Optimization of Airport Runway Configuration with Forecast-Augmented Offline Reinforcement Learning

Sumanth Nethi*, Milad Memarzadeh† and Krishna Kalyanam‡
NASA Ames Research Center, Moffett Field, CA 94035, USA

Runway configuration Management (RCM) governs the optimal utilization of runways based on variables such as traffic and meteorological conditions, making it a daunting task in air traffic management due to its dependency on volatile operational and environmental factors. This paper improves upon our previous work [1] on using offline model-free reinforcement learning for creating a Runway Configuration Assistance (RCA) decision-support tool. A novel integration of forecast data from LAMP (Localized Aviation Model Output Statistics Program) and TAF (Terminal Area Forecast) is introduced, enhancing the tool's accuracy and also its adaptability to quick wind changes. The performance is evaluated using two major US airports, Charlotte Douglas International Airport (CLT) and Denver International Airport (DEN). To counter scalability issues presented by the addition of discrete forecast variables, we transitioned to a continuous state space model, ensuring scalability and inclusion of longer forecast data. The results of our experiments reflect significant improvements in the RCA tool's prediction accuracy.

I. Introduction and Related Work

Runway Configuration Management (RCM), a critical component of Air Traffic Management (ATM), deals with the optimal selection of runways for arrivals and departures based on traffic, surface wind speed, wind direction, and other meteorological variables. Given the uncertain and complex nature of operational and meteorological variables, RCM relies on the local knowledge and experience of the Air Traffic Controllers (ATCo) and is invisible process for stakeholders and airspace users. Moreover, it takes time for a ATCo to build the local knowledge and skills to facilitate an efficient change in the runway configuration. The current practice sets the runway configuration by the ATCo based on relevant information available at the time. This makes the decision-making process subjective based on the accuracy of the available information, weather forecast models and ATCo experience and local knowledge. This process can result in delayed coordination and execution of a runway configuration change given the uncertainty in the forecast data and the complexity of the decision process. The search for an optimal policy might require evaluating exorbitant

*NASA OSTEM Intern, NASA Ames Research Center
†Universities Space Research Association, AIAA Member, Corresponding Author: milad.memarzadeh@nasa.gov
‡Aviation Systems Division, NASA Ames Research Center, AIAA Associate Fellow
This paper will be presented at the AIAA Scitech 2024 Forum in Orlando, FL, 8–12 January 2024.

number of possible scenarios that cannot be done by human reasoning alone. An automated approach based on Artificial Intelligence (AI) and Machine Learning (ML) can make use of historical data and search through all (or significant amount of) possible scenarios under forecast uncertainty and make well-informed decisions. Moreover, the automated technology will be able to facilitate the transitional periods between controllers shift changes.

AI/ML has been used previously to address the RCM problem. One popular approach uses different variants of model-based techniques, such as discrete choice modeling [2], dynamic programming [3], and its combination with queue modeling [4, 5]. Although model-based approaches are robust and interpretable, they suffer from a fundamental drawback. Their performance depends significantly on building an accurate model that can mimic the real-world operations and changes in the traffic and meteorological conditions. A potential remedy to this challenge is to use online model-free Reinforcement Learning (RL) approaches such as Monte Carlo Tree Search (MCTS) [6] for learning a good policy without relying on learning a specific model. However, such online approaches require a significant number of interactions in the operational setting to collect data and learn from the feedback. This interaction is impractical in safety-critical problems such as ATM, because data collection can be expensive. As a result, most of the recent literature is focused on data-driven approaches based on supervised ML to predict the best choice of runway configuration given all the independent factors such as weather and future traffic [7–9]. Although these approaches show great accuracy in predicting the runway configuration, such prediction is not backed up by any evidence that this would decrease the transition times or would alleviate a safety concern. The reason is that supervised learning only learns to mimic the ATCo with least amount of error and does not have an underlying mechanism (such as the reward/utility function in RL) to improve historical decisions.

This paper presents an innovative, automated approach using offline model-free RL to provide decision-support for RCM. The primary objective of this work is the development and optimization of a Runway Configuration Assistance (RCA) tool, specifically based on an algorithm called Conservative Q-Learning (CQL) [10]. The tool processes historical data about variables of interest, decisions made regarding RCM, and their subsequent outcome, to identify a policy that would encourage decision-making that optimizes airport rates and avoid the inappropriate decisions. The policy search is guided by an appropriately chosen weighted utility function, designed to improve efficiency of runway change timing and coordination, and mitigate delays and arrival go-arounds. Our previous work [1] demonstrated feasibility of the CQL algorithm in a simulated setup for Charlotte Douglas International Airport (CLT) and Denver International Airport (DEN). However, the developed model showed significant sensitivity in cases where the meteorological conditions (e.g., wind direction and/or speed) were rapidly changing. In practice, controllers tend to ignore rapid changes in the meteorological conditions due to the fact that switching the runway configuration takes a significant amount of time and is not efficient for such scenarios of rapid changes in the weather. As a result, the developed solution needs to be confident that the new configuration will be in place long enough that the expected accrued benefit exceeds the "cost" of making the change.

To address this, we propose to integrate weather forecast data from LAMP (Localized Aviation MOS Program) and TAF (Terminal Area Forecast) into the model. The forecast data for the upcoming one-hour period was incorporated into both the state space and the reward or utility function of the model. We hypothesize that this addition improves the model's ability to anticipate and adapt to rapid wind changes, leading to enhanced responsiveness and accuracy as well as stability in the decision-support tool for the ATCo. The model's state space and utility function is both augmented by encoding the forecast data into them. The state space is expanded to capture the forecast of wind speed, wind direction, and meteorological conditions for the next 1 hour. The utility function is also modified to adapt to sudden changes in wind direction or speed and penalize frequent configuration changes. In order to integrate multiple hours of forecast the state space is modified from discrete to a continuous setting and the new models' performance is evaluated.

In this paper, the performance of the enhanced model is analyzed and validated against the original benchmark model established in [1]. Two years of real-world data are used from two airports across the National Airspace System (NAS): CLT, as a representative of airports with simpler runway configuration, and DEN, as a representative of airports with complex runway configuration, thus demonstrating the model's potential performance on a variety of airports across the NAS. The results of our experiments from augmenting the model with one-hour forecast for CLT and DEN indicated significant improvements, especially in dealing with quick wind changes. By incorporating forecast data, the model became more adaptive and responsive, leading to better prediction performance.

II. Method

RCM is inherently complex due to the many dynamic variables that need to be considered. For this reason, we adopted the framework of a Markov Decision Process (MDP) to formulate the problem. By doing so, we are able to create a structured approach that systematically handles the various elements of RCM.

A. MDP Framework for RCM

In our study, we formalized the runway configuration prediction problem as an MDP, characterized by the tuple (S, A, T, U, γ):

- S: Represents the state, encapsulating all necessary information required for decision-making. This includes wind speed, wind direction, prevailing weather conditions, and any other pertinent operational information.
- A: Denotes the set of potential actions, corresponding to various runway configurations. Different configurations can have varying impacts on efficiency, safety, and operational flow. Number of runway configurations depends on the airport. Less complex airports like CLT has just two major configuration while more complex airport like DEN has 11 major runway configurations.
- $T : S \times A \rightarrow S$: The transition function describes how actions taken in the current state will influence the future state of the system.

- $U : S \times A \rightarrow \mathbb{R}$: The utility function, which provides feedback on the quality of decisions made in specific states. The long-term summation of utility is optimized to ensure the safety and efficiency of operations, reduce wait times, and increase the overall throughput of the airport.
- γ: The discount factor translating future utilities to their net present value.

The primary objective of using this MDP framework is to identify an optimal policy $\pi^* : S \rightarrow A$ that maximizes the expected long-term utility V^{π^*}. This relationship is mathematically expressed as:

$$V^{\pi^*}(s) = u(s, \pi^*(s)) + \gamma \sum_{s' \in S} p(s'|s, \pi^*(s)) V^{\pi^*}(s') \tag{1}$$

B. Integration of Forecast Data

One of the standout features of our approach is the integration of weather forecast data into the model. By incorporating data from LAMP or TAF into our model, we are able to enhance its predictive capability, especially in cases of rapidly changing meteorological conditions. Initially, the model's state representation was designed for a forecast duration of one hour, using a discrete representation. However, as we sought to incorporate longer forecasts, this discrete representation became limiting due to scalability concerns. To address this, we transitioned to a continuous state space representation. The refined state, S, now includes continuous variables that reflect the forecasted wind speed and direction, as well as other meteorological conditions across the desired forecast horizon. This change offers a richer data set for our model, allowing for more accurate predictions and decisions.

C. Refinement of the Utility Function

Based on the insights gained from our initial model and the challenges it faced with rapid meteorological changes, we made modifications to the utility function, U. The new utility function is not only responsive to sudden shifts in wind patterns but also penalizes frequent configuration changes. This ensures that the proposed configurations are not only optimal in terms of safety and efficiency but also feasible in terms of operational transitions. Details of the refinements are explained later in the *Results and Discussion* section.

D. Learning Algorithm - Conservative Q-Learning (CQL)

For the learning component, we turned to CQL, an advanced offline model-free RL technique [10]. This methodology was chosen due to its robustness in handling the distributional shift problem, which is prevalent in offline RL settings. CQL achieves this by regularizing Q-values for actions that are not well represented in the collected data, ensuring that the model remains grounded in reality. Our utilization of the CQL framework sets our approach apart, offering a more data-driven and adaptive solution to the RCM problem. Performance benchmarks were established in reference to our prior work in [1], allowing us to quantitatively evaluate the advancements made.

III. Results and Discussion

To evaluate the advancements made in our methodology, we employed multiple tests using real-world data for both CLT and DEN airports. The main contribution of our findings was the demonstrable value of incorporating forecast data, particularly in scenarios involving quick wind changes. In order to quantify and compare performance of the developed methodology with historical decisions made by the ATCo, we use the agreement metric. This metric quantifies on average, how often the developed RCA tool agreed with historical decisions made by ATCos. Although this metric should ideally be high, it should not be at 100%. The reason for this is that, the ideal tool should be able to identify sub-optimal decisions that have been made and correct them, which results in a lower agreement with historical decisions.

A. CLT and DEN Airport Setup

1. Charlotte Douglas International Airport (CLT) Setup

As depicted in Figure 1 (left panel), Charlotte Douglas International Airport (CLT) predominantly employs two principal runway configurations for both arrivals and departures: "North flow" (using runways 36L/C/R) and "South flow" (using runways 18L/C/R). The relatively simple configuration makes it a fitting case for testing the enhancements proposed in our method. We utilized hourly data from both 2018 and 2019, collating information from five principal sources:

- NASA's Sherlock Data Warehouse (`https://sherlock.opendata.arc.nasa.gov`)
- FAA's Aviation System Performance Metrics (ASPM) database (`https://aspm.faa.gov`)
- METeorological Aerodrome Reports (METAR) database (`https://www.aviationweather.gov/metar`)
- NOAA's Localized Aviation MOS Program (LAMP) database (`https://www.nws.noaa.gov/mdl/gfslamp/gfslamp.shtml`)
- Terminal Area Forecast (TAF) database (`https://www.faa.gov/data_research/aviation/taf`)

Our state space, structured in alignment with conventions in discrete state-action RL research, comprises five variables:

- Wind Direction: Split into 8 discrete states, segmenting the full 360° and centered at 0°, 45°, 90°, 135°, 180°, 215°, 270° and 315°.
- Wind Speed: binned into 4 groups in nautical miles per hour (knots): [0-5), [5-10), [10-15), and ≥ 15.
- Hour of the Day: 24 distinct states.
- 1-hour Forecast of Wind Direction (8 discrete states).
- 1-hour Forecast of Wind Speed (4 discrete states).

These variables culminate in a state space with 24,576 dimensions. As for the action space, we have the two aforementioned configurations for CLT: "North" and "South" flow.

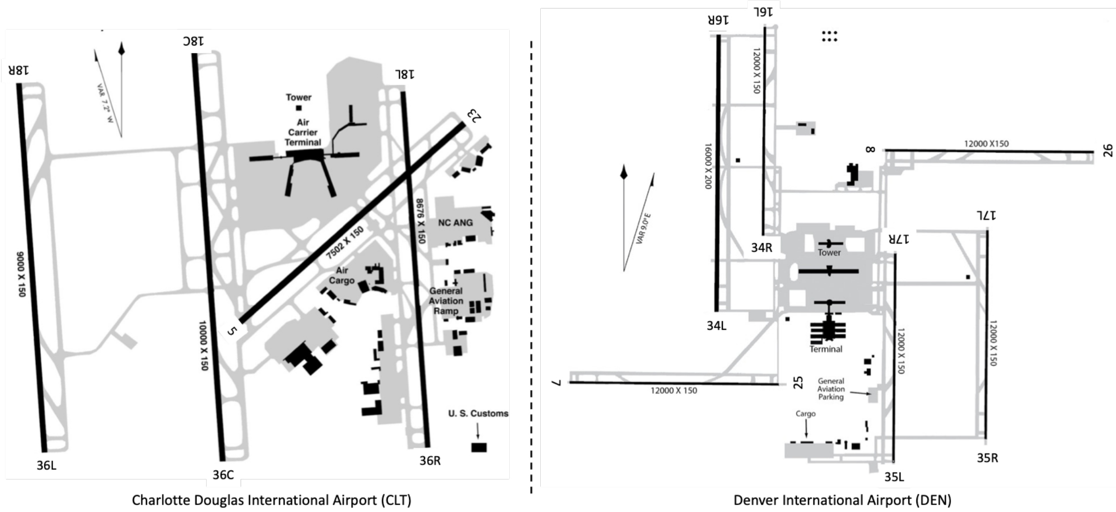

Fig. 1 This figure shows the airport surface (runway) diagrams for CLT (left) and DEN (right).

2. Denver International Airport (DEN) Setup

Denver International Airport (DEN), known for its intricate runway layout and configuration (shown in Figure 1, right panel), presents a much more complex setup. It has four parallel runways that operate North/South-bound (34/16 R/L, 35/17 R/L) and two parallel runways that operate East/West-bound (7/25, 8/26). Being one of the most complex airports, DEN poses unique challenges for runway configuration management. The data for DEN was sourced from the same repositories as CLT, spanning the years 2018 and 2019.

For DEN, the state space setup was similar to the CLT, except for the wind speed that is categorized into more categories as follows:

- Wind Speed: Given the intricate nature of DEN's environment, wind speed is more finely binned into 6 groups in nautical miles per hour (knots): [0-5), [5-10), [10-15), [15-20), [20-25), and $\geq$ 25.

The action space for DEN is larger, encompassing 11 distinct runway configurations to choose from. The bigger action space showcases the complexity inherent to DEN, requiring our model to make more difficult decisions.

B. Shortcomings of Prior Work in Quick Wind Change Scenarios

One of the pivotal challenges associated with the model in [1] was its restricted ability to respond to quick changes in wind conditions. Atmospheric conditions in aviation are full of complexities, and there exist instances wherein the wind undergoes a quick shift in both direction and speed, only to revert to its preceding state within a limited time span. Such rapid fluctuations in wind patterns present significant operational challenges, especially in the context of RCM.

It is worth noting that altering runway configurations is not a trivial task. In fact, from both an operational and efficiency standpoint, substantial resources and coordination efforts are required. Consequently, when faced with short-lived wind alterations, ATCos often opt for a more pragmatic approach. Instead of continuously adjusting the runways to momentarily changing winds, they often maintain the status quo, deeming it operationally more viable.

Figure 2 shows the wind patterns at specific intervals, emphasizing a key observation on 27 August 2018, around 5 PM. This span exhibited a significant fluctuation in wind speed and direction: transitioning swiftly from a Northwest orientation of 310° (NW) to a Southeast orientation at 150° (SE), while picking up intensity to 23 knots, before reverting to 330° (NW). This transient oscillation in the wind direction might seem negligible at first; however, it holds substantial implications for operational choices. ATCos displayed a preference for the 'North' configuration during this brief window, based on their expertise and considerations regarding the efficiency of the operations.

Yet, when we compare the configuration of preference for all similar wind direction and speed scenarios in the historical data in Figure 3, we can clearly see that the ATCo selected 'South' configuration for all other scenarios, except those two 15-minutes intervals depicted in Figure 2. It is also evident from the last two columns of Figure 3 that the RCA tool consistently assigned higher Q-values to the 'South' configuration across all analogous wind scenarios.

The example in Figure 2 shows a scenario which the ATCo refer to as a *quick wind change*. In this scenario, where the wind changes direction and/or speed in a short span of time, and then reverts back to the original values, the ATCo tend to not change the runway configuration. Changing a configuration for a short amount of time can be very expensive and it might cause more delays and operational inefficiencies. On the other hand, the RCA tool from [1], appears to be indifferent to the quick wind changes because it does not take wind forecast into account. This insensitivity to quick wind changes can be problematic.

Local Time (hrs)	flow	wind_dir	wind_spd
1630	N	310	4
1645	N	230	3
1700	N	150	23
1715	N	150	23
1730	N	180	12
1745	N	330	3
1800	N	330	3

Fig. 2 Wind patterns on 27 August 2018, highlighting rapid changes in wind direction and speed around 5PM.

The aforementioned observational analysis spotlights a vital shortcoming of the model presented in [1] — its limited adaptability to rapidly changing wind conditions. This paper endeavors to address this gap, aiming to offer a more synchronized and operationally aligned runway configuration decision-making framework.

Local Time (hrs)	flow	wind_dir	wind_spd	Q_N	Q_S
1215	S	150	22	701.53	705.67
1230	S	150	22	701.53	705.67
1945	S	150	20	455.09	457.41
2000	S	150	20	449.22	451.16
1815	S	140	22	614.50	614.68
1700	N	150	23	583.49	585.75
1715	N	150	23	583.49	585.75

Fig. 3 Comparison of Q-values for 'North' and 'South' configurations as suggested by the RCA tool.

C. Augmented 1-Hour Forecast Data and Modified Utility Function

With the intent to address the shortcoming highlighted in the previous section, we augmented our model with an additional 1-hour forecast data for wind direction, wind speed, and meteorological condition. This enhancement was based on the rationale that incorporating forecast data might provide the RCA tool with the ability to better adapt to rapid changes in wind patterns.

The 1-hour forecast data is sourced from either LAMP database. All forecasted data for the next hour are discretized and integrated as categorical variables similar to how the wind related variables for the current time are included, ensuring a seamless integration with the existing state space. The primary goal of this integration is to enhance the decision-making capability of the RCA tool. For instance, if an abrupt change in wind direction or speed is observed but the 1-hour forecast suggests a reversion to the original wind pattern, it can be deduced that the change is temporary. Hence, the RCA tool might decide against a runway configuration change and in alignment with what the ATCo would opt for.

In the realm of ATC, quick alterations in wind conditions can pose significant challenges for ensuring optimal runway configurations. Recognizing the imperative of making consistent decisions in the face of such rapid wind changes, our study aimed to refine the utility function used in previous work [1].

The hallmark of our updated utility function is its ability to incorporate a term specifically tailored to capture the nuances of rapid wind changes. This enhancement serves a dual purpose: it provides a theoretical foundation while also ensuring practical alignment with real-world ATCo practices, especially during transient wind conditions.

The refined utility function is represented as:

$$u_t = \lambda v_t - \mu \bar{\tau}_t - \beta c_{t,ga} - \eta c_{t,mga} - \zeta \mathbb{I}[a_t \neq a_{t-1}] - c d_{wind} \tag{2}$$

Here, the terms correspond to:

- v_t: traffic throughput

- $\bar{\tau}_t$: average transit times on the surface of the airport
- $\beta c_{t,ga} + \eta c_{t,mga}$: go-around and multiple go-arounds penalties
- cd_{wind}: quick wind penalty
- $\mathbb{I}[a_t \neq a_{t-1}]$: function that returns 1 when the configuration changes between time $t-1$ and t, and 0 if it remains the same.

Furthermore, the parameter d_{wind} is defined conditionally based on the threshold (thr) for wind direction (θ_{wdr}). Figure 4 illustrates the calculation of the quick wind change.

$$d_{wind} = \begin{cases} \theta_{wdr} & \text{if } \theta_{wdr} > thr \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

This formulation ensures only significant wind direction changes surpassing a predetermined threshold are taken into account, effectively filtering out minor fluctuations. In our simulations, the values for weights in Eq. (2) were finalized based on hyper-parameter tuning by dividing the training data into training and validation sets and performing grid search. For the CLT airport, we set the parameters as: $\lambda = \mu = 5$, $\beta = 10$, $\zeta = 1$, $\eta = 10$, and $c = 0.1$. Meanwhile, for the DEN airport, the weights were set to: $\lambda = \mu = 5$, $\beta = 10$, $\zeta = 1$, $\eta = 5$, and $c = 0.1$. These configurations were crucial in guiding the interpretations and insights derived from our plots.

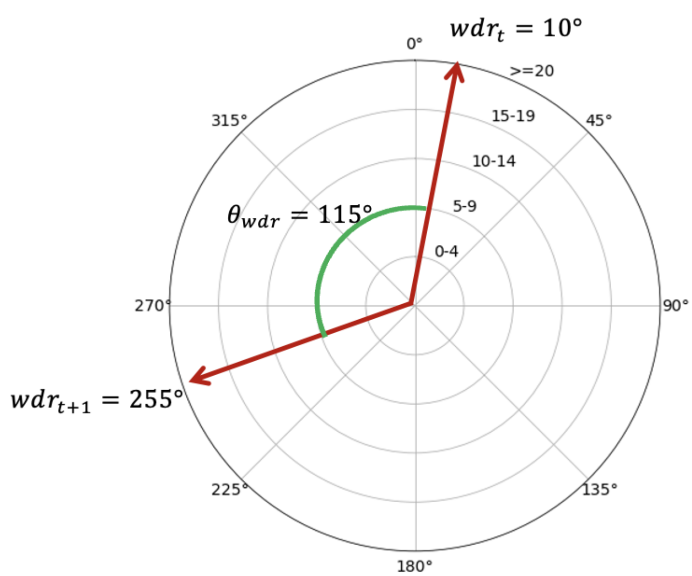

Fig. 4 Visual representation of θ_{wdr} in relation to wind direction changes.

The updated utility function combines the latest theoretical approaches with real-world needs, making sure the chosen runway configurations are not only the best choices but also realistic.

Building upon the RCA tool presented in [1], our modifications show noteworthy improvements in performance.

Referring to the Figure 5, it becomes evident that integrating the RCA tool with 1-hour forecast data from LAMP for CLT airport enhances its agreement with historical data by 3.2 percentage points (pp). This enhancement becomes even more pronounced during instances of quick wind changes, where the introduction of 1-hour forecast data improves the agreement by a significant 12.4pp. We quantify the quick wind change metric as the ratio of identified quick wind change instances to the instances where the model's decision was in alignment with actual decisions made by the ATCo.

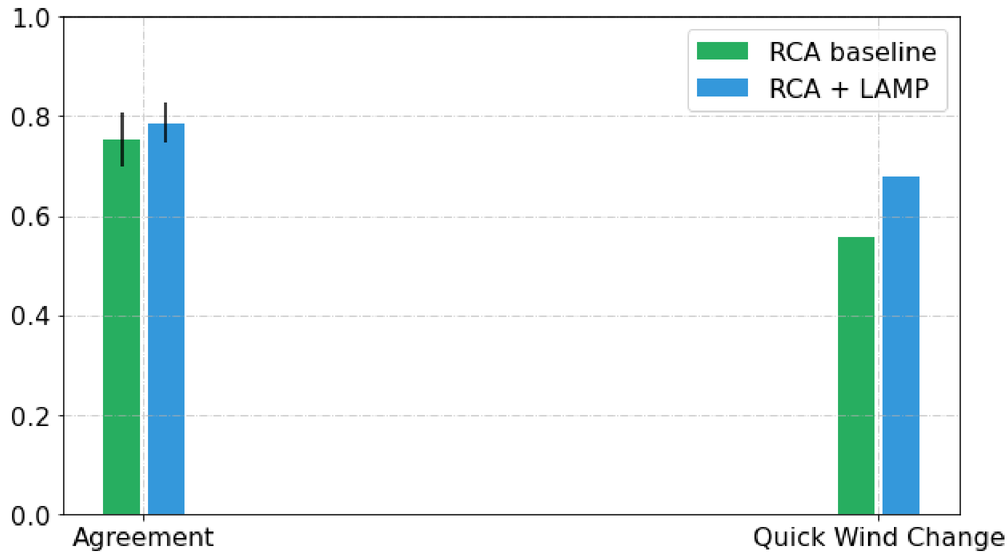

Fig. 5 Performance comparison for RCA with and without 1-hour forecast data for CLT.

As for DEN, discernible performance enhancement from the incorporation of forecast data is observed (refer to Figure 6). The forecast-augmented RCA tool outperforms the baseline by an appreciable margin of 4.2pp (with LAMP forecast) and 3.7pp (with TAF forecast). Moreover, during scenarios characterized by quick wind changes, the improvement in performance was about 4.9pp (with LAMP forecast) and 3.8pp (with TAF forecast). Such observations underscore the value of embedding forecast data into the RCA model, suggesting that its integration leads to a marked boost in the tool's efficiency.

In summary, the forecast-integrated model showcases a notable improvement in its operational efficacy for both the CLT and DEN airports, attesting to the merit of our proposed modifications.

D. Forecast Integration in Continuous State-Space Model

The inherent computational overhead arising from the inclusion of elongated forecast duration within the RCA tool necessitated a shift towards a more scalable solution. To address this computational challenge, we designed a model situated in a continuous state space. This not only promotes scalability but also facilitates the seamless augmentation of protracted forecast data. In this evolved model, the utility function remains unaltered. The principal transformation resides in the expansion of the state space. Traditionally, variables like wind speed and wind direction were discretized

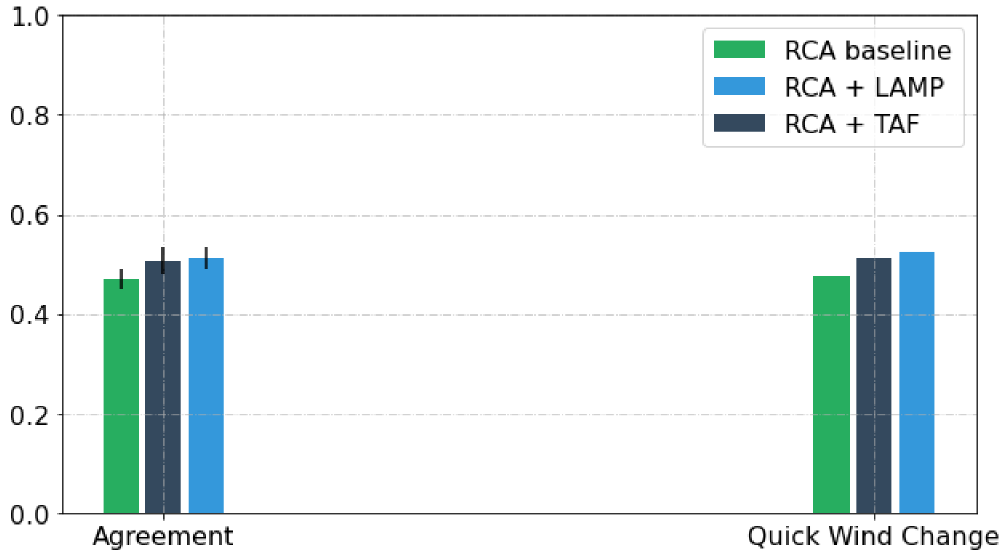

Fig. 6 Performance comparison for RCA with and without forecast data for DEN.

and binned, culminating in a categorical representation. However, in our continuous state space framework, these variables shed their categorical avatar and are introduced as normalized numbers, thereby justifying the "continuous" moniker. For context, in the conventional setup, the wind direction was delineated across eight discrete categories, while the wind speed was fragmented into four distinct bins for CLT and six for DEN. Incorporating merely an hour's worth of forecast data in this framework results in an addition of $8 + 4(6) = 12(14)$ for CLT (DEN) categorical divisions. This amplifies the total number of possible states exponentially, invoking computational challenges. However, under the continuous state space schema, we merely integrate two additional columns, symbolizing two extra variables, thereby maintaining computational efficiency. This will also improve the further augmentations of the state space to include further features related to traffic load, convective weather, etc.

Our exploratory findings demonstrate the distinct advantages of adopting a continuous state space approach. Even with the utility function maintained in its original form, without a dedicated quick wind penalty term, a marked enhancement in performance is witnessed by solely integrating forecast data into our state space. Especially in the case of DEN (Figure 8), where the decision-making is more complex, introducing forecast data can significantly refine our predictions.

Analyzing the plots provides several key observations. The expansion is specific to the state space, while the utility function is not changed (same utility function as used in [1]). In every instance within the continuous state space, there is an integration of six hours of forecast data. When considering CLT with LAMP forecast data, introducing a 1-hour forecast slightly improves the agreement with historical decisions made by the ATCo (i.e., agreement metric). However, further extending the forecast duration does not lead to significant performance improvement. On the other hand, for DEN, where both LAMP and TAF data sources were considered, performance using LAMP-augmented data shows a

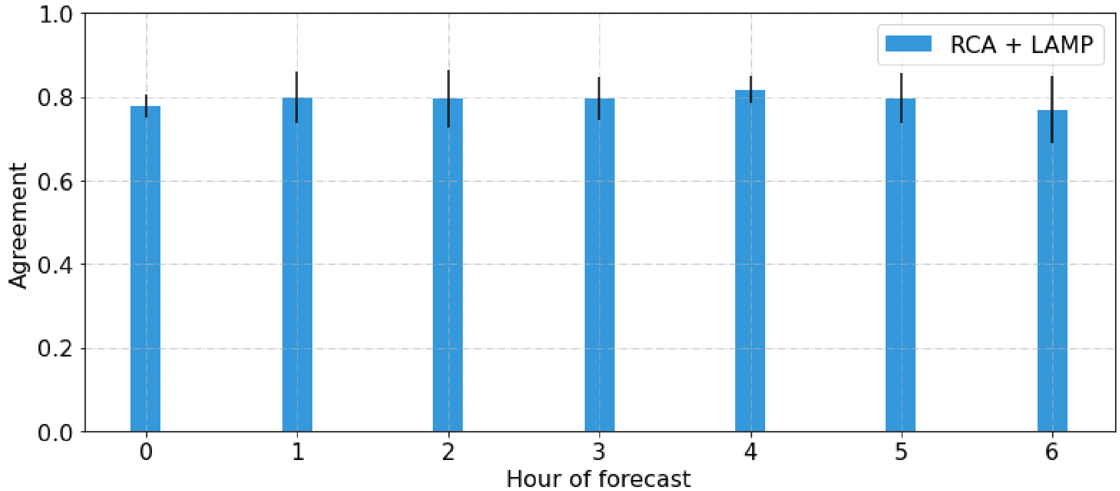

Fig. 7 Performance enhancement with continuous state space for CLT as forecast duration increases.

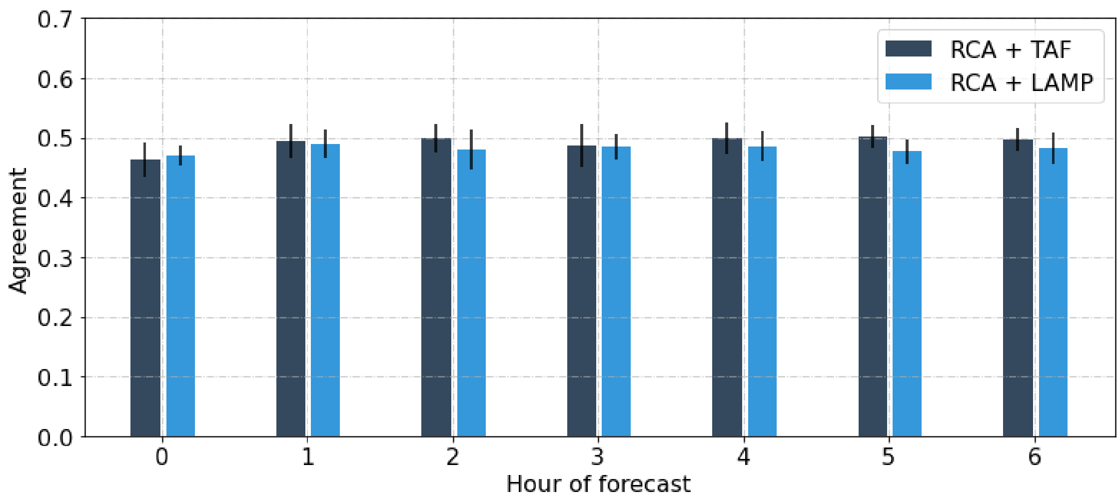

Fig. 8 Performance enhancement with continuous state space for DEN as forecast duration increases.

small improvement with a 1-hour forecast, stabilizing thereafter. More interesting, performance using TAF-augmented data improves steadily up to a 2-hour forecast, and stabilizes afterwards. The reason why the longer look-ahead forecast do not improve the runway configuration decision-making at the current time is that the longer horizon forecast become more uncertain and would influence the decision-making at the current time step less. Overall, forecast data from TAF is seen to perform better in the continuous state space model.

These observations suggest that there might be diminishing returns when incorporating longer forecast duration. The challenges of predicting further into the future could counteract the advantages of extended forecasts. However, the continuous state space model's ability to effectively utilize this forecast data indicates potential directions for enhancing the utility function, which warrants further investigation. In conclusion, the continuous state space RCA tool stands as a significant advancement in runway configuration decision-making, highlighting its effectiveness in terms of computational efficiency and prediction accuracy.

IV. Comparison Between Discrete and Continuous State Space

In this section, we aimed to evaluate the difference in performance when transitioning from a discrete to a continuous state space model. This comparison was only done with a 1-hour forecast integration due to computational complexity of including longer forecast in the discrete model.

Table 1 summarizes the performance comparison of the discrete and continuous state space models for both CLT and DEN. In the case of CLT, the forecast data was integrated from LAMP, while for DEN, both LAMP and TAF were taken into account and resulted in similar perforamnces.

Table 1 Comparison between discrete and continuous state space models.

Airport	Discrete	Continuous
CLT	0.79	0.8
DEN	0.51	0.49

Based on these findings, it's clear that shifting to a continuous state space with a 1-hour forecast generally offers performance similar to that of the discrete model. In the case of DEN, the continuous model's scores were just below those of the discrete model but this is within the margin of variations for the performance across different simulations. In essence, the continuous model provides an effective alternative to the discrete one, ensuring we make good runway decisions without losing out on accuracy.

V. Conclusions

Optimizing runway configuration management (RCM) is crucial for efficient air traffic management. Our research has focused on developing a state-of-the-art solution that is both effective and computationally efficient. One of our main achievements is the integration of forecast data into the RCM decision-making. This integration has improved the performance and stability of the model's predictive accuracy, particularly during quick and transient wind conditions, highlighting the importance of using forecast data. We also introduced a continuous state space model building upon our previously developed model [1] and have shown that such shift from a discrete to a continuous approach reduces computational burden and allows for easier integration of longer forecast periods, while achieving the same level of performance. To further advance the current tool, the continuous state space model can be adjusted, especially in the utility function, to improve its predictive accuracy.

In summary, this paper not only underscores the impact of judiciously integrating forecast data into RCM decision-making but also lays down the groundwork for future innovations in the field. Our findings serve as both a testament to progress made in RCM research and a call for continued exploration and refinement in the realm of ground based decision-support tools for air traffic management.

Acknowledgement

The authors acknowledge the invaluable support and feedback from collaborators and subject matter experts affiliated with the Federal Aviation Administration's (FAA) Office of NextGen. The material is based upon work supported by the National Aeronautics and Space Administration (NASA) under Contract Number NNA16BD14C, managed by the Universities Space Research Association (USRA).

References

[1] Memarzadeh, M., Puranik, T., Kalyanam, K., and Ryan, W., "Airport Runway Configuration Management with Offline Model-Free Reinforcement Learning," *AIAA SCITECH 2023 Forum, AIAA 2023-0504*, 2023. https://doi.org/10.2514/6.2023-0504.

[2] Avery, J., and Balakrishnan, H., "Data-Driven Modeling and Prediction of the Process for Selecting Runway Configurations," *Transportation Research Record*, Vol. 2600, 2016. https://doi.org/10.3141/2600-01.

[3] Li, L., Clarke, J.-P., Chien, H.-H., and Melconian, T., "A probabilistic decision-making model for runway configuration planning under stochastic wind conditions," *IEEE/AIAA 28th Digital Avionics Systems Conference*, 2009. https://doi.org/10.1109/DASC.2009.5347528.

[4] Jacquillat, A., Odoni, A. R., and Webster, M. D., "Dynamic Control of Runway Configurations and of Arrival and Departure Service Rates at JFK Airport Under Stochastic Queue Conditions," *Transportation Science*, Vol. 51, No. 1, 2016, pp. 155–176. https://doi.org/10.1287/trsc.2015.0644.

[5] Badrinath, S., Li, M. Z., and Balakrishnan, H., "Integrated Surface–Airspace Model of Airport Departures," *Journal of Guidance, Control, and Dynamics*, Vol. 42, No. 5, 2019, pp. 1049–1063.

[6] Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., and Colton, S., "A Survey of Monte Carlo Tree Search Methods," *IEEE Transactions on Computational Intelligence and AI in Games*, Vol. 4, No. 1, 2012. https://doi.org/10.1109/TCIAIG.2012.2186810.

[7] Khater, S., Rebollo, J., and Coupe, W. J., "A Recursive Multi-step Machine Learning Approach for Airport Configuration Prediction," *AIAA Aviation Forum*, 2021. https://doi.org/10.2514/6.2021-2406.

[8] Churchill, A., Coupe, W. J., and Jung, Y. C., "Predicting Arrival and Departure Runway Assignments with Machine Learning," *AIAA Aviation Forum*, 2021. https://doi.org/10.2514/6.2021-2400.

[9] Puranik, T., Memarzadeh, M., and Kalyanam, K., "Predicting Airport Runway Configurations for Decision-Support Using Supervised Learning," *42nd Digital Avionics Systems Conference, Barcelona, Spain*, 2023.

[10] Kumar, A., Zhou, A., Tucker, G., and Levine, S., "Conservative Q-Learning for Offline Reinforcement Learning," *arXiv:2006.04779*, 2020. URL https://arxiv.org/abs/2006.04779.