ECE 484: Principles of Safe Autonomy	Fall 2026
Written by: ECE 484 team	*MP0*
Website 	*Due Date: September 11, 2026*

1 HW0: Safety and Reachability

HW Problem 1 [10 points] [Individual] For an automaton $\mathcal{A} = \langle Q, Q_0, D \rangle$ and an unsafe set $U \subseteq Q$

(a) **[5 points]** Is Q an inductive invariant? Briefly explain your answer.

(b) **[5 points]** Write down the conditions for a set $I \subseteq Q$ to be an inductive invariant that helps prove safety with respect to U.

HW Problem 2 [10 points] [Individual] Let $I \subseteq Q$ be an invariant of an automaton $\mathcal{A} = \langle Q, Q_0, D \rangle$.

(a) **[5 points]** Must I necessarily satisfy $\text{Post}(I) \subseteq I$? Briefly explain your answer.

(b) **[5 points]** Explain why every inductive invariant is an invariant, but an invariant does not necessarily have to be an inductive invariant. You may use a small counterexample automaton in your explanation.

HW Problem 3 [10 points] [Individual] Suppose $I_1 \subseteq Q$ and $I_2 \subseteq Q$ are both inductive invariants of an automaton $\mathcal{A} = \langle Q, Q_0, D \rangle$, prove that $I_1 \cap I_2$ is also an inductive invariant of $\mathcal{A}$.

HW Problem 4 [35 points] [Individual] Consider the finite-state automaton shown below. The initial state is q_0.

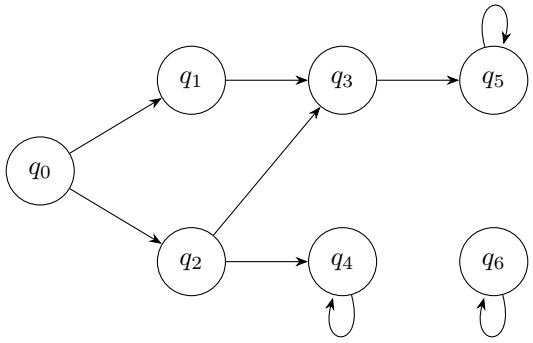

Define $Q = \{q_0, q_1, q_2, q_3, q_4, q_5, q_6\}$, $Q_0 = \{q_0\}$, $R_0 = Q_0$ and $R_{k+1} = R_k \cup \text{Post}(R_k)$.

(a) **[7 points]** Compute $\text{Post}(\{q_0\})$ and $\text{Post}(\{q_1, q_2\})$.

(b) **[10 points]** Compute $R_0, R_1, R_2, \ldots$ until the algorithm reaches a fixed point.

(c) **[6 points]** Write down the complete reachable set $Reach_A$.

(d) **[6 points]** Suppose the unsafe set is $U_1 = \{q_6\}$. Is the automaton safe with respect to U_1? Explain using $Reach_A$.

(e) **[6 points]** Now suppose $U_2 = \{q_4\}$. Is the automaton safe? If not, give a shortest counterexample execution.

HW Problem 5 [35 points] [Individual] Consider a simplified discrete-time emergency-braking system. The state is $q = (d, v)$, where

- $d \in \mathbb{R}_{\geq 0}$ is the distance from the vehicle to a stationary obstacle;
- $v \in \mathbb{Z}_{\geq 0}$ is the vehicle speed.

We use normalized units in which one time step is one unit of time and braking reduces the speed by one unit per step.

The system evolves according to the following discrete-time dynamics:

$$(d', v') = \begin{cases} (d - v,\ v - 1), & v \geq 1, \\ (d,\ 0), & v = 0. \end{cases}$$

The unsafe set is $U = \{(d, v) \mid d \leq 0\}$ and the initial state is $Q_0 = \{(d_0, v_0)\}$.

First consider the candidate invariant $I_1 = \{(d, v) \mid d > 0, v \geq 0\}$.

(a) **[7 points]** Is I_1 an inductive invariant? If not, give a specific state $q \in I_1$ whose successor satisfies $q' \notin I_1$ and briefly explain why the condition $d > 0$ alone is too weak.

(b) **[14 points]** Now consider the stronger candidate invariant

$$I_2 = \left\{ (d, v) \;\middle|\; d > \frac{v(v+1)}{2},\ v \geq 0 \right\}.$$

Prove that $\text{Post}(I_2) \subseteq I_2$.

(c) **[14 points]** Give a condition on the initial state (d_0, v_0) such that $Q_0 \subseteq I_2$. Then show that $I_2 \cap U = \emptyset$. Using your result from part (b), conclude that the system is safe under this initial-state condition.

ECE 484: Principles of Safe Autonomy	Fall 2026
Written by: ECE484 team	*MP0*
Website 	*Due Date: September 11, 2026*

2 MP0: Verification of air-traffic control via reachability

2.1 Overview

We don't have flying cars because traffic management and collision avoidance with untrained pilots would be a nightmare! Autonomous collision avoidance could solve this. See reports from NASA and AIAA on the challenges related to this topic [4]. In this MP, you will learn how to use reachability analysis to verify a simplified air collision avoidance protocol.

You are given three different scenarios inspired by ACAS-Xu (**A**irborne **C**ollision **A**voidance **S**ystem for Unmanned Aircraft)—a collision avoidance protocol for UAVs [3]. In each scenario, 2 airplanes are flying on the (x, y) plane—an *ownship* **O** (modelled with Dubin's plane model Link) and an *intruder* aircraft **I** that simply flies straight. **O** will follow *advisories* from ACAS-Xu to detect and avoid collisions with **I**. There are five possible advisories that the protocol can generate:

- **COC**: Fly straight, no turn
- **SL** : Turn a strong left
- **SR** : Turn a strong right
- **WL** : Turn a weak left
- **WR** : Turn a weak right

Each aircraft model has 4 different variables: **x,y**: are the (x,y) coordinates in meters (m), θ is the heading (measured in radians w.r.t. the x-axis), and **v** is the speed in m/s.

This controller will give an advisory to the ownship aircraft based on the following variables:

- ρ (m): The distance between the ownship and intruder aircraft
- Θ (radian): The angle made by the heading angle of the ownship and the line made connecting the position of ownship and intruder aircraft.

Given a particular initial condition in **R**, the complete system has a unique *execution*. An execution is *safe* if **O** is always greater than 25m distance from **I** within T_s. To put everything in perspective, you can take a look at Figure 1.

Your task is to prove the scenario is *safe* or *unsafe* from any initial condition in a given range **R**, in a given time T_s = 80, and with a given controller. In Python, the initial set is defined as [lowerBound, upperBound] where each bound is a state in the format of [x, y, θ, v].

Here are the following initial conditions that are included in the MP:

- $\mathbf{R_1}$: $x_{own} \in [-4100, -3100], y_{own} \in [-5000, -4950], \theta_{own} \in [0, 0], v_{own} \in [120, 120], x_{int} \in [-1450, -1250], y_{int} \in [1500, 2000], \theta_{int} \in [-\pi/2, -\pi/2], v_{int} \in [100, 100]$.
- $\mathbf{R_2}$: $x_{own} \in [-2200, -1200], y_{own} \in [-2000, -950], \theta_{own} \in [-\pi/2, -\pi/2], v_{own} \in [120, 120], x_{int} \in [-100, 100], y_{int} \in [1500, 2000], \theta_{int} \in [-3\pi/4 + 1.4, -\pi/12 + 1.4], v_{int} \in [100, 100]$.

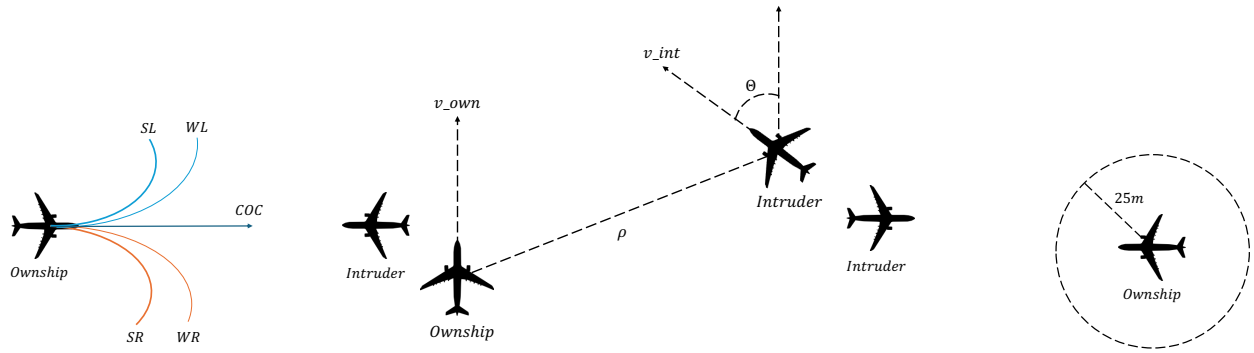

(a) (b) (c)

Figure 1: 1a) Illustration of possible advisories (Strong Left (SL), Weak Left (WL) , Clear of Conflict (COC), Strong Right (SR), Weak Right (WR)), 1b) Ownship's sensing based on intruder position, relative distance ρ, and relative angle Θ between 2 planes, 1c) Unsafe region when less than 25m separation between ownship and intruder.

- $\mathbf{R_3}$: $x_{own} \in [-7500, -7000], y_{own} \in [-6000, -5950], \theta_{own} \in [0, 0], v_{own} \in [120, 120], x_{int} \in [-100, 100], y_{int} \in [1500, 2000], \theta_{int} \in [-\pi/2, -\pi/2], v_{int} \in [100, 100]$.

2.2 Getting Started

Before you get started, learn about git starting here. Git is a version control system used in industrial development and we will use git throughout ECE484.

As we discussed in lectures, reachability analysis and $\overline{Post}$ computations can be hard and require careful engineering of over-approximations and set representations. There is a lot of ongoing research on this topic [1]. For this assignment, you will use a reachability analysis tool called Verse [2] that is actively being developed at Illinois. Verse is installed on all of the ECEB 5072 Ubuntu computers. Because of the large number of students, we recommend doing the MP on your own machine.

1. Visit this GitHub repo and follow the readme instructions for installing Verse. It is recommended to create a virtual environment to make package installation easier.
2. Get MP0 code with:

```
git clone git@github.com:safeautonomy-illinois-students/fa26-mp0-release.git
```

The relevant files are the following; you will only have to change the last one:

dubins_controller.py Contains the decision logic the ownship uses to avoid collision with the intruder.

dubins_agent.py Contains motion models for both aircraft.

dubins_sensor.py This file contains the implementation of processing the information of the initial states to return information that will be used to give advisory for the aircraft.

dubins_scenario.py This file contains the mechanism to create different aircraft scenarios. The three different initial ranges $\mathbf{R_1}$, $\mathbf{R_2}$ and $\mathbf{R_3}$ are provided in this file. **You will run and modify this file for the task.**

2.3 Executions: Verse Simulation

Verse provides simulation and reachability functions. The simulation function generates a random initial state within the specified ranges and computes an execution of the system (with 2 planes) from this initial state. If unsafety is detected from any point in the initial set, the scenario is sure to be unsafe. Therefore, it may be wise to run simulations in regions of the initial set likely to be unsafe.

This function runs in a few seconds by itself, but running it thousands of times will take much longer.

`trace = scenario.simulate(ownship_aircraft_init, intruder_aircraft_init, time_horizon, time_step, ax)`: Runs a single simulation from a randomly sampled initial point chosen from **R** and returns the trajectory. When the simulation is unsafe, you will get a red message `assert hit` in the terminal.

- `ownship_aircraft_init`: List, the initial state of the ownship (R1-R3).
- `intruder_aircraft_init`: List, the initial state of the intruder plane (R1-R3).
- `time_horizon`: Float, the total time for the scenario to evolve.
- `time_step`: Float, the time sampling period.
- `ax`: Plotter, the pyvista plotter for visualization
- `trace`: AnalysisTree object containing the simulated trajectory. The AnalysisTree object contains a list of AnalysisTreeNode that contains the execution of the scenario for each mode of **O** and **I**.

`traces = scenario.simulate_multi(ownship_aircraft_init, intruder_aircraft_init, time_horizon, time_step, init_dict_list, ax)` : simulate from all initial points in the init_dict_list instead of randomly choosing a point from the initial set. Unlike `simulate()`, The output will be a list of traces instead of just one trace.

- `ownship_aircraft_init`: List, the initial state of the ownship (R1-R3).
- `intruder_aircraft_init`: List, the initial state of the intruder plane (R1-R3).
- `time_horizon`: Float, the total time for the scenario to evolve.
- `time_step`: Float, the time sampling period.
- `init_dict_list`: List[Dictionary], list of initial points to simulate from. For example, an init_dict_list of [{'plane1': [-5,25,6,8], 'plane2':[170,-53,0,3]}] will simulate starting from plane1's state [-5,25,6,8] and plane2's state [170,-53,0,3].
- `ax`: Plotter, the pyvista plotter for visualization
- `traces`: List, list of traces. A trace is explained above in `scenario.simulate`

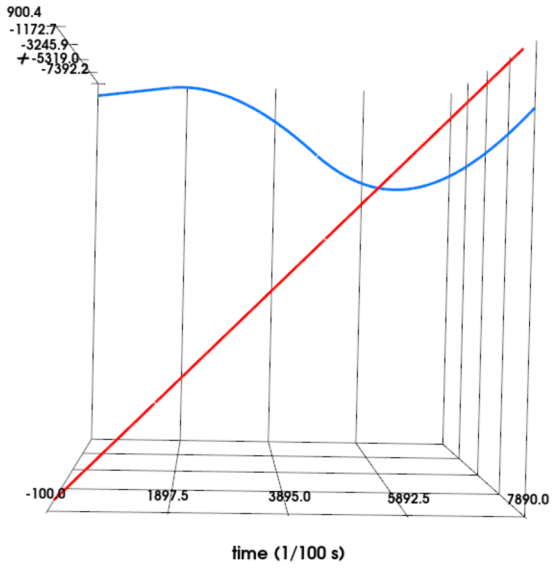

Figure 2: Simulation plot for a single simulation

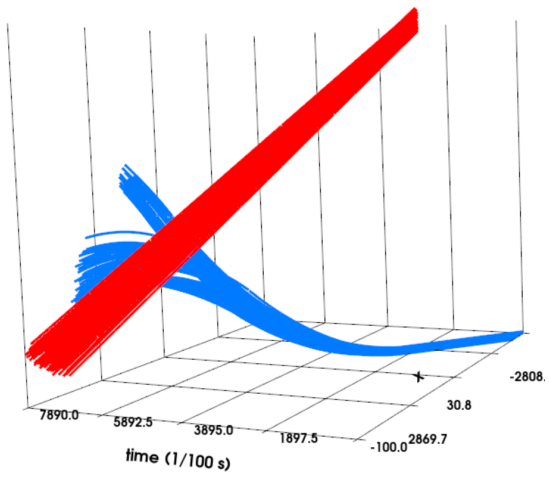

Figure 3: Simulation plot with multiple simulations run

2.4 Reachability

Recall, that reachability over-approximates all possible executions in the initial set. If the analysis determines that the system is safe, then safety is guaranteed. If unsafety is detected, that does not necessarily mean that the scenario is unsafe. The analysis usually generates an approximation of the reachable set larger than the actual reachable set. The reachability analysis could have over-approximated the reachable set such that it intersects with the unsafe set when the actual (smaller) reachable set would not. Internally, Verse uses repeated $\overline{Post}$ computations that you can learn more about by reading the paper, but for this MP you do not need to understand those internal workings.

`traces = scenario.verify(ownship_aircraft_init, intruder_aircraft_init, time_horizon, time_step, ax)`: Compute reachable set and check the safety of the scenario.

- `ownship_aircraft_init`: List, the initial state of the ownship (R1-R3).
- `intruder_aircraft_init`: List, the initial state of the intruder plane (R1-R3).
- `time_horizon`: Float, the total time to perform reachability analysis.
- `time_step`: Float, the time period that the reachable set is sampled.
- `ax`: Plotter, the pyvista plotter for visualization
- `traces`: AnalysisTree, the computed reachtube for the scenario. Has similar structure as that produced by `scenario.simulate()`. This reachtube will contain the reachable set for both planes across the entire time horizon

Figure 4 shows an example of performing reachability analysis and generating plots using above functions. If a safety assertion is violated, then a label will appear on the plot and a message will appear in the console.

2.5 Your Verification Task

Reachability analysis from a smaller initial set usually results in a tighter (less conservative) approximation. The more conservative the approximation, the less precise it is. Therefore, for a larger initial set, you

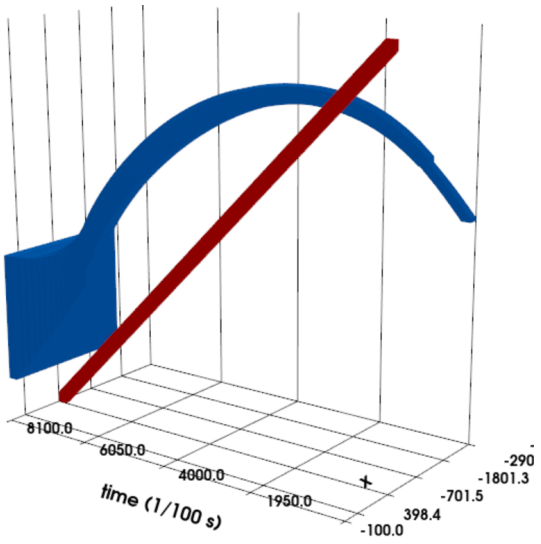

(a) Reachtube analysis that guarantees safety.

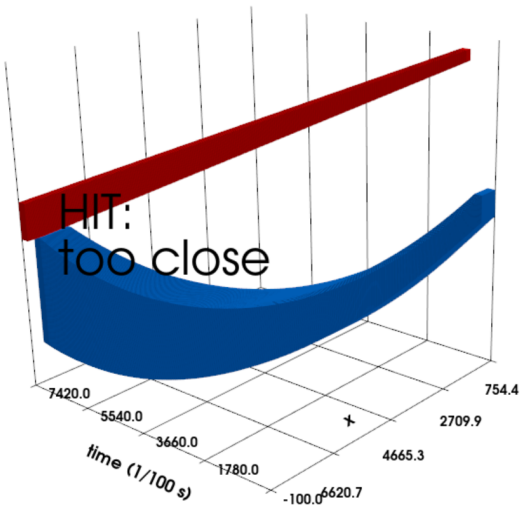

(b) Reachtube analysis that can't guarantee safety.

Figure 4: Reachability results with verify()

must determine a strategy to partition the initial set into smaller regions for a more precise reachability analysis.

Using the `verify()` function, you will write the function:

`traces = verify_refine(scenario, ownship_aircraft_init, intruder_aircraft_init, time_horizon, time_step, ax)`: Compute reachable set and check the safety of the scenario. `verify_refine()` must partition the initial ranges into smaller regions to get a more precise reachability analysis result.

- `scenario`: the scenario to verify.
- `ownship_aircraft_init`: List, the initial state of the ownship (R1-R3).
- `intruder_aircraft_init`: List, the initial state of the intruder plane (R1-R3).
- `time_horizon`: Float, the total time to perform reachability analysis. (Just pass this into `verify()`)
- `time_step`: Float, the time period that the reachable set is sampled. (Just pass this into `verify()`)
- `ax`: Plotter, the pyvista plotter for visualization (Just pass this into `verify()`)
- `traces`: List[`AnalysisTree`], the computed reachtube for the scenario. List of all traces returned by calls to `verify()`

One way to write this function is to use a DFS (depth first search) or BFS (breadth first search) style recursive traversal: continually partition the initial set into smaller and smaller regions until the result is safe. Each level of partition depth increases the number of calls to `verify()` exponentially.

Figure 5 shows partitioning on the dimension x. The final output is considered safe, since all leaf partitions are safe and together form the original initial set. In your scenarios, there are more dimensions than just x. The full state is defined in Section 2.1. You will need to decide which dimensions to partition and which order to partition. One method is to alternate splitting between dimensions. For example, you may split ownship's x dimension in half, then in the next layer of partitioning, split the intruder plane's y dimension into thirds.

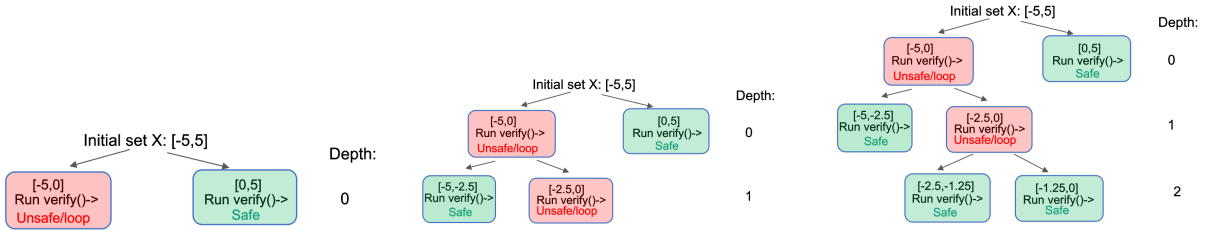

(a) Unsafety in some partitions (b) Partitioning on unsafe regions (c) All leaf partitions safe

Figure 5: Partitioning with tree method from initial set x: [-5,5]

You may also initially divide the initial set into some number of partitions in each dimension. (np.linspace function may be helpful here.) This may be faster than recursion if you try not run the `verify()` function more times than necessary. You may also combine the two approaches and partition the set in the beginning and then do recursive partitioning. We encourage you to try creative partitioning methods for this assignment.

Each call to `verify()` will generate a trace. All traces of partitions that run successfully without unsafety MUST be added to the traces list and returned to be visualized with the given plotting function. This plotting function simply visualizes all traces on one plot. Using this knowledge, you can answer Q9 in the report.

To conclusively prove safety, all of the partitions together MUST form the original initial set. That is, any point in the state space of the original initial set must be present in at least one of the partitions. You must run the `is_refine_complete()` function to validate this. To run this function, you must add your final partitions of the initial set to the partitions list.

We expect an optimal implementation of the function `verify_refine()` to run in 15 min or less for the scenarios provided. However, your `verify_refine()` does not need to be optimal, nor is it needed for all of the scenarios!
Implementation tips:

- Use the function `tree_safe()` to determine if a trace returned by `verify()` is safe.
- Reminder: In Python, lists are mutable objects, meaning that assigning one list to another variable will create a reference to the same object, not a copy. If you need to make a copy of an initial set (e.g., to preserve the original array), please use the `.copy()` method.
- It is also recommended to use a progress bar package such as alive-progress to keep track of what percent of the initial set is verified.
- If you see an "infinite loop detected" message, this essentially means that Verse's over-approximation was too big, and the reachability analysis cannot continue.

2.6 Reporting & Grading

Questions 6–8

For each R1–R3, answer the following:

6) [13 pts] Is the scenario safe or unsafe? Provide proof (images of the plot from multiple angles and console output) with either the `simulate/simulate_multi` or `verify/verify_refine` function that the scenario is either safe or unsafe.

7) [4 pts] Which function did you use to prove this and why?

8) [4 pts] If reachability was used, what was your strategy for partitioning (if partitioning is needed)? If simulation was used, what was your strategy for running simulations (number of simulations, regions to simulate in, etc.)?

Question 9 [7 pts]

Recall that running `verify()` will generate a plot and a reachability analysis result (safe or unsafe). If the different colored regions intersect on the plot, the analysis from `verify()` will always determine unsafety is present in the scenario. Is this necessarily true for a correct implementation of `verify_refine()`? That is, is there some way to partition the initial set such that there is visible intersection on the `verify_refine()` plot, but the analysis determines the scenario is safe? Explain why or why not.

Autograder [10 pts]

Your `verify_refine()` will be evaluated for correctness as well.

Demo Attendance [20 pts]

Attend your lab session on Sept 11th to demo your design logic. We will ask questions regarding invariance concepts related to the MP. Additionally, you will need to show your plots and the results for `verify_refine()`. This can be in the form of a video or the active window.

2.7 Submission

Upload the following files to **Gradescope**:

- Your homework report with questions 1-5 as `netid_hw0.pdf`
- Your code file `dubins_scenario.py`, renamed to `netid_mp0.py`. **Be sure to do this, otherwise you will not be graded!**
- Your MP0 report with questions 6-9 as `netid_mp0.pdf`

References

[1] Applied verification for continuous and hybrid systems, 2019. The workshop on applied verification for continuous and hybrid systems (ARCH).

[2] Yangge Li, Haoqing Zhu, Katherine Braught, Keyi Shen, and Sayan Mitra. Verse: A python library for reasoning about multi-agent hybrid system scenarios. In *International Conference on Computer Aided Verification*, pages 351–364. Springer, 2023.

[3] Michael P. Owen, Adam Panken, Robert Moss, Luis Alvarez, and Charles Leeper. Acas xu: Integrated collision avoidance and detect and avoid capability for uas. In *2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC)*, pages 1–10, 2019.

[4] Casey L. Smith and R Conrad Rorie. Helicopter pilot evaluations of the airborne collision avoidance system xr in a high-fidelity motion simulation - nasa technical reports server (ntrs). July 2024. [Online; accessed 2025-08-19].