

Analysis and prediction of VFR vs IFR traffic behavior to support uncrewed aircraft flight operations at regional airports

Vishwanath Bulusu*, Husni Idris† and Aastha Acharya‡
NASA Ames Research Center, Moffett Field, CA, USA

Uncrewed Aircraft flight operations at a regional airport will be affected by the uncertainty in the Visual Flight Rules traffic around. This paper analyzes Visual Flight Rules traffic behavior, compares it with Instrument Flight Rules traffic and develops traffic prediction methods to support uncrewed aircraft flight operations. The spatio-temporal distribution of traffic operating under Visual Flight Rules and Instrument Flight Rules was analyzed from one month of historical track data around Fort Worth Alliance airport as a representative regional airport. The traffic behavior was visualized as occupancy maps generated at different altitudes. The Instrument Flight Rules traffic was concentrated in fewer regions of the airspace along structured routes where the risk of interacting with one or more flights reached close to fifty percent in some areas. Visual Flight Rules traffic was spread over more regions, mostly segregated from the Instrument Flight Rules regions, and the risk of interaction was lower reaching up to twenty-five percent on average over the month in some regions. The interaction risk was predicted using predictive occupancy maps over multiple time horizons and conditioned on time and partial observation of traffic in the vicinity. The month-to-month predictability of Visual Flight Rules risk was lower than that of the Instrument Flight Rules traffic, consistently over all the conditions analyzed. However, the prediction and its accuracy were demonstrated to be sensitive to the conditions used. The predictive models generated can be used to support both strategic planning and in-flight decision-making by uncrewed aircraft during flight operations while maintaining an acceptable risk of interacting with other traffic.

I. Introduction

THE safety and efficiency of Uncrewed Aircraft (UA)-based cargo operations at scale [1] will be impacted by the presence of Instrument Flight Rules (IFR) and Visual Flight Rules (VFR) traffic in the airspace. Since VFR operators are not required to file a flight plan and often do not receive Air Traffic Control (ATC) separation services, VFR trajectories are harder to predict than IFR trajectories. How does VFR traffic behavior compare to IFR traffic behavior and can aggregate traffic behavior be predicted in a way as to inform and support UA flight operations? This paper investigates these questions in terminal airspace using historical VFR and IFR traffic data at a regional airport.

To enable seamless integration in all airspace environments, UA are expected to be operated under today's IFR [2–4]. At regional airports, unlike larger primary airports, the ratio of VFR traffic to IFR traffic is much higher. Since by design VFR traffic operates with minimal interaction with ATC and surveillance, the uncertainty in their location and intent is also much higher compared to IFR traffic. Therefore, it is necessary to characterize this uncertainty and demonstrate the differences in VFR and IFR traffic behavior. Furthermore, receiving separation services amidst this uncertainty could lead to missed ATC alerts and delays in UA maneuvering in response to controller commands, impacting flight safety [5]. Hence, it is also necessary to predict VFR traffic behavior based on the uncertainty characterization and use it for decision-making during UA flight operations to improve safety and performance.

The interaction of IFR traffic and VFR traffic has been structurally limited in aviation. Most VFR operations happen around non-primary airports or in Class G airspace entailing a natural segregation from IFR traffic while most of the operational inefficiencies occur at major primary airports where big jets operate. Hence, historically, research has focused on the operation of IFR traffic under strict ATC supervision to alleviate inefficiencies at major primary airports. On the other hand, pilots operating under VFR, flying in Visual Meteorological Conditions (VMC), primarily see-and-avoid other traffic by themselves without necessary ATC supervision. As a result, there is a lack of detailed characterization of VFR traffic in aviation literature.

*Aerospace Research Scientist, Crown Consulting Inc., AIAA Member

†Aerospace Research Engineer, Aviation Systems Division, AIAA Associate Fellow

‡Aerospace Engineer, Metis Technology Solutions

A few studies utilized VFR data to simulate conflicts with unmanned aircraft operations [6] or build encounter models [7]. Gariel et al. [8] analyzed six months of General Aviation (GA) traffic in the San Francisco bay terminal area and concluded that the traffic behavior is highly unpredictable. It would cause unacceptably high nuisance alarm rates on existing collision avoidance systems. However, they use this insight only as motivation for developing a novel collision alerting algorithm. UA-based cargo operations and other emerging operations, such as Advanced Air Mobility (AAM), are expected to grow at non-major airports, thereby increasing the frequency of interaction with VFR traffic. Characterizing the location and severity of such interactions could enable more informed flight planning and operations. Hence, characterizing VFR traffic and related uncertainty is timely and pertinent. Therefore, this continued research fills an important gap in the lack of detailed VFR traffic characterization and VFR traffic prediction in aviation literature.

To address the described research needs, the authors developed a method to characterize spatio-temporal distribution of VFR traffic and study how the runway operational capacity for UA is affected by different levels of VFR traffic uncertainty, based on empirical probabilities of traffic distribution near a regional airport [9]. The method was applied to Fort Worth Alliance (KAFW) as a representative regional airport. Heat maps, showing the distribution of traffic in a region of airspace, were developed as a means to characterize the uncertainty associated with VFR traffic density at different altitudes. VFR occupancy maps were derived from these showing the probability of interaction with VFR traffic in an airspace region. These provided further insight into runway occupancy by VFR traffic at the regional airport which was then used to derive the runway operational capacity as a function of VFR traffic interaction probability.

This paper serves as the next step in this research. To demonstrate the comparative variability of VFR traffic, first, the prior proposed method is applied to IFR traffic and used to compare VFR traffic and IFR traffic behavior. The empirical probabilities of traffic interaction are then used to generate predictive occupancy maps that predict the expected distribution of traffic conditioned on operational needs, for example detecting VFR aircraft at specific locations in flight. The predictive occupancy maps are validated using empirical data different from the data used to generate them. Finally, how such predictions can be utilized to support UA flight operations is discussed.

The paper is structured as follows. Section II describes the method to characterize and predict traffic behavior and contextualizes it with KAFW as an example of a regional airport for such analysis. Analysis results specific to the study region are presented in Section III. The difference between these two sections is that the main focus of Section II is on *the method*, while the main focus of Section III is on the detailed *application of the method*. Results are further divided into three parts. The first part shows the application to IFR traffic for comparison with VFR traffic and analysis using several metrics to demonstrate the difference between VFR traffic and IFR traffic distributions. The second part describes the use of occupancy maps for traffic prediction conditioned on observed traffic. Validation of these predictions is also discussed. The third part describes a simple use of predictive occupancy maps in UA flight operations.

II. Methodology

The first step is to characterize spatio-temporal density of traffic around a given regional airport. Owing to extensive surveillance and monitoring by multiple sources, historical IFR traffic data is quite abundant. Comparatively, historical VFR traffic data is derived primarily from VFR traffic squawking on transponder code 1200 and surface radar detection near an airport. Hence, the data is limited and often not available for every VFR flight. Since the proportion of VFR traffic at a regional airport is relatively higher, to accommodate UA and other novel operations, it is necessary to characterize VFR traffic behavior even with such limited data.

Extending prior work on this [9], the following sections describe the study region and the method to characterize uncertainty associated with VFR traffic density in space and time. This order of description is expected to make it easier for the reader to understand the methodology in the context of a regional airport.

A. Study Region

As a strong candidate for future UA-based cargo operations, KAFW is chosen to be the representative regional airport. Historical traffic data for flights at KAFW was gathered for September and October 2022. The data starts on August 31, 2022, and ends on Nov 1, 2022. The period was chosen because it had the highest monthly operations in the year in that region. Historical track data from the first month was used as test data to empirically model the uncertainty, and the second subsequent month of data was used for the validation of predictions made based on empirical uncertainty.

Fig. 1 shows the selected region of interest (RoI) around KAFW. Since the goal is to characterize traffic around a regional airport, a 100 x 100-mile region is selected (shown with a white box in the left figure), centered on the regional airport. All miles mentioned in this paper are statute miles unless otherwise specified. The right figure shows a

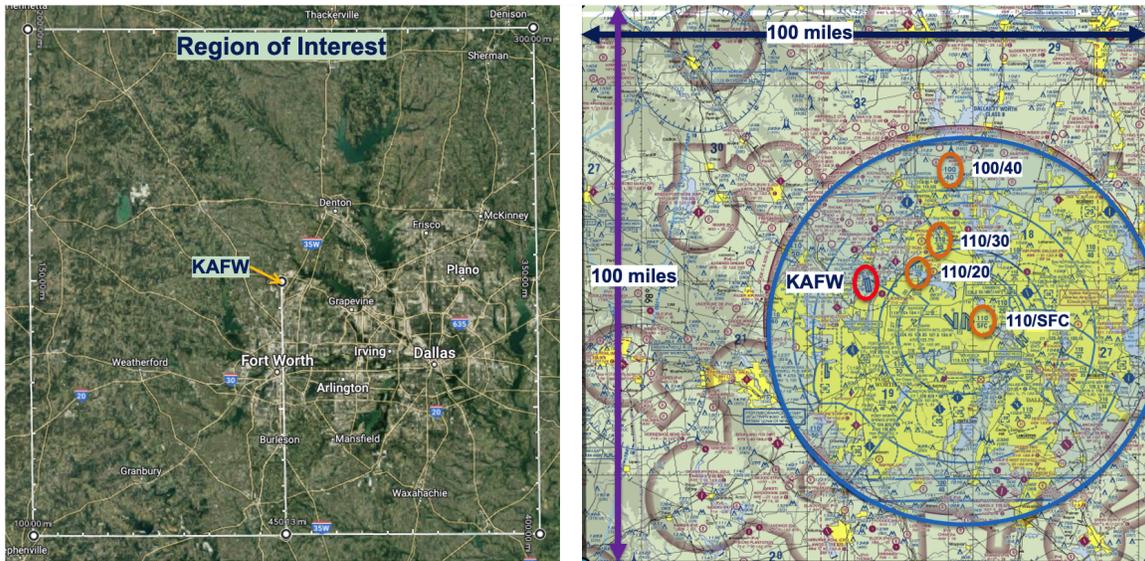


Fig. 1 Left: The region of interest (RoI) selected around KAFW demarcated by the white box (Background map credit: Imagery ©2023 Landsat/Copernicus, Imagery ©2023 TerraMetrics, Map data ©2023 Google). Right: RoI on a VFR sectional chart [10] showing controlled airspace. DFW class B airspace ceilings are highlighted.

VFR sectional chart for the same RoI. VFR operations primarily use Class E/G airspace. Operating in the vicinity of a Class B, C or D airport impacts the altitudes and patterns of operation. Controlled airspace complexities are therefore important to understand to give context to VFR traffic behavior. As an example, the right figure shows the external boundaries of the Dallas-Fort Worth (KDFW) airport’s Class B airspace with a translucent blue circle. The inverted wedding cake structure of this Class B airspace starts at different altitudes depending on the relative location from KDFW. For example, on top of KAFW, it starts at 4,000 feet above Mean Sea Level (ft MSL). At the north-eastern edge of the KAFW runways, it starts at 3,000 ft MSL. Roughly a couple miles southeast of KAFW it starts at 2,000 ft MSL and about six miles southeast, it starts from the surface. KAFW itself is located at an elevation of 722 ft MSL.

All geographically oriented figures in this paper have North pointing upwards and East pointing to the right.

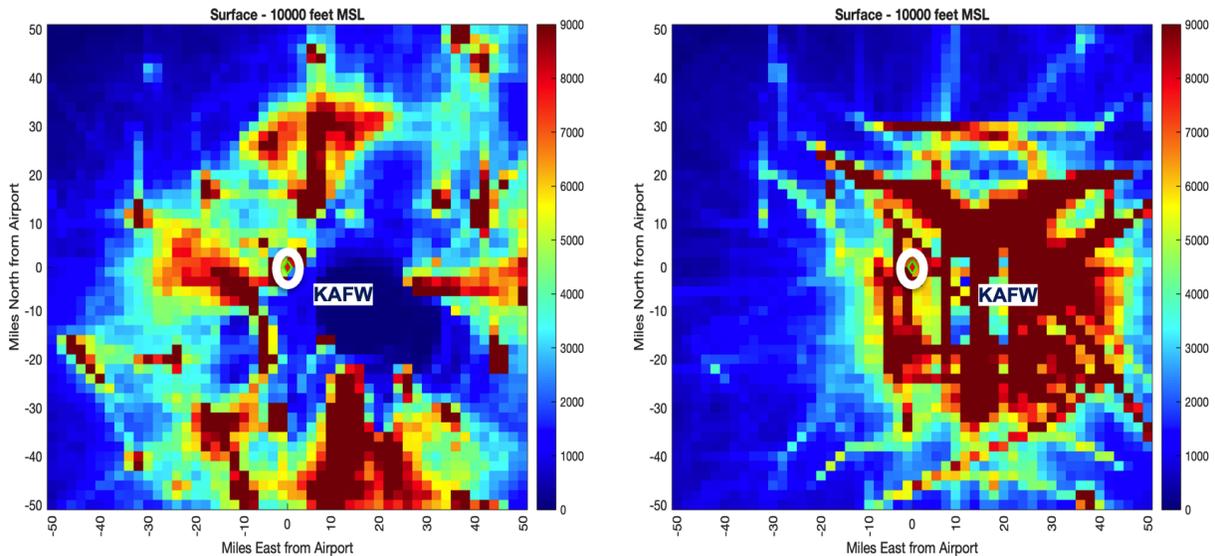


Fig. 2 Spatial density heat maps. Left: VFR Traffic. Right: IFR Traffic.

B. Spatial Density Heat Maps

The RoI is divided into square grid cells of 2-mile edge length. VFR aircraft are usually responsible for visual separation with other VFR aircraft and need to maintain a lateral separation of 1.5 nautical miles (nmi) or a vertical separation of 500 feet from IFR aircraft within controlled airspace. The 2-mile grid resolution was therefore chosen after repeated trials for the best traffic characterization without making the traffic in the grid cells excessively sparse. This same resolution is used for the rest of the analysis in the paper for both VFR and IFR traffic.

VFR traffic is restricted to fly under 18,000 ft MSL. Ninety-nine percent of VFR traffic in the RoI for the month of data used was under 10,000 ft MSL. The detailed distribution of VFR traffic and IFR traffic by altitude is discussed in more detail under Section III. Since the primary focus is on VFR traffic, the density of track points was plotted over an altitude band starting from KAFW surface to 10,000 ft MSL to generate the spatial density heat maps of VFR traffic and IFR traffic in the RoI around KAFW in Fig. 2. The spatial density depicted is the number of track points recorded within each cell volume for the month of data used.

The heat map in Fig. 2 shows regions of high track point density as redder and regions of low track point density as bluer. KAFW, in this and all further maps, is shown as a red diamond with a green outline in the center of the maps. The circular low-density blue region to the southeast of KAFW in the VFR traffic heat map on the left, closely corresponds to the Class B airspace that starts from the surface. The blue region implies that most VFR traffic in the region stays outside this airspace. On the contrary, most IFR traffic is concentrated within this airspace as seen in the IFR traffic heat map on the right. The redder regions on the IFR traffic heat map also correspond to the IFR approach and departure routes in the region. Further, the redder regions in the heat map on the left correspond to the Class D airports and Class E airspace starting at 700 ft Above Ground Level (AGL) (shown as thick and fuzzy magenta ring-shaped features on the VFR sectional chart [10]). UA flying through redder regions would be more likely to interact with traffic. To study the effect of the spatial distribution of the track data in more detail, the heatmap can be further partitioned vertically into altitude bands, as was demonstrated for VFR traffic in prior work [9].

C. Spatio-temporal Occupancy Maps

Using the spatial distribution of track data, the next step is to analyze the variation of traffic in the region over time. Times with higher concentrations of aircraft track points would entail a higher probability of interaction. To measure the high traffic concentration times, a distribution of track point data in time is generated over each grid cell volume. Since each track point is also associated with an aircraft identification, the probability of interaction with *at least one* aircraft over a given time-period is calculated.

$$P (>=1 \text{ Aircraft}) = \frac{\text{Number of time steps with at least 1 aircraft in the cell}}{\text{Total number of time steps}} \quad (1)$$

Plotted over all grid cells in the RoI this generates the spatio-temporal *Occupancy Maps* of the region at different altitude bands. These are also presented under results for altitude bands up to 7,000 ft MSL. The spatial aspect is captured by the location of the grid cell and the temporal aspect is captured by the probability calculation. Since each grid cell holds the temporal distribution of traffic in that cell's region, the probability of interaction with two or more aircraft can also be derived by this method. It is however noted that since the cell size is 2x2 square miles, the interaction with two or more aircraft is unlikely as most of the time there would be only one aircraft in such a cell size.

D. Predictive Occupancy Maps

The occupancy maps are an estimate of empirical probability derived from an aggregation of the temporal distribution of historic traffic in each cell, while the location of the cell captures the spatial distribution. In other words, the database stores a complete snapshot of traffic distribution in the airspace at every instant of time in the past covered by the underlying data. Predictive occupancy maps are occupancy maps that are regenerated with spatio-temporal conditions applied to traffic.

This can support UA flight operations in several ways. For strategic decision-making, suppose a UA flight needs to be planned for a weekend morning. Predictive VFR traffic occupancy maps could be generated conditioned on the time being a weekend and between 6 am to noon. These can then be used to plan a route avoiding areas of predicted VFR traffic interaction risk above a certain threshold considered acceptable by the operator. In a more tactical situation, suppose a UA detects a couple of VFR aircraft ahead of it at specific locations, again predictive occupancy maps could be generated on the fly by aggregating airspace snapshots conditioned on the presence of an aircraft in a cell (and its vicinity) at the locations detected. The UA can then tactically avoid regions of higher risk of VFR traffic interaction

on the fly. Predictive occupancy maps can also be used for contingency management. Suppose there is a lost link of communication and a UA needs to land, predictive occupancy maps based on spatial and temporal conditioning as described above, could provide the least risk alternatives to the nearest landing sites. Figures of predictive occupancy maps, validation of predictive occupancy maps, and a discussion of their sample applications are presented under results.

The authors note that prior work also developed a method to estimate runway occupancy and consequently UA runway operational capacity based on the occupancy maps approach. It is also extensible to IFR traffic but is not included in this paper for brevity.

III. Results

A. Comparison of VFR and IFR Traffic

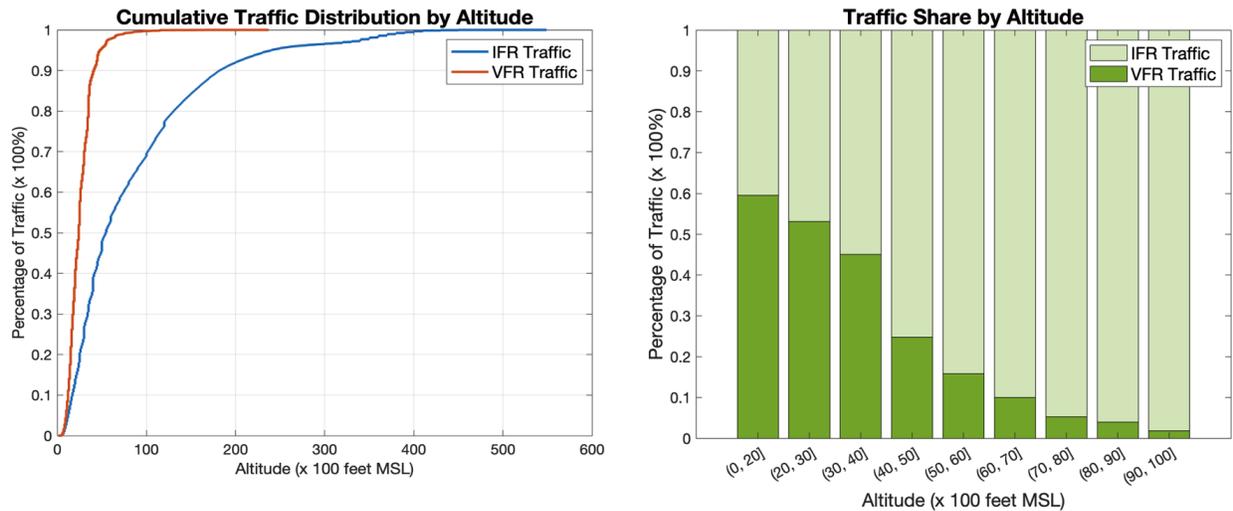


Fig. 3 Traffic distribution by altitude. Left: Cumulative distribution. Right: Traffic share

All VFR traffic and IFR traffic within the RoI was analyzed for distribution in space. Fig. 3 shows the vertical distribution analyzed in two ways. First, the figure on the left shows the cumulative distribution of each kind of traffic. VFR traffic was under 10,000 ft MSL almost all of the time while IFR traffic remained under 10,000 ft MSL only 70 percent of the time. Second, the figure on the right shows the share of track points in different 1,000-foot-wide altitude bands. The lowest band is wider and selected up to 2,000 ft MSL to account for surface elevation variation in the region. It is noteworthy that under 4,000 ft MSL, a major portion of the traffic flew under VFR varying between 45 to 60 percent.

It is noted that VFR traffic was under 7,400 ft MSL, 99 percent of the time while IFR traffic remained under the same altitude only 60 percent of the time. Since the primary focus of this work is VFR traffic, for further spatio-temporal analysis, altitude bands are studied only till 7,000 ft MSL.

Spatial density heat maps were generated for the KAFW RoI at different altitude bands. Since KAFW is located at an elevation of 722 ft MSL, the first band chosen was up to 2,000 ft MSL to coincide with the lowest Class B ceiling in the vicinity of KAFW. Above that, bands were chosen to be 1,000 ft wide to correspond with each shelf of the Class B airspace. These maps capture the whole month of the data used for analysis. Using the heatmaps, occupancy maps were generated at the same altitude bands to measure the probability of interaction with one or more aircraft. These are shown in Fig. 4 and Fig. 5 for VFR traffic and IFR traffic respectively. These quantify the temporal uncertainty of interacting with traffic in regions of high traffic density. The brighter regions imply a higher probability of interaction with traffic in that region. For example, with VFR traffic between the 3,000-4,000 ft MSL band, the highest interaction probability observed was 0.02 compared to around 0.25 under 2,000 ft MSL. That entails an order of growth in the risk of interaction. Both figures are normalized to a maximum interaction probability of 0.022 for comparison. This means all higher probabilities are shown in the figure as bright yellow. Using a higher value for normalizing the maps at different altitudes would make higher altitude occupancy maps much darker thereby limiting the observation of the patterns in the figure, for the reader. In contrast to VFR traffic, the maximum observed probability of interaction was

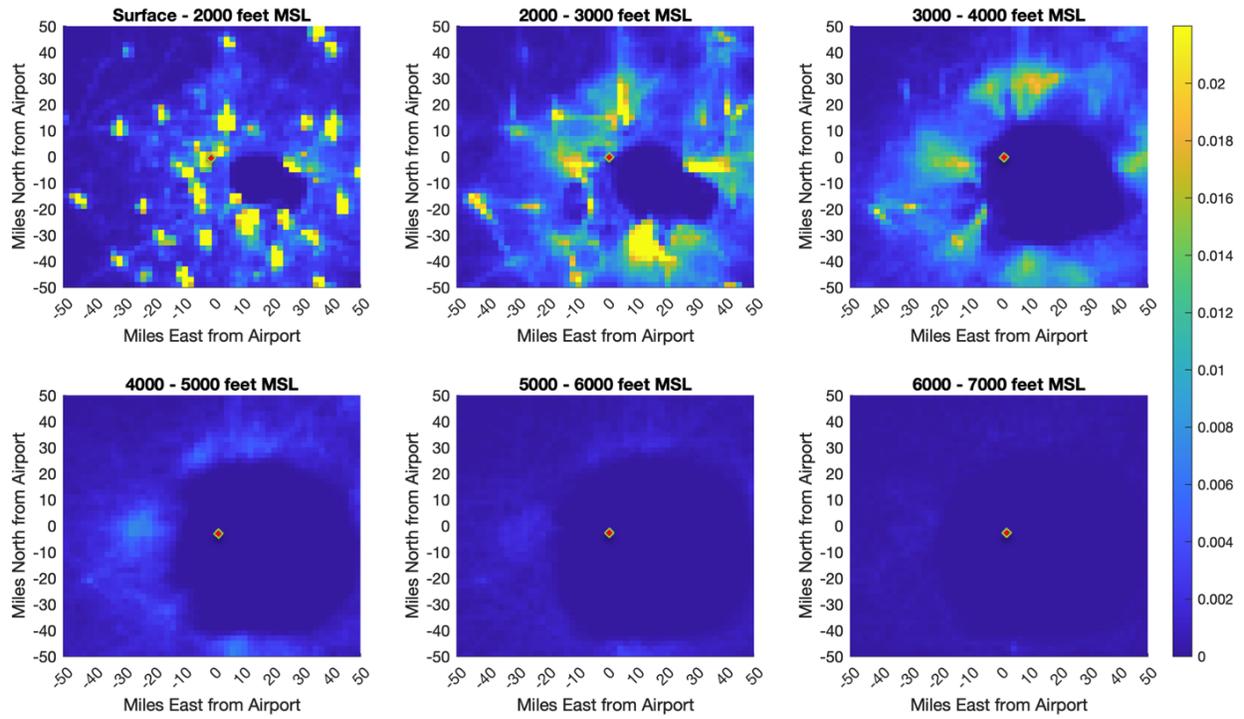


Fig. 4 Spatio-temporal occupancy maps for VFR traffic. Color indicates the probability of interaction with at least 1 aircraft.

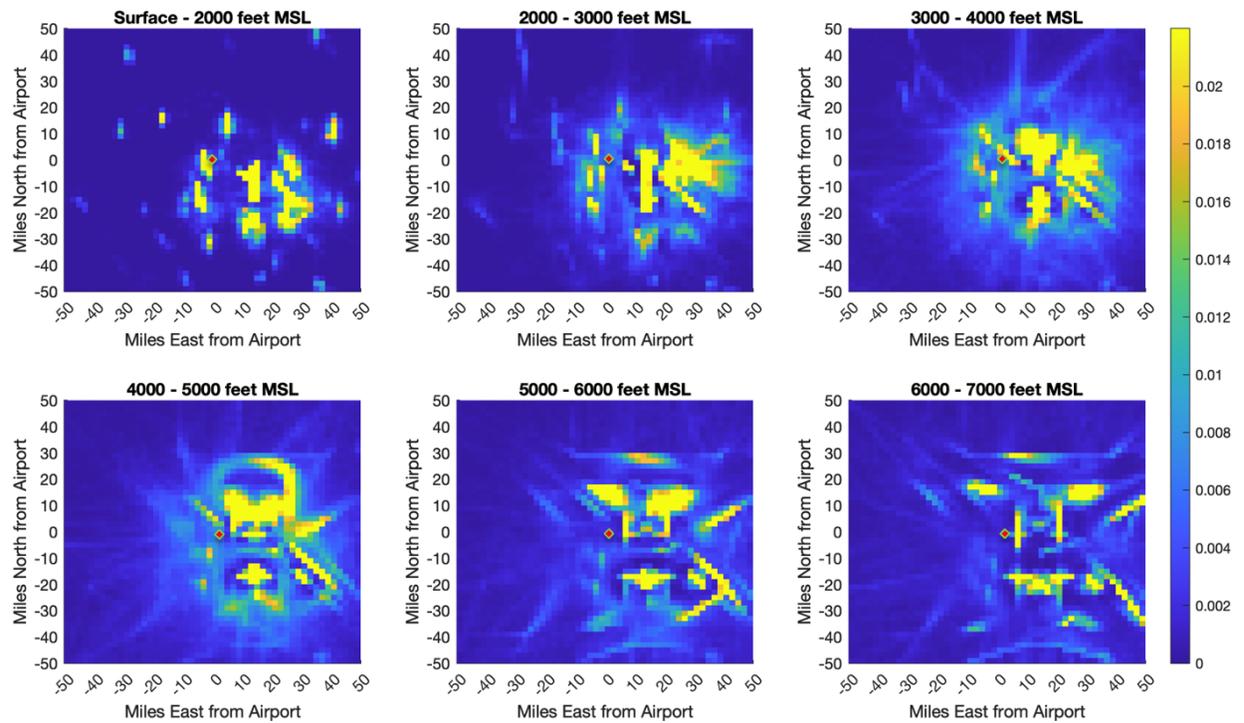


Fig. 5 Spatio-temporal occupancy maps for IFR traffic. Color indicates the probability of interaction with at least 1 aircraft.

close to 1 for IFR traffic. It is noteworthy that the dark blue hole in Fig. 4 (very low to no VFR traffic) to the east of KAFW closely corresponds to the Dallas-Fort-Worth (KDFW) Class B airspace and grows in size with altitude.

These maps bring out several insights into traffic behavior in the region and the differences between VFR traffic and IFR traffic. The brightest regions of VFR traffic coincide with the locations of the regional airports in the area such as Denton Enterprise (KDTO) to the north of KAFW and Fort Worth Meacham (KFTW) to the south. Hence, VFR traffic in the region is concentrated at or around these airports. On the contrary, the brightest regions of IFR traffic are concentrated inside the KDFW Class B airspace and other airports in the vicinity under the Class B airspace that receive IFR traffic. Next, the majority of VFR traffic was present between the surface to 4,000 ft MSL altitudes and spread across a wider portion of the RoI as shown by several dispersed swaths of yellow and green. This distribution is likely caused by the traffic staying out of the inverted wedding cake structure of the KDFW Class B airspace and flying primarily over low population density areas in the region. Meanwhile, the IFR traffic is primarily within the Class B airspace at these altitudes and above 4,000 ft MSL, concentrated within the arrival and departure routes in the region. Above 3,000 ft MSL, there are channeled gaps between VFR traffic-heavy regions that could be potentially utilized by UA and other novel entrants for flight planning and flow management decisions.

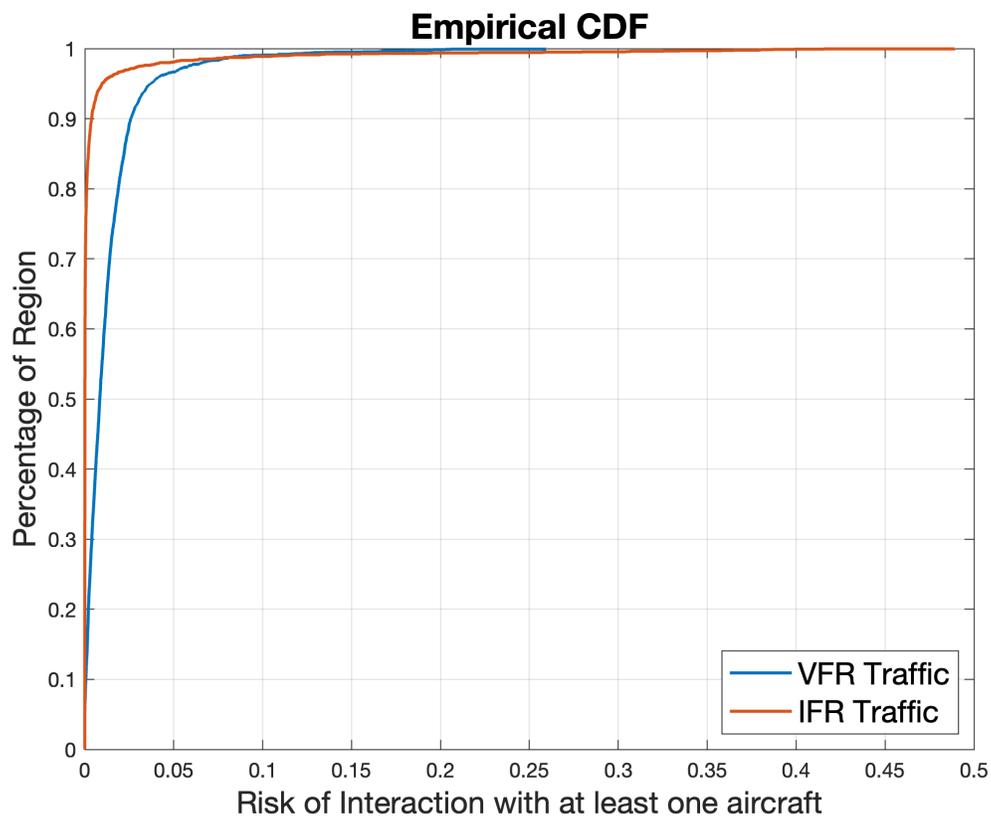


Fig. 6 Cumulative distribution of the percentage of the region as a function of the probability of interaction with at least one aircraft.

In general, VFR traffic was more dispersed and had less structure compared to IFR traffic. To demonstrate this, some sample metrics were computed using the surface-2,000 ft MSL occupancy maps. Fig. 6 shows the percentage of the region as a function of the probability of interaction with at least one aircraft. In that airspace, only 37 percent of the cells had a very low (under 0.001) VFR traffic interaction probability while almost 80 percent of the cells had a very low IFR traffic interaction probability. In the same airspace, the top 2 percent riskiest cells to traverse had an interaction probability of over 7.43 percent for VFR traffic and over 30 percent for IFR traffic. This implies VFR traffic-heavy regions are more likely to be encountered in this airspace but the risk of such encounters is lower, whereas IFR traffic-heavy regions are much more concentrated and the risk of interaction for encounters with IFR traffic in these

regions is very high.

B. Predictive Occupancy Maps

The prior set of results shows that the type of traffic to be avoided (VFR vs IFR) and the risk of interaction for each kind of traffic determines what airspace is available for UA traffic to traverse safely. The occupancy maps can therefore be used to make aggregate traffic behavior predictions, based on empirical probability, conditioned on flight operations requirements and observations. For example, the traffic behavior based on day of week and time of day variation was described by the authors in prior research [9]. Such predictions are useful for strategic flight planning. This paper explores more tactical, shorter time horizon predictions based on different kinds of conditions in flight.

For example, suppose a UA enters the airspace at 10 am. What is the probability of interaction with VFR traffic in the region after the next T minutes? To compute this, the empirically derived VFR traffic occupancy data for the region is conditioned to extract the temporal snapshots of the airspace " T minutes after 10 am". These temporal snapshots are then aggregated to derive Predictive Occupancy Maps showing the probability of interaction after the next T minutes. The UA operator chooses the prediction horizon of T minutes.

Another way of conditioning could be based on the partially observed state of the airspace, such as a few aircraft detected ahead while in flight. For example, suppose a UA flying in the region of interest, detects two VFR flights ahead at different positions (latitude, longitude and altitude). What is the probability of interaction with VFR traffic in the region after the next T minutes? To compute this, again the temporal snapshots of the airspace are extracted " T minutes after the time when aircraft were present in the vicinity of the detected locations in the empirical data". As before, these temporal snapshots are again aggregated to derive Predictive Occupancy Maps showing the probability of interaction after the next T minutes. The UA operator chooses both the size of the vicinity to be considered for the conditioning based on detected aircraft locations and the duration of T minutes.

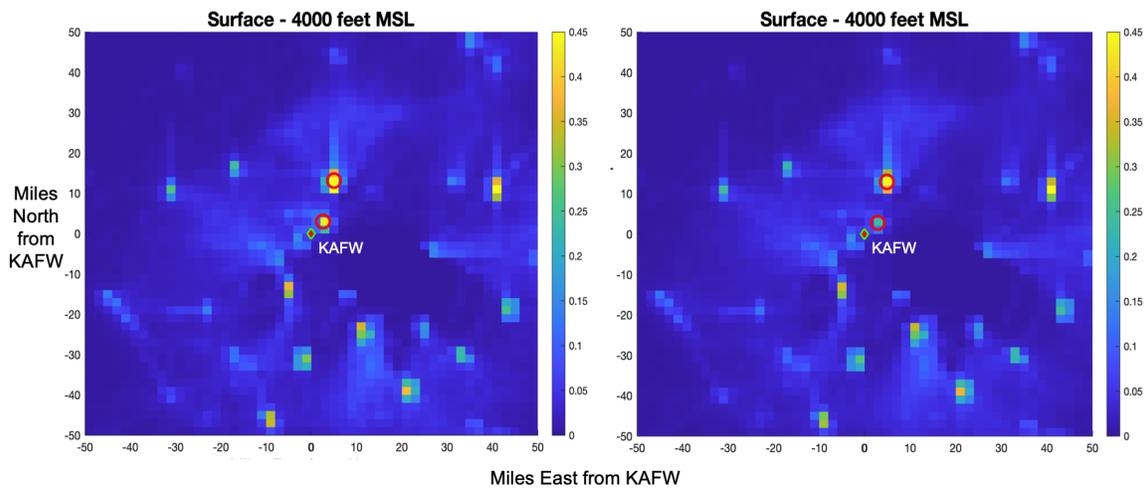


Fig. 7 Predictive occupancy maps. Left: After 1 min, Right: After 10 min. Color indicates the probability of interaction with one or more aircraft. Cells, where VFR aircraft were detected by UA, are shown in red circles.

A sample demonstration of the method based on detected aircraft is shown in Fig. 7. Suppose a UA entering the RoI at 3,800 ft MSL detects two VFR aircraft ahead. The cells where the VFR aircraft are detected, are marked by red circles. One VFR aircraft is detected at 1,500 ft MSL and another at 2,500 ft MSL. The operator is interested in the prediction of VFR traffic T minutes later and chooses to include the adjacent cells in the vicinity of the detected traffic. The figure on the left is a prediction after one minute, while the figure on the right is a prediction after ten minutes. As the time horizon is extended, the detected traffic is expected to move further away and therefore the probability of encountering it in the future at that location is expected to decrease. This is more noticeable in the cell close to KAFW as it becomes greener in the right figure.

To validate time-of-day-based predictions, predictive occupancy maps generated with a month of training data (September 2022) were compared against predictive occupancy maps generated under the same spatio-temporal conditions for the following month of test data (October 2022). This was done for both VFR traffic and IFR traffic at prediction time horizons varying between one and sixty minutes. Fig. 8 shows the conformance between training

and test data for VFR traffic (solid lines) and IFR traffic (dashed lines) when prediction is based on time of day. Only daylight hours with high traffic in the region are considered (8 am to 6 pm). A very short prediction time horizon would underfit to empirical data and exclude the effect of traffic propagation while a very long time horizon would overfit to traffic propagation in empirical data. To compare the predictions between training and test data, conformance is measured as the percentage of cells, with a difference in the predicted probability of interaction lower than one percent, between predictions based on training data and maps generated from the test data.

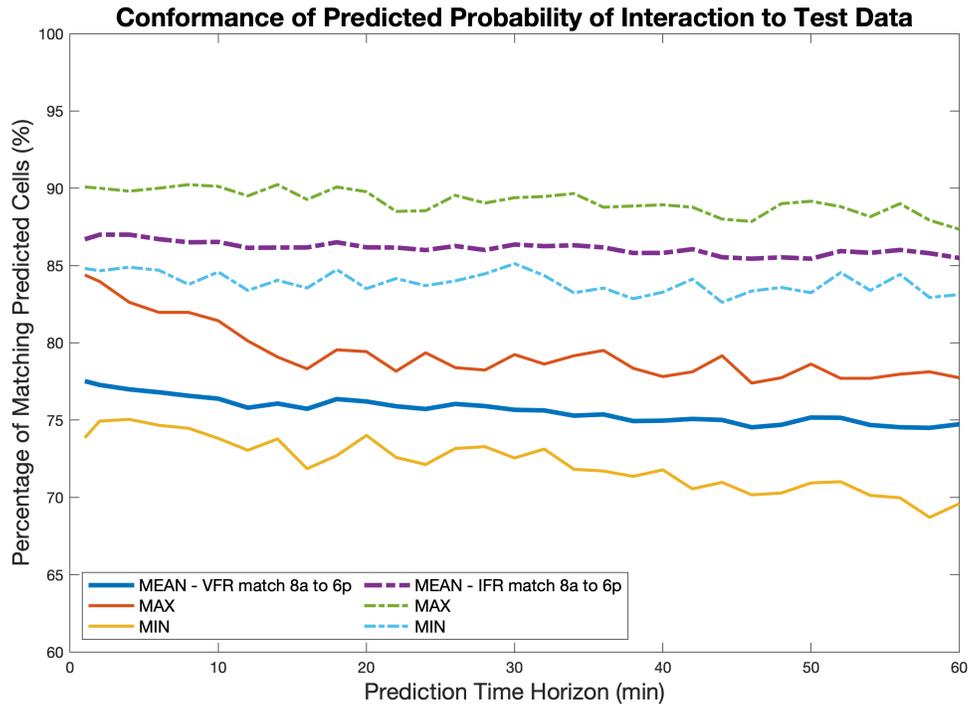


Fig. 8 Percentage of Region of Interest with the difference in the predicted probability of interaction between training and test data under 0.01.

The results show that there is 75 to 85 percent conformance for VFR predictions compared to 85 to 90 percent conformance for IFR predictions at best. Furthermore, the conformance decreases as the prediction time horizon increases. This trend is more pronounced for VFR traffic compared to IFR traffic. This is caused by the fact that IFR traffic is primarily concentrated in the approach and departure routes so predictions are more accurate. VFR traffic is more dispersed and therefore the quality of predictions decreases as the prediction time horizon increases.

A similar exercise is repeated to evaluate predictions based on aircraft detected in the airspace. The combinations of partially observed aircraft in the airspace can be infinite. Three combinations were chosen, each with two detected aircraft. The conformance between training and test data for these is shown in Fig. 9 with solid lines for VFR traffic and dashed lines for IFR traffic. Once again, in the chosen examples IFR traffic was more predictable than VFR traffic. There was 75 to 85 percent conformance for VFR predictions compared to 90 to 95 percent conformance for IFR predictions. It is noted that this trend is anecdotal. Since possible combinations of detected aircraft are infinite, it cannot be ascertained that no combination exists where VFR traffic is more predictable than IFR traffic. However, in practice, when the available information in the airspace is about some of the aircraft flying in the vicinity, this method of prediction would still be useful for predicting aggregate behavior. For the chosen sample conditions, the performance of the method was as shown in Fig. 9 but may vary for other conditions in practice. The quality of predictions is highly dependent on the observed state of the airspace used as conditions to predict based on the empirical data.

The general higher predictability of IFR traffic compared to VFR traffic is due to the fact that IFR traffic at these altitudes is highly structured and within the approach and departure routes. VFR traffic is highly dispersed and variable and therefore the predictability also highly varies.

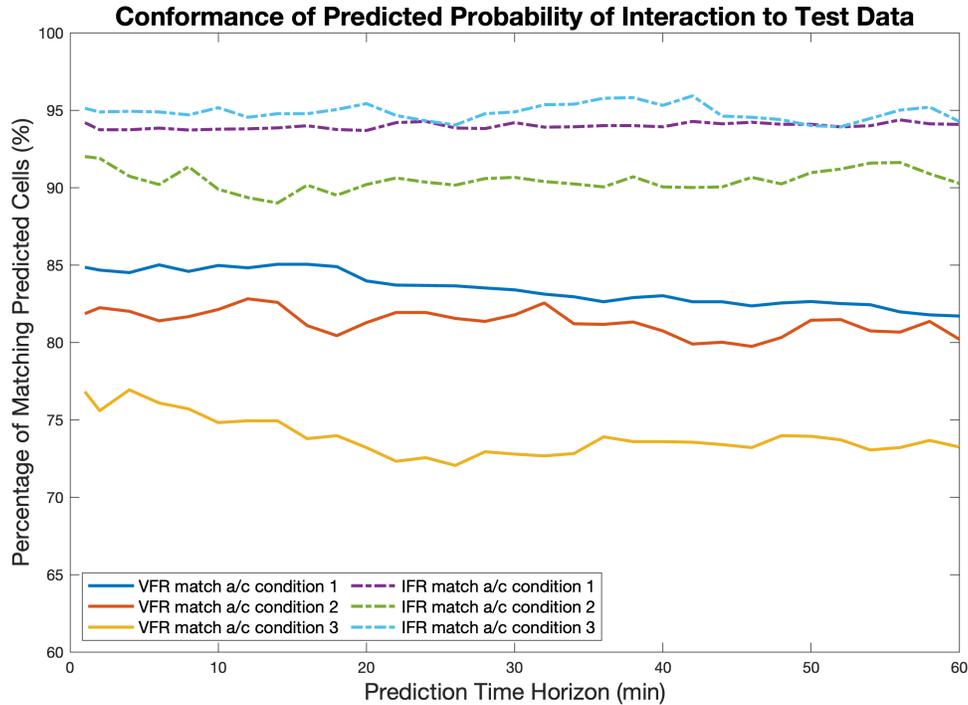


Fig. 9 Percentage of Region of Interest with the difference in the predicted probability of interaction between training and test data under 0.01.

Such predictive maps can be generated for detected IFR traffic also. These can be combined and/or compared with predictive maps for VFR traffic. The prediction method can also be used to set different vicinity and time horizon requirements for the two kinds of traffic. The underlying occupancy data can be continuously updated with the most recent VFR traffic data in the region, thereby making predictions more pertinent to the most recent state of the airspace. The authors also note that there are induced effects from monthly variations between test and training data. Furthermore, this method is limited by observed behavior in empirical data, and new or anomalous flight behavior would not be captured by this method.

Trajectory-based prediction is another commonly used method where aircraft physics is used to predict the expected behavior of the individual aircraft. Such prediction is useful for onboard conflict detection and resolution based on single (or few) aircraft trajectories. On the other hand, these predictive occupancy maps capture more of the aggregate behavior of similar traffic in the airspace and less so of the individual detected aircraft. They can inform longer time horizon decisions like contingency management which is described as a sample application next.

C. Sample Application - Contingency Management

Suppose a UA heading towards KAFW enters the RoI from the North, suffers a lost communication link or another kind of emergency, and needs to land within the next ten minutes. Following the example above, right before the emergency occurred, it detected the two VFR aircraft as described earlier. What are the available route alternatives if the UA operator can tolerate a VFR traffic interaction risk of up to 10 percent? To derive this we use the 10-minute predictive occupancy map from Fig. 7 and rescale it to show all cells with an interaction probability over 0.1 as bright yellow.

Fig. 10 shows the rescaled 10-minute prediction. The thick magenta circle is the location of the UA that needs to land. Based on the predictive occupancy map, three route alternatives are computed. Route 1 (R1), takes the UA to its intended destination KAFW in minimum time. Route 2 (R2) takes the UA to its first alternative landing airport Decatur Muni (KLUD) with minimal risk of VFR traffic interaction. Route 3 (R3) takes the UA to its second alternative landing airport Aero Country (T31) with minimal risk of VFR traffic interaction. The red dashed line shows the originally

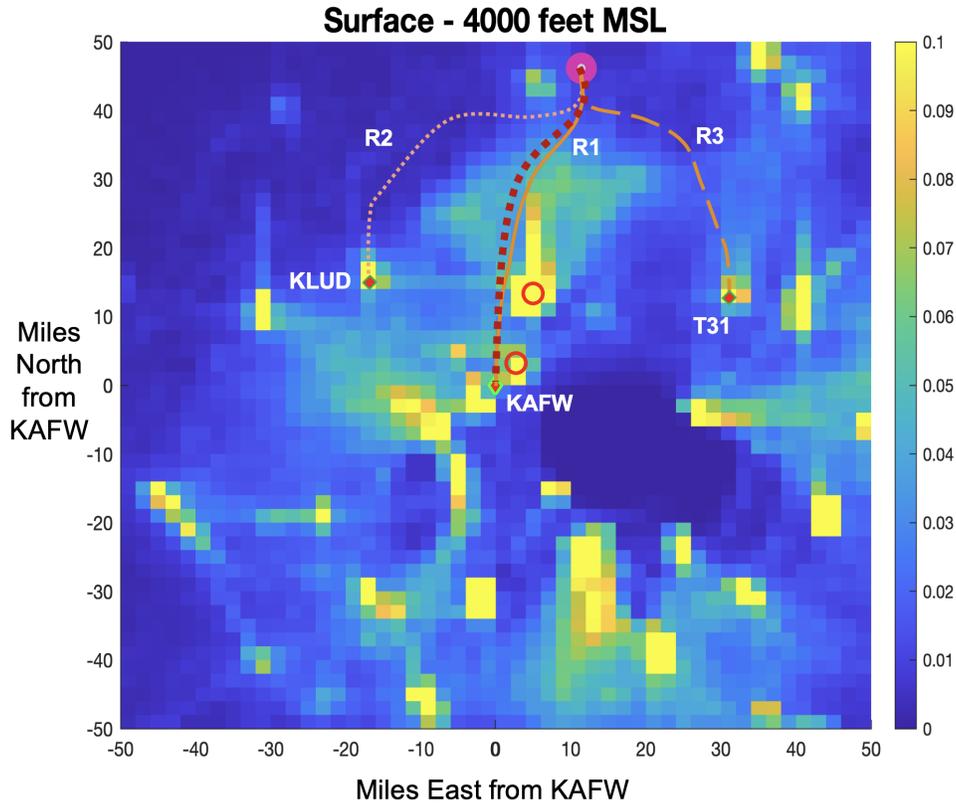


Fig. 10 Contingency routes for a UA avoiding VFR traffic with interaction risk above 0.1. R1: Minimum deviation route to destination airport KAFW. R2: Lowest risk route to the alternate airport KLUD. R3: Lowest risk route to the alternate airport T31. The nominal route based on arrival procedures at KAFW is shown as a red dashed route

intended trajectory of the UA following an arrival route into the airport for comparison.

IV. Conclusion

This paper extended a previously developed method to characterize the uncertainty associated with VFR traffic spatial and temporal density in the vicinity of a regional airport. The method was extended to IFR traffic and then used to compare the behavior of VFR traffic with IFR traffic. Historical traffic data was used to generate spatial density heatmaps around a regional airport at different altitude bands. These were used to develop spatio-temporal occupancy maps that capture the probability of interacting with VFR traffic or IFR traffic at a given place and time.

In the vicinity of KAFW regional airport, used as the study area, the comparison of both kinds of traffic showed several important insights. First, under 4,000 ft MSL, up to 60 percent of the traffic flew under VFR. Second, IFR traffic was more concentrated near Class B and C airports and tended to be structured within the approach and departure routes, whereas VFR traffic was more dispersed, stayed out of Class B airspace, and was closer to Class C and D airspace and low population density areas. Third, VFR traffic was more dispersed and had less structure compared to IFR traffic. In the airspace under 2,000 ft MSL, only 37 percent of the cells had a very low (under 0.001) VFR traffic interaction probability while almost 80 percent of the cells had a very low IFR traffic interaction probability. In the same airspace, the top 2 percent riskiest cells to traverse had an interaction probability of over 7.43 percent for VFR traffic, and over 30 percent for IFR traffic.

Next, conditioned on traffic states observed in the airspace, predictive occupancy maps were generated, to predict aggregate traffic behavior based on empirical probability. Predictions based on two types of conditions were demonstrated - time of day, and aircraft detected in the airspace. The predictive occupancy maps were validated by comparing predictions between the training month of data and the test month of data. When conditioned on the time of day,

predictions were accurate between 75 to 85 percent for VFR traffic and accuracy decreased with increasing prediction time horizon. For IFR traffic, predictions were accurate between 85 to 90 percent and showed a slight decrease with increasing prediction time horizon. The method was also shown to be useful in predicting aggregate traffic behavior based on few detected aircraft in the airspace. IFR traffic predictions conformed between 90 to 95 percent and VFR traffic predictions conformed between 75 to 85 percent between training and test data. However, these observations were only anecdotal in demonstrating the use of the method as the potential conditions to be tested would be infinite.

Finally, a sample application of such maps to contingency management path planning was discussed. Occupancy maps and predictive occupancy maps can be used to support both strategic planning and tactical decision-making by uncrewed aircraft during flight operations as shown by the sample application to contingency management. However, the developed prediction method is limited by the behavior observed in empirical data and its validation is limited by monthly and seasonal variations between training and test data.

The authors are currently developing a machine learning-based method to predict VFR traffic behavior which can eventually be compared with this empirical probability-based method and provide insights into what kinds of scenarios each method is more useful for. This continued research fills an important gap in the lack of detailed VFR traffic characterization and VFR traffic prediction in aviation literature. Future work will use this research for UA traffic scenario generation and derivation of capacity and other airspace metrics associated with safety and performance. Eventually, these methods can be released as a toolkit that would be equally useful for autonomous cargo and the broader AAM industry.

Acknowledgments

This work is funded by the Pathfinding for Airspace with Autonomous Vehicles sub-project at National Aeronautics and Space Administration Ames Research Center. The work was supported under Contract Number NNA16BD14C, managed by the Universities Space Research Association (USRA). The study was done in part by Crown Consulting, Inc. under a subcontract from USRA.

References

- [1] Hayashi, M., Idris, H., Sakakeeny, J., and Jack, D., "PAAV Concept Document," 2022.
- [2] Gigarjill, M., "Integration of Unmanned Aircraft Systems into the National Airspace System Concept of Operations," *Federal Aviation Administration*, 2012, pp. 1–120.
- [3] SC-228, R. F., *Minimum Operational Performance Standards (MOPS) for Detect and Avoid (DAA) Systems*, RTCA, Incorporated, 2017.
- [4] Johnson, W. C., and Shively, R. J., "UAS Detect and Avoid: Efforts from NASA's UAS Integration into the NAS Project," *2018 Aviation Technology, Integration, and Operations Conference*, 2018, p. 2871.
- [5] Bulusu, V., Chatterji, G. B., Lauderdale, T. A., Sakakeeny, J., and Idris, H. R., "Impact of Latency and Reliability on Separation Assurance with Remotely Piloted Aircraft in Terminal Operations," *AIAA AVIATION 2022 Forum*, 2022, p. 3704.
- [6] Johnson, M., Santiago, C., and Mueller, E., "Characteristics of a well clear definition and alerting criteria for encounters between uas and manned aircraft in class e airspace," *Air Traffic Management (ATM) Research and Development Seminar*, 2015.
- [7] Adaska, J. W., "Computing risk for unmanned aircraft self separation with maneuvering intruders," *2012 IEEE/AIAA 31st Digital Avionics Systems Conference (DASC)*, IEEE, 2012, pp. 8A4–1.
- [8] Gariel, M., Hansman, R. J., and Frazzoli, E., "Impact of GPS and ADS-B reported accuracy on conflict detection performance in dense traffic," *11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, including the AIAA Balloon Systems Conference and 19th AIAA Lighter-Than*, 2011, p. 6893.
- [9] Bulusu, V., Idris, H. R., and Chatterji, G., "Analysis of VFR traffic uncertainty and its impact on uncrewed aircraft operational capacity at regional airports," *AIAA AVIATION 2023 Forum*, 2023, p. 3553.
- [10] "FAA Sectional Chart for Dallas-Fort Worth Area," https://www.faa.gov/air_traffic/flight_info/aeronav/productcatalog/VFRCharts/Sectional/, 2023. Images produced by the U.S. Government and in the public domain.