

Safety Optimization of a Layered Airspace Structure with Supervised Learning

Caranti, L.; Ribeiro, M.J.; Ellerbroek, J.; Hoekstra, J.M.

Publication date
2021

Document Version
Final published version

Published in
11th SESAR Innovation Days

Citation (APA)

Caranti, L., Ribeiro, M. J., Ellerbroek, J., & Hoekstra, J. M. (2021). Safety Optimization of a Layered Airspace Structure with Supervised Learning. In *11th SESAR Innovation Days*

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Safety Optimization of a Layered Airspace Structure with Supervised Learning

Leonardo Caranti, Marta Ribeiro, Joost Ellerbroek and Jacco Hoekstra
Control and Simulation, Faculty of Aerospace Engineering
Delft University of Technology, The Netherlands

Abstract—The capacity of the current system of air traffic is rapidly reaching a limit with the increasing demand for air transportation. Expected future traffic densities not only make automated conflict detection and resolution a necessity, but also force a re-evaluation of coordination elements to decrease conflict rate and severity. It has been acknowledged that airspace structure plays a positive role by acting as a first layer of conflict protection, reducing the likelihood of aircraft meeting and, consequently, the likelihood of conflicts and losses of minimum separation. In the recent past, different airspace structures have been explored. Research shows that the layered airspace concept, where groups of aircraft with similar headings remain separated by cruising at different altitudes, increases airspace capacity. However, implementation of this concept often employs an evenly distributed heading range per vertical layer, which is not optimal for all traffic scenarios, since it may lead to unevenly distributed numbers of aircraft per layer. In this work, we use supervised learning to determine a heading range distribution per layer adapted to the current traffic. This method resulted in a reduction of both conflicts and losses of minimum separation when compared to an evenly distributed layers concept. Results show that conflicts prevention, with a structure which efficiently segments aircraft through the airspace, may have a greater impact on safety than applying conflict resolution methods.

Keywords—Airspace Design, Conflict Detection & Resolution (CD&R), Supervised Learning, Modified Voltage Potential (MVP), BlueSky ATC Simulator

I. INTRODUCTION

The current air traffic control system must evolve in order to sustain the increasing traffic demands of future operations [1]. Existing research, such as the Free Flight project [2], advocates that higher traffic densities can be achieved through a reduction of current traffic flow constraints [3]. Researchers have proposed that transferring traffic separation responsibility to each aircraft (“self-separation”), may enable an increase of the traffic allowed into airspace [4], [5]. This redirects attention towards elements which reduce conflict probability. In particular, the airspace structure which is known to decrease this probability by directly affecting the likelihood of aircraft meeting during their flights.

The Metropolis project explored different types of distributed structures for a (hypothetical) high-density urban airspace [6]. Its most effective airspace structuring concept (the so-called “layers” concept) increases airspace capacity, and reduces conflicts and losses of minimum separation (LoSs), by dividing aircraft over several layers of airspace. This creates different groups of aircraft that remain separated from each other (segmentation effect). Moreover, within each layer,

heading limitations are introduced that enforce a degree of alignment between aircraft, thereby reducing the relative speed between aircraft cruising at the same altitude, which in turn reduces the likelihood of conflicts within a layer of airspace (alignment effect) [7]. However, only evenly distributed heading ranges per layer have been investigated. This is not optimal when headings of the current traffic are not uniformly distributed. In reality, this is often not the case. Take, for instance, a common example from manned aviation: over the Atlantic, eastbound flights mostly occur during night-time, whereas westbound flights mostly operate during daytime. In both instances, there is no uniform heading distribution.

When the airspace structure does not align with the current traffic scenario, aircraft will not be equally divided through the available airspace as expected. In situations where aircraft predominantly adopt a certain heading range, one layer will have higher traffic density than the others. In a worst-case scenario, uneven aircraft distribution per layer will significantly reduce segmentation over the airspace, eliminating the benefit of a layered structure. Herein, we defend that the heading range per layer must be set according to the real expected traffic distribution, guaranteeing that traffic is fully segmented over the available airspace. Safety-wise, this is expected to prevent congestion, collisions, and to reduce travel time. Moreover, an automated control is preferable in order to guarantee fast response times and higher structure variability.

In this work, a machine-learning approach is used to obtain an optimal heading distribution per layer given a non-uniform distribution of aircraft headings. This will be done by resorting to two neural networks, in parallel with the open source, multi-agent ATC simulation tool BlueSky [8]. In this context, the optimal heading distribution per layer is the one resulting in the lowest number of conflicts. The improvement on conflict prevention, obtained by having structures catered to the traffic scenario, will be directly compared to a fixed, evenly distributed heading range structure. In both cases, remaining conflicts will be solved using a tactical conflict resolution method. This work employs the Modified Voltage Potential (MVP) method [2], which has proven to be efficient in reducing the effect of resolution manoeuvres on flight efficiency while improving safety [9].

II. METHOD - SUPERVISED LEARNING

Supervised learning is used to identify the best heading range distribution in a layered airspace structure in function

of the traffic distribution. A system of two neural networks is used to approach this problem. The first neural network, the Conflict Estimator Network (CEN), is trained to be able to estimate the resulting total number of conflicts given two inputs: (1) traffic scenario, (2) heading range distribution per layer of the airspace structure. Validation of the estimated number of conflicts is done by direct comparison with the final conflict count in simulations with the ATC simulation tool Bluesky [8]. The objective of the second neural network, the Layer Heading Network (LHN), is to output a heading range distribution which results in a minimal number of conflicts for the given traffic scenario. The number of conflicts output by CEN is used by the LHN to train towards the optimal heading range distribution. The training process for every episode for the two neural networks is represented in Fig. 1.




Figure 1. Training process for every episode for the Conflict Estimator (CEN) and Layer Heading (LHN) networks.

One of the inputs for both networks, is the expected traffic demand scenario. The heading of all aircraft must be known to the network. For simplification, an array of fixed dimension is used to represent the headings of all aircraft. This way, the dimension of the state array is uniform and not dependent on the number of aircraft. The dimension of the array was set to 12. The total aircraft heading range, 0° to 360° , is divided into 12 bins matching the size of the state array. Each position on the array represents the number of aircraft whose heading is included in the heading range for each bin. For example, the first bin will have the number of aircraft travelling between 0° to 30° . Naturally, a bigger array may be used, allowing for better representation of the heading differences between aircraft. However, increasing the dimension of the state array also increases the number of possible states and state-action combinations. As the size of the problem's solution space grows, so does training time.

The CEN receives the traffic scenario and the heading range distribution per vertical layer as input. During training and validation, the latter is a pre-defined distribution. The

output indicates the expected total number of conflicts for the given traffic scenario when operating in an airspace divided according to the given heading range distribution. This traffic scenario is then run on Bluesky, and the total observed number of conflicts is used to train the CEN. The Mean Squared Error between the two conflict values is used as a loss metric. Once the CEN has been validated, its output is used to train the LHN. This network trains towards finding heading range distributions that result in a minimum amount of conflicts for every traffic scenario.

It is clear that the CEN imitates the behaviour of having a Bluesky simulation during the training of LHN and it may be seen as a duplicate. This was found the best solution as the output of the CEN can be set as the necessary data type to train the LHN. Additionally, this allows for a faster training of the LHN, as it does not require a complete run of the simulation scenario in order to obtain the total number of conflicts. However, naturally this may result in limited accuracy for the LHN if the CEN is not capable of correctly assessing the number of conflicts.

Details on both the CEN and LHN are presented in Table I.. The LHN has a final activation function (a Softmax function) which ensures that the sum of the LHN output is equal to one. When multiplied by 360, each element of the output of the LHN resembles the heading range for that specific layer. For example, if the output vector starts with $[0.128, 0.223]$, the first two layers would have a heading range of 0° to 46.08° and 46.08° to 126.36° . A fixed number of 8 vertical layers is used in this study. Naturally, in a real-world scenario, a different or even a non-fixed number of layers may apply.

TABLE I. NEURAL NETWORK ARCHITECTURE PARAMETERS.

Parameter	CEN	LHN
Input layer size	16	12
Hidden layer 1 size	82	128
Activation function 1	ReLU	ReLU
Hidden layer 2 size	82	128
Activation function 2	Sigmoid	ReLU
Output layer size	1	8
Activation function 3	-	Softmax

III. EXPERIMENT: OPTIMIZED AIRSPACE STRUCTURE WITH SUPERVISED LEARNING

A. Simulation Scenarios

To train and test the two networks, an airspace is defined with two different areas: (1) a measurement area ($10\,404\text{ NM}^2$), and (2) an experiment area ($14\,400\text{ NM}^2$). The first is a square area where aircraft spawn locations (origins) are placed. Ideally, aircraft would only operate within the measurement bound, thereby ensuring a constant density of aircraft within that area. However, aircraft may temporarily leave the area during the resolution of a conflict and should not be deleted in this case. Therefore, a second, larger square area encompassing the measurement area is considered: the experiment area. As a result, aircraft in a conflict situation close to their origin or destination are not deleted incorrectly from the

simulation. Ultimately, an aircraft is deleted once it leaves the experiment area. Note that we assume a no-boundary setting, with sufficient flight space around the measurement area, in order to avoid edge effects from influencing the results.

Aircraft fly a straight line towards their destination. The heading of each aircraft is computed in a random, non-uniform way. At every episode, a random number of normal heading distributions (between 1 and 4) is picked. The mean and standard deviation of each heading distribution are randomly picked between 0° to 360° and 18° to 72° , respectively. For every aircraft, its heading is sampled randomly from one of the heading distributions. Even though aircraft fly straight trajectories through the measurement area, still several waypoints are created along each trajectory, to prevent aircraft from leaving the measurement area entirely in an attempt to avoid conflicts. The cruise altitude, or cruising layer, for each aircraft is set once the heading range distribution per each layer is determined, which happens before the start of the simulation. All aircraft within the same layer fly at exactly the same altitude, in the middle of the layer. 8 layers are defined, between altitudes 1500 ft to 6500 ft.

Logging of conflicts is restricted to the cruise phase of the flight. Conflicts resulting from aircraft climbing and descending to their cruising altitude are not taken into consideration. Considering only conflicts during cruise leads to optimal convergence of the machine learning model for this phase. It is likely that reducing conflicts during climb and descent phases requires a different approach, and thus might result in non-convergence of the machine learning model if they were to be modelled together.

B. Apparatus and Aircraft Model

The open Air Traffic Simulator BlueSky [8] is used for simulation of an airspace environment. All vehicles in the simulation used a Boeing-737 performance model. Bluesky uses a kinematic aircraft performance model based entirely on open data [10]. An average speed of 140 kts was considered, resulting in an average flight time of roughly 850 seconds.

C. Conflict Detection and Minimum Separation

We consider a horizontal separation of 5 NM. Vertical separation is set in function of the dimension of a vertical layer, thus aircraft cruising at adjacent vertical layers are not in conflict. The experiment employs state-based conflict detection. A look-ahead time of five minutes is used. The conflict evaluation interval is set to one second; each second, the current conflicts and LoSs are detected, and the necessary conflict avoidance manoeuvres are calculated.

D. Conflict Resolution

The geometric derivation of a resolution with the MVP model is displayed in Fig. 2, as previously defined [2]. When a conflict is detected, MVP uses the predicted future positions of both ownship and intruder at the closest point of approach (CPA). These calculated positions “repel” each other, and this “repelling force” is converted to a displacement of the

predicted position at CPA. The avoidance vector is calculated as the vector starting at the future position of the ownship and ending at the edge of the intruder’s protected zone, in the direction of the minimum distance vector. This displacement is thus the shortest way out of the intruder’s protected zone. Dividing the avoidance vector by the time left to CPA yields a new speed, which can be added to the ownship’s current speed vector resulting in a new advised speed vector. From the latter, a new advised heading and speed can be retrieved. In a multi-conflict situation, the final avoidance vector is determined by summing the repulsive forces with all intruders. As it is assumed that both aircraft in a conflict will take (opposite) measures to evade the other: MVP is implicitly coordinated.




Figure 2. MVP geometric resolution. CPA represents the closest point of approach between the two aircraft. $PZ_{intruder}$ represents the protected zone of the intruder. Adapted from Hoekstra [2].

E. Independent Variables

The heading range distribution per layer, and traffic density are set as independent variables.

The heading range distribution per layer is set to pre-defined values during training and validation of the CEN. Once the CEN has been validated, its output (expected number of conflicts for the given traffic scenario and heading range distribution per layer) is then used to train the LHN. The latter network outputs a heading range distribution per layer which is directly compared with evenly distributed heading ranges per layer, by running the same traffic scenarios with both.

During training, a medium traffic density is used. During testing, lower and higher traffic densities, as per Table II, are introduced in order to analyse how the supervised learning model performs at traffic densities it was not trained in.

TABLE II. TRAFFIC VOLUME USED IN THE EXPERIMENTAL SIMULATIONS.

	Low	Medium	High
Traffic density [$ac/10\,000\,NM^2$]	0.005	0.01	0.015
Number of spawned aircraft [-]	50	100	150

F. Dependent Measures

Two different categories of measures are used to evaluate the effect of different airspace structures: *safety*, and *efficiency*.

Safety is defined in terms of the total number conflicts and LoSs. Naturally, fewer conflicts and LoSs are safer.

Efficiency is evaluated in terms of distance and time travelled; significantly increasing the path or time travelled is considered inefficient.

IV. EXPERIMENT: HYPOTHESES

It was hypothesized that a supervised learning model would be able to identify an airspace structure which minimizes the number of conflicts for the given traffic scenario. Through segmentation of the traffic per the available airspace according to their real heading distribution, more conflicts can be prevented versus a structure that assumes an even traffic distribution. Moreover, although the total of number of conflicts is not directly proportional to the total number of LoSs, fewer conflicts tend to lead to fewer LoSs [11]. Reducing the number of conflicts also leads to fewer conflict resolution manoeuvres. This leads to higher efficiency as aircraft can follow their planned trajectory, instead of adopting a conflict resolution path to avoid intruders.

During testing, the supervised learning model is run in the traffic density it was trained in, as well as with lower and higher traffic densities. The objective is to verify if the performance of the supervised learning model declines when used with different traffic densities. It was hypothesized that the model would deteriorate in performance with higher traffic densities. A higher number of aircraft increases complexity of the environment, and in turn, the difficulty of determining an optimal airspace structure.

V. EXPERIMENT: RESULTS

The performed experiment involves a training and a testing phase. First, the supervised learning model is trained continuously with a set of traffic scenarios for medium traffic density. Second, it is tested with unknown traffic scenarios for low, medium, and high traffic density. Performance with these new scenarios is directly compared to a baseline employing an evenly distributed heading range per vertical layer.

A. Training the Model

The first step is to train the CEN. This network is repeatedly trained on a set of 16 scenarios. Then it is validated with only one scenario. All scenarios run for 500 seconds. A termination rule was imposed as follows: when the 60-episode training moving average (MA) does not decrease for more than 30 episodes, the network finishes the training. Fig. 3 displays the evolution of the loss during training. Results show that it converges towards a minimal average loss (roughly 0.18) within 300 episodes. There are, however, occasional large spikes (i.e., episodes where the CEN estimated number of conflicts diverged significantly from the value observed with the Bluesky simulation). This is considered a result of the formulation of the traffic scenario not identifying the exact direction of the traffic (i.e., all traffic is discretised into heading bins). As a result, the CEN is not capable of identifying “hotspots” of conflicts from intercepting routes within a heading bin. The larger the heading range in a layer is, the more likely it is that flight routes within this layer will intersect. The CEN does not account for these.

Fig. 4 shows the evolution of the “real” and “estimated” number of conflicts during training and validation of the CEN. The “real” number of conflicts is obtained through simulation




Figure 3. Evolution of the loss during training of the CEN. The Mean Squared Error between the “real” and the “estimated” number of conflicts is used as a loss metric.




Figure 4. Evolution of the estimated number of conflicts during testing and validation of the CEN.

with Bluesky; “estimated” is the value output by the CEN. Throughout the training of the CEN, the precision of the prediction of conflicts evolves up to an average error of only 5.3 conflicts.

The training of the LHN then follows. This network is repeatedly trained on a set of 48 scenarios. Training stops once it is detected that the network reaches a stable value, and is no longer learning. There are no validation steps, due to the fact that the outcome of this training will be validated only later, by means of a comparison with evenly distributed heading ranges per layer. Fig. 5 shows the evolution of the total number of conflicts through training of the LHN. After 6000 episodes, the LHN is capable of finding heading range distributions which minimize the total number of conflicts.

An example of the structure output by the supervised learning model is presented in Fig. 6. Here, a traffic scenario is outlined in blue (top graph), and the supervised learning model assigns heading ranges per vertical layer according to the gray graph (bottom). When comparing the two graphs it can be seen that the supervised learning model attempts to distribute aircraft evenly over the vertical layers by selecting




Figure 5. Evolution of the loss during training of the LHN.

smaller heading ranges per layer within directions with higher incidence of traffic. This results in an optimized segmentation of the traffic per the available vertical space. Although a similar number of aircraft in all layers may be expected, this may not be optimal safety-wise. A wide heading range in a layer results in considerable heading differences between aircraft travelling in the same layer, leading to intercepting routes and large conflict angles. Adding more aircraft to a layer with a wide heading range comes at a higher cost than adding aircraft in a layer with a smaller range. Thus, an equal number of aircraft in all layers may not be desirable.




Figure 6. Example of a heading range distribution (bottom) output by the supervised learning model for a given traffic scenario (top).

B. Testing the Model

The effect on safety and efficiency of the structures output by the supervised model is directly compared to employing evenly distributed heading ranges. The former is denominated as “Adaptive” and the latter as “Even”, in the following graphs. 800 traffic scenarios are run for both situations, for each traffic density. Each scenario runs for 1500 seconds.

1) Safety Analysis

Fig. 7 displays the total number of conflicts for all traffic densities. The supervised learning model was able to decrease




Figure 7. Total number of conflicts with a heading range distribution per layer output by the supervised learning model (Adaptive) and an even heading range distribution (Even), for all traffic densities with conflict resolution.

the number of conflicts for all traffic densities, compared to the baseline model. Such was expected for the medium traffic density, since this resembles the environment the model trained in. The fact that the latter was also able to reduce conflicts at lower and higher traffic densities, shows that its behaviour is expandable to different numbers of aircraft.

Figs. 8, 9, and 10 further explore how the supervised learning model structures traffic based on demand. For simplification, we focus on the output structures for a medium traffic density.

Fig. 8 displays the relationship between the heading range and the number of aircraft per layer, both for the supervised learning model and for an evenly distributed structure. The latter assigns the same heading range to all the layers; it does not defend against a high incidence of aircraft in one layer. Consequently, it has a higher maximum for the number of aircraft per layer. Naturally, as the number of aircraft per layer increases, so does the number of conflicts. The results with the supervised learning model show a more optimal segmentation of aircraft. By reducing the heading range in each layer, the model limits the number of aircraft per layer, and consequently the number of conflicts. Additionally, it is interesting to note some linearity between the number of aircraft and the heading range. This is probably due to the linearity of the hidden layers in the neural network, and the partial linearity of the Rectified Linear Unit (ReLU) activation function.

Fig. 9 plots the probability density of the number of aircraft per heading range in each layer for the supervised learning model. It shows that the model tries to exploit two properties: (1) a small heading range, or (2) a low number of aircraft. However, the frequency with which the model assigns many aircraft to a layer with a small heading range is higher than the frequency with which it assigns few aircraft to a layer with a larger heading range.

Figs. 10a and 10b display the probability density function between the number of aircraft and the total number of conflicts in a layer for the supervised learning model and for an even distribution, respectively. The evenly distributed heading range concept has a higher average and a higher maximum number of conflicts. Additionally, the contours of Fig. 10b are




Figure 8. Relationship between the heading range and the number of aircraft per layer. A point is plotted for each layer for each tested traffic scenario.





Figure 9. Probability density between heading range and the number of aircraft in each layer for the output of the supervised learning model.


approximately symmetrical with respect to the mean number of conflicts. This can be seen as a benefit, since the probability of having an episode with a great number of conflicts is nearly none; the system is more dependable.

Fig. 11 displays the total number of LoSs for all traffic densities. Although the objective of the supervised learning model was to reduce the total number of conflicts and not of LoSs, as hypothesized, a reduction in the number of conflicts leads to fewer LoSs.

Fig. 12 compares distributions of the occurrence of losses of separation, for the four experiment conditions. The data for these graphs is based on the medium traffic density scenarios. Comparison of these distributions enables a comparison of the relative influence of (1) adaptive airspace structuring, and (2) conflict resolution, on the total number of losses of separation. The top two graphs represent the simulations run without conflict resolution, and the bottom two run using the MVP model for resolving conflicts. As previously seen in Figs. 7 and 11, the scenarios with the supervised learning model have fewer conflicts and fewer LoSs. As expected, within the same airspace structure, applying conflict resolution further reduces the number of LoSs. However, it also increases the disparity



(a) With an evenly distributed heading range.



(b) With the supervised learning model output.

Figure 10. Probability density function of the number of conflicts against the number of aircraft in each layer, both for output of the supervised learning model and for an evenly distributed heading range per layer.




Figure 11. Total number of LoSs resulting from employing the heading ranges per layer output by the supervised learning model (Adaptive) and an even heading range distribution per layer (Even), for all traffic densities with conflict resolution.

between the mean and the median values, albeit with a smaller impact when the supervised learning model is used. Finally, it is interesting to note that the layer distribution output by the supervised learning model without MVP (blue) appears to perform even better than the evenly distributed layers with




Figure 12. Total number of LoSs resulting from employing the heading ranges per layer output by the supervised learning model and an even heading range distribution per layer. The top two results were run without conflict resolution, and the bottom two were run using the MVP model for resolving conflicts.

MVP (brown). Conflict prevention resorting to airspace design may be more efficient at improving safety than just applying conflict resolution.

2) Efficiency Analysis

To analyse efficiency, we focus on the effect of airspace structuring and conflict resolution on the flight routes. With the Bluesky simulator, this information is gathered once aircraft finish their flights. We denominate aircraft which have finished their flight as “landed” aircraft. In order to have a direct comparison between the structures output by the supervised learning model and an evenly distributed structure, all episodes run for the same amount of time. As a result, a different number of aircraft might have landed for different structures according to its effect on flight path/time.

Fig. 13 displays the total number of landings throughout each scenario for all traffic densities. On average, with the structures output by the supervised learning model, more aircraft landed. This signifies that, with an even heading distribution per layer, aircraft take longer to finish their routes. As hypothesized, a higher number of conflicts leads to aircraft having to adopt more conflict resolution manoeuvres which alter their otherwise linear trajectory. The more aircraft have to diverge from their linear path to avoid intruders, the longer their path will be.

Fig. 14 displays the total distance travelled by each aircraft for all traffic densities. Note that these values are only from the aircraft that landed (i.e., that finished their path within the running time of each scenario). Even though only the flights which finished their trajectory were included, there was already a strong effect visible. These values confirm that, with an even heading range distribution per layer, aircraft tend to travel slightly longer.

Fig. 15 displays the total time travelled throughout each scenario for all traffic densities for “landed” aircraft. Here, no significant differences are visible. Comparing with Fig. 14,




Figure 13. Total number of landings resulting from employing the heading ranges per layer output by the supervised learning model (Adaptive) and an even heading range distribution per layer (Even), for all traffic densities with conflict resolution.




Figure 14. Total distance travelled resulting from employing the heading ranges per layer output by the supervised learning model (Adaptive) and an even heading range distribution per layer (Even), for all traffic densities with conflict resolution.




Figure 15. Total time travelled resulting from employing the heading ranges per layer output by the supervised learning model (Adaptive) and an even heading range distribution per layer (Even), for all traffic densities with conflict resolution.

it is likely that, with an even heading range distribution per layer, aircraft adopt higher speeds for conflict resolution, thus compensating for the extra travelled distance.

VI. DISCUSSION

Results show that having a dynamic airspace structure capable of adapting to the current traffic scenario leads to an increase in the airspace capacity, by reducing conflicts and LoSs. The supervised learning model optimized segmentation

of aircraft; the heading range allowed in each vertical layer was reduced when it is detected that a high number of aircraft will travel within that range. Limiting the traffic density at each vertical layer had a positive effect of limiting conflicts and LoSs, thus increasing safety. Moreover, the scenarios ran with the supervised learning model, without conflict resolution, had fewer LoSs than the scenarios run with an evenly distributed heading range per layer with conflict resolution activated. This shows that conflict prevention may be the best form of conflict “resolution”; in a situation with several multi-actor conflicts, conflict resolution algorithms can encounter deadlocks, which prevent them from resolving all conflicts.

However, the supervised learning model can still be further improved. First, the number of aircraft assigned to a layer is occasionally zero. This is a misallocation of useful vertical space, causing more stress on the other layers. It is likely that a more complex discretization of the current traffic scenario, or even focusing on the number of conflicts on each layer instead of the total number of conflicts in the airspace, may mitigate this issue. Second, performance of the LHN is dependent on the performance of CEN, as the number of conflicts estimated by the latter is used to train the former. Optimizing conflict information, will also optimize the output of the LHN. Third, training focused solely on the number of conflicts. Considering also LoSs, or even efficiency factors such as increased flight path or flight time, may further optimize the structures found by the supervised learning model. Lastly, since the objective is to minimize the number of conflicts, this can be seen as improving a cost function as typically used with reinforcement learning. This represents a different approach from the work herein performed; it is of interest to compare both.

This work assumed a fixed number of vertical layers. This is a simplification that favoured optimal convergence of the supervised learning model. Although it is fair to assume that a certain range of flight levels is allocated for air transportation and that aircraft must adhere to these limits, it may also be that controlling the number of vertical layers may further optimize capacity of the airspace. Naturally, adopting fewer or more layers has an impact on the segmentation of traffic. For the cruising phase, more layers are expected to improve segmentation and thus potentially decrease the number of conflicts. However, in future work, the effects of climbing and descending towards the correct layer must be considered. Climb and descent phases account for a large portion of conflicts and LoSs in environments with non-linear routes [11], [12]. Considering vertical deviations will likely help improve safety in the airspace. However, it will also add complexity to the training of a machine learning model. It is likely that reducing conflicts during cruise, climb, and descent phases requires different approaches, and consequently, models with different learning policies.

Future work should explore fuel consumption, and resulting environmental impact, of climbing, descending, and allocating aircraft to sub-optimal altitudes. Moreover, before a real-world implementation, the method must be further tested with different traffic densities and trajectories, improving its

capability to generalise. There is also potential for this method to be applied to unmanned aviation, where it may have a bigger impact given the higher variability of trajectories and traffic types. Finally, no machine learning application can be blindly implemented into a real-life scenario. Further examination is necessary to explain the choices made by the model, as well as safeguards for potential bad decisions when applied to aircraft densities/trajectories not previously seen.

VII. CONCLUSION

This work focused on using neural networks to create a safer, dynamic version of the layered airspace concept adapted to the current traffic scenario. Results showed that a supervised learning model is capable of optimally dividing aircraft per the available airspace in function of their heading distribution, thus increasing safety. Proper segmentation of traffic even had a greater effect on safety than employing a conflict resolution model. Multi-conflict situations are extremely difficult to resolve; preventive action towards limiting the occurrence of these situations may be the only way to resolve them.

Future research should consider non-linear trajectories, which will likely create density “hotspots”, heavily increasing the number of conflicts as well. This will likely require a more complex representation of the environment, as to identify heading changes. Finally, the research presented herein can be extended towards more competitive operational environments with a different number of layers, differences in the performance limits, as well as preference for efficiency over safety.

REFERENCES

- [1] “Performance review report covering the calendar year 2018,” EUROCONTROL, Tech. Rep., 2018.
- [2] J. Hoekstra, R. van Gent, and R. Ruigrok, “Designing for safety: the ‘free flight’ air traffic management concept,” *Reliability Engineering & System Safety*, vol. 75, no. 2, pp. 215–232, feb 2002.
- [3] J. M. Hoekstra, “Free flight in a crowded airspace?” 2000.
- [4] K. Bilimoria, H. Lee, Z.-H. Mao, and E. Feron, “Comparison of centralized and decentralized conflict resolution strategies for multiple-aircraft problems,” 08 2000.
- [5] J. Shortle and Y. Zhang, “Safety comparison of centralized and distributed aircraft separation assurance concepts,” *Reliability, IEEE Transactions on*, vol. 63, pp. 259–269, 03 2014.
- [6] E. Sunil, J. Hoekstra, J. Ellerbroek, F. Bussink, D. Nieuwenhuisen, A. Vidosavljevic, and S. Kern, “Metropolis: Relating Airspace Structure and Capacity for Extreme Traffic Densities,” in *ATM seminar 2015, 11th USA/EUROPE Air Traffic Management R&D Seminar*, Portugal, 2015.
- [7] M. Tra, E. Sunil, J. Ellerbroek, and J. Hoekstra, “Modeling the intrinsic safety of unstructured and layered airspace designs,” in *Twelfth USA/Europe Air Traffic Management Research and Development Seminar*, 2017.
- [8] J. Hoekstra and J. Ellerbroek, “Bluesky ATC simulator project: an open data and open source approach,” in *Conference: International Conference for Research on Air Transportation*, 2016.
- [9] M. Ribeiro, J. Ellerbroek, and J. Hoekstra, “Review of conflict resolution methods for manned and unmanned aviation,” *Aerospace*, vol. 7, no. 6, p. 79, jun 2020.
- [10] J. Sun, J. M. Hoekstra, and J. Ellerbroek, “Open Aircraft Performance Modeling: Based on an Analysis of Aircraft Surveillance Data,” Ph.D. dissertation, Delft University of Technology, 2019.
- [11] M. Ribeiro, J. Ellerbroek, and J. Hoekstra, “Velocity obstacle based conflict avoidance in urban environment with variable speed limit,” *Aerospace*, vol. 8, no. 4, 2021.
- [12] M. Doole, J. Ellerbroek, V. L. Knoop, and J. M. Hoekstra, “Constrained urban airspace design for large-scale drone-based delivery traffic,” *Aerospace*, vol. 8, no. 2, 2021.