

# Formal Verification of a Machine Learning Tool for Runway Configuration Assistance

Pouria Razzaghi<sup>1,\*</sup>, Milad Memarzadeh<sup>2</sup> and Krishna Kalyanam<sup>2</sup>

<sup>1</sup>*Metis Technology Solutions, Inc, NASA Ames Research Center, Moffett Field, CA 94035, USA*

<sup>2</sup>*NASA Ames Research Center, Moffett Field, CA 94035, USA*

Correspondence\*:

Pouria Razzaghi

pouria.razzaghi@nasa.gov

## 2 ABSTRACT

3 This study explores the use of formal verification techniques to evaluate the efficacy of  
4 suggestions made by the Runway Configuration Assistance (RCA) tool, a machine learning-based  
5 decision support system that our group developed independently Memarzadeh et al. (2023). By  
6 using model-checking approaches, in particular Computation Tree Logic (CTL), this study verifies  
7 the compliance of the RCA tool with predefined safety regulations under different conditions  
8 of surface winds. By simulating a range of scenarios at three major US airports, Charlotte  
9 Douglas International Airport (CLT), Denver International Airport (DEN), and Dallas-Fort Worth  
10 International Airport (DFW), we thoroughly test the predictions of the tool to ensure that they  
11 meet strict safety margins with respect to crosswind and tailwind. The application of formal  
12 verification methods provides a strict analysis of the RCA tool, enhancing its validity and utility  
13 for possible implementation in an operational environment. Initially, a Monte Carlo simulation is  
14 carried out to analyze all possible wind conditions both velocity-wise and direction-wise. This part  
15 is intended to rigorously test the model against extreme, worst-case conditions to evaluate its  
16 performance. Second, we improve our methodology by performing simulations driven by realistic  
17 scenarios informed by actual historical data. This approach allows for a more accurate reflection  
18 of typical wind conditions (seen in the test airport) and provides a robust assessment of the  
19 model's effectiveness in maintaining safety standards under realistic environmental conditions.  
20 The model-checking reveals that overall 70% and 94% of the predictions satisfy the safety criteria  
21 in worst-case and realistic wind scenarios, respectively.

22 **Keywords:** Formal Verification, model-checking, Air Traffic Management, Machine Learning, Safety Criteria, Runway Configuration  
23 Management

## 1 INTRODUCTION

24 Runway Configuration Management (RCM) is the task of selecting appropriate runways for arriving  
25 and departing aircraft at an airport. It is a complex task that involves multiple stakeholders and must  
26 take into consideration air traffic services (ATS), weather, and other contributing factors at an airport.  
27 Typically, it involves either switching the runway direction for takeoff and landing or switching between  
28 different (combinations of) runways available on the airport surface. Most airports, depending on the  
29 surface geometry, capacity, and local weather patterns, have multiple configurations that can be used. Many  
30 factors, including incoming/outgoing traffic load, wind direction and speed, convective weather, cloud  
31 ceiling, and other environmental or noise-related factors, can affect the choice of runway configuration.  
32 In current practice, air traffic controllers (ATC) select the runway configuration based on the information  
33 (weather, traffic, and other forecast) available to them at the time of decision making. As a general rule, it  
34 is preferred that the aircraft take off and land into the wind (for maximal lift and braking, respectively). So,  
35 the surface wind (speed and direction) is the dominant feature that determines the runway configuration.  
36 However, the human decision making process is subjective based on training and expertise, which can  
37 be impacted by bias and can sometimes lead to poor outcomes. In particular, for airports with multiple

38 available configurations, dynamic traffic and weather patterns can make it difficult for humans to compute  
39 optimal solutions in near real time.

40 In recent years, automated approaches based on Artificial Intelligence (AI) and Machine Learning (ML)  
41 have been proposed to solve the aviation problem (Razzaghi et al., 2024). A popular approach used in this  
42 context is online Reinforcement Learning (RL) and variants such as discrete choice modeling (Avery and  
43 Balakrishnan, 2016), dynamic programming (Li et al., 2009), and queueing theory (Jacquillat et al., 2017;  
44 Badrinath et al., 2019) have been proposed. The goal is to model the dynamics of surface operations at an  
45 airport and use a simulator to learn a near-optimal policy for the runway configuration selection problem.  
46 However, building an accurate simulator (and/or learning accurate models for the surface dynamics) is hard.  
47 On the other hand, model-free RL approaches such as Monte Carlo Tree Search (MCTS) (Browne et al.,  
48 2012) can be used to learn a near-optimal policy by interacting directly in the operational environment.  
49 However, this interaction is not possible in safety-critical environments, since the model tends to make  
50 mistakes at the early stages of learning, which we cannot afford to do. So, in the absence of a good model  
51 (simulator) and difficulties in learning in the real-world environment, we turn our attention to offline  
52 model-free RL.

53 To go into the evaluation of our approach, this paragraph introduces the verification method applied  
54 to the RL model. Model-checking is a formal verification method used to assess whether a system's  
55 finite-state model satisfies specified requirements. The underlying verification mechanism is based on the  
56 popular model-checking technique called temporal logic (Bérard et al., 2013). Specifications, typically  
57 articulated through temporal logic, that is, Linear Temporal Logic (Vardi, 2005) or Signal Temporal Logic  
58 (Baheri et al., 2022), or Computation Tree Logic (CTL) (Li et al., 2015), define properties over time,  
59 including safety properties. The algorithm then systematically explores the state space of the model to  
60 check if the specification is valid. Upon detecting a specification breach, a counterexample showcasing the  
61 fault is provided. Although model-checking streamlines the verification process through automation and  
62 ensures thorough system examination, it faces the challenge of managing the rapidly expanding state space  
63 associated with complex systems. Also, a recent trend in formal verification of AI systems is including  
64 hybrid system verification and runtime monitoring (Sobeh (2024); Paul et al. (2023)), which makes this  
65 more challenging. Despite these challenges, our study employs the CTL process to validate the RCM  
66 model, which does not have a large and complex state space and can be deployed in an offline manner.

67 In previous work, we adopted an offline model-free RL methodology, known as Conservative Q-Learning  
68 (CQL) (Kumar et al., 2020), to successfully tackle the RCM problem (Memarzadeh and Kalyanam, 2025;  
69 Memarzadeh et al., 2023; Nethi et al., 2024). The tool we developed to provide runway recommendations is  
70 referred to as the Runway Configuration Assistance (RCA) tool - for details, see (Memarzadeh et al., 2023;  
71 Nethi et al., 2024). The offline nature of the tool removes the need for interactions with the operational  
72 environment. Instead, the RCA tool relies only on historical data which includes all relevant system state,  
73 input (human decisions), and output (traffic flow) to identify a near-optimal policy.

74 Given the safety critical nature of the problem, it is crucial that a thorough verification and validation of  
75 the RCA tool is performed. As a first step in this direction, we investigate a verification process to ensure  
76 that the outcomes of the RCA tool adhere to predefined safety criteria. It is important to note that the safety  
77 criteria discussed here relate only to the wind conditions under which runway operations are regarded  
78 safe. The verification method used ensures that the recommendations made by the RCA tool comply with  
79 these conditions. Ultimately, airport controllers guarantee the overall safety of the system by applying their  
80 expertise to ensure that all operational guidelines and safety standards are rigorously upheld.

81 In the first part of the verification process, we randomly sample wind conditions (both speed and direction)  
82 and pipe them into the RL model (RCA tool) and assess whether the output, i.e., recommended runway  
83 configuration, meets the safety standards. Specifically, we are performing an experiment as follows: for the  
84 selected wind speed and direction and the recommended runway configuration (RCA tool output), is it  
85 safe for aircraft to takeoff and land as stipulated by the tailwind and crosswind safety criteria? We then  
86 record the response to this question to compute the validation performance metrics. In the second part,  
87 the verification process will be repeated by more realistic scenarios coming from real data. The process  
88 will identify the boundaries within the input ranges where the system remains safe. The approach will be  
89 performed using data from three major US airports, Charlotte Douglas International Airport (CLT), Denver  
90 International Airport (DEN), and Dallas-Fort Worth International Airport (DFW).

## 2 MODEL-CHECKING

91 Model-checking, a formal and automated verification method, is widely used in various fields,  
 92 including computer software, hardware systems, communication protocols, control systems, and security  
 93 authentication protocols (Baier and Katoen, 2008). When verifying complex concurrent systems, it is  
 94 typical to come across uncertain and inconsistent information. For example, intelligent autonomous  
 95 transport systems often generate complex computing tasks for autonomous vehicles (Gao et al., 2022).  
 96 Model-checking is a technique to verify whether a given system model satisfies a specified property,  
 97 typically expressed in temporal logic such as CTL logic, which is branching-time logic. This means that it  
 98 allows one to express properties over trees of possible execution paths (states) rather than linear paths.

### 99 2.1 System Model

100 First, let us define the system model as a tuple  $M = (S, R, L)$ , where:

- 101 •  $S$  is a set of states.
- 102 •  $R \subseteq S \times S$  is a state transition relation.
- 103 •  $L : S \rightarrow 2^{|AP|}$  is a labeling function that maps each state to a set of Atomic Propositions ( $AP$ ) that  
 104 are true in that state.

### 105 2.2 CTL Logic Syntax

106 The CTL logic formulas are built from  $AP$ s, boolean operators, and path quantifiers along with temporal  
 107 operators. Some common CTL logic operators include:

- 108 • **EX** $\phi$  : There exists a next state such that  $\phi$  as the required condition, holds.
- 109 • **AX** $\phi$  : For all next states,  $\phi$  holds.
- 110 • **EF** $\phi$  : There exists a path where  $\phi$  eventually holds.
- 111 • **AF** $\phi$  : On all paths,  $\phi$  eventually holds.
- 112 • **EG** $\phi$  : There exists a path where  $\phi$  always holds.
- 113 • **AG** $\phi$  : On all paths,  $\phi$  always holds.

### 114 2.3 Basic Model-Checking Algorithm

115 This section includes previously developed material, reintroduced here to enhance clarity and ensure  
 116 completeness of the discussion. A high-level approach to the CTL logic model-checking algorithm is  
 117 provided below:

- 118 1. **Labeling states with  $AP$ s** : Each state in  $S$  is labeled with the  $AP$ s that are true in that state based on  
 119 the function  $L$ .
- 120 2. **Recursive Check** : For each subformula  $\phi$  of the CTL logic formula:
  - 121 • If  $\phi$  is an  $AP$ , return the set of states for which  $\phi$  is true.
  - 122 • For Boolean operations (AND, OR, NOT), compute the result based on the results of the operands.
  - 123 • For temporal operations involving path quantifiers (E or A combined with X, F, G), compute the  
 124 set of states that satisfy these formulas by:
    - 125  **EX** $\phi$  : Return states for which there exists a transition to a state satisfying  $\phi$ .
    - 126  **AX** $\phi$  : Return states for which all transitions lead to states satisfying  $\phi$ .
    - 127  **EF, EG, AF, AG** : Use fixpoint computations. For instance: **EF** $\phi$  starts with the set of states  
 128 satisfying  $\phi$  and iteratively adds states that can reach this set until no more states can be added.
- 129 3. **Evaluate the main formula** : The root of the CTL logic formula (or the main property to check) is  
 130 evaluated last, utilizing the results from the evaluations of its subformulas.
- 131 4. **Interpret results** : The algorithm returns whether the initial state (or any specified state) of the model  
 132 satisfies the CTL logic formula.

## 133 2.4 Model-Checking Algorithm for an ML System

134 Using model-checking to verify properties of an ML model involves a fairly abstract and theoretical  
 135 approach, as CTLLogic is traditionally used to check logical properties in systems described by state-  
 136 transition models. However, we can conceptualize how this might be approached by considering the ML  
 137 model as a dynamic system where each state represents a set of parameters or decisions, and transitions  
 138 reflect changes or iterations in the learning process, or the final outcomes of the model.

### 139 2.4.1 Steps to Conceptualize CTLLogic model-checking for ML Models

#### 140 1. Define the System Model:

141 • **States:** In the context of an ML model, states could represent specific configurations or snapshots  
 142 of the model during training (e.g., after each epoch), testing, or/and validating.

143 • **Transitions:** Transitions between states could represent the update steps of model parameters.

144 2. **Specify Properties in CTLLogic:** You would need to define the properties you want to verify in  
 145 CTLLogic. For an ML model, these might involve convergence, stability over epochs, or fairness  
 146 metrics, depending on the interpretability of the ML model's operations as logical transitions.

147 3. **Labeling States:** Each state must be labeled with APs that are true in that state. For an ML model,  
 148 these labels could be derived from the performance metrics, error rates, or other measurable output of  
 149 the model at each point.

150 4. **Build the Transition System:** Construct a transition system where each node corresponds to a state of  
 151 the ML model at a given point, and directed edges represent transitions due to steps.

152 5. **Run CTLLogic model-checking:** Using a model-checking tool or library that supports CTLLogic (e.g.,  
 153 NuSMV, SPIN), run the model-checking process on the constructed transition system with the CTLLogic  
 154 properties defined in Step 2.

155 6. **Interpret Results:** Analyze the results from the model checker to understand whether the ML model  
 156 satisfies the specified properties.

157 We should note that the formal model-checking is conducted after training to verify the safety constraint  
 158 satisfaction of the RCA tool for both synthetic and real wind condition scenarios. Training in offline  
 159 reinforcement learning is separate from verification. The RCA tool is first trained using historical data via  
 160 CQL. Then, the verification module uses CTLLogic-based model-checking to check if the output of the  
 161 trained model satisfies operational safety constraints for a range of wind conditions.

## 162 2.5 Model-Checking Algorithm on RCA Tool

163 In this section, we will present the model-checking algorithm for the RCA tool and check the basic  
 164 CTLLogic formulas, i.e.  $(\mathbf{EF}, \mathbf{EG})\phi$ , where  $\phi$  is the safety criteria for selecting the runway configuration. It  
 165 is important to note that we employ these formulas because the computation of the criteria occurs at a fixed  
 166 point within the system, where states are fixed and the criteria are checked at this point. This means that  
 167 in the preceding nodes in the search tree, there are fixed-in-time states. The formulas can be described as  
 168 follows.

169 •  $\mathbf{EF}\phi$  : There is at least one path in the model such that the safety criteria are always satisfied.

170 •  $\mathbf{EG}\phi$  : For each path, the safety criteria are always satisfied.

171 We should note that here model-checking was performed using the NuSMV tool due to its support for  
 172 CTL logic. Simulations and policy inference are executed using Python 3.9 and NumPy/Pandas for data  
 173 manipulation. Historical wind transitions are modeled using custom-built probabilistic transition matrices  
 174 derived from datasets.

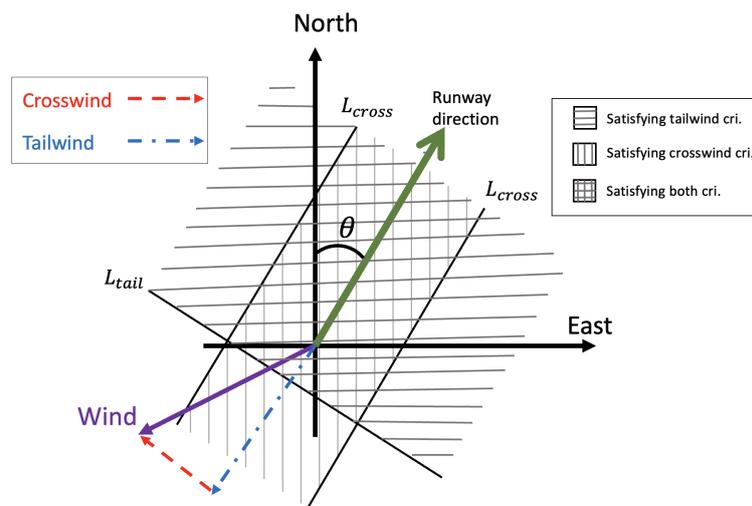
### 175 2.5.1 Safety Criteria

176 The criteria that ensure safe runway selection are based on crosswind and tailwind limits. These limits  
 177 are determined according to the active runway and state that wind components must not exceed specified  
 178 thresholds. These thresholds are typically derived from aircraft performance manuals or operational  
 179 regulations, which stipulate the maximum allowable values for crosswind and tailwind. For instance,  
 180 for a Boeing 737-300 on a dry runway, the maximum permitted crosswind and tailwind speeds are 29

181 and 10 knots, respectively. If the surface wind speed exceeds these limits during takeoff or landing, the  
 182 structural safety of the aircraft could be compromised. Both the Federal Aviation Administration (FAA)  
 183 and the International Civil Aviation Organization (ICAO) recommend these maximum limits of the wind  
 184 component to ensure safe flight operations. According to FAA Order 8400.9 (FAA, 1981-11-09), the safety  
 185 standards are as follows:

- 186 • Maximum crosswind component (including gust)
  - 187 1. Dry Runway: 25 kts
  - 188 2. Wet Runway: 15 kts
  - 189 3. Contaminated Runway: 15 kts
- 190 • Maximum tailwind component (including gust)
  - 191 1. Dry Runway: 10 kts
  - 192 2. Wet Runway: 10 kts
  - 193 3. Contaminated Runway (< 8000 ft): <3 kts (reported as calm)
  - 194 4. Contaminated Runway (> 8000 ft): 5 kts

195 A contaminated runway is one that has standing water, ice, snow, slush, or any material that will reduce  
 196 braking ability. The FAA defines certain values of the wind component for contaminated surfaces due  
 197 to increased risk during takeoff and landing. Before establishing the crosswind and tailwind criteria, it  
 198 is essential to define the coordinate system used in aviation. This system, as illustrated in Fig. 1, is a  
 199 right-hand Cartesian system where the x-axis points north and the y-axis points east, and the z-axis points  
 down into the earth as the NED (north-east-down) convention.



**Figure 1.** Tailwind and crosswind components and their corresponding criterias in the coordinate system.  $\alpha$  is the angle between the wind vector and the north axis.

200

201 In this framework, variables  $\theta$  and  $\alpha$  represent the runway direction, and the wind direction respect to  
 202 the north axis, respectively. The direction of the runway is determined by the direction of the aircraft  
 203 during landing or takeoff. In contrast, the direction of the wind indicates where the wind originates. Both  
 204 angles are measured from the true north, with their values ranging from 0 to 360 degrees. In Fig. 1, the  
 205 runway is represented by a solid green line marked with an arrow that indicates the direction of the runway.  
 206 Additionally, a solid red line segment is used to represent the wind vector, with its length and orientation  
 207 indicating the wind's strength and direction, respectively.

208 Let  $V$  represent the wind magnitude, with  $L_{\text{cross}}$  and  $L_{\text{tail}}$  being the limits for the crosswind and tailwind,  
 209 respectively. The criteria for crosswind and tailwind are then defined on the basis of these parameters as

210 follows:

$$\begin{aligned} \begin{bmatrix} V_{\text{tail}} \\ V_{\text{cross}} \end{bmatrix} &= \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} V \cos \alpha \\ V \sin \alpha \end{bmatrix} \\ &= \begin{bmatrix} V \cos \theta \cos \alpha - V \sin \theta \sin \alpha \\ V \sin \theta \cos \alpha + V \cos \theta \sin \alpha \end{bmatrix} = \begin{bmatrix} V \cos(\theta + \alpha) \\ V \sin(\theta + \alpha) \end{bmatrix} \end{aligned} \quad (1)$$

$$211 \quad V_{\text{tail}} < L_{\text{tail}} \quad (2)$$

$$V_{\text{cross}} < L_{\text{cross}} \quad (3)$$

212 where,  $V_{\text{tail}}$  and  $V_{\text{cross}}$  are the tailwind and crosswind components of the wind in the runway direction.  
 213 In Fig. 1, the area to the right of the line perpendicular to the runway encompasses all wind vectors that  
 214 comply with Eq. (2), related to tailwind criteria. Likewise, the space between the two parallel lines with  
 215 runway direction visually represents all wind velocities that meet the criteria outlined in Eq. (3), which  
 216 pertains to crosswind. Therefore, the intersecting region of these two areas indicates all wind vectors  
 217 that fulfill both sets of criteria, indicating that the runway in question is operable under wind conditions  
 218 depicted by the overlapping area.

## 219 2.5.2 Definitions of Model-Checking Steps

220 Building on the concept introduced in Section 2.4.1, we now define wind states (wind speed and direction  
 221  $(V, \alpha)$ ) for an airport with multiple runways. For an airport with  $N_r$  runways, we introduce an indicator  
 222 function for each runway. This function determines whether a runway is operational under specific wind  
 223 conditions.

$$I_i(V, \alpha, \theta_i) = \begin{cases} 1 & \text{runway } r_i \text{ is active for wind condition set } (V, \alpha) \\ 0 & \text{runway } r_i \text{ is inactive} \end{cases} \quad (4)$$

224 All wind conditions that produce identical values for all indicator functions collectively form a wind state,  
 225 denoted  $\omega_j$ . Essentially, a wind state corresponds to a specific combination of active runways  $R_j \subseteq R$ ,  
 226 where  $R$  represents the complete set of runways at the airport and  $N_W$  is the length of the wind state vector.

$$R_j \quad : \{r_j \mid I_i(V, \alpha, \theta_i) = 1, \forall i = 1, \dots, N_r\} \quad (5)$$

$$\omega_j \quad : \{(V, \alpha) \mid I_i(V, \alpha, \theta_i) = 1, \forall r_i \in R_j \text{ and } I_i = 0, \forall r_i \notin R_j\} \quad (6)$$

227 Transitions, in the context of this research, are changes over time in environmental conditions, namely  
 228 wind speed and direction, rather than alterations to the static set of runway configurations. The system  
 229 transitions through various wind states and, at any given point, the RCA tool determines an appropriate  
 230 configuration. Verification checks whether these configurations remain compliant with all temporal changes  
 231 in wind conditions.

232 By these definitions, we establish a one-to-one mapping transition between runway configuration and  
 233 wind state. We select two CTLogic properties to verify the safe selection decision made by the RCA tool.  
 234 These properties are  $(\mathbf{EF}, \mathbf{EG})\phi$ , where  $\phi$  is the safety criteria to select the runway configuration defined by  
 235 Eqs. (2) and (3). We calculate the properties in each state transition explained above and check the safety at  
 236 each point.

237 To place the use of CTLogic in the context of this research, it is important to redefine the terms "path"  
 238 and "transition" as they apply to the dynamics of changing runway configurations. Although runway  
 239 configurations are static by definition (in that each configuration represents a stable operating state), the  
 240 transitions of interest in this research are defined as changes in environmental conditions, namely changes  
 241 in wind speed and direction, over a particular temporal frame. These wind conditions influence the RCA  
 242 tool's decisions at different time steps. Thus, a path in CTLogic corresponds to a trajectory of wind state

243 changes throughout a day or simulation time interval, with the RCA tool making a configuration choice at  
244 every time step.

245 We utilize two temporal logic formulas to check the safety of model outputs:

- 246 • **EF** $\phi$ : This expression calculates if there exists at least one path (i.e., series of wind conditions) in  
247 which the RCA tool suggests a configuration that satisfies the safety constraints  $\phi$  (i.e., the crosswind  
248 and tailwind limits).
- 249 • **EG** $\phi$ : This finer property demands that throughout the entire trajectory, for each discrete time step, the  
250 configuration of the RCA tool satisfies  $\phi$ . With this condition, the safety retention can be assessed in  
251 the face of changing dynamic conditions.

252 Accordingly, while a static logic verification might ensure that a particular configuration is safe under  
253 given wind conditions, the CTLogic-based model-checking framework utilized here allows us to assess  
254 safety temporally—ensuring that the RCA tool always produces safe choices under varying operational  
255 conditions. This is a requirement in safety-critical systems, where ongoing compliance over time, not just  
256 at a moment, is a basic requirement. Let us consider a smaller example: an airport having two runways, A  
257 and B. Suppose that under wind direction  $\alpha = 30^\circ$  and velocity  $V = 20$  knots, Runway A is safe, Runway  
258 B is not. For **EF** $\phi$ , the model checker verifies if there is at least one such safe wind condition under which  
259 RCA selects Runway A. For **EG** $\phi$ , it verifies that for all hourly wind conditions during the day, RCA never  
260 recommends Runway B. Experimental results are presented in the next section. For a formal summary of  
261 the RCA-to-CTLogic model-checking procedure, please see Appendix A.

### 3 RESULTS

262 The results section is divided into two main parts that examine the verification process in three different  
263 airports CLT, DEN, and DFW. The first part of our analysis involves a Monte Carlo simulation, which  
264 comprehensively evaluates the tool under all possible wind conditions. In these simulations, we rigorously  
265 assess the safety criteria to determine how the model performs under worst-case conditions. In the second  
266 part, we refine our approach by conducting simulations based on more realistic scenarios that are informed  
267 by actual historical data. This two-pronged strategy allows us not only to test against worst-case wind  
268 profiles (which may or may not occur in the real world for the specified airport), but also validate model  
269 performance against real-world conditions. All runway configurations and the usage frequency (how often  
270 each configuration is used historically) for all three airports are reported in Table 1. Here,  $N/N$  means that  
271 the north configuration is used for both arrivals and departures, and an example runway identifier: 36R/C/L  
272 means runway 36 (which is oriented 360 degrees from true North) and R/C/L stands for Right/Center/Left.  
273 Figure 2 shows the CLT airport diagram FAA (2024) with pictorial details of the runways with the  
274 corresponding labels.

#### 275 3.1 Monte Carlo Random Wind Condition Simulation

276 In the first part of our analysis, we randomly generate wind conditions to serve as input for the RCA tool.  
277 Wind conditions are sampled uniformly within the ranges of 0 knots to the maximum wind speed specified  
278 in Table 2 and 0 to 360 degrees for wind direction. For each wind condition, we utilize the RCA tool to  
279 determine the recommended runway configuration. We then evaluate whether the tool output satisfied  
280 the established safety criteria, specifically focusing on crosswind and tailwind limits. This procedure is  
281 repeated 100,000 times for each airport (CLT, DEN, and DFW) to ensure a comprehensive analysis.

282 To clarify this procedure, we look at three examples of runway configurations in different airports  
283 with random wind conditions. Figure 3 represents different wind conditions in three selected runways of  
284 different airports. The colored boxes indicate the safety criteria for the chosen runway configurations, each  
285 labeled in a format divided by dashes. The first segment of the label displays the takeoff configuration— $N$   
286 representing north for CLT and DEN, and  $SSE$  indicates south-southeast for DFW. The second segment  
287 details the landing configuration. If the wind vector is shown in green, it indicates that the runway selection  
288 is safe for both takeoff and landing. If the wind vector appears red, it signifies that the runway selection is  
289 unsafe and should not be used for aircraft operations for this wind condition, and when the wind vector is  
290 yellow, it denotes that the runway is safe for either takeoff or landing but not both. A safe (feasible) runway  
291 configuration selection indicates that the operating runway complies with the safety criteria for the current



Figure 2. Charlotte Douglas International Airport Surface Diagram

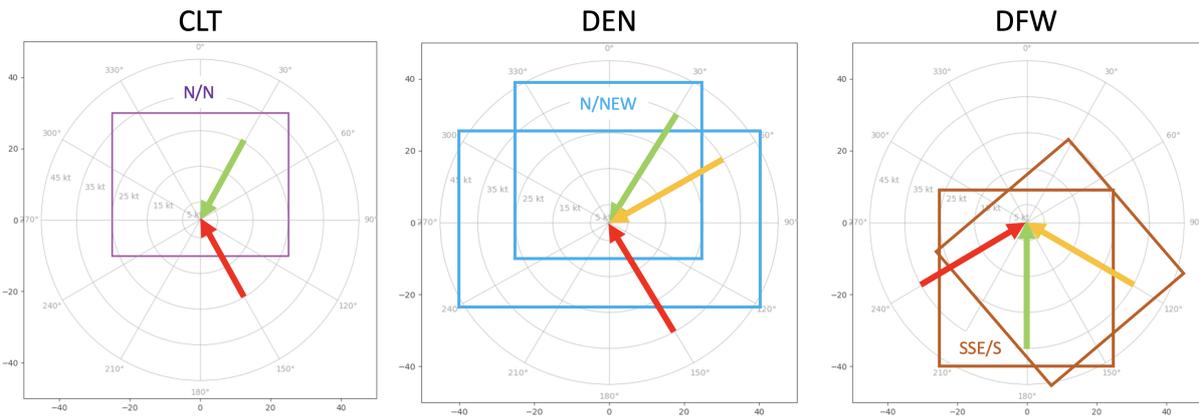


Figure 3. Safe and unsafe wind conditions on selected runway configurations for different airports. Each colored boxes represents the area satisfying the safety criteria. Two boxes in the middle and right subfigures show two different configurations for takeoff and landing. The first part in the box’s label represents the takeoff configuration and second one is for the landing. If the wind vector is shown in green, the selected runway configuration is safe for both takeoff and landing. A red vector indicates that the configuration is unsafe for both operations. A yellow vector signifies partial safety—safe for either takeoff or landing, but not both.

292 wind conditions. Also, the tool’s efficiency is assessed by how frequently each configuration is deemed  
 293 optimal under varying wind conditions. In the following, we show the validation results of the RCA tool  
 294 in the Monte Carlo simulations. Additionally, as detailed in Table 1, the most efficient preferred runway  
 295 configurations based on their frequency of use are *N/N* for CLT and *SSE/S* for DFW.

296 Table 2 shows the results of randomly generated wind condition simulation procedure. The results are  
 297 calculated based on  $EG\phi$  values on all configurations. The RCA tool demonstrates an overall effectiveness  
 298 of above 70% in selecting safe runway configurations under simulated wind conditions for all airports.

**Table 1.** The runway configurations for all three airports

Config. [Arr/Dep]	Arrival Runways	Departure Runways	Usage Frequency (%)
<b>CLT</b>			
N/N	36R/C/L	36R/C	60.8
S/S	18R/C/L	18C/L	39.2
<b>DEN</b>			
SE/SE	16R/L, 17R/L 7, 8	16R/L, 17R/L 7, 8	18.8
S/S	16R/L, 17R/L	16R/L, 17R/L	15
N/NEW	34R/L, 35 R/L	34R/L, 35 R/L 8, 25	14.5
S/SEW	16R/L, 17R/L	16R/L, 17R/L 8, 25	12.6
N/N	34R/L, 35R/L	34R/L, 35R/L	12.3
NE/NE	34R/L, 35R/L 7, 8	34R/L, 35R/L 7, 8	11.7
NW/NW	34R/L, 35R/L 25,26	34R/L, 35R/L 25,26	8.6
SW/SW	16R/L, 17R/L 25, 26	16R/L, 17R/L 25, 26	3.4
E/E	7,8	7,8	1.6
NS/EW	34R/L, 35R/L 16R/L, 17R/L	8,25	1.2
W/W	25, 26	25, 26	0.3
<b>DFW</b>			
SSE/S	13R, 17R/C/L, 18R	17R, 18R/L	61.5
NNW/NNW	31R, 35R/C/L, 36R/L	31L, 35C/L, 36R/L	21.3
S/S	17R/C/L, 18R	17R, 18R/L	7.6
N/NNW	35R/C/L, 36R/L	31R, 35R/C/L, 36R/L	5.1
NNW/N	31R, 35R/C/L, 36R/L	35R/C/L, 36R/L	3
N/N	35R/C/L, 36R/L	35R/C/L, 36R/L	1.1
SSE/NNW	13R, 17C/L, 18R	31R, 35R/C/L, 36R/L	0.2
NNW/S	31R, 35R/C/L, 36R/L	17R, 18L	0.1
NW/NW	31R/L	31R/L	0.1

299 In both CLT and DFW, the tool identifies efficient and preferred runway configurations. In scenarios  
300 where multiple runways are viable, the RCA tool successfully selects the optimal configuration 63% of  
301 the time for CLT and 74% for DFW. These figures underscore the tool's capability to effectively prioritize  
302 safety while accommodating varying airport layouts and conditions. The results highlight the RCA tool's  
303 robust performance in ensuring compliance with established safety criteria. This performance is especially  
304 noteworthy given the complexity of managing multiple variables, including varying wind conditions and  
305 airport-specific runway configurations. The tool's success in these areas suggests it could serve as a valuable  
306 asset in enhancing operational safety and efficiency in dynamic airport environments. The objective of this  
307 analysis was to assess how well the model's predictions aligned with safety standards in a wide variety

**Table 2.** Safety criteria of RCA tool prediction through randomly generated wind conditions

<b>Airport</b>	<b>CLT</b>	<b>DEN</b>	<b>DFW</b>
Max wind speed (kts)	30	40	40
Tailwind violation (%)	0	15	20.3
Crosswind violation (%)	4.7	16.7	6.2
Safe RCA prediction (%)	100	70.5	84.5
Safe arrival prediction (%)	-	69	86.6
Safe departure prediction (%)	-	99	78.7
Efficient runway prediction (%)	63	-	74

308 of wind scenarios. The results were analyzed to identify the safe versus unsafe predictions faced with the  
309 worst-case scenarios.

### 310 **3.2 Simulated Daily Wind Transitions**

311 In the second part of the analysis, we utilize actual weather data for years 2018 and 2019 to construct  
312 a transition matrix that describes changes in wind conditions throughout a day. This transition matrix  
313 provides a probabilistic framework for simulating realistic changes in wind conditions over time.

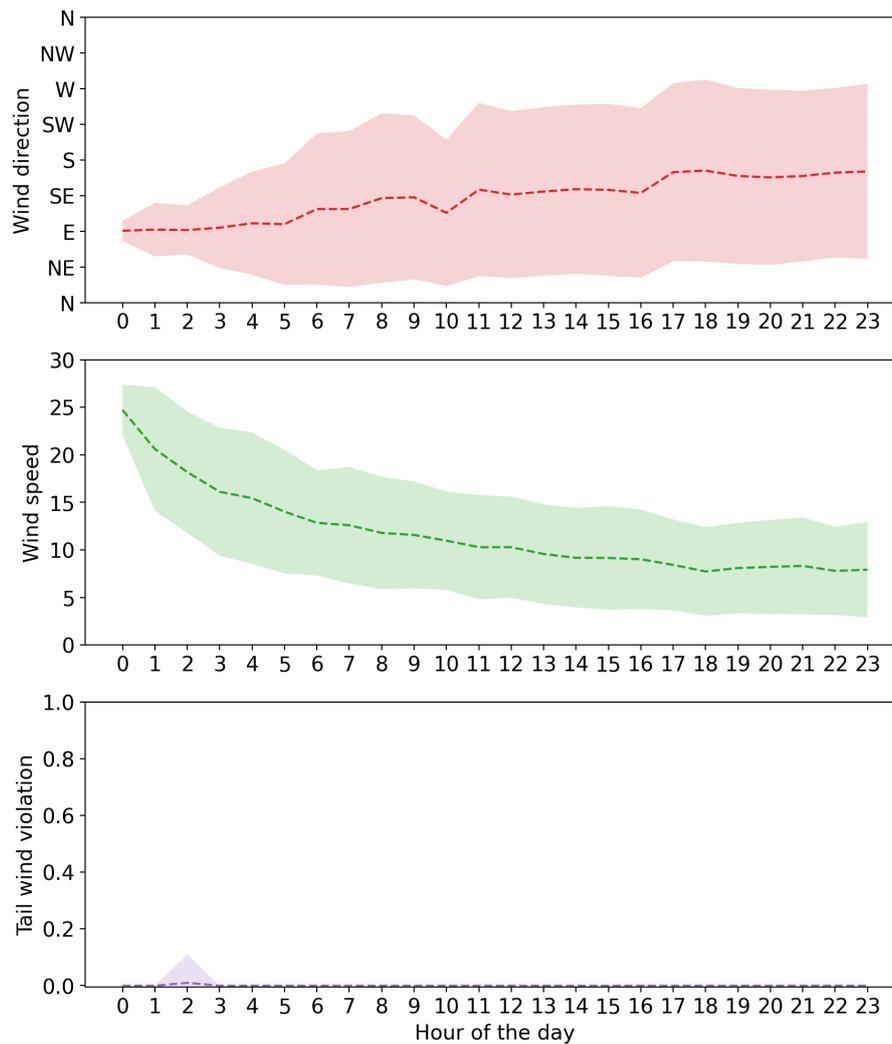
314 We estimate reasonable changes on an hourly basis in wind and meteorological conditions and simulate  
315 realistic operations (based on historical data), each multiple times to consider random variations. In each  
316 episode, the wind speed and direction, the hour of the day, and the meteorological conditions are sampled,  
317 and the RCA tool predicts the runway configuration.

#### 318 **3.2.1 Procedure**

- 319 1. Random Day Selection: A day is randomly selected from the dataset.
- 320 2. Initial Condition: A random wind condition is chosen as the starting point for the simulation.
- 321 3. Simulation Process:
  - 322 • Using the transition matrix, we simulate the evolution of the wind conditions for the entire day.
  - 323 • This process is repeated 100 times for each selected day to capture a broad spectrum of possible
  - 324 wind condition transitions.
  - 325 • The tool predicts the runway configuration
- 326 4. Safety Evaluation:
  - 327 • For each simulated day, the output of the ML model is evaluated against the safety criteria.
  - 328 • This entire procedure was repeated for 100 different days to ensure the robustness and reliability
  - 329 of the results.

330 We showcase a series of figures depicting the simulated days in various scenarios. For CLT, we chose  
331 a randomly selected day. Similarly, for DEN and DFW, we selected and analyzed the most challenging  
332 scenarios involving random cases of tailwind and crosswind violations. We then compared these realistic  
333 scenario outcomes with those generated by Monte Carlo simulations. Each figure provides detailed  
334 information, displaying the mean and standard deviations calculated from 100 simulation iterations.  
335 This approach offers a comprehensive view of the variability of each scenario and the robustness of our  
336 simulation methodology in capturing the dynamics of runway safety under different wind conditions.  
337 Figure 4 illustrates a realistic simulation of a whole day in CLT. It captures fluctuations in wind speed and  
338 direction throughout the day, along with instances of tailwind violations. The depicted tailwind violation  
339 data represent the mean of 100 values, each coded as 0 or 1, where 0 indicates a safe scenario and 1 denotes  
340 an unsafe scenario. This provides a clear visual representation of the frequency and distribution of tailwind  
341 violations throughout the day.

342 Figures 5 and 6 show three different scenarios, worst-case scenarios of tailwind and crosswind violation,  
343 and a random one for DEN and DFW, respectively. The worst-case scenarios highlight instances of  
344 maximum safety violations within the system. By adjusting the predictions made by the RCA tool, we  
345 can effectively address and mitigate these violations. The tailwind and crosswind violations sub figures



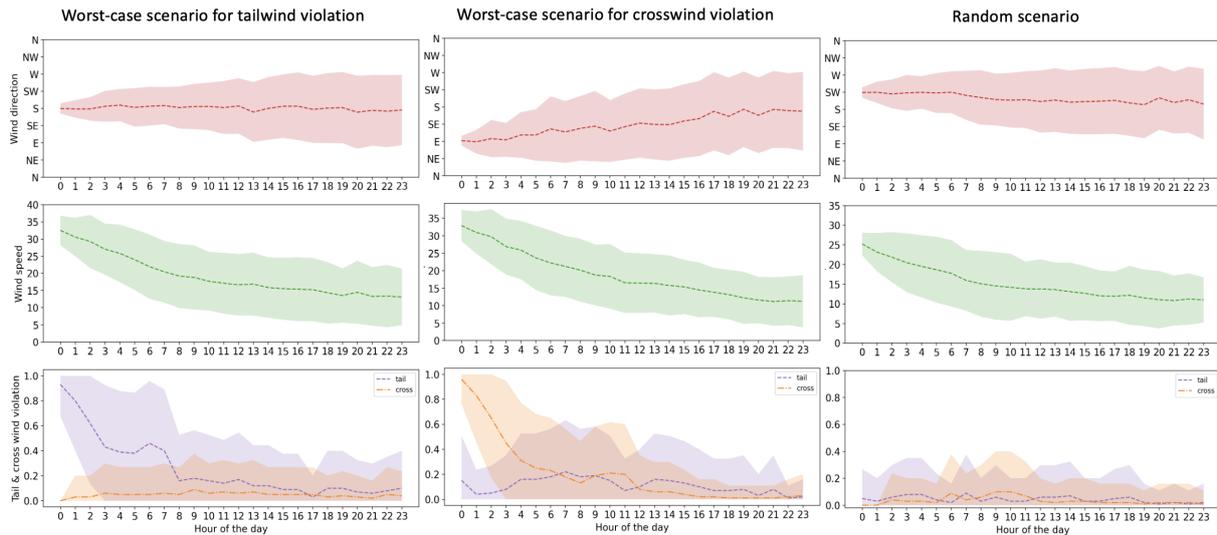
**Figure 4.** Realistic scenario simulation of CLT airport

**Table 3.** Comparison between Monte Carlo (MC) and realistic simulations for all airports

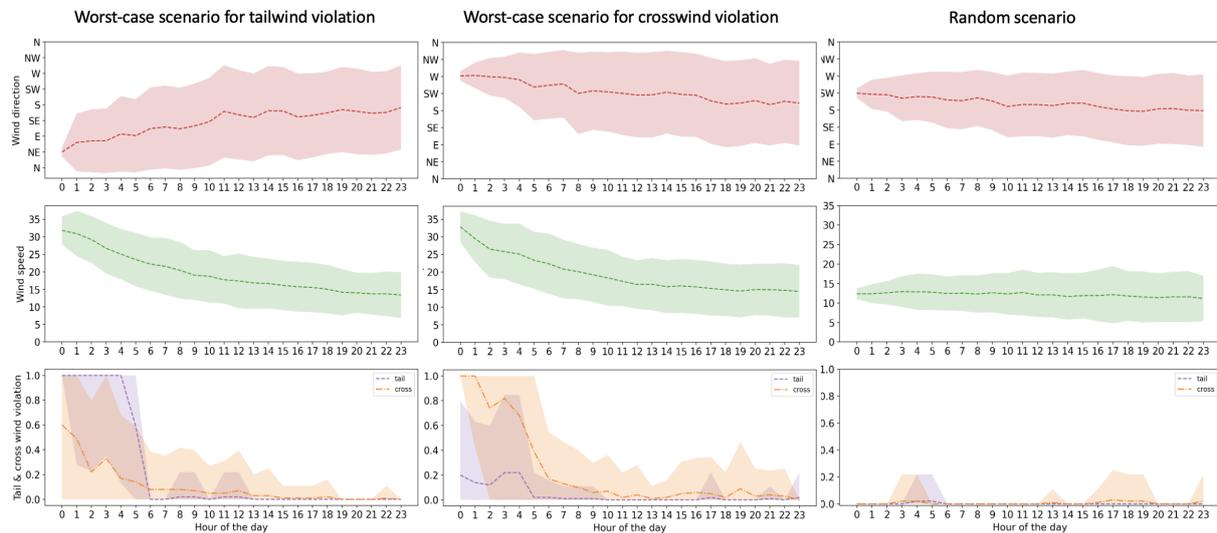
Airport	CLT		DEN		DFW	
	MC	Realistic	MC	Realistic	MC	Realistic
Crosswind violation (%)	4.7	0.4	16.7	3.7	6.2	1.3
Tailwind violation (%)	0	0	15	5.8	20.3	2.6

346 show that the model predictions decrease the safety violation throughout the day. The random scenario  
 347 sub-figures also show significant low violations among 100 simulations during a day. The dashed line  
 348 indicates the mean value, and the colored boundary shows the standard deviation.

349 Table 3 shows the instances in which safety criteria are violated during both types of simulation. Notably,  
 350 there is a significant reduction in the percentage of violations when comparing the results of Monte Carlo  
 351 simulations to those from realistic scenarios. This indicates a marked improvement in the adherence to  
 352 safety standards under more realistic operating conditions. This notable discrepancy between the two  
 353 simulation scenarios, Monte Carlo and realistic, underscores the robustness and adaptability of our realistic  
 354 simulation method, particularly in its superior ability to capture and respond to dynamic environmental  
 355 variables when compared to the Monte Carlo simulations. This suggests that realistic simulation provides a  
 356 more effective framework for understanding and managing complex real-world scenarios. The goal was to



**Figure 5.** Realistic scenario simulation of DEN airport



**Figure 6.** Realistic scenario simulation of DFW airport

357 evaluate the performance of the model in predicting safe outputs over extended periods, reflecting realistic  
 358 daily variations in wind conditions. This approach helps us understand how temporal changes in wind  
 359 conditions impact the safety and reliability of the model's predictions.

## 4 CONCLUSION

360 The application of formal verification methods to the RCA tool helps us validate the tool's output in  
 361 adhering to crucial safety criteria in air traffic control. This study underscores the value of using formal  
 362 methods, such as model-checking, to rigorously assess the safety of machine learning algorithms within a  
 363 safety critical operational setting. By validating the RCA tool's compliance with safety criteria in various  
 364 simulated wind conditions, the research highlights the potential of formal verification to enhance the  
 365 trustworthiness of automated decision-support systems. The results of the Monte Carlo random wind  
 366 condition simulations provided insight into the model's ability to handle a wide range of wind speeds and  
 367 directions, highlighting potential biases or limitations in handling extreme conditions. In addition, the  
 368 simulations based on the (historical data-based) transition matrix offered a detailed view of the model's

369 performance in realistic and dynamic wind scenarios, revealing how well the model maintains safety  
370 standards throughout typical daily wind fluctuations.

371 It is worth to mention that while the RCA tool currently uses CQL with tabular representation and  
372 function approximators, future extensions to deep RL architectures could leverage abstractions or symbolic  
373 encodings (e.g., decision trees, BDDs) to preserve tractability in model-checking. Existing approaches in  
374 neural-symbolic verification (e.g., abstraction-refinement) could support scalability.

## ACKNOWLEDGMENTS

375 We thank Subject Matter Experts (SMEs) affiliated with the FAA office of NextGen for their guidance and  
376 support in performing the validations.

## APPENDIX A. RCA-CTLOGIC VERIFICATION ALGORITHM

377 To complement the formal definitions and theoretical model-checking framework described above,  
378 Algorithm 1 outlines the complete verification procedure used to assess the safety of RCA tool outputs  
379 under varying wind conditions. The algorithm receives historical or simulated wind data as input, applies  
380 the RCA model to generate runway configuration decisions, and evaluates each decision against established  
381 crosswind and tailwind safety thresholds. The results are then analyzed using CTLogic to determine  
382 whether the model satisfies the temporal safety properties  $EF\varphi$  and  $EG\varphi$ .

---

### Algorithm 1 RCA-CTLogic Model Checking Algorithm

---

**Require:** RCA\_model, Wind\_Data, Runway\_Set, Safety\_Limits, Transition\_Matrix (optional)

**Ensure:**  $EF\varphi$  and  $EG\varphi$  satisfaction results

```

1: Initialize CTL_State_Set  $\leftarrow []$ 
2: for each  $(V, \alpha)$  in Wind_Data do
3:   config  $\leftarrow$  RCA_model.predict( $V, \alpha$ )
4:    $(V_{tail}, V_{cross}) \leftarrow$  compute_components( $V, \alpha, config$ )
5:   is_safe  $\leftarrow (V_{tail} < L_{tail}) \wedge (V_{cross} < L_{cross})$ 
6:   Append  $(V, \alpha, config, is\_safe)$  to CTL_State_Set
7: end for
8: Build Transition_Graph from CTL_State_Set using temporal ordering
9:  $EF\varphi \leftarrow$  EXISTS_PATH(CTL_State_Set,  $\lambda s: s.is\_safe$ )
10:  $EG\varphi \leftarrow$  ALL_PATHS_ALWAYS(CTL_State_Set,  $\lambda s: s.is\_safe$ )
11: return  $EF\varphi, EG\varphi$ 

```

---

## REFERENCES

- 383 AlSobeh, A. (2024). Osm: Leveraging model checking for observing dynamic 1 behaviors in aspect-  
384 oriented applications. *arXiv preprint arXiv:2403.01349*
- 385 Avery, J. and Balakrishnan, H. (2016). Data-driven modeling and prediction of the process for selecting  
386 runway configurations. *Transportation research record* 2600, 1–11
- 387 Badrinath, S., Li, M. Z., and Balakrishnan, H. (2019). Integrated surface–airspace model of airport  
388 departures. *Journal of guidance, control, and dynamics* 42, 1049–1063
- 389 Baheri, A., Ren, H., Johnson, B., Razzaghi, P., and Wei, P. (2022). A verification framework for certifying  
390 learning-based safety-critical aviation systems. *AIAA AVIATION 2022 Forum*, 3965
- 391 Baier, C. and Katoen, J.-P. (2008). *Principles of model checking* (MIT press)
- 392 Bérard, B., Bidoit, M., Finkel, A., Laroussinie, F., Petit, A., Petrucci, L., et al. (2013). *Systems and*  
393 *software verification: model-checking techniques and tools* (Springer Science & Business Media)
- 394 Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., et al. (2012). A  
395 survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in*  
396 *games* 4, 1–43
- 397 [Dataset] FAA (2024). Airport diagram of CLT

- 398 [Dataset] FAA, O. A.-. (1981-11-09). National safety and operational criteria for runway use programs  
399 Gao, H., Huang, W., Liu, T., Yin, Y., and Li, Y. (2022). PPO2: Location privacy-oriented task offloading  
400 to edge computing using reinforcement learning for intelligent autonomous transport systems. *IEEE*  
401 *transactions on intelligent transportation systems*
- 402 Jacquillat, A., Odoni, A. R., and Webster, M. D. (2017). Dynamic control of runway configurations and  
403 of arrival and departure service rates at JFK airport under stochastic queue conditions. *Transportation*  
404 *Science* 51, 155–176
- 405 Kumar, A., Zhou, A., Tucker, G., and Levine, S. (2020). Conservative q-learning for offline reinforcement  
406 learning. *Advances in neural information processing systems* 33, 1179–1191
- 407 Li, L., Clarke, J.-P., Chien, H.-H. C., and Melconian, T. (2009). A probabilistic decision-making model  
408 for runway configuration planning under stochastic wind conditions. In *2009 IEEE/AIAA 28th Digital*  
409 *Avionics Systems Conference (IEEE)*, 3–A
- 410 Li, Y., Li, Y., and Ma, Z. (2015). Computation tree logic model checking based on possibility measures.  
411 *Fuzzy Sets and Systems* 262, 44–59
- 412 Memarzadeh, M. and Kalyanam, K. (2025). Runway configuration assistance: Offline reinforcement  
413 learning method for air traffic management. *Journal of Aerospace Information Systems* 22, 275–287
- 414 Memarzadeh, M., Puranik, T. G., Kalyanam, K. M., and Ryan, W. (2023). Airport runway configuration  
415 management with offline model-free reinforcement learning. In *AIAA SCITECH 2023 Forum*. 0504
- 416 Nethi, S., Memarzadeh, M., and Kalyanam, K. (2024). Optimization of runway configurations with  
417 forecast-augmented offline reinforcement learning. In *AIAA SCITECH 2024 Forum*. 0533
- 418 Paul, S., Cruz, E., Dutta, A., Bhaumik, A., Blasch, E., Agha, G., et al. (2023). Formal verification of  
419 safety-critical aerospace systems. *IEEE Aerospace and Electronic Systems Magazine* 38, 72–88
- 420 Razzaghi, P., Tabrizian, A., Guo, W., Chen, S., Taye, A., Thompson, E., et al. (2024). A survey on  
421 reinforcement learning in aviation applications. *Engineering Applications of Artificial Intelligence* 136,  
422 108911
- 423 Vardi, M. Y. (2005). An automata-theoretic approach to linear temporal logic. *Logics for concurrency:*  
424 *structure versus automata* , 238–266