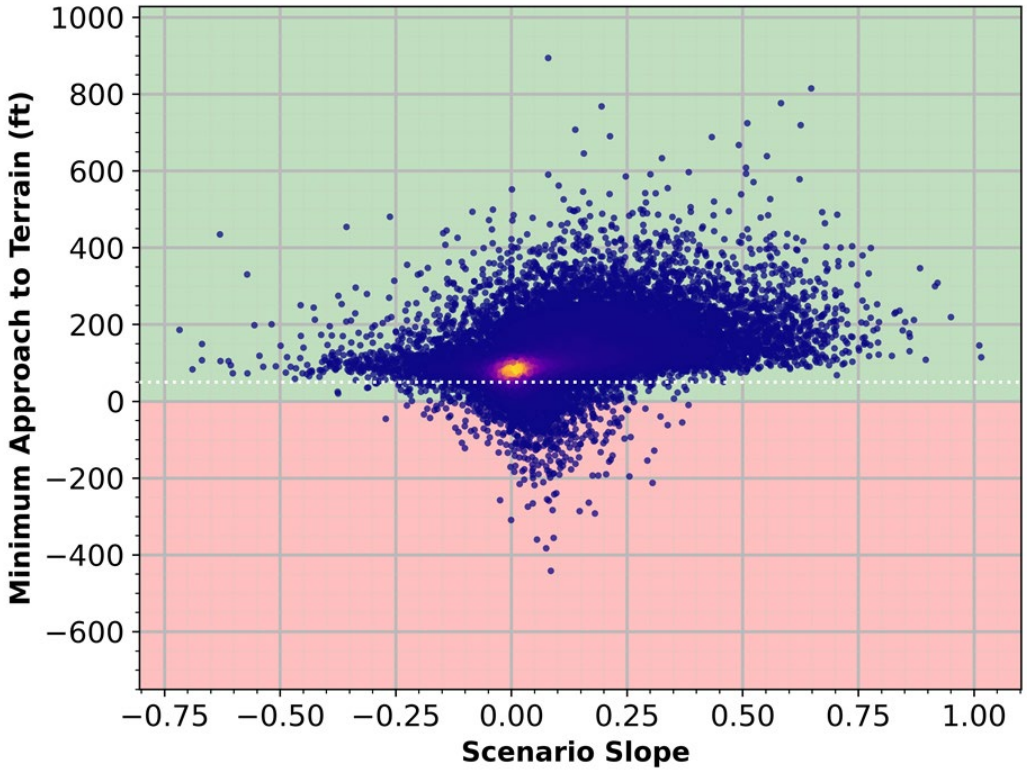

Figure 55. MAT as a function of scenario slope

5.6.2.2 Calibrated airspeed

Figure 56 highlights the system's safety ability with variation in KCAS. Given that most activations occurred with a nose-down configuration, as shown in Figure 58, a relatively high airspeed is expected as a consequence. Figure 56 reflects this with a high-density area of activations around 90 KCAS.

If you follow the cloud of activations, lower-airspeed scenarios are met with a lower MAT, while higher-airspeed trials result in a higher MAT. This observation makes clear that airspeed is positively correlated with GCAS protection.

There are two key takeaways from this experiment. Firstly, the GCAS demonstrates a strong performance across a majority of activations, providing a significant safety net during most flight conditions. Secondly, there is a notable decrease in system protection at lower airspeeds. This trend highlights the aeronautical understanding that the aircraft may lack the sufficient power to maneuver away from terrain at lower airspeeds. This behavior is natural and anticipated in the design and functionality of any GCAS. Still, the majority of activations experienced a comfortable (and perhaps nuisance-prone) MAT > VCB.

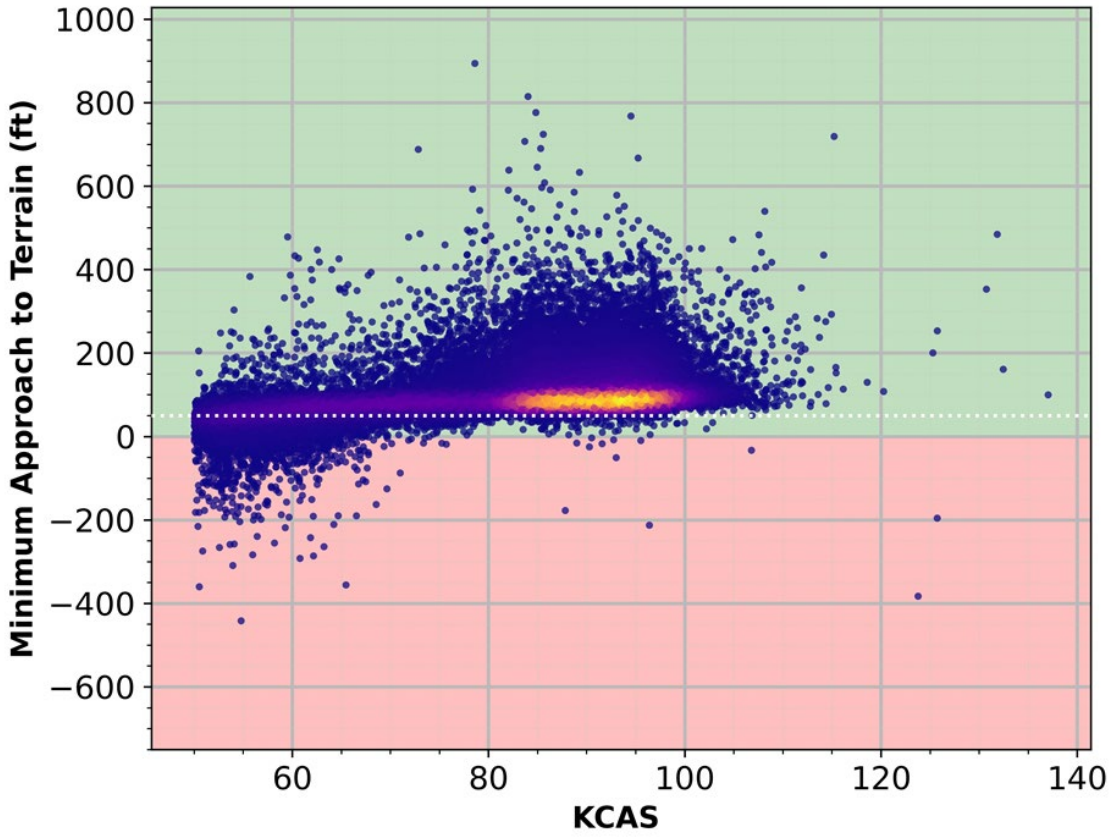

Figure 56. MAT as a function of KCAS

5.6.2.3 Density altitude

Figure 57 reveals the system's performance as a function of density altitude. A very dense region of activations can be observed at an altitude of approximately 3000 ft. This anomaly can be attributed to the fact that most of the terrain is around 3000 ft in elevation (Figure 52). The figure depicts sparsely distributed data points with no apparent clustering or grouping. This implies that the protection abilities of the system are not inhibited at any particular altitude. A larger testing area with more diverse terrain may be necessary to obtain a better representation of the system's performance with varying altitude. In this research, most activations occurred around 3000 ft, with higher-altitude activations being sparser.

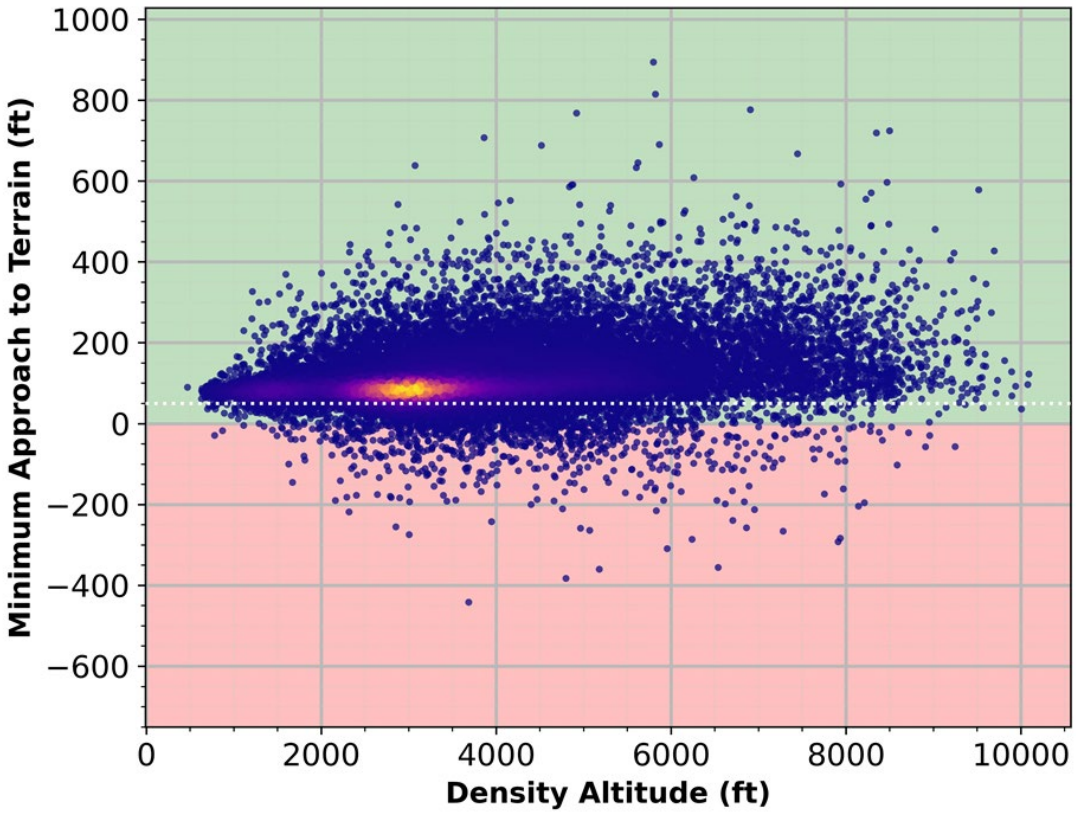

Figure 57. MAT as a function of density altitude

5.6.2.4 Flight path angle

Figure 58 highlights the system's protection performance as a function of the aircraft's flight-path angle γ at activation. Overall, non-extreme descent values of gamma (e.g. $-10° \leq \gamma \leq 0°$) do not appear to influence the MAT. Instead, the outer edge outliers (like $\gamma < -10°$ and $\gamma > 10°$) appear contribute to a stronger positive slope. From the figure, infer that "nominal" values of γ have little influence on the protection abilities of the system, while more extreme values of γ do.

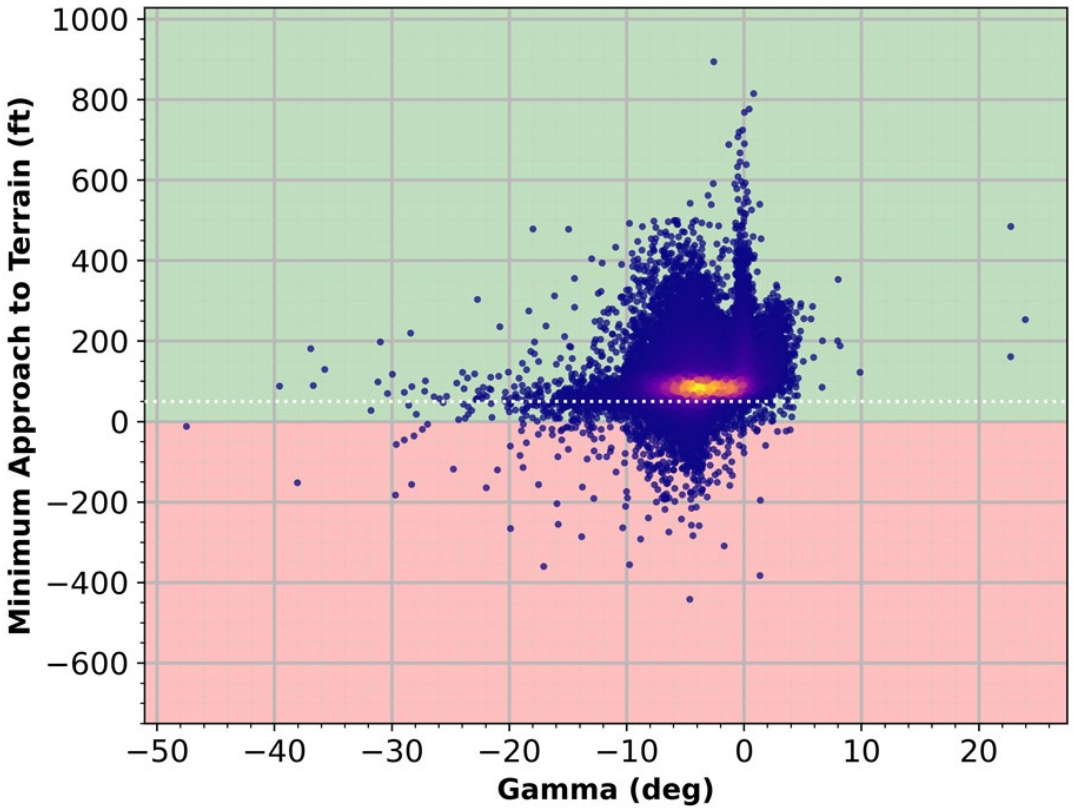

Figure 58. MAT as a function of gamma

5.6.2.5 Terrain database discrepancy

Figure 59 reveals the system's protection abilities as a function of the terrain database error. For context, the compressed digital terrain map (CDTM) was developed by NASA for Auto GCAS applications (Suplisson, 2015). The same technology is used in EVAA's GCAS and is a variable under investigation. A separate database provided by the USGS features a high-fidelity digital map of the 1° × 1° southern California region specified in the ICD. On the other hand, EVAA's CDTM is a coarser database which uses binary-tree tip-tilt compression techniques (Suplisson, 2015). This experiment aimed to study the relationship between local terrain database uncertainties and the resulting MAT.

From the figure, most of the activations occur left of the y-axis. This indicates that there is a slight positive-up elevation bias on the CDTM database compared with the USGS database. This is an inherent consequence of the binary-tree tip-tilt compression technique as it vertically elevates each tile above all of its source raster data. This is desirable for a GCAS because errors are only allowed in the conservative direction (Suplisson, 2015). The figure shows a high-density

region around x = −15 ft. This implies most of the CDTM map had a positive-up bias of ≈ 15 ft compared to the high-fidelity USGS terrain database for the southern California region tested, in accordance with the compression algorithm. The largest terrain discrepancy occurred with CDTM = USGS −150 ft. In this trial, the 150-ft difference still resulted in a successful recovery. Overall, small errors in the GCAS terrain database do not appear to largely influence the protection abilities of the system.

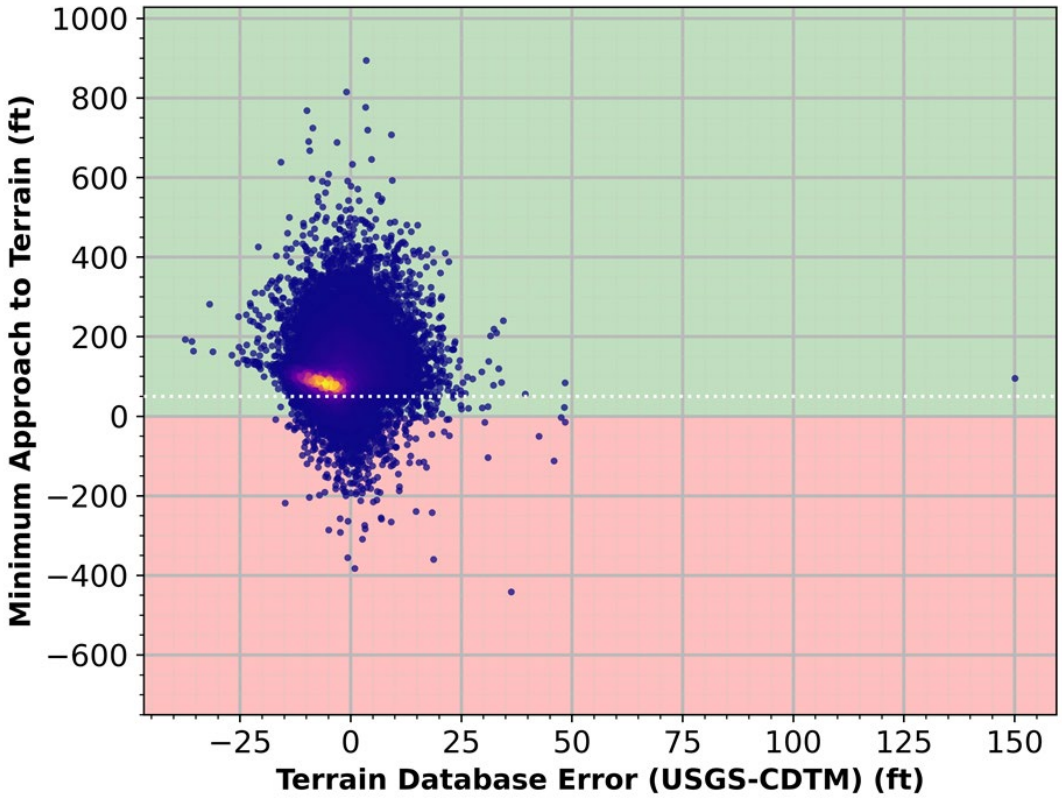

Figure 59. MAT as a function of terrain database error

5.6.3 Non-influential variables

There are aircraft state variables which have little to no impact on the MAT. For the non-influential variables, one of two things can be assumed. Either the system models that particular variable well or the variable is inherently uniform. Neither classification will influence the protection abilities of the system.

5.6.3.1 Roll angle

Figure 60 reveals the system's performance as a function of roll (ϕ). In theory, we expect ϕ to have a non-zero influence on the system's protection ability. However, the TPA roll modelling capabilities appear to be quite good, as the high-density cloud of activations appears consistently linear from $-60^{\circ} \leq \phi \leq 60^{\circ}$. Across a wide range of roll angles, no roll-dependent anomalies or biases appear. This uniformity implies that GCAS has a nominal model of aircraft roll, as it has no impact on the protection rate of the system, suggesting a well-functioning system in terms of roll angle variability.

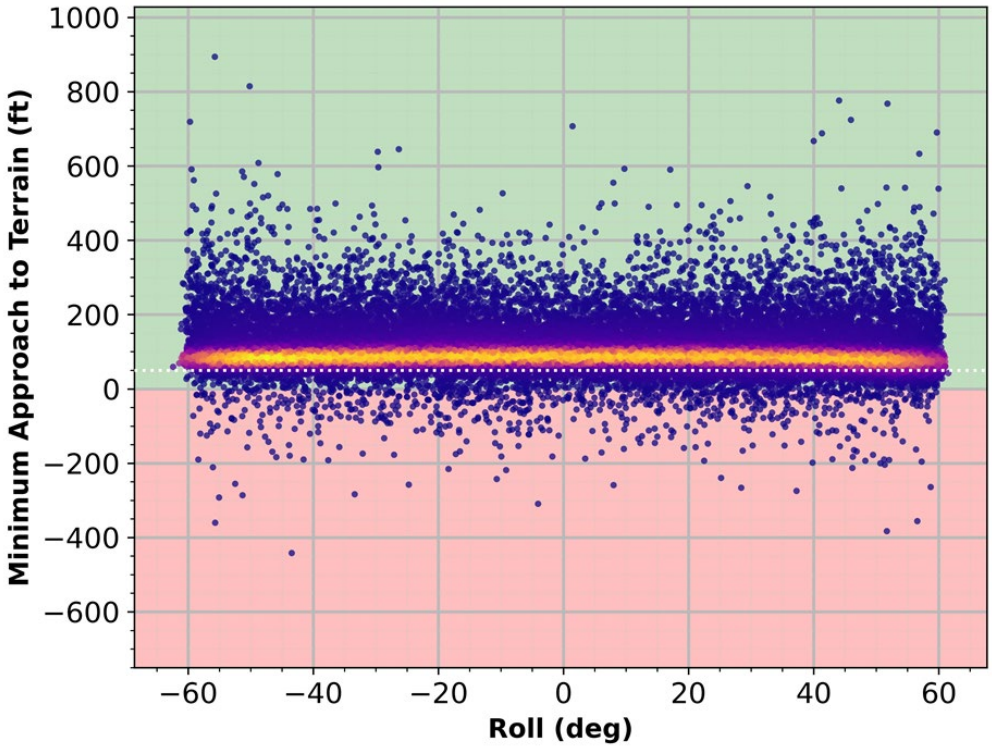

Figure 60. MAT as a function of roll

5.6.3.2 Winds

Assuming uniform and consistent terrain, the direction of wind should have no impact on the protection abilities of the GCAS. Figure 61 reflects this expectation, as the distribution is uniform and is not influenced by the direction of wind. Figure 62 highlights the system's performance with different wind speeds. An χ^2 distribution with six degrees of freedom was used to generate the wind speed for each trial. To prevent extreme outliers, a saturation wind speed of 35 knots was imposed. Equation 21 shows the logic for environmental wind speed creation.

$$v_{kts}^{wind} = \min(2\chi_6^2, 35) \tag{21}$$

As expected by the χ^2 distribution, the majority of activations occur in lower wind speeds, around five to ten knots. The activations become thinner with increasing wind speeds. All outlier wind speeds are caught and saturated to 35 knots, which projects a vertical wall at $x = 35$.

The activation density cloud for this figure strongly implies that the GCAS wind modelling capabilities are operating nominally because the cloud resides above the VCB = 50ft barrier. From this, it is safe to assume that the GCAS functions well in nominal flight conditions regardless of the wind speed.

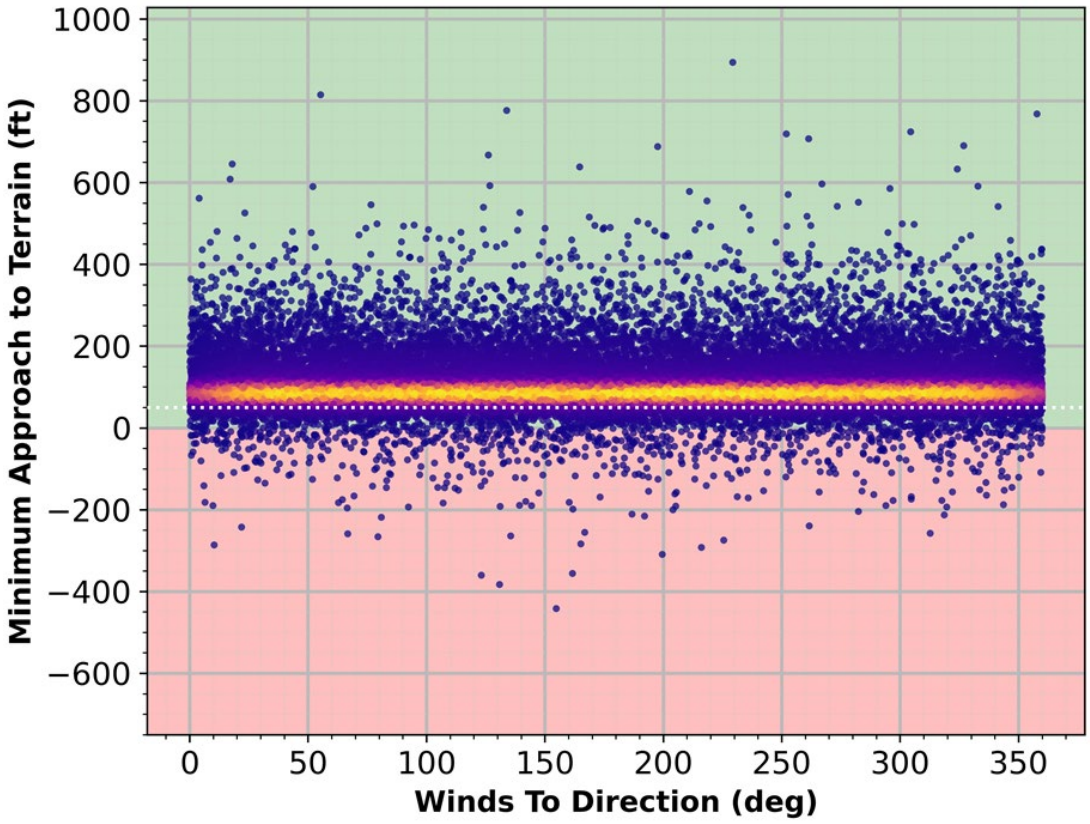

Figure 61. MAT as a function of wind direction

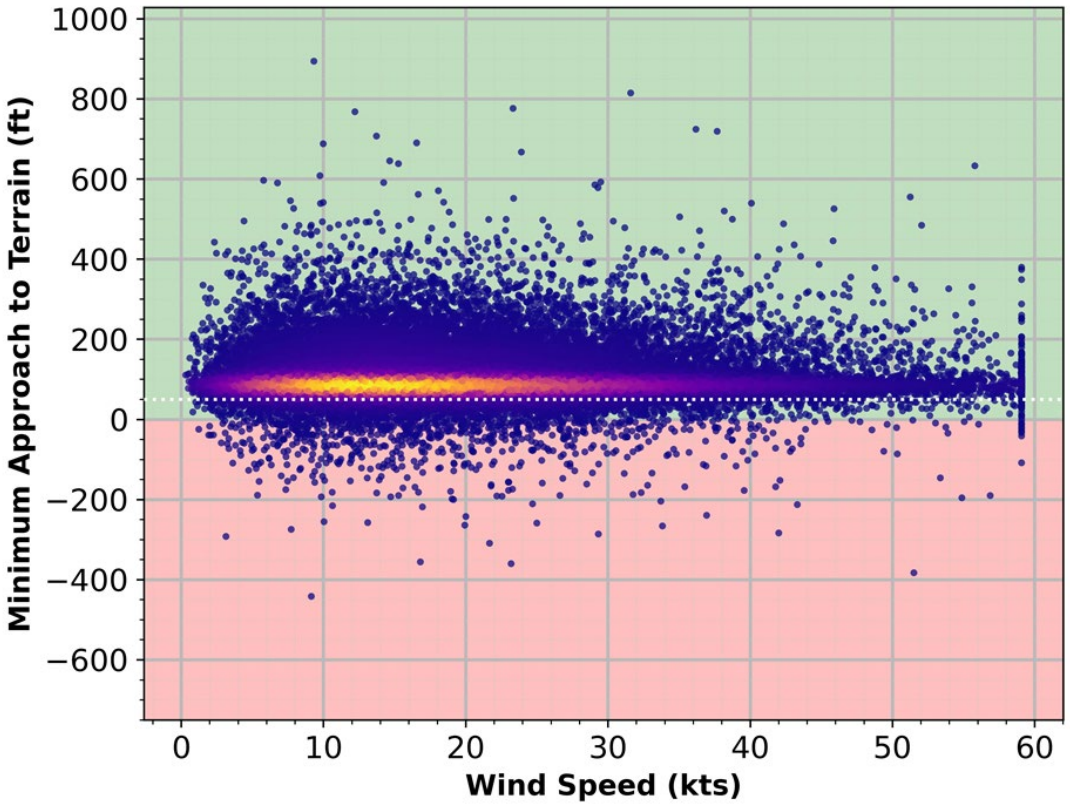

Figure 62. MAT as a function of wind speed

5.6.3.3 Normal acceleration

Figure 63 illustrates the system's performance as a function of normal acceleration n_z. Though a large variety of normal accelerations occurred at activation time, no observable anomalies were found. This implies that the protection abilities of the GCAS are not affected by n_z.

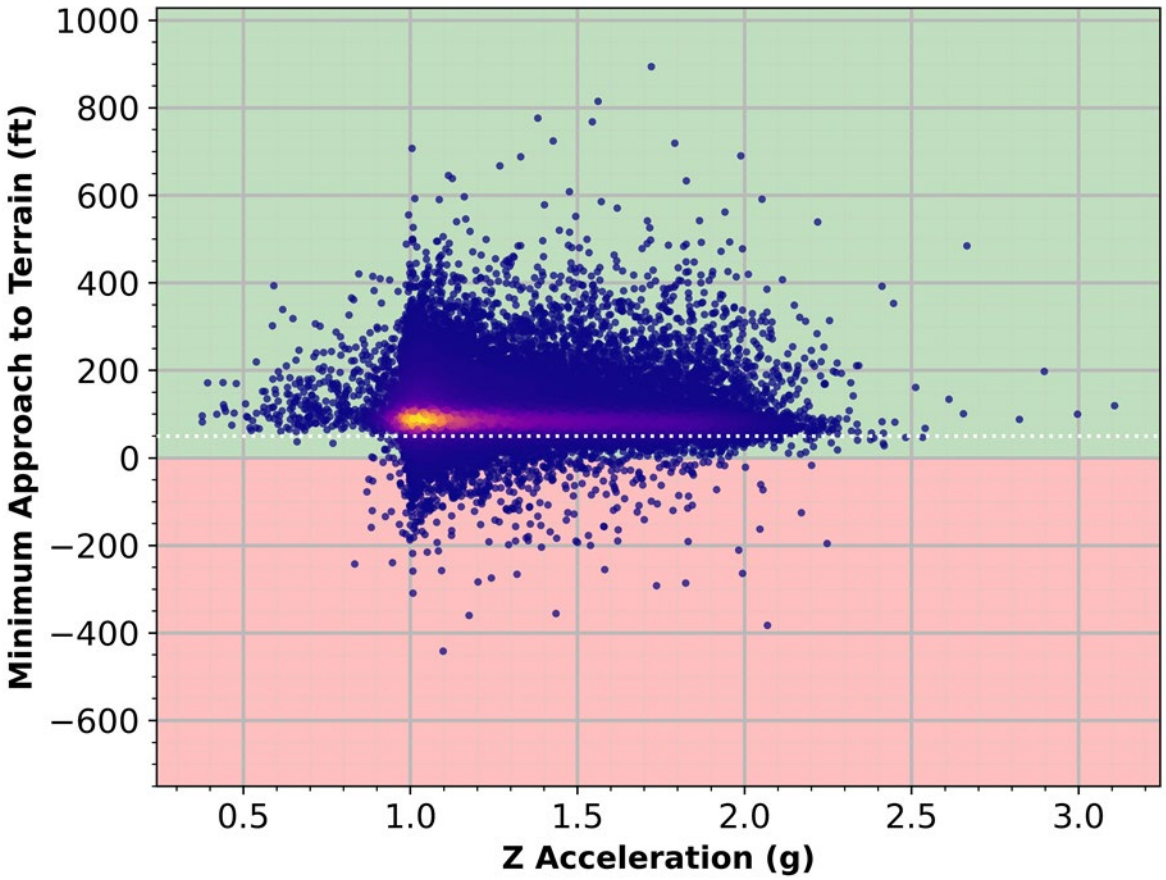

Figure 63. MAT as a function of normal acceleration

5.6.4 Negatively correlated variables

Aircraft state variables exist which have a negative relationship with the MAT. Large values within this category will negatively influence the MAT, decreasing the protection ability of the system and increasing the likelihood of a collision.

5.6.4.1 Angle of attack

Figure 64 reveals the system's performance as a function of angle of attack α. This research discovered that α is the only variable that has a negative correlation with the MAT. This means that a higher angle of attack is associated with a smaller MAT. Intuitively, a large α is more prone to stall the aircraft, leading to a collision and a MAT ≤ 0. This natural aircraft behavior may contribute to this uncovered anomaly.

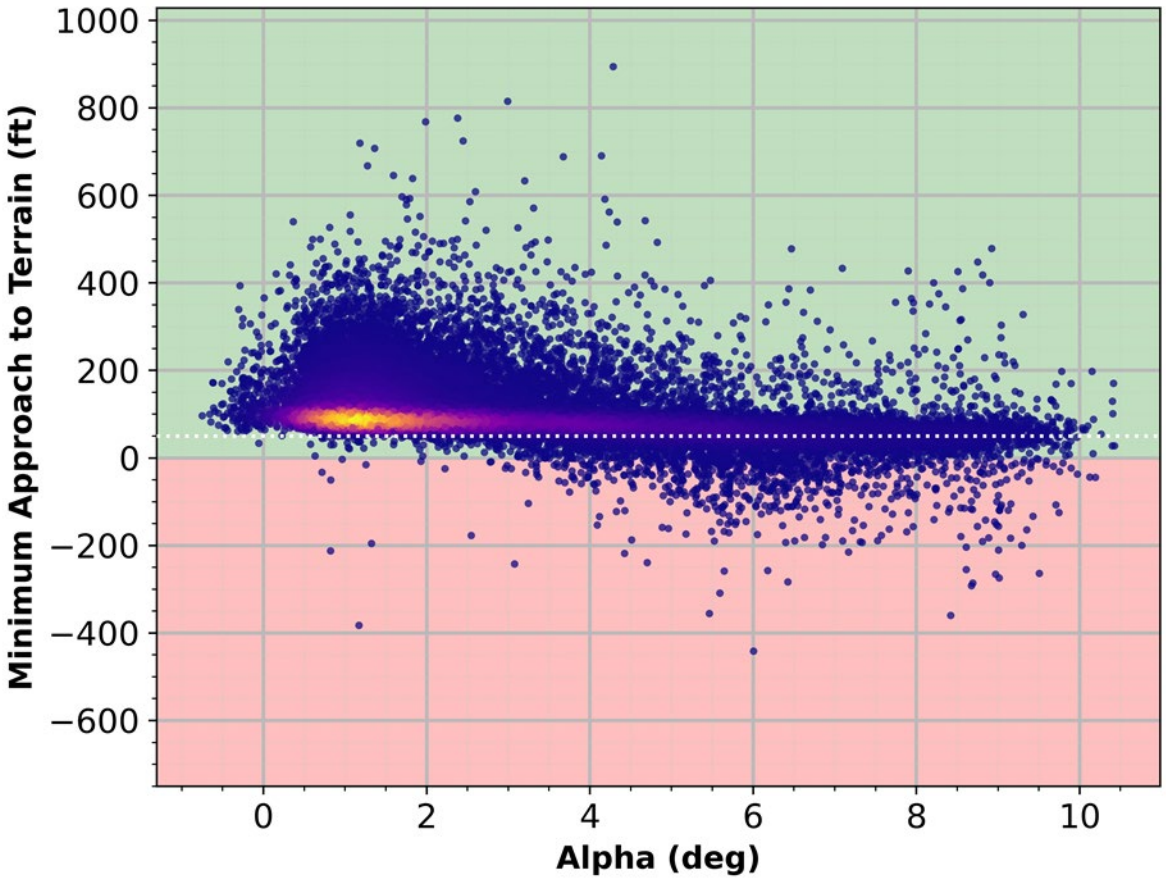

Figure 64. MAT as a function of angle of attack

5.7 Final remarks

Overall, EVAA's GCAS was found to have a 98.45% protection rate in over 49,804 activations. This highlights the system's ability to protect the pilot in many different conditions, both nominal and extreme.

5.7.1 Terrain

GCAS works better over flat terrain than in mountains, a natural consequence of slope diversity and other uncertainties, as shown in Figure 52 indicate that the system has inhibited performance amidst mountainous terrain compared to flat terrain. This is intuitive, a natural consequence of imperfect modeling capabilities and other uncertainties.

5.7.2 Variable correlation

GCAS has certain variables that may or may not have correlation with the MAT, indicative of incident severity. These results, though intuitively expected, can be used to further tune the GCAS, increasing protection abilities in more unfortunate circumstances.

- The variables which have a **positive** influence on the MAT are scenario slope, airspeed, density altitude, flight path angle, terrain database discrepancy, vertical velocity.
- The variables which have **no influence** on the MAT are winds, pitch, normal acceleration, yaw, course angle, roll, and roll rate.
- The variables which have a **negative** influence on the MAT are pitch rate and angle of attack.

This paper presented a Monte Carlo analysis-based evaluation of a GCAS to assess its protective capabilities. This research provides simulation and analysis techniques to evaluate a GCAS by providing insights into its performance across a large number of scenarios. The methodology developed during this research contributes to aviation safety by offering a robust framework for analyzing technologies that can save lives. By improving the evaluation process, this paper supports the ongoing advancement of safe flight operations and the development of reliable life-saving systems in aviation.

5.8 Future work

The correlation matrix, while meaningful, does not remove interdependencies between variables. To perform a more thorough analysis on a GCAS, a model of the system could be created using statistical or machine learning techniques. This approach could help identify which variables most significantly influence the effectiveness of the GCAS. By employing these techniques, the relationships, interactions, and interdependencies between variables could be explored. These insights could be very beneficial in the creation and tuning of GCAS parameters for optimal performance and safety.

Our sincere thanks are extended to the FAA for their substantial support and funding, which were pivotal in the successful completion of this project. Additionally, we acknowledge and appreciate the contributions of those involved in our study, including NASA and Blackbird Aerospace, whose cooperation and support were invaluable in achieving the results of this work.

6 Conclusions

6.1 Final remarks in the evaluation of NASA's EVAA GCAS

Our findings show that EVAA is a very capable system in a large domain of flight envelopes. It provides extensive protection in both nominal and more extreme scenarios.

6.1.1 Nuisance

The investigation into nuisance revealed that the EVAA's GCAS tends to be particularly susceptible around terminal areas. This was determined using real-flight ADS-B data to assess inappropriate activations, finding marginally acceptable results. In all processed flights, research found that GCAS may have activated roughly one-third of the time, showing pronounced nuisance (Table 3). This is significant because these were real flights that did not result in a collision, yet GCAS likely would have activated if it found itself in that scenario. Leveraging ADS-B data from actual flights proved to be an effective method for assessing the GCAS's susceptibility to unwanted interference.

6.1.2 Protection

For the assessment of system protection, the GCAS was integrated with a proprietary recovery autopilot. In every analysis focused on protection, this recovery autopilot represented a real world aircraft, which the GCAS sought to safeguard. The GCAS TPA was tuned to this particular recovery autopilot to align with the simulated aircraft's flight path. Following the creation of the recovery autopilot and tuning of the TPA, our research ventured into two distinct dimensions of protection evaluation: TPA analysis and a holistic system analysis. In the TPA analysis, we compared the actual simulated flight path against the TPA's anticipated path. This comparison is crucial as terrain not included in the predicted path is not considered for ground collision risks. Consequently, an inaccurate projected trajectory undermines the GCAS decision-making accuracy. Our findings indicate that the TPA generally performs exceptionally well. Further tuning for left-turning maneuvers may be beneficial to system performance (see Paper 3: TPA evaluation). In the system-wide analysis, Monte Carlo analysis methods were used to simulate a myriad of unique scenarios, offering a macroscopic view of the system's efficacy. These simulations, totaling over 60,000 distinct scenarios, revealed that EVAA's GCAS maintains a high level of safety, achieving an overall protection rate of 98.45% (Table 18).

6.1.3 Results

The system-wide analysis corroborates the findings from the nuisance and protection studies, finding both a high protection rate and potential nuisance. The ADS-B-based nuisance analysis indicated a higher-than-ideal activation rate. The system-wide evaluation revealed that most activations occurred prematurely which results in a MAT > VCB, thus ensuring a significant buffer between the aircraft and terrain. Additionally, the TPA study highlighted potential weaknesses in left-turn recoveries. The system-wide analysis showed a left-turning climb protection rate of 96.61%, slightly lower than the 97.65% for straight climbs and 98.78% for right-turning climbs Table 19. In any GCAS, there is an inherent trade-off between protection and nuisance. Our studies suggest that EVAA's GCAS leans towards high-nuisance and high-protection. Further discussions to attenuate some nuisance while sacrificing some protection may be required to make the system more well balanced. Overall, the GCAS demonstrates strong protective capabilities, but additional fine-tuning could further reduce nuisance issues. This project has successfully met and surpassed the expected benchmarks in evaluating both protection and nuisance of the GCAS under investigation, advancing the goal of using technology to decrease aviation fatalities.

All tasks outlined as a part of the University of Tulsa's role in the evaluation of GCAS were completed successfully. Using real-world ADS-B to evaluate nuisance of GCAS proved extremely effective as did Monte Carlo analysis of the protection capabilities. Future work to expand these technologies to additional safety systems for general aviation should be performed.

7 References

Amin, S., Clark, T., Offutt, R., & Serenko, K. (2014). Design of a cyber security framework for ADS-B based surveillance systems. *2014 Systems and Information Engineering Design Symposium (SIEDS)* (pp. 304-309). IEEE. doi:10.1109/SIEDS.2014.6829910

Anderson, T., Jones, W., & Beamon, K. (2011). Design and implementation of TAWS for rotary wing aircraft. *2011 Aerospace Conference* (pp. 1-7). IEEE. doi:10.1109/AERO.2011.5747519

ASTM International. (2017). ASTM F3269-17, Standard practice for methods to safely bound flight behavior of unmanned aircraft systems containing complex functions.

ASTM International. (2021). ASTM 3269-21 standard practice for methods to safely bound behavior of aircraft systems containing complex functions using run-time assurance.

Beard, R. W., & McLain, T. W. (2012). *Small unmanned aircraft: Theory and practice.* Princeton University Press.

Berndt, J. (2004). JSBSIM: An open source flight dynamics model in C++. *AIAA Modeling and Simulation Technologies Conference and Exhibit*, (p. 4923). Retrieved from https://jsbsim.sourceforge.net/JSBSimAnOpenSourceFlightDynamicsModelinCPP.pdf

Breen, B. C. (1999). Controlled flight into terrain and the enhanced ground proximity warning system. *IEEE Aerospace and Electronic Systems Magazine, 14*(1), pp. 19-24. doi:10.1109/62.738350

CAST/ICAO Common Taxonomy Team (CICTT). (2011). Aviation occurrence categories. p. 13. Retrieved from https://www.icao.int/APAC/Meetings/2012_APRAST/OccurrenceCategoryDefinitions.pdf

Cessna 172N pilot's operating handbook. (1977). Cessna Aircraft Company.

Connor, M. (2020, October). *Resilient autonomy project develops EVAA software*. Retrieved from Nasa.gov: https://www.nasa.gov/centers/armstrong/features/ resilient-autonomy-project-develops-evaa-software.html

Federal Aviation Administration. (2009). *Pilot's handbook of aeronautical knowledge.* Skyhorse Publishing , Inc.

Franklin, G. F., Powell, J. D., Emami-Naeini, A., & Powell, J. D. (2002). *Feedback control of dynamic systems* (Vol. 4). Upper Saddle River: Prentice Hall.

Fuller, J. G., & Hook, L. R. (2020). Understanding general aviation accidents in terms of safety systems. *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, (pp. 1-9). doi:10.1109/DASC50938.2020.9256778

Fuller, J. G., & Hook, L. R. (2023). Effectiveness of automatic safety systems in reducing general aviation fatalities. *Journal of Aerosace Inforamtion Systems, 20*(5), 240-250. doi:10.2514/1.I011069

Hook, L. R., Clark, M., Sizoo, D., Skoog, M. A., & Brady, J. (2016). Certification strategies using runtime safety assurance for part 23 autopilot systems. *2016 IEEE Aerospace Conference* (pp. 1-10). IEEE. doi:10.1109/AERO.2016.7500817

Hook, L. R., Ryan, W., Skoog, M. A., & Fuller, J. (2023). Exploring the potential of automatic safety systems in general aviation. *2023 IEEE Aerospace Conference* (pp. 1-10). IEEE. doi:10.1109/AERO55745.2023.10115864

Hook, L. R., Sizoo, D., & Fuller, J. G. (2022). How digital safety systems could revolutionize aviation safety. *2022 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)* (pp. 1-9). IEEE. doi:10.1109/DASC55683.2022.9925863

Hook, L. R., Skoog, M., Garland, M., Ryan, W., Sizoo, D., & VanHoudt, J. (2018). Initial considerations of a multi-layered run time assurance approach to enable unpiloted aircraft. *2018 IEEE Aerospace Conference* (pp. 1-11). IEEE. doi:10.1109/AERO.2018.8396622

International Civil Aviation Organization. (2009). Review of the classification and definitions used for civil aviation activities. *Tenth Session of the Statistics Division.* Montreal. Retrieved from https://www.icao.int/meetings/sta10/documents/sta10_wp007_en.pdf

Maley, P. D., Hubbard, A. M., Urban, J. M., & Hook, L. R. (2023). Recovery autopilot analysis for a general aviation ground collision avoidance system. *2023 IEEE Aerospace Conference* (pp. 1-12). IEEE. doi:10.1109/AERO55745.2023.10115694

Nagarajan, P., Kannan, S. K., Torens, C., Vukas, M. E., & Wilber, G. F. (2021). ASTM F3269-AN industry standard on run time assurance for aircraft systems. *AIAA Scitech 2021 Forum*, p. 0525. doi:doi.org/10.2514/6.2021-0525

National Transportation Safety Board. (2021). *Accidents, Fatalities, and Rates, 2020 Statistics, US Aviation.* [Excel Spreadsheet]. Retrieved from https://www.ntsb.gov/safety/data/Documents/AviationAccidentRates_ 2001-2020_20211014.xlsx

Ostroumov, I., KUzmenko, N., & KYzymchuk, O. (2022). Automatic dependent surveillence-broadcast trajetory data processing. *2022 IEEE 16th International Conference on Advanced Trends in Radioelectronics, Telecommunications, and Computer Engineering (TCSET)* (pp. 43-47). IEEE. doi:10.1109/TCSET55632.2022.9767058

Skoog, M. A., Hook, L. R., & Ryan, W. (2020). Leveraging ASTM industry standard F3269-17 for providing safe operations of a highly autonomous aircraft. *2020 IEEE Aerospace Conference* (pp. 1-7). IEEE. doi:10.1109/AERO47225.2020.9172434

Sorokowski, P., Skoog, M., Burrows, C. S., & Thomas, S. (2015). *Small UAV automatic ground collision avoidance system design considerations and flight test results.* National Aeronautics and Space Administration. Retrieved from https://ntrs.nasa.gov/api/citations/20150014106/downloads/20150014106.pdf

Stevens, B. L., Lewis, F. L., & Johnson, E. N. (2015). *Aircraft control and simulation: dynamics, controls design, and autonomous systems.* John Wiley & Sons.

Sun, J., Ellerbroek, J., & Hoekstra, J. (2016). Large-scale flight phase identification from ADS-B data using machine learning methods. In D. Lovell, & H. Fricke (Ed.), *7th International Conference on Research in Air Transportation*, (pp. 1-8).

Suplisson, A. W. (2015). *Optimal recovery trajectories for automatic ground collision avoidance systems (AUTO GCAS).* [Doctoral dissertation], Air Force Institute of Technology, Ohio. Retrieved from https://scholar.afit.edu/etd/183/

Swihart, D. E., Barfield, A. F., Griffin, E. M., Lehmann, R. C., Whitcomb, S. C., Flynn, B., . . . Processor, K. E. (2011). Automatic ground collision avoidance system design, integration & flight test. *IEEE Aerospace and Electronic Systems Magazine*, pp. 4-11. doi:10.1109/MAES.2011.5871385

T. Pandas Development Team. (2020, February). *pandas-dev/pandas:Pandas.* doi:https://doi.org/10.5281/zenodo.3509134

U.S. Geological Survey. (2023). *TIFF terrain map of N35W118.* Retrieved from https://prd-tnm.s3.amazonaws.com/index.html?prefix=StagedProducts/Elevation/13/TIFF/current/n35w118/USGS_13_n35w118_20221019.tiff

Urban, J. M., Hubbard, A. M., Hook, L. R., Mark, S., & Sizoo, D. (2023). Evaluation of a trajectory prediction algoritm within a ground collision avoidance system. *2023 IEEE/AIAA 42nd Digital Avionics Systems Conference (DASC).* doi:10.1109/DASC58513.2023.10311267

Wetsig, W. (2021). *Congressional report commends AFRL for life-saving collision avoidance technology; integrated air and ground system remains near transition.* Air Force Research Laboratory Public Affairs. Retrieved from https://www.afrl.af.mil/News/Article/2678754/congressional-report-commends-afrl-for-life-saving-collision-avoidance-technolo/