NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

ENHANCING AIR TRAFFIC CONFLICT RESOLUTION THROUGH MACHINE LEARNING, CONFORMAL AUTOMATION, AND FLOW-CENTRIC PARADIGMS

YASH GULERIA

SCHOOL OF MECHANICAL AND AEROSPACE ENGINEERING

2024

ENHANCING AIR TRAFFIC CONFLICT RESOLUTION THROUGH MACHINE LEARNING, CONFORMAL AUTOMATION, AND FLOW-CENTRIC PARADIGMS

YASH GULERIA

School of Mechanical and Aerospace Engineering

A thesis submitted to the Nanyang Technological University
in partial fulfilment of the requirement for the degree of
Doctor of Philosophy

2024

Statement of Originality

I hereby certify that the work embodied in this thesis is the result of original research, is free of plagiarised materials, and has not been submitted for a higher degree to any other University or Institution.

25/01/2024

Date

Yash Guleria

Supervisor Declaration Statement

I have reviewed the content and presentation style of this thesis and declare it is free of plagiarism and of sufficient grammatical clarity to be examined. To the best of my knowledge, the research and writing are those of the candidate except as acknowledged in the Author Attribution Statement. I confirm that the investigations were conducted in accord with the ethics policies and integrity standards of Nanyang Technological University and that the research data are presented honestly and without prejudice.

25/01/2024

. . . .

Date   ` Sameer Alam

Authorship Attribution Statement

This thesis contains materials from three papers published in the following peer-reviewed conferences and journals in which I am listed as an author.

Chapter 3 is published with material from:

(1) Guleria, Y., Tran, P.N., Pham, D.T., Alam, S. and Durand, N., 2021, December. A machine learning framework for predicting ATC conflict resolution strategies for conformal automation. In 11^{th} SESAR Innovation Days.

AND

(2) Guleria, Y., Pham, D.T., Alam, S., Tran, P.N. and Durand, N., 2024. Towards conformal automation in air traffic control: Learning conflict resolution strategies through behavior cloning. Advanced Engineering Informatics, 59, p.102273.

The contributions of the authors are as follows:

- Prof. Sameer Alam proposed the research idea and provided the research direction of the project.
- Prof. Sameer Alam and I developed the methodology for the research.
- Dr. Thinh and I developed the machine learning model.
- Dr. Phu and Dr. Thinh provided constant feedback and discussion on the model development and experimental design.
- I implemented the source code, conducted the experiments and prepared the manuscript draft.
- Prof. Nicolas Durand provided feedback on the research and reviewed the manuscript.

Chapter 4 and **Chapter 5** are published with material from:

(1) Guleria, Y., Pham, D.T., Alam, S., Tran, P.N. and Durand, N., 2024. Towards conformal automation in air traffic control: Learning conflict resolution strategies through behavior cloning. Advanced Engineering Informatics, 59, p.102273.

The contributions of the authors are as follows:

- Prof. Sameer Alam proposed the research idea and provided the research direction of the project.
- Prof. Sameer Alam and I developed the methodology and the scope of the research.
- I had constant discussions with Dr. Thinh on the performance test of the machine learning model and validation exercises. Dr. Thinh and Dr. Phu provided constant feedback and discussion on the model development and experimental design.
- I implemented the source code, conducted the experiments, and prepared the manuscript draft.
- Prof. Nicolas Durand provided feedback on the research, hosted me in his lab to conduct the experiments, and reviewed the manuscript.
- Prof.Sameer Alam and Dr. Thinh reviewed the manuscript.

Chapter 6 is published as:

(1) Guleria, Y., Pham, D.T. and Alam, S. (2023). An agent-based approach for air traffic conflict resolution in a flow-centric airspace. 26^{th} IEEE International Conference on Intelligent Transportation Systems (ITSC 2023).

The contributions of the authors are as follows:

- Prof. Sameer Alam proposed the research idea and provided the research direction of the project.
- Dr. Thinh and I developed the experimental design, decided the scope, and the methodology of the research.

- I implemented the source code, conducted the experiments, and prepared the manuscript draft.
- Prof. Sameer Alam and Dr.Thinh reviewed the manuscript.

25/01/2024

Date

Yash Guleria

Acknowledgements

First, I wish to express my sincere gratitude to my advisor, Professor Sameer Alam for giving me the opportunity to pursue research at ATMRI, providing guidance in research, and lessons for life as well, throughout my time at ATMRI. His positive attitude, exceptional work ethic, and vision have inspired me to be the best that I can be in the situations that present themselves. He has been extremely patient with me, and has been a source of constant support throughout my Ph.D. journey.

I would like to thank Professor Vu Duong for agreeing to be my thesis advisor. He has been a source of constant motivation with his vision for the future of air traffic management, and a caring attitude towards all of us at ATMRI. My sincere thanks to Professor Michael Schultz for agreeing to be my thesis advisor, providing valuable inputs on my work, and engaging in discussion on research and life. I am grateful to Professor Nicolas Durand for accommodating me in his lab and giving the opportunity to perform experiments there. The entire team at ENAC has been very warm and accommodating. A heartfelt thanks to Dr. Thinh for mentoring me, and being a friend when I needed support. His valuable support made the Ph.D. journey relatively less stressful. I would also extend my gratitude to Dr.Phu for mentoring me in the initial phase of my journey.

I am grateful for having met the current and past lab mates at ATMRI, Dr. Ma, Dr. Arif, Dr. Ali, Phuoc, Dr. Cai, Dr. Goh, Dr. Kiran, Greg, Zhi Jun, Duy Anh, Nam and Duy. Special thanks to Narendra Sir and Luke for guiding me and enriching my operational knowledge. I am also grateful to Suan, Yati, and Sunny for their help and support.

I must acknowledge a few great friends back home, who have been with me through the

thick and thin.

Finally, I take this opportunity to express how blessed I am to have my family: my parents, my sister, and my wife, without whom, my current life is unimaginable. Their unconditional love, support, and motivation has brought me this far.

Contents

Acknowledgement **viii**

Abstract **xiii**

List of Figures **xvii**

List of Tables **xx**

Institutional Review Board Approvals **xxi**

List Of Abbreviations **xxii**

1 Introduction **1**

- 1.1 Air Traffic Conflict Resolution 2
  - 1.1.1 Background 2
  - 1.1.2 Air Traffic Conflict Resolution Strategies 5
  - 1.1.3 Associated Challenges 7
- 1.2 Research Motivation 11
  - 1.2.1 ATCO-Automation Conformance 11
  - 1.2.2 Scalability Constraints of the Airspace 12
- 1.3 Research Objectives and Scope 13
- 1.4 Thesis Contributions 14
  - 1.4.1 Identification of ATCOs' Conflict Resolution Strategies 15
  - 1.4.2 Machine Learning Models for ATCO-Conformal Conflict Resolution with Behavior Cloning 16

1.4.3 Air Traffic Conflict Resolution for Flow-Centric Operations 17
1.5 Thesis Organization 18

2 Literature Review 19
2.1 Air traffic conflict resolution Methods: An overview 20
2.2 Air Traffic Conflict Resolution Using Conventional Methods 23
2.3 Air traffic Conflict Resolution Using Learning Based Methods 25
2.3.1 Conventional Machine Learning Approaches 25
2.3.2 Learning from Expert's Demonstrations 29
2.4 Air Traffic Conflict Resolution in Flow-Centric Operations 34
2.5 Research Gaps 36

3 Air Traffic Conflict Resolution Strategies: Identification Methodology and Analysis 38
3.1 Overview 39
3.2 Conflict Resolution Maneuver Representation 39
3.3 Experiment Design 41
3.3.1 Simulation Environment 41
3.3.2 Conflict Scenarios and Data Generation 42
3.3.3 Evaluation metrics 43
3.3.4 Participant Demographics 44
3.4 Strategy Identification and Analysis 45
3.4.1 ATCOs' Conflict Resolution Maneuvers 48
3.4.2 ATCOs' Conflict Resolution Strategies 51
3.5 Summary 52

4 Air Traffic Conflict Resolution Strategies: Learning and Prediction 54
4.1 Learning ATCOs' strategies: Supervised Learning Approach 55
4.1.1 Overview 55
4.1.2 Problem Formulation 56

4.1.3 Conflict Resolution Dataset 56
4.1.4 Learning Architecture 57
4.1.5 Feature Engineering 59
4.1.6 Model Selection And Performance Evaluation 60
4.2 Learning ATCOs' Strategies: Deep Reinforcement Learning Approach . . . 71
4.2.1 Overview 71
4.2.2 Problem Formulation 71
4.2.3 Conflict Resolution Dataset 72
4.2.4 Learning Mechanism 72
4.2.5 Results and Discussions 77
4.3 Operational Implications 79
4.4 Summary 80

5 Air Traffic Conflict Resolution Strategies:
Conformance and Evaluation **83**
5.1 ATCO Validation and Conformance Evaluation 84
5.1.1 Experiment Objectives 84
5.1.2 Experiment Design and Interface 84
5.1.3 Results and Discussions 85
5.2 Conflict Resolution Preference Evaluation 88
5.2.1 Experiment Objectives 88
5.2.2 Experiment Design and Interactive Interface 88
5.2.3 Results and Discussions 92
5.3 Summary 95

6 Agent-Based Conflict Resolution Strategies: Flow-Centric Paradigms **97**
6.1 Flow-Centric Operations: Introduction 98
6.2 Flow Centric Conflict Resolution as a Reinforcement Learning Problem . . 101
6.3 Methodology 102
6.3.1 Conflict Scnearios 102

6.3.2 Markov Decision Process Framework 104
6.3.3 Learning Algorithm 107
6.4 Experiments and Results 109
6.4.1 Experimental Setting 109
6.4.2 Safe Separation Analysis 110
6.4.3 Maneuver Efficiency Analysis 111
6.5 Summary 114

7 Conclusions and Future Work 116

7.1 Thesis Summary 117
7.2 Revisiting the Research Questions 118
7.3 Challenges in Integration of ML Based Conflict Resolution Models in Air Traffic Control 121
7.3.1 Challenges in Conflict Resolution Model Development 122
7.3.2 Challenges in Validation and Certification of the ML Based Conflict Resolution Models 125
7.3.3 Challenges in Adoption of ML Models by ATCOs 127
7.4 Opportunities for Future Research 128
7.4.1 Models for Increased Conflict Scenario Complexity 128
7.4.2 Investigation of Generalizability and Model Robustness 129
7.4.3 Development of Methods for Adaptive Conformal Conflict Resolution Advisories 130
7.4.4 Integration With Other Air Traffic Control Subsystems 130
7.5 Concluding Remarks 131

A Post Experiment Questionnaire 148

A.1 Workload 148
A.2 Situational Awareness 149
A.3 Trust 150
A.4 Additional Questions 151

Abstract

Air traffic conflict resolution is a dynamic, time-sensitive, and safety-critical aspect of air traffic control, which involves a complex interaction of humans, machines, and procedures. In current sector-based operations where the airspace is subdivided into smaller geographical regions, ensuring safe separation between the aircraft by resolving potential conflicts and maintaining an efficient flow of traffic is the primary responsibility of air traffic controllers (ATCOs). The task being safety-critical, demands high cognitive effort. Researchers have proposed several computational and learning-based methods for air traffic conflict resolution. However, the translation of such methods to operations and their acceptance by ATCOs remains a challenge because of a mismatch between how the ATCOs perceive the conflicts and resolve them, and the resolutions proposed by the computational methods. In other words, the ATCOs' conflict resolution maneuver preferences may differ from the optimized solutions generated by such computational and learning-based methods. In this regard, behavior cloning is an established methodology to encapsulate the preferences of a human in performing complex tasks. Moreover, sector-based operations also present inherent scalability constraints that hinder meeting future demands. New concepts of operations such as flow-centric operations (FCOs), which involve managing air traffic from an aggregated perspective instead of individual flights (as in sector-based operations) based on the formation and evolution of major traffic flows, are being proposed to accommodate future traffic. This poses an additional challenge in ensuring safe separation between air traffic flows in flow-centric operations. Thus, there are two-fold challenges in meeting future traffic demands for sustained growth while ensuring safe operations. First, the development of methods that incorporate ATCOs'

conflict resolution preferences, and second, the development of novel conflict resolution techniques to accommodate future air traffic management concepts.

Research approaches in the development of methods for air traffic conflict resolution have transitioned from mathematical models and optimization approaches to learning-based methods. This shift has been triggered to better address the limitations of mathematical models in accommodating uncertainties and stochasticity associated with the environment, and generalizability to non-nominal scenarios with non-standard model inputs. Nonetheless, the proposed approaches have the following fundamental limitations. Firstly, air traffic conflict resolution is a sequential decision-making problem, where the ATCOs must make a series of sequential decisions to ensure safe separation between aircraft. Such sequential decisions are a result of the inherent preferences or *strategies* that the ATCOs develop over time and constitute an end-to-end conflict resolution maneuver of the aircraft. The machine learning methods proposed in the literature do not generate an end-to-end conflict resolution maneuver but rather, only the initial vectoring segment. ATCOs' conflict resolution strategies are also not captured by such approaches. In this regard, behavior cloning is an established methodology to encapsulate the preferences of a human in performing complex tasks. Furthermore, in the design of such methods, researchers have employed simplistic simulation environments, which fail to emulate the complexity of an end-to-end conflict resolution maneuver. Secondly, the existing literature lacks analysis and discussion on how current conflict resolution methods can be adapted to function in flow-centric operations. In this context, the thesis addresses two key research questions. First, can the ATCOs' conflict resolution strategies be learned through behavior cloning and incorporated into a learning-based model? This research question encompasses developing a methodology to identify the ATCOs' conflict resolution strategies, and then, developing suitable machine learning models that can learn these strategies. Second, how can air traffic conflicts be resolved in a flow-centric paradigm where traffic is modeled as intersecting flows? The approaches proposed to address these research questions have resulted in three research contributions.

The first contribution is based on developing a methodology to identify the conflict resolution strategies of the ATCOs. To investigate the ATCOs' conflict resolution strategies and the factors affecting them, high-quality data that is representative of the ATCOs' behaviors is required. Therefore, a conflict resolution data collection experiment was designed with the objective of collecting ATCOs' conflict resolution maneuvers for the corresponding conflict scenarios. In this experiment, conflict scenarios were simulated in a high-fidelity simulation environment that emulates an air traffic control radar interface. Eight ATCOs were involved in the experiments. The participants were shown real-time air traffic conflicts in the simulation environment representing a sector of the Singapore flight information region and their interactions with the simulation environment to resolve the conflicts were recorded. Individual analysis and comparison of each ATCO's data demonstrated that ATCOs have distinct conflict resolution strategies, specifically in terms of maneuver direction and the magnitude of the heading deviations. The metrics used to identify these strategies were the selection of the aircraft to be maneuvered, maneuver initiation time, maneuver direction, magnitude of deviations, and preferred safe separations between the conflicting aircraft. It was also observed that the ATCOs' conflict resolution strategies were influenced by factors such as the proximity of the conflict point from the sector boundary and which aircraft arrived first at the conflict point.

The second contribution developed machine learning-based ATCO conformal conflict resolution models using behavior cloning. Here, 'conformal' implies resolutions that are similar to the ATCOs' conflict resolution strategies. This work formulated conflict resolution as a sequential decision-making task and used a sequence of regressor and classifier-supervised machine learning models to map the environment states to low-level actions and clone the behavior of the ATCOs. Based on the ATCOs' strategies identified previously, personalized (matching individual ATCO's preference) and group conformal (matching the preference of a group of ATCOs) models were developed using data collected during the conflict resolution experiments with the ATCOs. The prediction results demonstrated the proposed models are able to generate ATCO conformal predictions on the test dataset. The trained models were also tested for added Gaussian noise to the data, to evaluate the models' robustness. The results demonstrated that the models are

robust up to 7.5% added noise, with low mean absolute errors and high accuracy for the predictions (for example, the classification accuracy was $> 92.7\%$ for predicting the maneuvering aircraft, MAE for maneuver initiation distance was < 5.3 NM and MAE for predicting the heading angle was $< 5.3°$ for the prediction models). The work was further evaluated by human-in-loop experiments involving professional ATCOs, to evaluate their choices in the selection of their own conflict resolution strategies, and the conflict resolution strategies obtained from the prediction models. The results demonstrated that the ATCOs selected solutions matching their conflict resolution strategies for over 70% of the scenario, which reinstates that conformal conflict resolution advisories receive greater acceptance by the ATCOs.

The third contribution focused on the emerging concepts of operations involving the flow-centric paradigm and the development of inter-flow (between the flows) and intra-flow(within each flow) conflict resolution models for flow-centric operations. Crossing conflict scenarios under uncertainty in a flow-centric paradigm were investigated in this work. The research was modeled as a sequential decision-making problem using the Markov Decision Process (MDP). This choice was made to align with the inherent sequential decision-making nature of the conflict resolution process. Each flow, which consists of a varying number of aircraft with different cruise speeds and the associated location uncertainties, was modeled as a self-stabilizing graph structure. The two-stage conflict resolution process involved training a conflict resolution policy to ensure inter-flow safe separation along with the use of a self-stabilizing graph structure to ensure intra-flow safe separation. The trained policy can ensure both intra and inter-flow safe separation between the aircraft for 100% of the scenarios. The performance was further analyzed in terms of maneuver efficiency and deviations from the flight plans. Inspite of the uncertainties and dynamics associated with the flows' size, speed, and evolution over time, the absolute delays for the flow were 2.53 and 9.49 minutes respectively.

In its entirety, this thesis addresses both immediate and future strategies for tackling the challenge of escalating air traffic by proposing innovative approaches for safe separation of air traffic. It presents methods for the creation of learning-based conformal models by cloning the behavior of ATCOs and the formulation of a conflict resolution model within a flow-centric airspace paradigm, which holds significance for the evolving concepts of operations.

List of Figures

1.1 Modes to achieve safe separation 3
1.2 Safe separation in procedural airspace 4
1.3 Safe separation in controlled airspace 4
1.4 Vectoring to achieve safe separation 5
1.5 Representation of a strategy which is a mapping from the system states to corresponding actions. 6
1.6 IFR movement forecast for Europe: 2050 7
1.7 Relationship between air traffic and perceived conflict rate. 8
1.8 Graphical representation: ATCO-automation resolution mismatch 12
1.9 Graphical representation: ATCO-automation solution space. 12
1.10 Thesis organization 15

2.1 Graphical representation of ATCO tasks and their sub-categories, along with the corresponding automation methods. 21
2.2 Agent-environment interaction in a Markov decision process. 27
2.3 Interaction between the learning environment and the AI agent, Pham et al. [1] 28
2.4 Strategic conformance: Overall idea 31
2.5 Conformal automation in air traffic conflict resolution 32
2.6 Conflict resolution with the Interactive conflict solver Tran et al. [2] 34
2.7 Major flows identified in the current sectored airspace over the French airspace, Ma et al. [3] 35

3.1 Concept diagram: ATCOs' strategies identification 39

3.2 End-to-end conflict resolution maneuver representation 41
3.3 ESCAPE Light simulation interface 42
3.4 Conflicting flight plan data generation pipeline 43
3.6 Individual dataset size: Visual representation 45
3.7 Combined data distribution of all ATCOs 46
3.8 Data distribution of ATCOs' maneuvers 47
3.9 Visualization of the resolved flight trajectories. 47
3.10 Intra-ATCO consistency in merging waypoint selection 50

4.1 Concept diagram: ATCO conformal machine learning model 55
4.2 Visual representation of supervised machine learning framework 56
4.3 Chained-prediction architecture 58
4.4 A comparison of the original resolutions and the trajectories predicted through chained predictions 63
4.5 Estimation of the degree of conformance through measurement of similarity between the aircraft trajectories using Frechet distance. 66
4.6 Permutation feature importance analysis for the random forest models developed for ATCOs A, ATCO B and ATCO X. 68
4.7 Sensitivity analysis for the ML models developed for ATCO A, ATCO B and ATCO X, over increasing test set noise. 70
4.8 Behavior cloning with DRL: Concept diagram 72
4.9 ATCO A's CPA and CTD distribution 73
4.10 Model convergence: episode length and the average reward. 77
4.11 Comparison of the closest point of approach (CPA) distribution of the ATCO and the DRL policy. 78
4.12 Comparison of the cross-track deviation (CTD) distribution of the ATCO and the DRL policy. 78
4.13 DRL agent's maneuver direction for conflict point MUMSO. 79

5.1 Conflict resolution maneuver acceptance evaluation interface 85

5.2 Conflict resolution interactive interface 91

5.3 Conflict resolution preferences of the ATCOs. 94

6.1 Flow-centric operations: An illustration 100

6.2 A concept diagram for the interaction between the agent and the learning environment 102

6.3 Distribution of aircraft cruise speed from the ADS-B data. 103

6.4 An illustration of the conflict scenarios with two flows in a 90° crossing conflict. 104

6.5 Model convergence: episode length (left) and the average reward (right). . 110

6.6 Number of intra-flow actions at each time step. 111

6.7 Delay distribution in terms of the difference between the original flow plan and the actual flow plan. 112

6.8 The maximum cross-track deviation distribution (NM) for Flow 1 and Flow 2. 112

6.9 CPA distribution between Flow 1 and Flow 2. 113

6.10 Flow1 size and its variation with scenario evolution 113

List of Tables

2.1 Conflict resolution methodologies discussed in the literature review. 22

3.1 Resolution maneuver statistics (averaged values) with their standard deviation. 45

4.1 Description of the features extracted from ATCOs' data and their abbreviations used in the figures. 59
4.2 Range of hyperparameters used in grid search to find the optimal values. . 60
4.3 Performance of random forests (R.F), XGBoost, and SVM on the dataset of ATCOs *A*, *B* and X, averaged for comparison 61
4.4 Average values of performance metrics on train and test sets of the ATCOs *A* and *B* 62
4.5 Average values of performance metrics on train and test sets of the ATCO X. 64
4.6 Parameters for PPO training. 76

5.1 Conflict resolution maneuver selection percentage for ATCO *A* and ATCO *B*. 85
5.2 Overall analysis of ATCOs' maneuver selection 93
5.3 ATCOs' post-experiment feedback. 95

6.1 Encoding of the agent's action space. 106
6.2 Parameters for PPO training. 109

Institutional Review Board Approvals

The thesis documents the results from two human-in-the-loop experiments. These are included in **Chapter 3** and **Chapter 5** of the thesis. The experiments have been performed after obtaining approval from the Nanyang Technological University Institutional Review Board, with reference numbers **IRB-2020-11-030** and **IRB-2023-046**.

List Of Abbreviations

ADSB Automatic Dependent Surveillance Broadcast

AI Artificial Intelligence

ANSP Air Navigation Service Provider

ATCO Air traffic Control Officer

ATM Air Traffic Management

ATS Air Traffic Service

CPA Closest Point Of Approach

CR Conflict Resolution

CDR Conflict Detection and Resolution

CTD Cross-track Deviation

DRL Deep Reinforcement Learning

EASA European Union Aviation Safety Agency

FCO Flow-Centric Operations

FIR Flight Information Region

FRA Free Route Airspace

IATA International Air Transport Association

ICAO	International Civil Aviation Organisation
IFR	Instrument Flight Rules
LOS	Loss of Separation
MARL	Multi-Agent Reinforcement Learning
MIT	Maneuver Initiation Time
ML	Machine Learning
MTCD	Medium Term Conflict Detection
NM	Nautical Miles
SCTA	Short Term Conflict Alert
TCP	Trajectory Change Point

Chapter 1

Introduction

Safety and efficiency in present-day air traffic control are constrained by the air traffic controllers' cognitive capabilities and their ability to handle the increasing traffic. With the current air traffic control system already reaching its operational limits, ensuring safety while meeting the increasing future air traffic demands requires novel solution approaches. Numerous automation methods have been proposed to assist air traffic controllers in safety-critical operations such as air traffic conflict resolution. Nevertheless, there are limitations of such methods in terms of their acceptance by the air traffic controllers and adaptation to future concepts of operations.

This chapter presents an overview of air traffic conflict resolution, the present and future traffic trends, and establishes the motivation and research objective of the thesis. The key contributions and the structure of the thesis are presented at the end of the chapter.

1.1 Air Traffic Conflict Resolution

Air traffic conflict resolution is a dynamic, time-sensitive, and safety-critical task. Air traffic controllers (ATCOs) use different methods for aircraft conflict resolution, which are influenced by various factors. This section introduces the concept of air traffic conflict resolution, different methods used by ATCOs for conflict resolution, ATCOs' conflict resolution strategies, and the challenges associated with conflict resolution.

1.1.1 Background

The primary responsibility of the Air traffic controllers (ATCOs) is to ensure a safe and expeditious flow of traffic, with responsibilities across domains such as conflict detection, conflict resolution, air traffic flow management, and search and rescue operations. ATCOs are a critical and indispensable component of the air traffic control system because of their ability to access, analyze, integrate information, and make judgments and efficient decisions to ensure flight safety. While in the initial years of air traffic control, the ATCOs' responsibilities were limited to the departure and approach phases of flights, over the years that followed, their responsibilities have grown to cover all the segments of a flight, with a significant increase in their workload.

Ensuring safety in an airspace primarily comprises of two key components- conflict detection and conflict resolution. A conflict is an event where two or more aircraft experience a loss of minimum separation (LOS). For example, in a radar environment, a loss of separation occurs if the horizontal separation of 5 NM and the vertical separation of 1000 feet(ft) is breached [4, 5]. In such situations, ATCOs must take preventive actions to avoid any possibility of a loss of separation between the aircraft. Figure 1.1 shows the categorization of the safe separation standards along with their sub-categories. The requirements for a safe vertical separation depend on whether the airspace and the aircraft are approved to operate under reduced vertical separation minima (RVSM) or not. In an en-route RVSM airspace (above FL290 to FL410, inclusive), the vertical separation requirement is reduced from 2000 ft to 1000 ft.

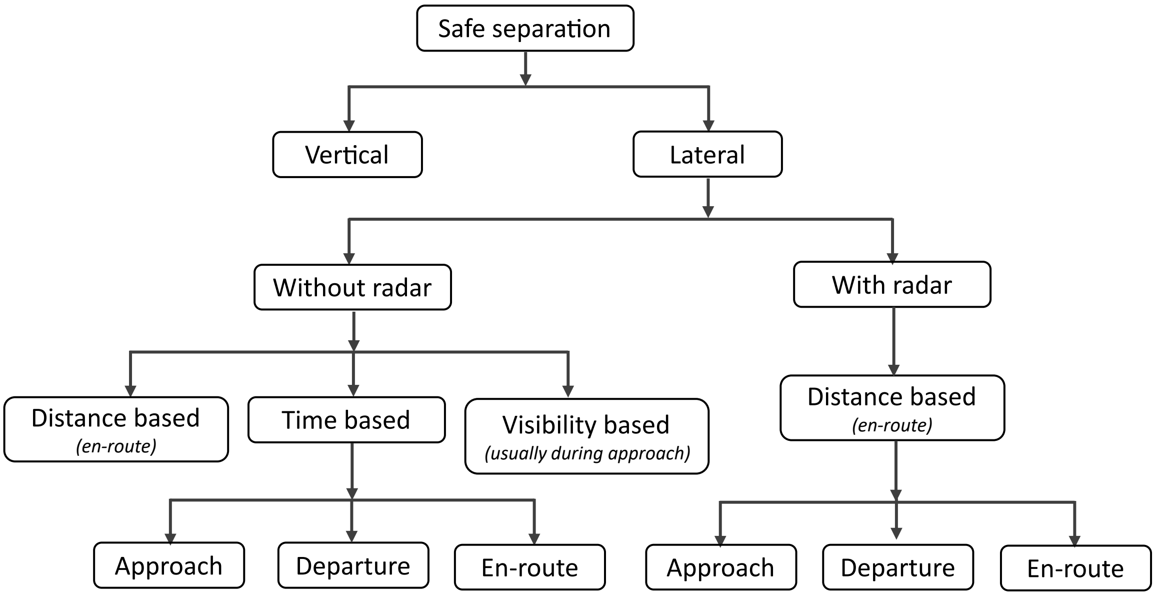

Figure 1.1: Modes to achieve safe separation in non-radar and radar airspace.

Lateral separation requirements vary based on procedural airspace (without radar) or controlled airspace (with radar). Procedural airspace includes areas where surveillance cover is not available, such as oceanic airspace or areas with sparse flight traffic. In procedural airspace, if the aircraft are at the same flight level, time-based or distance-based separation is utilized, for which the standard separation requirements are depicted in Figure 1.2, with reference to the ICAO document 4444: Procedures for Air Navigation Services [5].

In a controlled airspace, distance-based separation standards are employed. The minimum vertical safe separation between two aircraft is 1000 ft and the minimum lateral separation is 5 NM, as shown in Figure 1.3.

The ATCOs may employ different methods to achieve the required minimum separation, which include speed change, vectoring (heading change), flight level change, or a combination of these. Flight level change for achieving safe separation is usually employed in tactical scenarios, where an immediate resolution of a potential conflict is required. Nonetheless, vertical maneuvers cause more fuel burn and further, introduce an additional dimension of conflict detection for the maneuvered aircraft. Speed change is usually preferred when a potential conflict is sufficiently ahead in time and an immediate resolution is not required. This is because the effect of such a maneuver is fairly delayed.

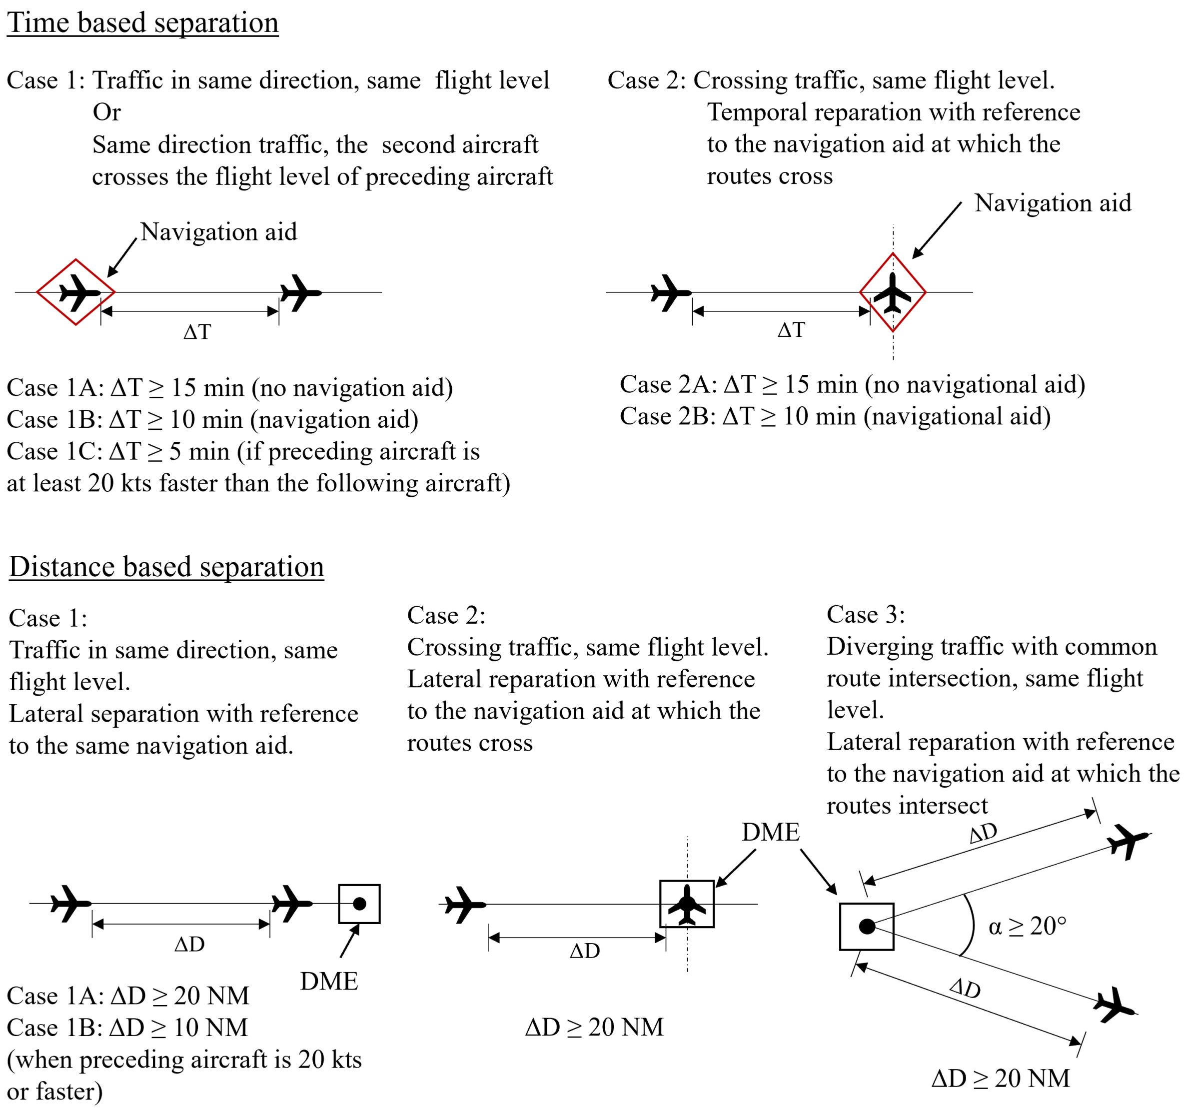

Time based separation

Case 1: Traffic in same direction, same flight level
Or
Same direction traffic, the second aircraft crosses the flight level of preceding aircraft

Case 1A: ΔT ≥ 15 min (no navigation aid)
Case 1B: ΔT ≥ 10 min (navigation aid)
Case 1C: ΔT ≥ 5 min (if preceding aircraft is at least 20 kts faster than the following aircraft)

Case 2: Crossing traffic, same flight level.
Temporal reparation with reference to the navigation aid at which the routes cross

Case 2A: ΔT ≥ 15 min (no navigational aid)
Case 2B: ΔT ≥ 10 min (navigational aid)

Distance based separation

Case 1:
Traffic in same direction, same flight level.
Lateral separation with reference to the same navigation aid.

Case 1A: ΔD ≥ 20 NM
Case 1B: ΔD ≥ 10 NM
(when preceding aircraft is 20 kts or faster)

Case 2:
Crossing traffic, same flight level.
Lateral reparation with reference to the navigation aid at which the routes cross

ΔD ≥ 20 NM

Case 3:
Diverging traffic with common route intersection, same flight level.
Lateral reparation with reference to the navigation aid at which the routes intersect

ΔD ≥ 20 NM

Figure 1.2: Time-based and distance-based separation requirement in procedural airspace(non-radar environment). Here, DME refers to the distance measuring equipment. α is the angle between the two intersecting flight paths.

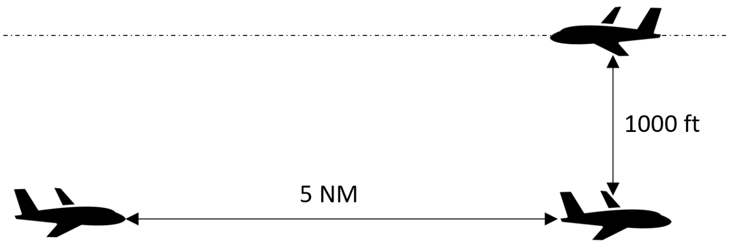

Figure 1.3: Vertical and lateral separation requirement in a controlled airspace(radar environment).

Most often, ATCOs prefer vectoring to achieve the desired safe separation between aircraft in head-on and crossing conflicts, and sequencing operations as well (Figure 1.4).

This is because vectoring is relatively fuel efficient, and does not disturb the vertically stratified structure of the airspace. It involves taking a sequence of actions, from selection of the aircraft to maneuver, initiation of a maneuver at a preferred time, deciding the magnitude and direction of the maneuver, and finally, reverting the aircraft to its initial path. This sequence of actions constitutes an *end-to-end* conflict resolution maneuver for the concerned aircraft. As in Figure 1.4, two common methods to resolve crossing conflict are either to maneuver one aircraft towards the other aircraft (Figure 1.4 (a.1)) or direct one aircraft to a subsequent waypoint (Figure 1.4 (a.2)) to achieve the desired separation, depending on the conflict configuration.

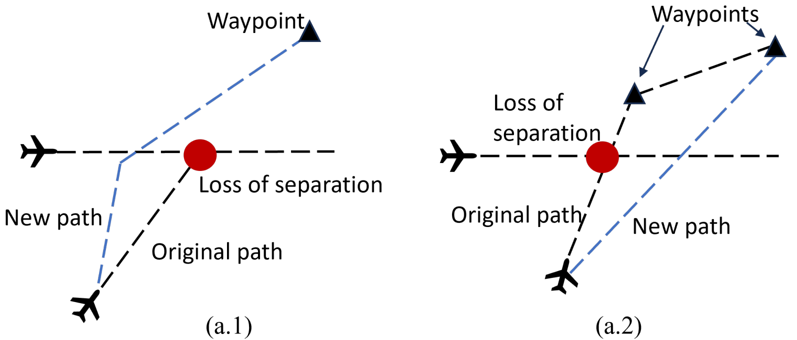

(a.1) (a.2)

(a) Crossing conflicts

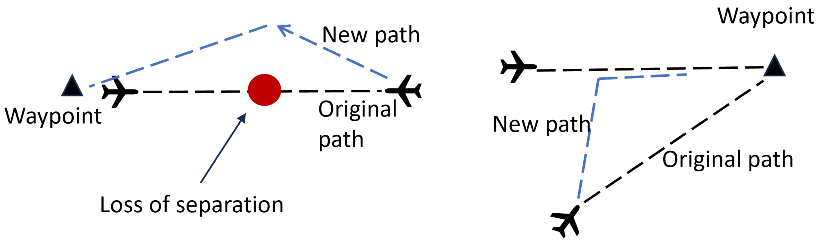

(b) Head-on conflict (c) Sequencing

Figure 1.4: Vectoring one of the two aircraft to avoid a loss of separation or sequence the aircraft.

1.1.2 Air Traffic Conflict Resolution Strategies

Researchers have defined a strategy as a high-level decision-making process, or a method or plan towards a goal [6]. A strategy can also be defined as a function that maps the states of a system to the corresponding actions, as shown in Figure 1.5. In other words, a function that governs the actions that should be taken in a particular state of the system. Strategies in air traffic control have been defined as a working method or a specific group of air traffic control activities that achieve one or more objectives such as safety, expeditiousness, and orderliness of the traffic within a certain investment of time [7]. They may be used for tasks such as situation assessment and perception, attention and workload management, solution aircraft conflict situations, planning, and taking decisions [8]. ATCOs' strategies reduce the likelihood of overall task performance being compromised. From the context of air traffic conflict resolution, an ATCO's conflict resolution strategy can be defined as a state-action mapping that is intrinsic to the ATCO, that the ATCO uses to avoid a potential conflict. Thus, the strategy of different ATCOs to resolve a given conflict may be different.

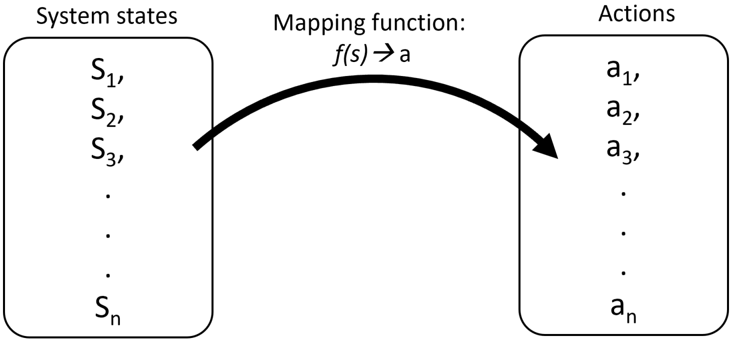

Figure 1.5: Representation of a strategy which is a mapping from the system states to corresponding actions.

These strategies exist at different levels of decision-making. For instance, the strategy to select the type of separation method (speed change, vectoring, flight level change) to avoid a conflict may depend on multiple influencing factors such as the distance of the aircraft from the conflict, the flight level of the aircraft, etc. In further detail, if vectoring is selected, the subsequent strategy of selecting the type of vectoring maneuver (*path stretch* (1.4.a.1) or *direct-to* 1.4.a.2)) may be affected by multiple other factors, which may include the leading and the trailing aircraft, speed of the aircraft, the flight paths of the aircraft etc.

Certain basic strategies are available at the ATCOs' disposal to ensure safe separation. For example, selecting the trailing aircraft to maneuver in case of same aircraft speeds, directing the maneuvering aircraft towards the other aircraft in case of crossing conflicts (Figure 1.4.a.1) and *direct-to* maneuvers for situations where these are advantageous (Figure 1.4.a.2). Over time, ATCOs develop certain inherent preferences in managing air traffic conflicts, which can be termed *conflict resolution strategies*. ATCOs use these strategies to manage air traffic while preserving their cognitive resources.

1.1.3 Associated Challenges

Managing air traffic conflicts presents ATCOs with various challenges, from handling increased air traffic to balancing cognitive resources. Understanding and addressing these challenges are vital for enhancing air traffic control efficiency and safety.

Increasing Air traffic

The civil aviation sector has witnessed phenomenal growth in the past decade. The pre-COVID-19 forecasts by the International Civil Aviation Organisation (ICAO) suggested that the global passenger traffic would double by the year 2035 if the annual growth rate of 4.3% was sustained [9]. In 2018 alone, the civil aviation sector served 8 billion people and operated over 90 million flights worldwide.

With clear indications of a strong recovery from COVID-19, the air traffic growth forecast by EUROCONTROL is shown in Figure 1.6. In the most likely base scenario, the year 2050 will witness 16 million flights in Europe, which is a growth of 44% from the year

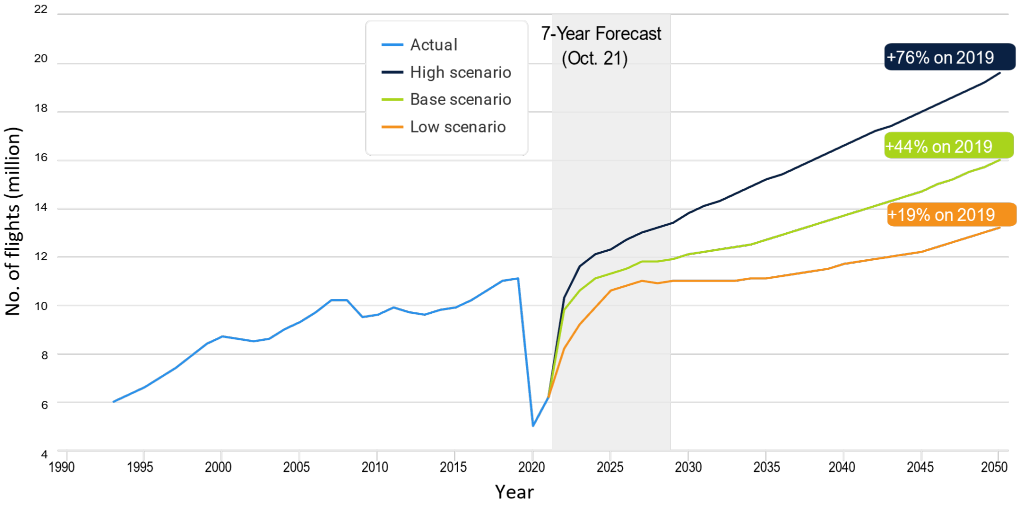

Figure 1.6: IFR movement forecast for Europe, with total growth between 2019 and 2050. [10]

2019, at an average growth of 1.2% per year [10]. Similarly, the global airline passengers grew by 21.8% post-COVID-19 and demonstrated significant growth in the year 2022 as well [11].

With the growth of air traffic, the present-day airspace is already reaching its operational limits. Direct effects of increased air traffic include airspace congestion leading to flight delays, negative economic effects, increased passenger inconvenience, and most significantly, reduced safety and efficiency due to increased ATCO workload. As shown in Figure 1.7, researchers have identified that with the increase in the number of aircraft in a sector, the perceived conflict rate as experienced by the ATCO increases quadratically [12]. Furthermore, with the increase in air traffic, the ability of ATCOs to detect

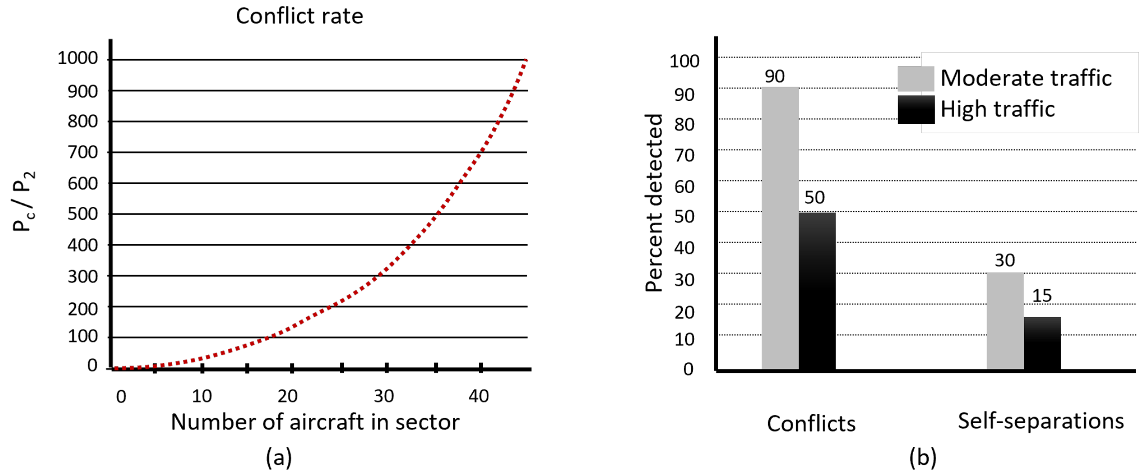

Figure 1.7: (a) Relation between the number of aircraft in a sector and the perceived conflict rate, from the perspective of the ATCOs. Here P_2 refers to the probability of two aircraft having a conflict flying a direct route and P_c is the global conflict probability (adopted from [12]). (b) Reduction in conflict detection and self separations detection with increasing traffic (adopted from [13]).

conflict drops by 40%, and detecting self-separation maneuvers by the aircraft decreases by 15% as the traffic increases [13]. Thus, while the ATCOs are a critical component of the air traffic control system, their cognitive capabilities are likely to remain the main functional limitation in safely meeting the future traffic demand. [14, 15].

Complexity of the conflict resolution task

Air traffic conflict resolution is a time-sensitive, dynamic, and safety-critical task that is influenced by factors such as the sector's characteristics (its geometry, no-go-zones, etc), surrounding traffic, location of the conflict, weather-related uncertainties, and human factors such as ATCOs in the current, and the preceding and subsequent sector, the conflict resolution strategies of the ATCOs, and their hand-over/take-over practices. Furthermore, due to a large action space available to resolve a conflict, different ATCOs may select different conflict resolution actions. As discussed earlier, air traffic conflict resolution is not a one-step process but rather a sequential decision-making problem, that is also influenced by the inherent stochasticity and uncertainty associated with the air traffic control system as a whole. Thus, a combination of multiple factors affects and governs the conflict resolution process, making it a significantly complex research problem.

Advanced concepts of operations

To ensure sustainable air traffic growth and meet the projected traffic demands, advanced concepts of operations such as free route airspace (FRA), flight-centric, and flow-centric operations (FCOs) are being investigated. These operational concepts are significantly different from the existing sector-based operations. For instance, in FRA the user may freely plan their route between a defined entry and exit point, subject to only a few limitations of avoiding no-go zones (danger areas, temporary reserved areas, etc) [16], in contrast to the current airspace where standard airways must be used. On similar lines, FCOs propose grouping the aircraft operating on multiple airways with distinct geographical characteristics, including flight trajectory orientation, proximity to their current geographic area, flights originating from and heading towards the same region, and flights proceeding to destinations within the same area/region [17]. Such novel operational concepts pose challenges such as the emergence of new and unexpected conflict birth points in the airspace, and difficulty in the identification and resolution of conflicts.

Human-centric control

Unlike other tasks such as conflict detection, for which there are advisory tools such as medium-term conflict detection (MTCD) and short-term conflict alert (STCA), ATCOs perform the task of conflict resolution by either vectoring (providing heading change), speed change, flight level change, or a combination of these with no decision support. The primary reason is that conflict resolution is a human-centric task wherein inputs from an ATCO in the form of resolution actions are required. As previously discussed, different ATCOs may take different conflict resolution actions based on their inherent strategies. Thus, air traffic conflict resolution is not just a mathematical problem but involves human-centric aspects as well, which the advisory tools based on computational methods do not capture.

Aviation sustainability

The aviation industry faces significant sustainability challenges, primarily driven by its environmental impact. Aviation contributes approximately 2-3% of global CO2 emissions [18], with forecasts suggesting a potential tripling by 2050 if current trends continue [19]. The reliance on fossil fuels and the emissions of carbon dioxide and nitrogen oxides present substantial hurdles to reducing the industry's carbon footprint. Technological advancements such as more efficient aircraft designs and alternative fuels like biofuels and synthetic fuels are essential but come with high development costs and significant implementation challenges. Furthermore, regulatory discrepancies across regions complicate the global standardization of sustainable practices. Addressing these challenges is crucial not only for environmental protection but also for ensuring the long-term viability of the industry in an increasingly eco-conscious global market.

The pursuit of sustainability in aviation is therefore not just an environmental imperative but also a strategic economic necessity. This challenge must be addressed with innovations for cleaner fuels, efficient and quieter aircraft, and smarter operations to improve the system efficiency.

1.2 Research Motivation

Conflict resolution is influenced by a multitude of environmental, technical, and human factors. As such, the motivation for this thesis lies in addressing the challenges associated with meeting the increasing air traffic demand, specifically from the perspective of the ATCOs in ensuring safe and efficient air traffic operations. The key motivating factors are as follows:

1.2.1 ATCO-Automation Conformance

Several automation methods have been proposed for fully autonomous air traffic conflict resolution as well as with an aim to assist the ATCOs in this task. These approaches encompass mathematical algorithms [20], modeling, and optimization approaches [21], to new perspectives using learning-based methods [1]. As discussed, there exists a large action space to resolve a given conflict and different ATCOs may perceive a conflict scenario differently. Over time, ATCOs develop certain inherent preferences in managing air traffic conflicts, which can be termed '*conflict resolution strategies*'. ATCOs use these strategies to manage air traffic while preserving their cognitive resources. The existing automation methods typically prioritize immediate conflict resolution through vectoring (horizontal, vertical, or both) and do not incorporate the conflict resolution strategies of the ATCOs. Furthermore, most of the existing methods do not provide an end-to-end conflict resolution, but rather the initial maneuver only. As a result, the low acceptance of such methods by the ATCOs can be attributed to a mismatch or a lack of conformance in how ATCOs perceive the conflict scenarios and how such automation tools provide the advisories [22], as depicted in Figure 1.8.

As shown in Figure 1.9, the aim should be to increase the existing overlap between the perceived solution space of the ATCOs and the proposed solutions by the automation methods, to increase their acceptance in operations.

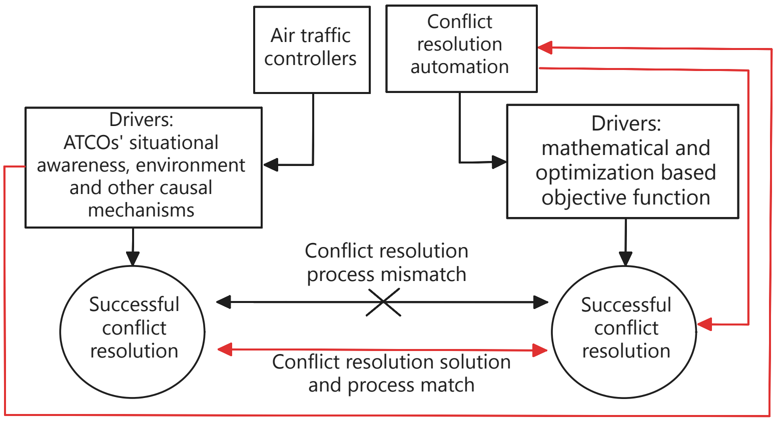

Figure 1.8: Graphical representation of the ATCO-automation solution mismatch. The key idea is that an ATCO's conflict resolution decision is governed by a multitude of factors, and thus the solutions proposed by the existing automation methods do not conform to the ATCOs' conflict resolution strategies.

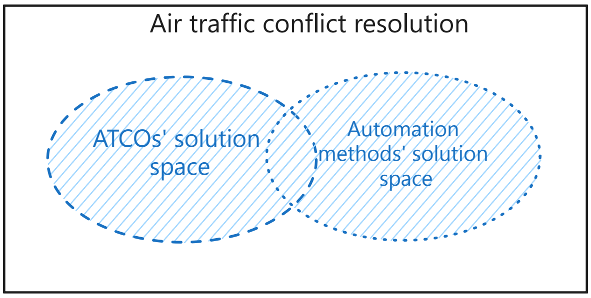

Figure 1.9: Graphical representation of the ATCO-automation solution space.

1.2.2 Scalability Constraints of the Airspace

In the existing airspace operations, the airspace is subdivided into smaller geographical regions called sectors. In day-to-day operations, the airspace is dynamically reconfigured according to the ATCOs' workload [2]. Under-loaded sectors are collapsed to form larger sectors, and overloaded sectors are split into smaller sectors, which are then operated by different ATCOs. This sectorization, especially splitting the sectors, becomes inefficient beyond a certain threshold, rendering sector sizes impractical for operational purposes. Further, considering the ICAO and EUROCONTROL forecasts, air traffic will continue to grow. Organizations such as the Single European Sky ATM Research (SESAR) in Europe and the Federal Aviation Administration (FAA) in the United States are working towards significantly increasing the sector capacities without adding in more air traffic controllers [23, 24]. In this regard, the development of automation methods to assist

ATCOs in their tasks is a promising solution to meeting the near future air traffic demands. Addressing the long-term future air traffic demand also requires a paradigm shift from the current concepts of operations. As such, alternative concepts, such as flight-centric and flow-centric operations (FCOs), are being explored. These concepts introduce novel challenges in adapting the existing conflict resolution methodologies to such emerging paradigms.

1.3 Research Objectives and Scope

With increasing air traffic, methods to assist ATCOs in operations such as air traffic conflict resolution are required. This is because air traffic conflict resolution is a safety-critical task and demands high cognitive effort. Consequently, the ATCOs' workload and cognitive capabilities are one of the factors that limit meeting future airspace demands. The existing automation methods proposed to assist ATCOs in cognitively demanding and safety-critical tasks such as air traffic conflict resolution hold significant limitations. These include the inability of such methods to incorporate ATCOs strategies in the proposed solutions which is the primary factor limiting their acceptance by the ATCOs. Further, the existing automation methods do not provide end-to-end conflict resolution advisories for aircraft conflicts.

Future concepts of operations such as the flow-centric paradigm offer to meet the long-term increasing air traffic demand. Currently, research on flow-centric operations and how automation assistance should assist the ATCOs, in still in infancy. The development of air traffic conflict resolution assistance methods may help the ATCOs in better management of airspace traffic in the current as well as future airspace operations. On this premise, this thesis proposes to leverage the experts' knowledge of operating in complex systems and use the advances in machine learning to develop methods for air traffic conflict resolution for the current and future concepts of operations. Specific research objectives include the following:

- Investigating methodologies for identifying ATCOs' conflict resolution strategies, developing effective representations of these strategies, and investigating the factors affecting the identified strategies.

- Investigating and developing robust machine learning models that can efficiently learn the ATCOs' strategies.
- Investigating, conceptualizing, and proposing innovative models specifically tailored for air traffic conflict resolution in a flow-centric airspace.

This thesis focuses on the cruise flight phase of an aircraft and the ATCOs' conflict resolution strategies in this flight phase. A cruise flight phase refers to the fight segment between the top of the climb and the top of the descent of an aircraft. The time horizon involves a pre-tactical time frame where an immediate, reactive conflict resolution action by the ATCO is not required. As such there is sufficient take for the ATCOs to analyze the situation, evaluate the available solution space, and opt for a '*strategic conflict resolution*' maneuver. Researchers have also described strategic conflict resolution as one where efficiency of the maneuver is also a consideration along with safety as a constraint [25,26]. The action space considered in this thesis is focused on the lateral direction due to the following reasons. Lateral maneuvers are more appropriate in the context of strategic conflict resolution [26, 27], because such maneuvers cause significantly less discomfort to the passengers, cause less fuel burn, and do not distort the vertically stratified structure of the airspace. Furthermore, while adding a vertical dimension to the action space allows for an increased variety of maneuvers, this extra dimension also increases the complexity of an optimal route calculation. Further, the research on flow-centric operations considers lateral and speed change-based action space.

1.4 Thesis Contributions

This thesis presents a data-driven approach to identify the ATCOs' conflict resolution strategies, develops machine learning-based ATCO-conformal models using the ATCOs' conflict resolution data, and proposes a two-stage conflict resolution approach for the flow-centric concept of operations where the traffic is modeled as intersecting flows.

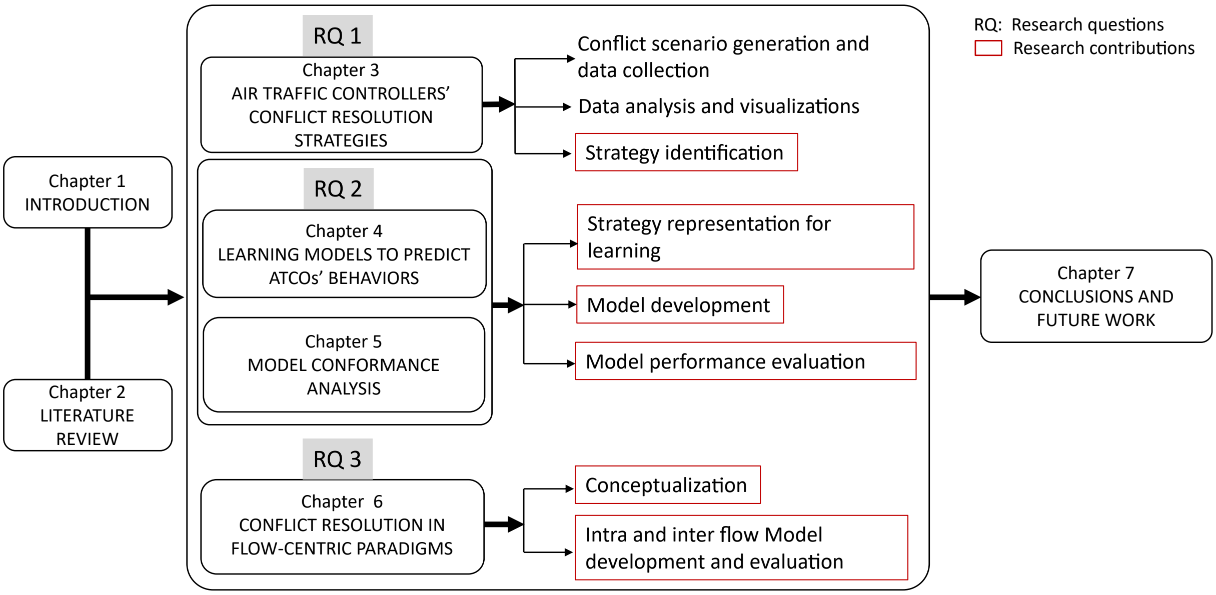

Figure 1.10: Overview of the thesis structure with contents covered in each chapter. Research Questions (RQ) and the key contributions are indicated.

A graphical representation of the thesis layout is provided in Figure 1.10 with the chapters that address the corresponding research questions and the research contributions highlighted in red. In concrete, this thesis makes three novel contributions which are discussed in the following subsections.

1.4.1 Identification of ATCOs' Conflict Resolution Strategies

The first contribution develops a methodology to identify the conflict resolution strategies of the ATCOs. To investigate the ATCOs' conflict resolution strategies and the factors affecting them, high-quality data that is representative of the ATCOs' behaviors is required. It has been identified in the literature that the likelihood of accurately measuring ATCOs' performance increases with the extent to which one is able to place the ATCOs in a standardized and realistic environment [28]. Therefore, a data collection experiment is designed with the objective of collecting ATCOs' conflict resolution maneuvers for the corresponding conflict scenarios. In this experiment, conflict scenarios are simulated in a high-fidelity simulation environment that emulates an air traffic control radar interface. Eight ATCOs are involved in the experiments. The participants are shown real-time air traffic conflicts in the simulation environment representing a sector of the Singapore flight information region (FIR) and their interactions with the simulation environment to resolve the conflicts are recorded. Individual analysis and comparison of each ATCO's data demonstrated that ATCOs have distinct conflict resolution strategies, specifically in terms of maneuver direction and the magnitude of the heading deviations. The metrics used to identify these strategies are the selection of the aircraft to be maneuvered, maneuver initiation time, maneuver direction, magnitude of deviations, and preferred safe separations between the conflicting aircraft. It is also observed that the ATCOs' conflict resolution strategies are influenced by factors such as the proximity of the conflict point from the sector boundary and trailing and the leading aircraft involved in the conflict.

1.4.2 Machine Learning Models for ATCO-Conformal Conflict Resolution with Behavior Cloning

The second contribution involves the development of machine learning models for ATCO-conformal conflict resolution using behavior cloning. In this context, 'conformal' denotes the resolutions aligning with the conflict resolution strategies employed by ATCOs. This research frames conflict resolution as a sequential decision-making task and employs a series of supervised machine learning models, including regressors and classifiers, to map environmental states to specific low-level actions, effectively emulating the behavior of ATCOs. Leveraging previously identified ATCO strategies, personalized (matching individual ATCO strategies) and group conformal (aligning with the strategies of a group of ATCO) models are constructed using data collected during conflict resolution experiments with ATCOs.

Results from predictions on the test dataset indicate that the proposed models can generate ATCO-conformal conflict resolutions, in line with ATCOs' strategies. The trained models are subjected to testing with added Gaussian noise to assess their robustness. The results demonstrate that the models are robust up to 7.5% added noise, characterized by low mean absolute errors and high prediction accuracy. Additionally, the research is further evaluated through human-in-loop experiments involving professional ATCOs to analyze their choices in conflict resolution strategy selection. Results highlight that

ATCOs opted for solutions aligning with their conflict resolution strategies for over 70% of scenarios. Furthermore, the collected conflict resolution data is utilized to train a reinforcement learning agent to conform to ATCO conflict resolution maneuvers. To achieve this, ATCOs' data is used as a reference to guide the maneuvers of the agent. The agent is able to propose resolution maneuvers with comparable distributions for the closest point of approach (CPA) and cross-track deviations, with safe separation for 99.8% of the scenarios.

1.4.3 Air Traffic Conflict Resolution for Flow-Centric Operations

The third contribution pertains to the evolving operational concepts centered around flows and the formulation of a two-stage model for resolving air traffic conflicts within such a framework. This study delves into scenarios involving crossing conflicts amid uncertainties inherent in a flow-centric paradigm. The investigation employs a Markov Decision Process (MDP) framework to formulate the research as a sequential decision-making problem. Each flow, comprising a varying number of aircraft aircraft with different cruise speeds and associated location uncertainties, is represented as a self-stabilizing graph structure. The two-stage conflict resolution process encompasses training a conflict resolution policy to ensure safe separation between flows and utilizing a self-stabilizing graph structure for ensuring safe separation within each flow. The trained policy guarantees both intra and inter-flow safe separation for 100% of the scenarios. Performance evaluation includes an analysis of maneuver efficiency and deviations from flight plans. The absolute delays for the flow stand at 2.53 and 9.49 minutes, respectively, falling within acceptable limits considering uncertainties and dynamic characteristics of flow size, speed, and temporal evolution

1.5 Thesis Organization

The subsequent chapters of this thesis are organized as follows.

- Chapter 2 provides a detailed review of the literature pertaining to the aforementioned contributions and highlights the research questions addressed in the thesis.
- Chapter 3 presents a methodology for identifying ATCOs' conflict resolution strategies. It discusses the details of the data generation process, participant demographics, simulation environment and the experiments, and the details of the identified ATCOs' conflict resolution strategies.
- Chapter 4 presents the work on the development and performance evaluation of machine learning models to predict the ATCOs' conflict resolution strategies.
- Chapter 5 presents the details of two human-in-the-loop (HITL) experiments that were conducted to evaluate the consistency of the ATCOs in selecting their conflict resolution strategies and acceptance of the machine learning predictions by the ATCOs.
- Chapter 6 presents the research on air traffic conflict resolution in flow-centric airspace, with a two-stage conflict resolution process.
- Chapter 7 presents the conclusion of the thesis, discusses the challenges in the integration of ML and AI-based models in air traffic control, and the open research directions for future work.

Chapter 2

Literature Review

The present air traffic control relies on ATCOs to ensure aircraft safety. Several methods have been proposed to automate the conflict resolution task with the aim of alleviating the ATCOs' workload. Initial work in this direction was based on mathematical models, exact optimization methods, and heuristics. Subsequently, with the advent of machine learning and the availability of data, research on data-driven approaches for air traffic conflict resolution has gained momentum. Methods to incorporate the ATCOs' conflict resolution strategies in the automation methods have also been proposed. Nonetheless, the existing methods have limitations that limit their feasibility and operational use. This chapter provides a detailed review of the existing work on air traffic conflict resolution methods, with their limitations, and the identified research questions.

2.1 Air traffic conflict resolution Methods: An overview

Air traffic conflict resolution is an indispensable component in all airspace architectures - sectored, free-route, and flow-centric (refer Figure 2.1). As discussed in Chapter 1, for radar operations, the minimum lateral separation between two aircraft must be 5 NM, and based on the reduced vertical minimum separation (RVSM) regulations, the minimum vertical separation between two aircraft must be 1000 ft. In the event of an identified potential LOS, ATCOs must take a sequence of actions to maneuver one or more of the concerned aircraft to ensure safe separation. For radar-based en-route operations, speed change, flight level change, vectoring, or a combination of these may be used to achieve the desired separation. Nonetheless, vectoring (lateral maneuvers) is preferred for conflicts with longer looks ahead (pre-tactical or strategic conflict resolution, as defined in Chapter 1.3.) This is because such maneuvers cause significantly less discomfort to the passengers and do not distort the vertically stratified structure of the airspace. Furthermore, ATCOs usually do not prefer to change the flight level of the maneuvering aircraft because this adds an additional dimension to the conflict detection process and increases the complexity of an optimal route calculation. Lateral maneuvers are also efficient in terms of fuel burn and emissions when compared to vertical maneuvers.

Researchers have proposed numerous computational and learning-based methods for air traffic conflict resolution in sectored airspace. Computational methods consist of approaches such as exact optimization and heuristics, while learning-based methods use either the experts' knowledge or an empirical approach that involves self-exploration and trial and error to learn from the collected data. Both supervised learning and reinforcement learning have been used in the literature for learning-based conflict resolution methods. On similar lines, there has been significant research progress on conflict resolution in free-route airspace. As of now, no significant research has been performed on conflict resolution for flow-centric airspace.

Figure 2.1 provides a graphical representation of the above discussion and highlights the research focus of the research in this thesis. It involves using ATCOs' knowledge to

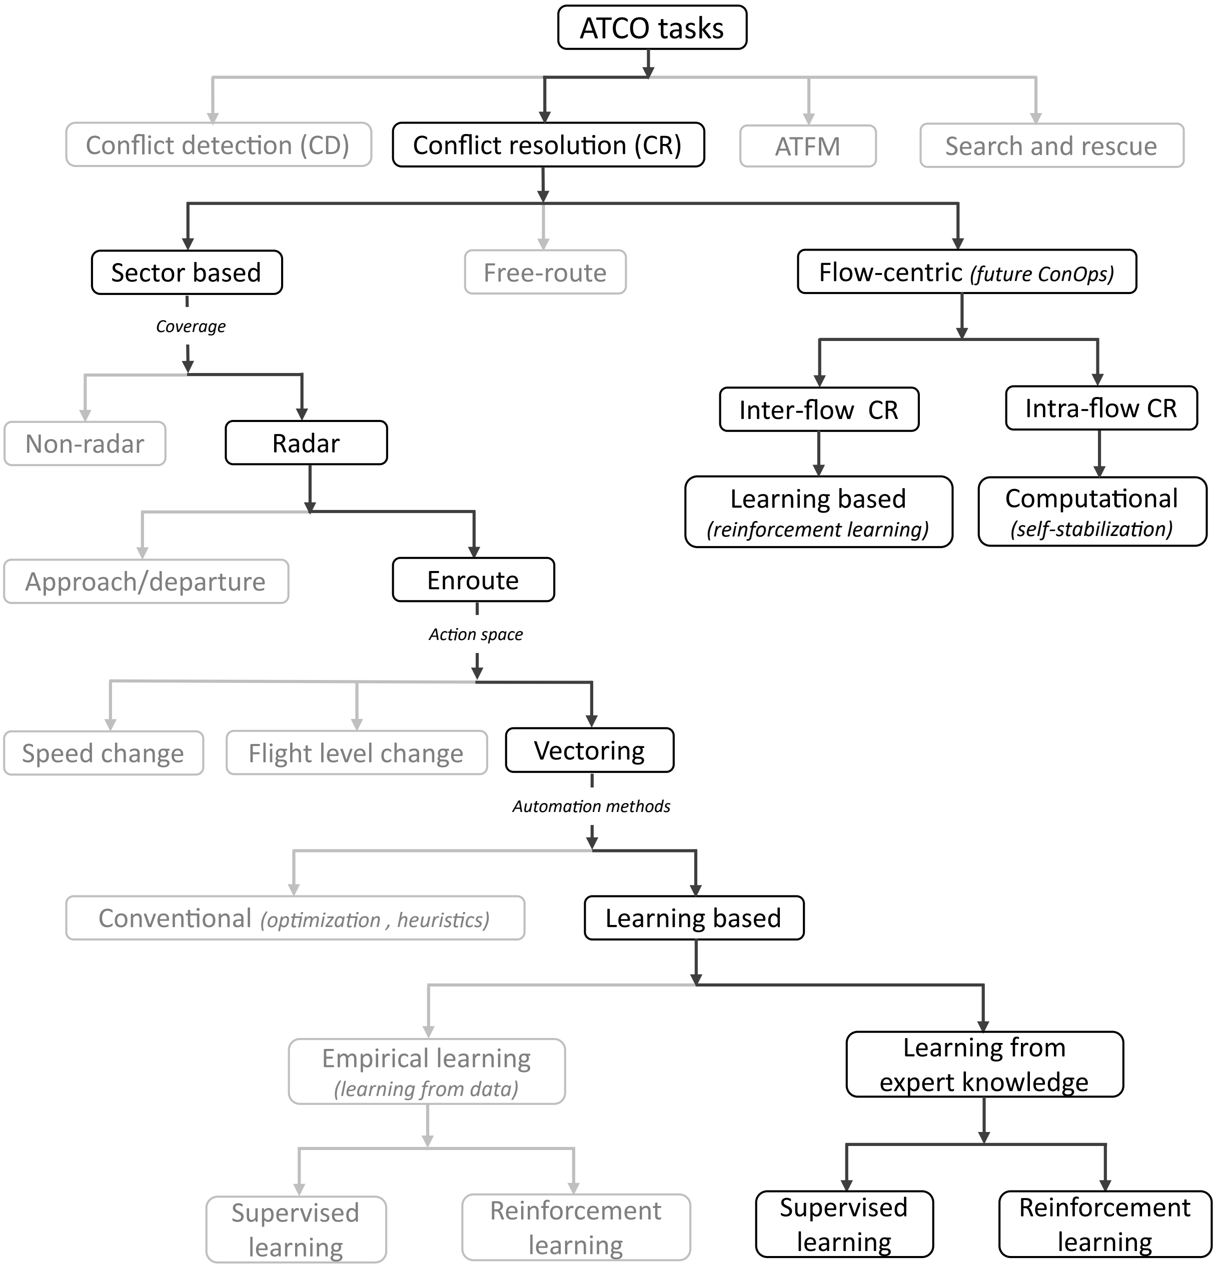

Figure 2.1: Graphical representation of ATCO tasks and their sub-categories, along with the corresponding automation methods.

identify their conflict resolution strategies, developing ATCO-conformal machine learning models using behavior cloning, and the development of a conflict resolution model for flow-centric operations where the traffic is modeled as intersecting flows. The literature review in the subsequent sections will focus on the computational and learning-based methods for air traffic conflict resolution, and the concept of air traffic conflict resolution in flow-centric operations. Subsequently, the associated limitations that led to the corresponding research questions for the thesis will be highlighted. Table 2.1 presents an overview of the methodologies discussed in the literature review along with the authors.

Table 2.1: Conflict resolution methodologies discussed in the literature review.

Methodology	Authors
Optimization and heuristics	Durand et al. [20] Durand et al. [21] Durand et al. [29] Bilimoria et al. [30] Hoekstra et al. [31] Kosecka et al . [32] Zaghal et al. [33] Christodoulou et al. [34] Delahaye et al. [35] Feron et al. [36] Pallotinno et al. [37] Alonso-Ayuso et al. [38] Allignol et al. [39] Allignol et al. [40] Stollenwerk et al. [41] Durand et al [42]
Learning based: ATCO action identification and prediction	Rantenan and Wickens [43] Shin et al. [44] Kim et al [45] Olive et al. [46] Malakis et al. [47] Dang et al. [48] Liang [49] Pham et al. [1]
Learning based: Reinforcement learning	Pham et al. [50] Brittain and Wei [51] Hochreiter and Schmidhuber [52] Brittain and Wei [53] Brittain and Wei [54] Mollinga and hoof [55] Ribeiro et al. [56] Zhao and Liu [57] Ghosh et al. [58] Mukherjee et al. [59]
Learning based: From expert demonstrations	Westin et al. [22] Karikawa et al. [60] Xiaotian and Zhang [61] Kirwan et al. [62] Nunes and Mogford [63] Corver and Grote [64] Bekier et al. [65] Bekier et al. [66] Lee et al. [67] Westin et al. [68] Klien et al. [69] Svensson et al. [70] Westin et al. [71] Westin et al. [72] Westin et al. [73] Rooijen et al. [74] Regtuit et al. [75] Tran et al. [2]

Methodology	Authors
	Goh et al. [76]
	Durand et al. [77]
FLow-centric operations	Delahaye et al. [78]
	Ma et al. [3]
	SESAR [17]

2.2 Air Traffic Conflict Resolution Using Conventional Methods

Computational methods for air traffic conflict resolution can be categorized based on centralized or decentralized approaches, the type of method used (prescribed, force field, or optimized), the maneuver dimensions used (speed, lateral, vertical, or combined), the time horizon (tactical or strategic), and the management of conflicts (pairwise or global). In a centralized approach, just like the current air traffic control operations, a single entity or agent is responsible for controlling all the aircraft. On the other hand, a decentralized approach is more applicable for free route operations [42]. Prescribed methods fix the resolution maneuver during system design based on a set of predefined procedures. These include the work of Bilimoria et al [30] who analyzed 3 types of maneuvers: airspeed change, heading change, and altitude change for conflict resolution in free flight operations. Force field approaches consider each aircraft as a charged particle and use modified electrostatic equations to generate conflict resolution maneuvers, wherein the repulsive force between the aircraft is used to define the maneuver used to avoid a conflict. These include prominent works of Hoekstra et al. [31] and Kosecka et al. [32] for FRA, and Zeghal [33], who propose a decentralized force field based approach for conflict resolution with cooperative intruders. In terms of the management of conflicts, a pairwise approach aims to solve multiple conflicts sequentially. In other words, if a conflict resolution induces a new conflict, the original solution can be modified until a conflict-free solution is found. On the other hand, a global approach considers all the potential conflicts (primary and secondary) simultaneously to identify a resolution maneuver.

The majority of the work has been focused on optimization methods (exact methods and heuristics) with horizontal or a combination of horizontal and vertical maneuvers for air traffic conflict resolution. Exact methods include approaches such as mixed integer linear programming (MILP) [36,37] with a CPLEX solver to achieve global optimality of the solution, and mixed integer non-linear programming approaches as well [34]. However, such exact methods require long computational time which makes them impractical for operational applications. On these lines, heuristics such as Evolutionary Algorithms [29,35], Ant Colony Optimization [21] and Variable Neighborhood Search [38] have been commonly used to achieve good, but not necessarily optimal solutions for the conflict resolution problem. In other works, Allignol et al. [39] propose a heuristic combining Genetic Algorithm with Tabu Search to allow for 3D conflict resolution. Stollenwerk et al. [41] propose the use of quantum heuristics to solve a simplified version of the conflict-resolution problem, focusing on wind-optimal trajectories and minimizing trajectory adjustments. An extensive review of computational methods for air traffic conflict resolution has been provided by Kuchar and Yang [79].

While such computational methods have demonstrated suitability for the air traffic conflict resolution problem to an extent, such methods for conflict resolution suffer from notable limitations. Firstly, they require complete knowledge of the mapping between conflict scenarios and appropriate maneuvers, making them highly complex and less effective in situations with high uncertainty where comprehensive knowledge of environmental factors cannot be achieved. Secondly, such methods heavily rely on standardized input scenarios and do not possess the ability to adapt and learn when confronted with novel or non-standard situations. High computational time is also a limitation when considering operational implications. While these methods exhibit competence in controlled environments, their adaptability to the inherent uncertainties and dynamic nature of real-world airspace operations prove to be a formidable challenge. The inherent uncertainties and stochasticity associated with the air traffic control environment thus, make it a challenging problem for such methods.

2.3 Air traffic Conflict Resolution Using Learning Based Methods

In recent years, the approach to air traffic conflict resolution has undergone a transformative shift from the traditional mathematical and optimization-based approaches towards the application of machine learning methods. Multiple factors such as the increased availability of data, improved computational power and parallel processing, improved and robust algorithms, interdisciplinary collaboration between aviation and machine learning experts, and an increased human-centric approach to model development have led to this transformation.

Different machine learning methodologies such as supervised learning, reinforcement learning, and even unsupervised learning have been used for conflict resolution and identifying patterns in the aircraft trajectory data to detect the trajectory deviations that might be potentially due to a conflict resolution maneuver. The following subsections discuss the learning-based methods in detail.

2.3.1 Conventional Machine Learning Approaches

The conventional approaches that use machine learning for air traffic conflict resolution can be divided into two categories: first, the identification and prediction of conflict resolution actions using supervised, unsupervised learning and other data-driven methodologies, and second, the use of reinforcement learning to train single or multiple agents for conflict resolution.

ATCOs' Action Identification and Prediction

Researchers have attempted to identify ATCO actions and the reasons underlying these actions, and attempted to group the identified actions on the basis of different similarity metrics. In order to identify the factors affecting ATCO conflict resolution actions, Rantanen and Wickens [43] analyzed 256 cases of conflict resolution from radar track data of five U.S ATC centers. They identified expediency, preservation of airspace structure, and visualizations as factors affecting ATCOs' conflict avoidance actions. To understand the types of resolutions performed by ATCOs to maintain aircraft separation standards, Shin et al. [44] proposed an algorithm to detect off-track flights and then identify the resolution types. It consisted of two sub-components - conformance monitoring to detect off-track flights and identify those due to aircraft conflict resolutions among them, and an inference algorithm to identify the conflict resolution type. The data were simulated using the Airspace Concept Evaluation System (ACES). In another research, Kim et al. [45] used the ACES data and attempted to classify conflict resolution maneuvers into 9 classes, based on resolution maneuvers used by ATCOs in the past using supervised learning algorithms (Neural networks and support vector machines). This was based on a list of 16 features that they used to train their models. Researchers have also used machine learning algorithms such as auto-encoders with the historical Mode-S data [46] to infer ATCOs' actions, tree-based models such as decision trees on simulated data in terminal approach radar environment to classify the air traffic scenarios to understand the ATCOs' response to such scenarios [47], and the use of isolation forests and density-based clustering on ADS-B data to detect aircraft maneuvers [48]. Man Liang [49] proposed a similarity measure-based approach to identify the similarity between the target conflict case and the source cases, to identify the most suitable resolution for the target case. In another work, Pham et al [80] have used historic ADS-B data to train a tree-based ensemble learning model to predict the ATCOs' actions given a list of input features extracted from the data.

While the supervised machine learning models trained using historical data to predict ATCOs' actions show high prediction accuracy as in [45, 80], there are two primary limitations of such methods, especially from the scope of conflict resolution. First, in an operational setting, the ATCOs may maneuver an aircraft due to a multitude of reasons which include expediting the aircraft travel time to manage the air traffic flow, avoiding no-go zones such as temporary military use airspace, avoiding weather cells, maneuvering the aircraft to avoid a LOS or in some cases, maneuvers based on pilots' request for climb or descend. Thus, it is difficult to estimate which maneuvers were performed for the purpose of conflict resolution. Also, air traffic conflicts are a rare event and scarce in the historical data. Second, a typical day in air traffic control consists of multiple ATCOs working in shifts, with usually no more than 2 hours on one position at a time without a break. In other words, historical data such as the ADS-B is a collection of actions of multiple ATCOs. Thus, it is not possible to identify how different ATCOs perform the conflict resolution task.

Reinforcement Learning for Conflict Resolution

Researchers have also explored the use of reinforcement learning (RL) for air traffic conflict resolution. RL is based on the fundamental idea of Markov Decision Process (MDP), which is a mathematical framework used to model sequential decision-making tasks. The main components of an MDP include states, actions, transition probabilities, rewards, and the discount factor. In most cases, the transition probabilities or the environment dynamics are difficult to estimate. A simple representation of the agent and its environment is shown in Figure 2.2. At each time step, an agent receives a representation of the environment and takes an action. Based on the action taken, the state of the environment changes and the agent receives a corresponding reward. The agent now finds itself in a new state, S_{t+1}. The agent's objective is to learn a policy, i.e. a mapping from states to the corresponding actions (in other words, a strategy for selecting actions), that maximizes the expected cumulative reward over time [81]. All the work on RL is based on this agent-environment interaction. Unlike supervised learning methods that rely on historical data to predict ATCO actions, reinforcement learning agents learn optimal actions through trial and error in simulated environments. Furthermore, in contrast

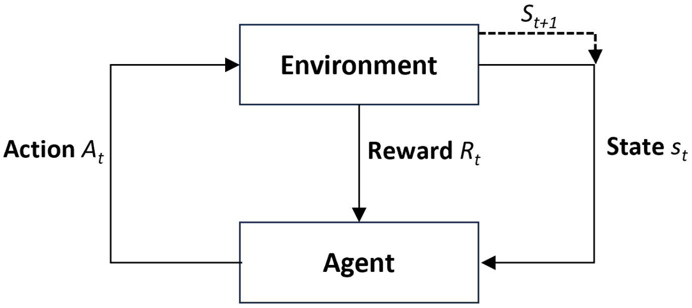

Figure 2.2: Agent-environment interaction in a Markov decision process.

to the rule-based or optimization algorithms, RL agents continuously refine their policies based on the feedback received during simulated interactions. This adaptability makes them promising candidates for addressing the inherent uncertainties and complexity of air traffic conflict scenarios.

One of the initial research for conflict resolution using RL was performed by Pham et al. [1] who used a Deep Deterministic Policy Gradient (DDPG) algorithm for conflict resolution with varying degrees of uncertainty. Figure 2.3 shows the underlying idea of reinforcement learning, as visualized by [1]. In the works that have followed, Brittain et al. have proposed multiple methods for conflict resolution in dense air traffic scenarios with uncertainties. The research includes a multi-agent reinforcement learning approach for centralized learning and a decentralized execution approach using PPO for safe-separation [51], use of Long Short-Term Memory Network (LSTM) [52] to encode a variable number of aircraft into a fixed length vector for conflict resolution [54], and the use of attention mechanism [82] to allow the RL agents to access variable aircraft information in the sector [53]. Mollinga and van Hoof [55] introduced an RL algorithm capable of steering an arbitrary number of aircraft through unstructured three-dimensional airspace every 5 seconds while avoiding conflicts and collisions. In other works, Ribeiro et al. [56] use a combination of a geometric approach (Modified Voltage Potential) and deep reinforcement learning with DDPG to achieve aircraft safe separation in a hybrid environment with

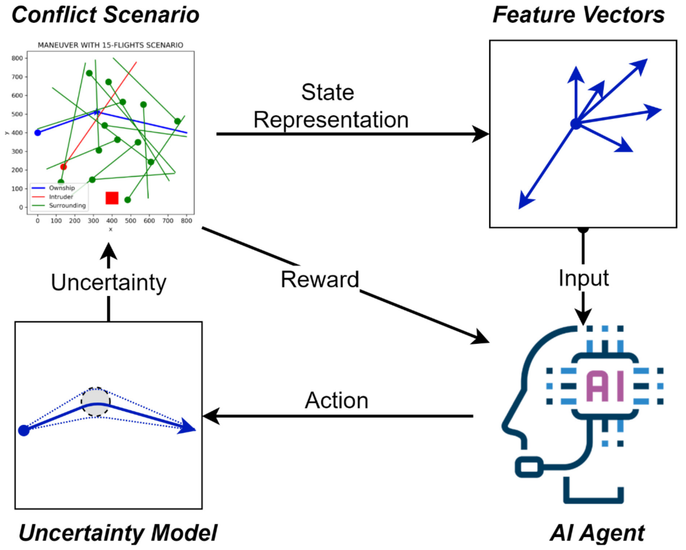

Figure 2.3: Interaction between the learning environment and the AI agent, Pham et al. [1].

manned and unmanned aircraft. Zhao and Liu [57] integrate the information of a conflict scenario into a solution space diagram (SSD) which is then passed to a convolutional neural network (CNN) with an actor-critic network to learn a conflict resolution policy. On similar lines, Mukherjee et al. [?] use a CNN-based Deep Q-Network to achieve aircraft safe separation under aircraft position uncertainty using heading maneuvers with 100% safe separation.

Although the current reinforcement learning (RL) and deep reinforcement learning (DRL) methods have gained attention in terms of their ability to better address complex decision-making problems in air traffic control, and provide optimal solutions, such projects are related to fully autonomous air traffic control systems and do not incorporate the ATCOs' conflict resolution strategies in their resolution advisories. Furthermore, the presented works incorporate actions taken at each time step, and in some cases after a certain number of simulation steps [51]. This is also a limitation from the operational perspective where the number of interactions of the ATCOs with the aircraft should be limited, due to the already high workload of the ATCOs.

2.3.2 Learning from Expert's Demonstrations

Automation tools that align with the behavior and preferences of the ATCOs, or are 'conformal' to their preferences, have a higher likelihood of being accepted by the ATCOs. Strategic conformance, in this context (Figure 2.4), refers to the degree to which the solution provided by the automation and the problem-solving approach matches that of the human operators [22]. This concept shares similarities with behavior cloning (BC), a method that seeks to learn to perform tasks based on an expert's demonstrations [83,84].

In general, a strategy is a high-level decision-making process or the art of employing a method or a plan towards a goal [6]. Autonomous agents' strategies in games have been discussed in many benchmark papers like playing ATARI with DRL [85] and Alpha GO [86], where the agents opt for strategies that maximize the overall reward, even accepting penalties with a potential of higher delayed rewards. In contrast to such games, strategies in air traffic conflict resolution are strikingly different, where the immediate tasks of conflict resolution must carry the highest reward. Within the scope of air traffic conflict resolution, a strategy is a high-level decision-making process to avoid a potential conflict. As discussed earlier, ATCOs' strategies may also be defined as a sophisticated planning skill, which is an essential element of the ATCOs' skills allowing them to handle a large amount of traffic while reserving their cognitive resources [60, 61].

EUROCONTROL's Conflict Resolution Assistant Project (CORA) [62] was one of the initial attempts to identify formal and informal rules and principles used by ATCOs during conflict resolution. The research used a standardized set of 12 static (image-based) conflict scenarios shown to controllers from seven countries in the European Union. Nunes and Mogford [63] also documented an interesting research on identifying controller strategies for conflict detection based on the mental representation of the current traffic scenario. According to their hypothesis, ATCOs build a comprehensive representation of the existing traffic which helps them to ensure that the aircraft are safely separated as they travel in the concerned airspace. The ability of the ATCO to ensure aircraft safety is dependent on the maintenance of this representation, which they called the 'picture'. There have been research attempts to identify strategies used by ATCOs' in arrival sequencing and coping with uncertainties in air traffic control [60, 64].

The extent to which automation should be involved to ensure operational acceptability by human operators in air traffic control has also been discussed in the literature. [65–67]. Bekier et al. [66] revealed a 'tipping point' in automation acceptance or rejection, as a point where the level of automation shifted the decision-making away from the ATCOs. They also claimed that automation methods may be more successful in tasks related to efficiency improvements and as support systems. The Multidimensional Framework for Advanced SESAR Automation (MUFASA) project hypothesized that the ATCOs would more readily accept the automation tool's advice if the tool behaved like the ATCOs [68], as demonstrated in Figure 2.4.

Researchers have discussed the main challenges such as interpretability, collaboration, goal negotiation, mutual performance monitoring etc for the automation to be a 'team player' in joint human-agent activities [69, 70]. On these lines, the concept of strate-

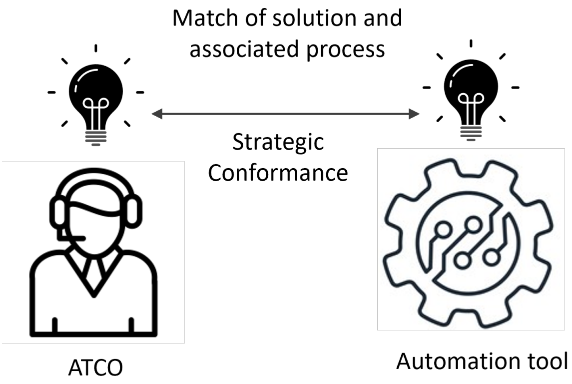

Figure 2.4: The underlying idea of strategic conformance where both, the solution and the process used to achieve the solution by the automation tool matches that of the ATCO.

gic conformance has been discussed as a measure to estimate how close the automation method's problem-solving style is to that of a human operator, as depicted in Figure 2.4.

Westin et al. [22] have argued that strategic conformance can not only improve the initial acceptance of the automation tools but also improve the overall system performance since the operators are more likely to use such tools in operations. Researchers have shown that ATCOs commonly use decision-making heuristics in air traffic conflict detection and resolution (CDR). Subsequently, heuristic form of automation, which is common in situations where operators' behavior is homogeneous, has been used in the past to develop automation methods in air traffic conflict detection and resolution (CDR) [27, 87, 88]. These research attempts were based on the hypothesis that ATCOs develop and maintain conflict resolution heuristics in a 'mental library' and scan through this library for a suitable solution for detecting a conflict [89]. One of the fundamental limitations of the heuristic approach is that in case of heterogeneity in ATCOs' behavior and conflict resolution preferences, it becomes difficult to identify a single strategy that applies to all [90].

This paved the way for a relatively newer approach in strategic conformal air traffic conflict resolution called individual sensitive automation, which basically implies that the conflict resolution advisories are personalized according to the individual ATCOs' strategies (refer Figure 2.6). Westin et al. [68, 73] have demonstrated the benefits of conformal resolution advisories in terms of increased acceptance by the ATCOs and have explored factors such as consistency, source bias, and transparency on automation acceptance. Re-

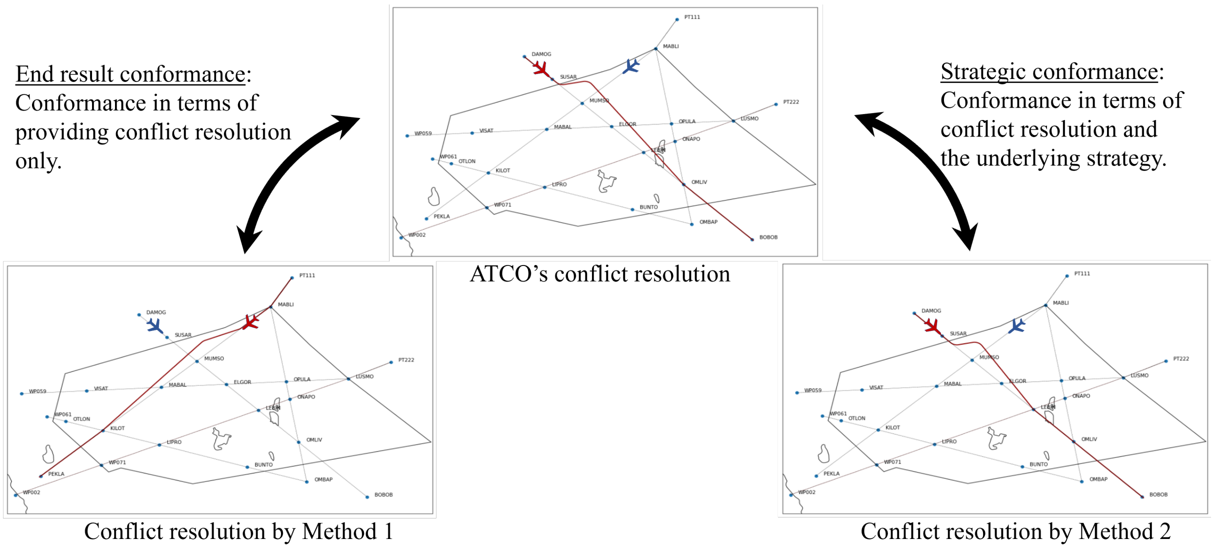

Figure 2.5: An ATCO-conformal conflict resolution advisory generated by Method 2 that matches the strategy of the ATCO.

searchers have also discussed the degree of acceptance of personalized models (individual sensitive advisories), group-conformal models (trained on combined data of ATCOs), and transparent conflict resolution advisories by the ATCOs [72].

As mentioned earlier, this concept is similar to behavior cloning or learning from the demonstrations of experts, in this case, the ATCOs. The general idea behind behavior cloning (BC) is that if a human can perform a task, rather than asking the human to explain how it is performed, the human is asked to perform it. Machine learning is then used to produce a symbolic description of the skill [91]. This symbolic description refers to a mapping between the state of the task environment and the corresponding action that should be taken. In one of the first works on behavior cloning, Sammut et al. [92] demonstrated the feasibility of controlling a system by watching a human operator at work. A flight simulator program was modified to log the actions of a human subject while flying an aircraft. These logs were used to create input features for the program and a set of 4 decision trees was used to predict actions to control the elevator, rollers, thrust, and flaps, based on the attributes of the training examples. In other words, the work aimed at developing a mapping of the respective states to the low-level actions of the operator. Along similar lines, K-nearest neighbors (KNN) has been used to train robots to avoid obstacles and perform desired tasks based on human inputs [93]. Researchers have also used probabilistic approaches such as Gaussian mixture models [94] and Bayesian networks [95] for mappings of the states to low-level actions to achieve behavior cloning. BC is associated with sequential decisions that are based on the state of the system [96]. Therefore, it is a suitable approach for air traffic conflict resolution, where the ATCOs must make a series of decisions to ensure safe separation between aircraft in a dynamic environment.

For air traffic conflict resolution, supervised learning and reinforcement learning approaches have been used with the aim of learning and replicating ATCOs' conflict resolution strategies [2, 74, 75]. Regtuit et al. [75] proposed conformal automation for air traffic control by learning conflict resolution maneuvers through operator demonstrations with reinforcement learning (RL) by creating a reward function to replicate the strategies. The scenarios do not include any variation and it appears to be an overfit of the RL model on 3 scenarios. The scenarios are two aircraft conflicts with conflict angles of 45°, 90°, and 135°, and the ATCO can only maneuver a pre-decided aircraft, which limits the scope of strategy analysis. Rooijen et al. [74] used a convolutional neural network architecture (CNN) that utilized solution space diagram (SSD) images to learn the ATCOs' resolution maneuvers. In one of the recent works on behavior cloning in air traffic conflict resolution, researchers have attempted to estimate the trajectory change point (TCP) of the maneuvered aircraft based on 500 data samples collected from the ATCOs [2]. While this approach accurately reflects the concept of behavior cloning, only the TCP is predicted by the agent, in an abstract simulation environment. Researchers have also used error-related potential generated during electroencephalography (EEG) experiments with ATCOs, to capture the ATCOs' preferences for conflict resolution [76].

Overall, the current research on understanding and incorporating ATCOs' strategies in automation methods for conflict resolution does not provide end-to-end conflict resolution maneuver predictions and does not capture the sequential decision-making characteristic of the conflict resolution task. Furthermore, these works use simulation environments that are different (low and medium fidelity) from the actual radar interface used by the ATCOs in operations. Researchers have highlighted this as a limitation in capturing

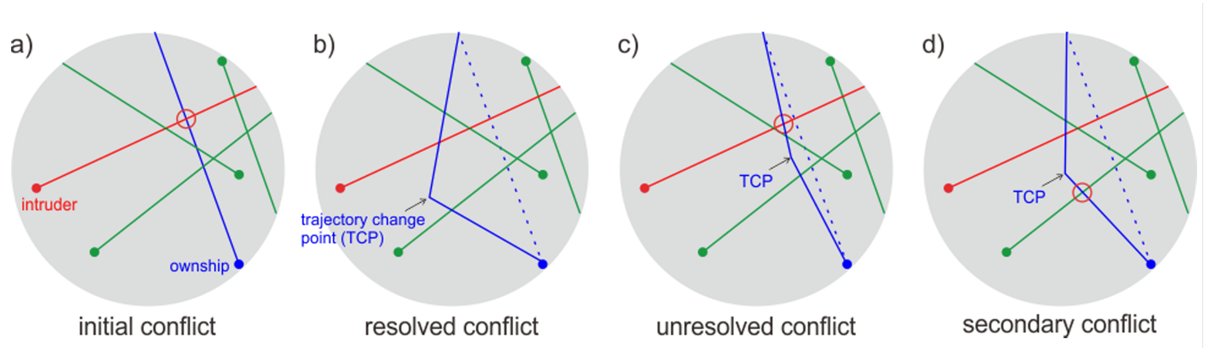

Figure 2.6: The interactive conflict resolution interface developed by Tran et al. [2] to collect the ATCOs' preferences for the trajectory change point of the maneuvered aircraft. The conflict scenarios were generated in a circular, abstract environment as shown.

the ATCOs' maneuver preferences [77], and discussed that the likelihood of accurately measuring controller performance increases with the extent to which one is able to place controllers in a standardized and realistic environment [28].

2.4 Air Traffic Conflict Resolution in Flow-Centric Operations

The existing sector-based air traffic control system faces serious limitations in terms of its inherent scalability constraints that hinder meeting future air traffic demands and ensuring sustainable growth [97]. In routine operations, the airspace is dynamically reconfigured according to the controllers' workload [98], wherein under-loaded sectors are collapsed to form larger sectors and overloaded sectors are split into several smaller sectors. These reconfigured sectors are then operated by different air traffic controllers (ATCOs). This sectorization, especially splitting the sectors, becomes inefficient beyond a certain threshold, rendering sector sizes impractical for operational purposes.

Hence, there is a need for a fundamental change in current operational concepts to effectively plan for and accommodate the anticipated growth in air traffic. Emerging paradigms, such as flight-centric and flow-centric operations (FCOs), are under consideration. While flight-centric operations are associated with a free route airspace, FCOs are proposed as a hybrid of flight-centric and structured airspace operations. In this concept, a traffic flow refers to a group of aircraft operating on multiple airways with distinct geo-

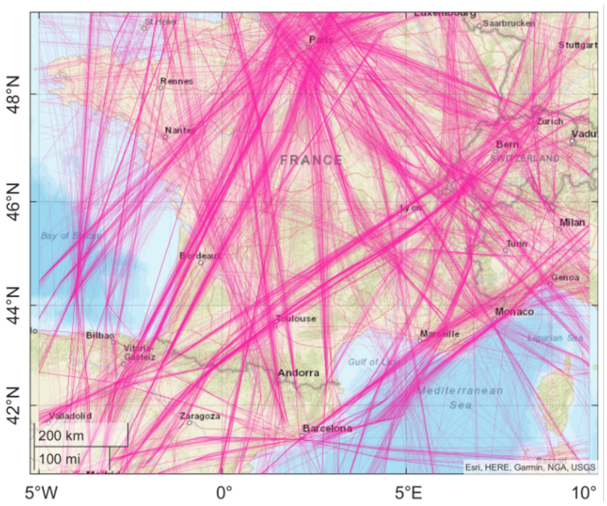

Figure 2.7: Major flows identified in the current sectored airspace over the French airspace. The flight trajectories are clustered as flows connecting the major hub airports [3].

graphical characteristics, including flight trajectory orientation, proximity to their current geographic area, flights originating from and heading towards the same region, and flights proceeding to destinations within the same area/region. [17]. FCOs focus on managing air traffic from an aggregated perspective, centered around the formation and evolution of major traffic flows rather than individual flights. This aggregated traffic management positions FCOs as an evolution of trajectory-based operations (TBOs), which prioritize flight-centric operations to address capacity demand imbalances. The primary goals of FCOs are to optimize regional traffic, distribute workload efficiently across ATC units, whether under or over-loaded, and overcome the scalability constraints inherent in current geographic sectors. Consequently, this approach shifts the responsibility of ATCOs from managing all traffic within a given sector to overseeing a specified number of aircraft throughout their flight segment within an airspace. The research on FCOs is in its infancy with recent works on traffic flow coordination at major flow intersections [3], and flow prediction in FCOs using the ADS-B data of the French airspace [78]. These works are based on the air traffic flows connecting major hub airports and based on the current sectored airspace, as shown in Figure 2.7.

Though FCOs hold significant potential as the future concept of operations given the rate of traffic increase, they induce significant challenges in terms of adapting safety-critical tasks such as conflict resolution to these new paradigms. This is because in

FCOs safe separation must be ensured within the flows(between the flows) as well as intra-flow (within each flow). Moreover, a cooperative strategy among ATCOs overseeing distinct flows must be maintained consistently. In addressing conflict resolution, essential considerations include the nature of maneuvers, their frequency, and the exchange of information, encompassing both intra-flow and inter-flow perspectives. As of now, there are no existing methods for air traffic conflict resolution in FCOs.

2.5 Research Gaps

It is evident from the literature that numerous air traffic conflict resolution methods have been proposed for autonomous operations and as assistance tools for ATCOs. The proposed methods use exact optimization and heuristics, supervised learning, unsupervised learning, and reinforcement learning. While the performance of these methods is promising, there are fundamental limitations of such methods which are summarized as follows: For computational methods that use optimization and heuristics, a complete knowledge of the mapping between conflict scenarios and appropriate maneuvers is required, making them complex and less effective in capturing the uncertainty and stochasticity associated with air traffic control. Additionally, they struggle with adaptability to novel scenarios and have high computational demands, posing challenges in real-world airspace operations with inherent uncertainties and dynamics.

Learning ATCOs' actions from historical data poses two limitations. Firstly, in real-world operations, ATCOs execute aircraft maneuvers for diverse purposes, complicating the isolation of maneuvers specifically intended for conflict resolution. Secondly, due to the operational nature of air traffic control, the identification of distinct conflict resolution strategies employed by individual ATCOs from the historical data is difficult. Reinforcement learning offers another solution for air traffic conflict resolution in scenarios with high uncertainty and traffic density (similar to the characteristics of air traffic control). Nonetheless, most of the work is inclined towards fully autonomous operations. Another limitation is the integration of such RL-based approaches into real-world operations where limited interaction between the ATCO and the maneuvering aircraft would be preferred due to the already high workload of the ATCOs.

Current research on incorporating ATCOs' strategies into conflict resolution automation lacks the ability to predict complete end-to-end maneuvers and capture the sequential decision-making characteristics of the conflict resolution task. Additionally, these studies use simulation environments of varying fidelity, deviating from ATCOs' actual radar interfaces. This has been recognized as a limitation in capturing ATCOs' strategies and there is a need for standardized, realistic environments to accurately capture these strategies. Further, the existing literature lacks research on air traffic conflict resolution in flow-centric paradigms.

Based on the research gaps in the literature, the following research questions have been identified, which this thesis aims to answer.

Research Question 1

What is a suitable representation of the ATCOs' conflict resolution strategies? How can we identify the ATCOs' strategies based on this representation and analyze the factors affecting the strategies?

Research Question 2

What type of machine learning models are suitable for generating ATCO-conformal conflict resolution advisories?

Research Question 3

How can air traffic conflicts be resolved in a flow-centric paradigm where traffic is modeled as intersecting flows?

Chapter 3

Air Traffic Conflict Resolution Strategies: Identification Methodology and Analysis

It is evident from the literature review that the existing methods for air traffic conflict resolution automation have inherent limitations in terms of incorporating the ATCOs' strategies in the proposed solutions. The first challenge in achieving conformity of automation methods in performing experts' tasks such as air traffic conflict resolution is the identification of the ATCOs' conflict resolution strategies. In order to identify the ATCOs' strategies and the factors affecting these strategies, human-in-loop experiments must be performed to collect the relevant data from the ATCOs in a high-fidelity simulation environment. The collected data then needs to be processed and analyzed to answer the research question.

This chapter provides details regarding the human-in-loop experiments, the simulation environment, participant demographics, and the identified strategies. Further, factors affecting the identified strategies based on data analysis and participants' feedback are analyzed.

3.1 Overview

Figure 3.1 illustrates the process used to collect the conflict resolution data and identify the ATCOs' conflict resolution strategies. A conflict scenario generator was developed which generated conflicting flight plans in Sector 6 of the Singapore Flight Information Region (FIR). These generated flight plans served as input for the ESCAPE Light simulation platform, a high-fidelity environment Experienced ATCOs were invited to participate in the conflict resolution exercise where they were presented with the conflict scenarios in the simulation environment. Subsequently, the conflict resolution data of the ATCOs was systematically collected and subjected to analysis to ascertain the strategies employed in conflict resolution. Additionally, ATCOs' feedback on the factors that influenced their decision was also collected. The subsequent sections delve into the specifics of each step involved in identifying and analyzing these conflict resolution strategies.

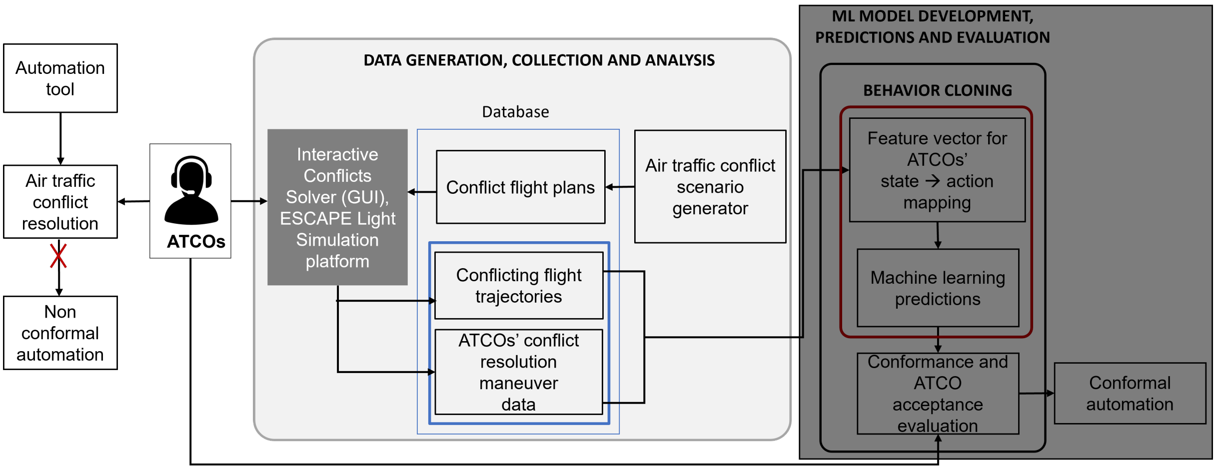

Figure 3.1: The overall concept diagram highlights the process of data generation, collection, and analysis.

3.2 Conflict Resolution Maneuver Representation

A conflict resolution maneuver involves a series of sequential decisions. These decisions encompass selecting the aircraft to maneuver (C), determining the maneuver initiation time (T), maneuver direction (D), cross-track deviation (D_c), and choosing the merging waypoint (M_{wp}). In this context, farther and closer are the two aircraft in conflict, T_{sim}

is the current simulation time, T_{los} represents the simulation time at the loss of separation and D_{bound} indicates the distance of the maneuvered aircraft from sector boundary. Farther aircraft refers to the aircraft that reaches the conflict waypoint later. On similar lines, the closer aircraft researches the conflict point sooner. We represent an end-to-end conflict resolution, P, as a sequenced tuple of the governing elements, such that $P = (C, T, D, D_c, M_{wp})$. Each element in P takes value from the feasible actions available in its domain. In particular,

$$C \subset \{farther, closer\} \tag{3.1}$$

$$T \subset \{T_{sim}, ..., T_{los}\} \tag{3.2}$$

$$D \subset \{left, right\} \tag{3.3}$$

$$D_c \subset \{0, ..., D_{bound}\} \tag{3.4}$$

$$M_{wp} \subset \{\textit{list of waypoints available after } T\} \tag{3.5}$$

In the discussed representation of an end-to-end conflict resolution maneuver, while the direction of the maneuver (left or right) is represented as D, the magnitude of this component is the heading angle. An ATCO's conflict resolution strategies can be identified by analyzing the data distributions of these actions. This is because a strategy, which is defined as a high-level decision-making process, is a generalization of the actions taken during a conflict resolution maneuver. Based on the sequence of actions discussed, different conflict resolution strategies may be possible, depending on the preference of the ATCOs and a combination of other external factors. For instance, ATCOs' preferences in the selection of the farther or the leading aircraft *(C)* for the maneuver. This selection may also be affected by the aircraft that enters the concerned sector earlier or later. ATCOs' strategies may also differ in terms of when *(T)* they prefer to start the maneuver of the selected aircraft i.e. early maneuver, or a relatively delayed maneuver. This may subsequently affect the direction (D) and magnitude of maneuver (heading deviation) required to ensure safe separation. Subsequently, the decision on the preferred cross-track deviation (D_c) (lower or higher) may be influenced by the ATCOs' preferences on the

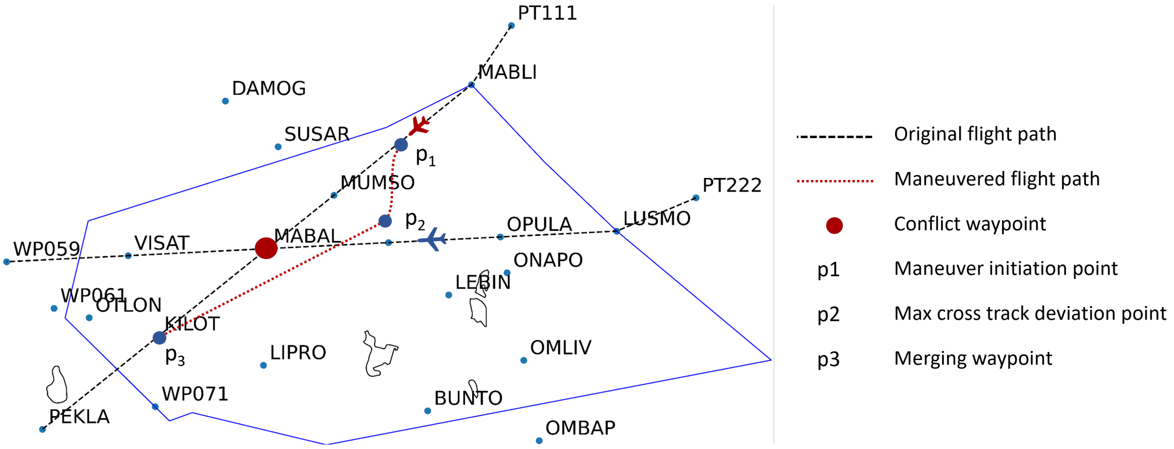

Figure 3.2: Representation of an end-to-end conflict resolution maneuver for the maneuvered aircraft, for a pair of aircraft bound to witness a loss of separation. For a conflict resolution maneuver, one aircraft initiates a maneuver at point p_1 (maneuver initiation point). This aircraft is deviated to point p_2 (point of maximum cross-track deviation) and then merged into its initial path at point p_3(merging waypoint). In our work, ATCOs provide conflict resolution maneuvers based on such path stretch vectors.

level of safe separation, orderliness, and efficiency of the traffic flow. The direction and magnitude of these deviations may also influenced by the proximity of the conflict to the sector boundary. Finally, the ATCOs may have different preferences in terms of where they merge the maneuvered aircraft into its original flight path (M_{wp}).

3.3 Experiment Design

3.3.1 Simulation Environment

To accurately observe the actual behavior of ATCOs in conflict situations, it is necessary to conduct experiments in simulation environments that can replicate the ATCOs' radar interface. Previous studies have shown that presenting conflict scenarios to ATCOs in environments that differ from their actual radar interface tends to be ineffective in capturing their true resolution preferences [77]. It has also been identified that standardized and realistic environments may be able to capture the ATCOs' performance better [28]. Therefore, in this research, we utilized the ESCAPE Light Simulation Platform, provided by EUROCONTROL, to perform experiments. This platform was chosen as it offers a suitable environment for replicating the ATCOs' radar interface and conducting realistic simulations [99]. ESCAPE Light is a high-fidelity, real-time, simulation environment. Along with key features such as en-route and terminal airspace design, data input and extraction, and Base of Aircraft Data information, ESCAPE Light provides a visual interface similar to the ATCO radar display (Figure 3.3).

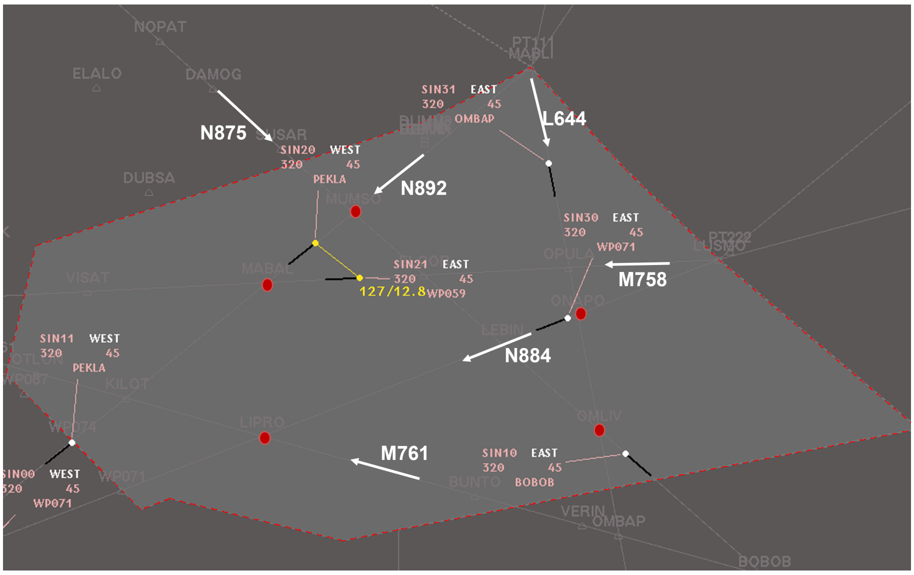

Figure 3.3: A snapshot of ESCAPE Light's ATCO radar screen interface with Sector 6 of the Singapore FIR. The red markers indicate the conflict points in the sector. The conflict points are the waypoints common to the intersecting airways. The aircraft in conflict have yellow icons. The white arrows show the direction of the aircraft on the airways along with the airway names.

3.3.2 Conflict Scenarios and Data Generation

The experiments to collect ATCO's conflict resolution data were performed in Sector 6 of the Singapore FIR. This is because the airway structure of Sector 6 enables a detailed analysis of crossing conflicts as compared to other sectors. Co-ordinate information of the sector and the waypoints were input to the simulator and conflicting flight routes were identified (refer to Figure 3.3). The following conflict angles and airways were used in the experiments based on the airway design: 36.2°(N892 and M758), 36.6°(N884 and M761), 37.29°(N875 and L644), 79.82°(L644 and N884) and 99.94°(N875 and N892). The aircraft appeared approximately 20 NM outside the sector boundaries so that the ATCOs develop situational awareness about the aircraft entering their sector. Time offsets were added to ensure that the aircraft experienced a loss of separation on the identified routes. Further, the start time of the aircraft was perturbed by adding a randomized noise in the form of time so that though the conflicts were ensured, they were not exactly the same. The flight plans were then uploaded into the simulation environment and the ATCOs' inputs in the form of conflict resolutions were collected. Figure 3.4 shows the conflict scenario data generation process.

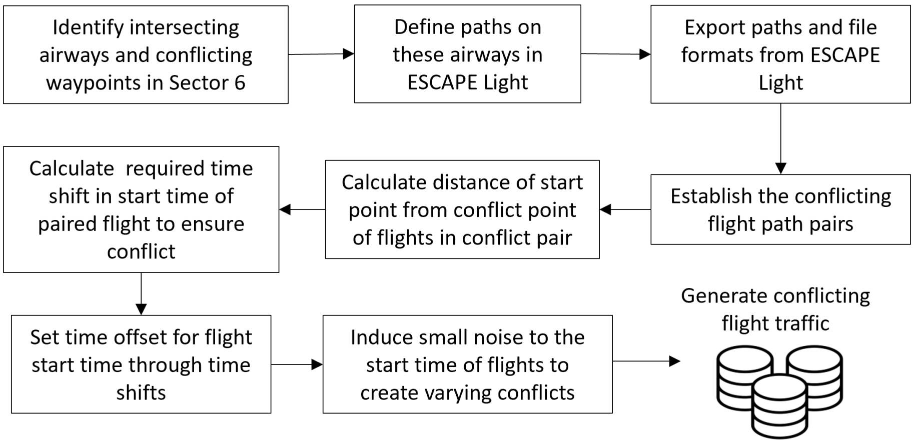

Figure 3.4: Conflicting flight plan data generation pipeline which was used to generate conflict flight trajectories in the simulation environment.

The interaction between ATCOs and the simulation environment was facilitated through mouse clicks. In this setup, the simulation environment recorded the positions of aircraft at regular 5-second intervals. This allowed the system to effectively log the updated locations of aircraft, reflecting the instructions provided by ATCOs for conflict resolution.

3.3.3 Evaluation metrics

After collection of the conflict resolution trajectories from each ATCO, these trajectories were compared to the original unmaneuvered trajectories. This comparison at each time step of the simulation gave a quantitative measure of the metrics such as maneuver initiation time (T), direction and magnitude of the maneuvered aircraft (D) i.e the heading angle, the cross-track deviation (D_c), and the merging waypoint (M_{wp}). The safe separation achieved between the two conflicting aircraft was also extracted. These metrics were used to identify the patterns and the strategies from the ATCOs' data. Figure 3.5 shows the feature extraction process used to obtain the aforementioned evaluation metrics, that constitute an end-to-end conflict resolution maneuver.

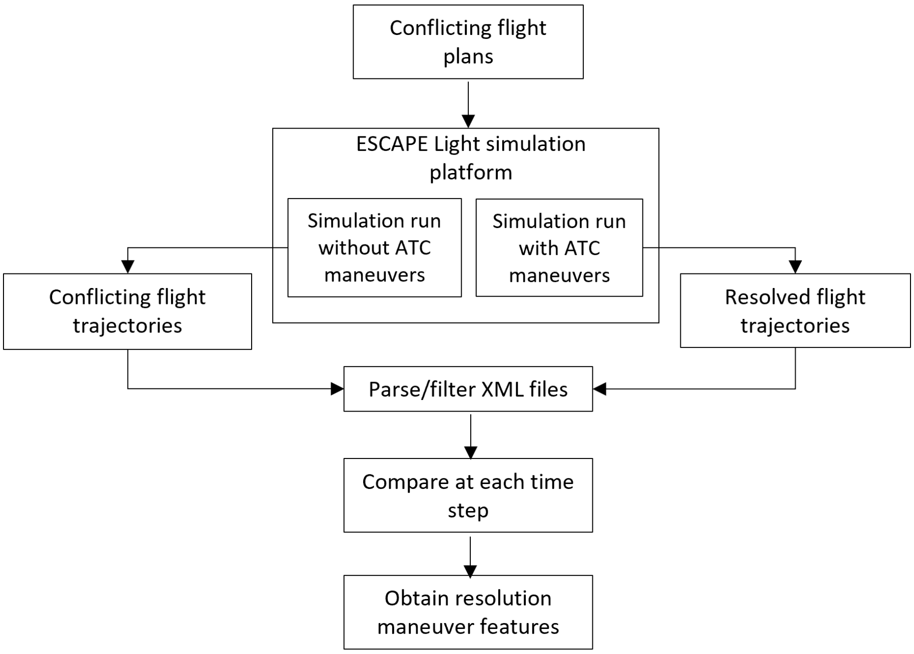

Figure 3.5: Comparing maneuvered flight trajectories with the unmaneuvered flight trajectories to extract conflict resolution maneuver features i.e. the components of the end-to-end conflict resolution maneuver.

3.3.4 Participant Demographics

A total of 8 experienced ATCOs participated in the experiments to collect the conflict resolution data. The experiments were performed after obtaining approval from the Institutional Review Board (IRB) with the approval reference number: IRB-2020-11-030. The participants were divided into two groups. One group was familiar with the airspace, while the other group was unfamiliar with the airspace and the traffic flows within it. The ATCOs' data collection varied according to their availability to participate in the experiments. The simulation run-time was 2 hours. As a consequence, each ATCO was engaged for multiple sessions to collect the data. Before beginning the experiments, the ATCOs were briefed about the traffic flows and the airspace structure. They were also given time to familiarize themselves with the simulation environment. In the experiments, the same conflict scenario dataset was used for all the ATCOs, to enable a comparison between their conflict resolution maneuver preferences (refer to Figure 3.6). For example, the 30 conflict scenarios presented to ATCO H were also presented to all other ATCOs while ATCO A was presented with a total of 612 unique conflict scenarios.

Figure 3.6: A Venn diagram representing the number of conflict scenarios shown to each ATCO (within brackets).

3.4 Strategy Identification and Analysis

Table 3.1 shows the dataset size and summarizes the average values of different maneuver parameters for all the ATCOs.

Table 3.1: Resolution maneuver statistics (averaged values) with their standard deviation.

ATCO	Experience (years)	Dataset size	MIT (mins)	Heading Change (°)	Max CTD (NM)	Separation achieved (NM)
A	23	612	6.6 ± 1.8	35.8 ± 20.8	10.9 ± 4.5	8.5 ± 2.4
B	2	571	8.7 ± 2.9	31.5 ± 13.6	16.1 ± 7.0	8.8 ± 2.6
C	15	117	8.1 ± 2.1	11.4 ± 4.8	9.9 ± 3.6	7.7 ± 1.3
D	21	120	7.3 ± 1.9	16.1 ± 7.5	11.7 ± 4.2	8.9 ± 1.4
E	22	49	6.5 ± 2.1	14.8 ± 6.1	9.8 ± 2.7	8.3 ± 1.4
F	30	46	6.5 ± 1.7	16.1 ± 7.4	11.1 ± 3.6	8.3 ± 2.0
G	27	36	7.7 ± 2.3	15.6 ± 5.0	12.0 ± 4.2	8.2 ± 1.9
H	20	30	7.4 ± 2.3	17.8 ± 7.1	11.5 ± 2.9	9.6 ± 1.6

MIT: Maneuver initiation time prior to loss of separation
CTD: Cross-track deviation

Figure 3.7 shows the overall average values for maneuver initiation time, absolute heading change, maximum cross-track deviation, and the safe separation achieved on a combined dataset of all ATCOs. It is visible that except for maneuver initiation time, the data for the other plots is skewed to the right. The average values for the maneuver initiation time, absolute heading change, maximum cross-track deviation, and the safe separation achieved are 7.39 minutes, 28.90°, 12.81 NM, and 8.58 NM respectively.

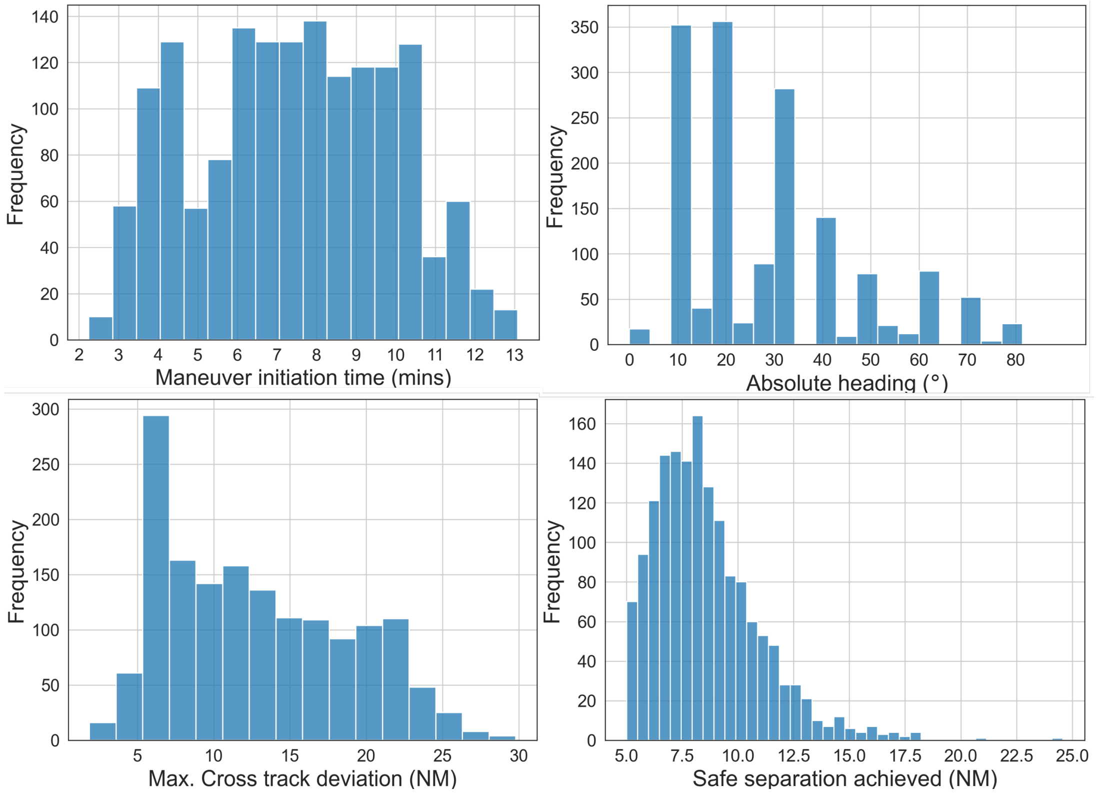

Figure 3.7: Combined data distribution of the maneuver parameters for all ATCOs.

Here, ATCOs A and B represent the group that was familiar with the airspace and ATCOs C to H represent the group unfamiliar with the airspace. Data distribution for the maneuver parameters is shown in Figure 3.8. The maneuver preferences of the ATCOs based on the data analysis is discussed in the following subsections. The conflict resolution data collected from each ATCO was analyzed in terms of the maneuver initiation time, absolute heading angle, maximum cross-track deviation, and the safe separation achieved (Figure 3.8. The following subsections discuss the ATCOs' conflict resolution maneuvers based on these parameters and the other components of the maneuver tuple (refer Section 3.2) representing an end-to-end conflict resolution maneuver.

Only a qualitative measure of ATCOs' maneuvers can be documented when discussing them in isolation. Hence, the conflict resolution maneuvers of all the ATCOs are discussed in parallel to establish quantitative measures of the identified strategies. The data collected from ATCOs A and B was the highest at 612 and 571 resolution maneuvers for the conflict scenarios presented. Their respective maneuvers for the conflict are visualized in

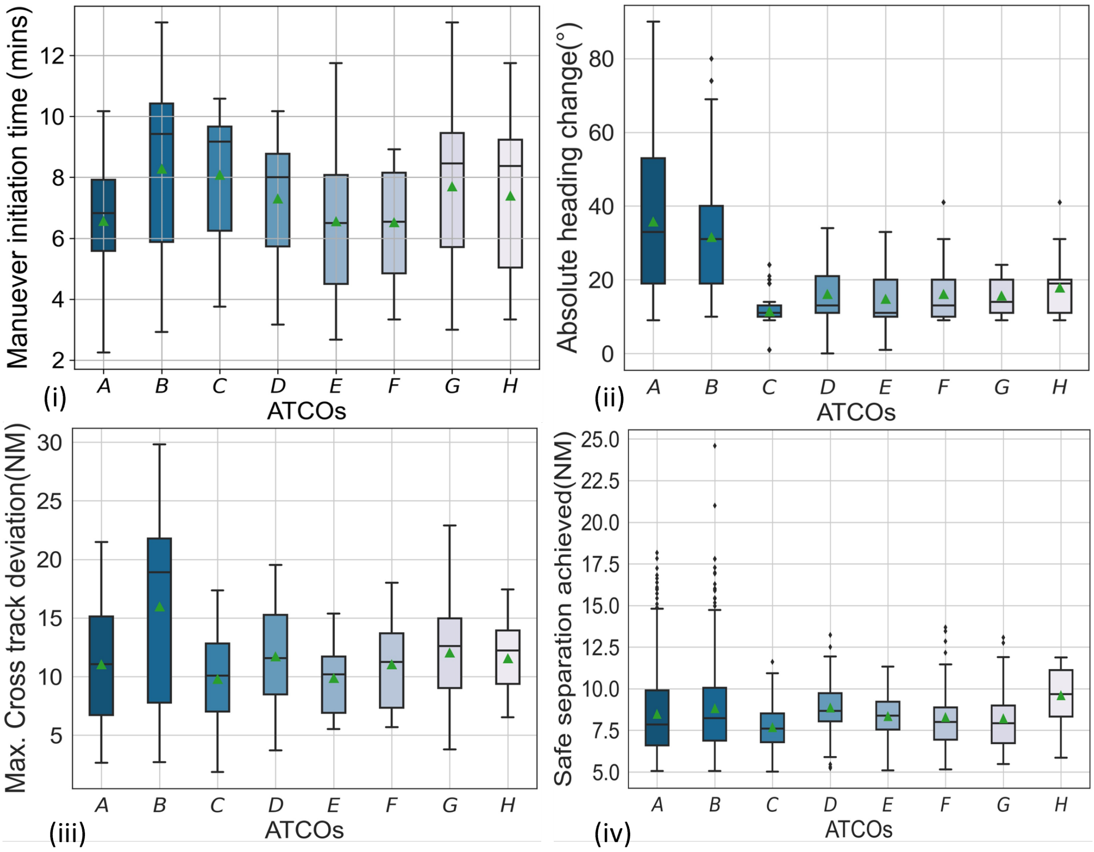

Figure 3.8: Distribution of maneuver parameters for the ATCOs. The triangles within the boxplots represent mean values.

Figure 3.9.

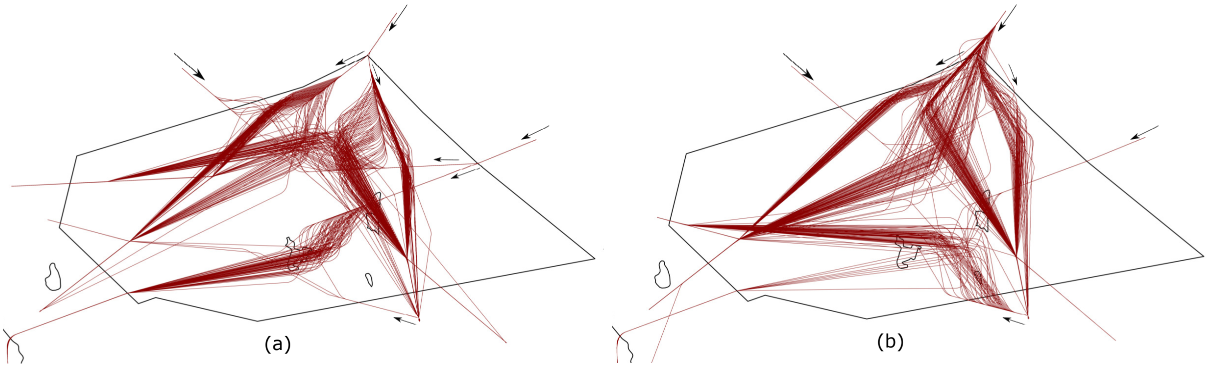

Figure 3.9: Resolved flight trajectories for ATCO *A* (a) and ATCO *B* (b) for the same dataset. A qualitative measure of the differing conflict resolution strategies can be seen by comparing the two images. The arrows indicate the direction of aircraft coming into Sector 6.

3.4.1 ATCOs' Conflict Resolution Maneuvers

Choice of aircraft to maneuver

Analysis of the conflict scenarios and the maneuvering aircraft preferences revealed that except ATCO B, all the ATCOs preferred the farther aircraft for the majority of conflict resolution maneuvers. This choice is effective because extending the path of the farther aircraft always ensures the necessary safe separation. The closer aircraft was chosen in cases where it was difficult to visually identify which aircraft was closer or farther from the conflict point. In these cases, choosing between a farther or closer aircraft to maneuver was equally likely. ATCO B on the contrary had mixed preferences in terms of choice of aircraft to maneuver. Out of the 571 conflict pairs presented to ATCO B, the aircraft closer to the conflict point was maneuvered 323 times, while the aircraft further away was maneuvered 248 times. Furthermore, feedback from the ATCOs revealed that the majority preferred to maneuver the west-bound aircraft in this conflict scenario involving waypoint MUMSO (Figure 3.10) because it is the first to arrive in their sector and for a longer duration before reaching the conflict point. Maneuvering the east-bound aircraft was not a common choice, because it would necessitate maneuvering the aircraft immediately, or even before it entered the sector (as done by ATCO D and E in some scenarios (Figure 3.10)).

Maneuver initiation time

Maneuver initiation time (MIT) is the amount of time remaining prior to the loss of separation when the resolution maneuver is initiated. Analysis of the maneuvered trajectories revealed that ATCOs A, E and F preferred to perform relatively delayed maneuvers, with mean values less than 7 minutes. Strategies of other ATCOs involved performing relatively early maneuvers, with mean values above 7 minutes (Table 3.1). The distribution of data (Figure 3.8.(i)) shows that MIT has a relatively higher spread of data for all the ATCOs except C, D, and F, and overall, there is less consistency in the preference of MIT for all the ATCOs.

Direction of the resolution maneuver

The majority of ATCOs consistently directed the maneuvering aircraft to the tail of the leading aircraft. In other words, the maneuvers were always directed within the region bound between the conflicting aircraft trajectories. It was also observed that for ATCOs B, D, and E, the direction of the maneuver was influenced by the location of the conflict point and when the maneuver was initiated. As shown in Figure 3.10, conflict point MUMSO is close to the sector boundary. ATCOs D and E preferred to keep the maneuvered aircraft at a sufficient distance from the sector boundary. ATCO B employed two kinds of distinct strategies. One strategy was similar to the other ATCOs, but in some cases, ATCO B directed the aircraft away from the aircraft in conflict with the maneuvering aircraft. These were the cases where the maneuver initiation was delayed (average MIT 4.07 minutes), thus requiring large CTD. In terms of the absolute headings given to the maneuvered aircraft, ATCOs A and B preferred larger magnitudes of headings as compared to other ATCOs (Figure 3.8 and Table 3.1). ATCOs C to H preferred relatively smaller heading changes with average values between 11.4°(ATCO C) to 17.8°(ATCO H). Further, standard deviations of the heading angles show that there was relatively higher consistency in the magnitude of headings given to the aircraft for ATCOs C to H.

Cross-track deviation

For all ATCOs except ATCO B, CTD values were equal to or less than 12 NM. ATCOs C and E had the lowest values for CTD with ATCO E being the most consistent (Figure 3.8). The CTD values for ATCO B ranged between 2.7 NM to 29.8 NM with an average CTD of 16.07 NM. These deviations can also be attributed to the ATCOs' level of confidence in the assigned maneuver and the extent of their training.

Selection of merging waypoint

Merging waypoint selection depends on the choice of waypoints available and the ATCOs' preference to redirect the flight to its original path. A key observable pattern is the consistency of waypoint selection by the ATCOs. In Figure 3.10 (iii), for example, even

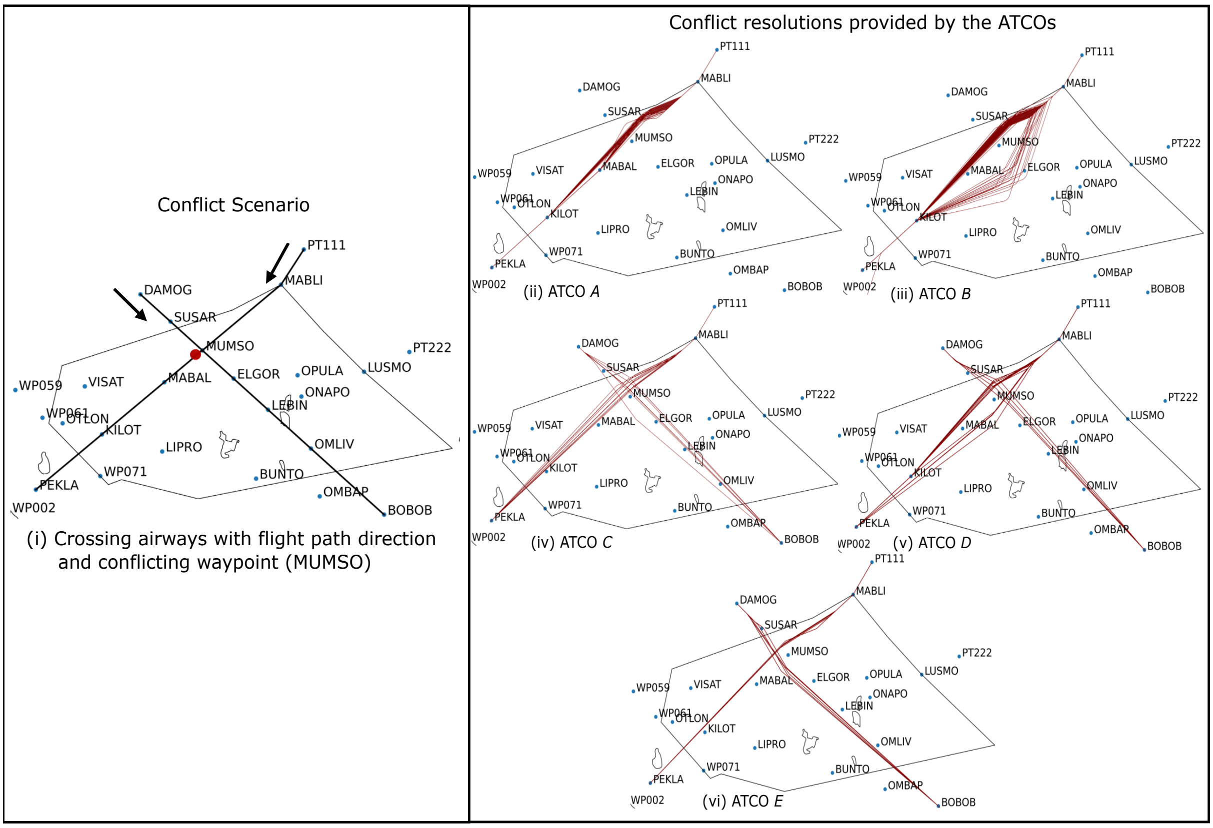

(i) Crossing airways with flight path direction and conflicting waypoint (MUMSO)

(ii) ATCO *A*

(iii) ATCO *B*

(iv) ATCO *C*

(v) ATCO *D*

(vi) ATCO *E*

Figure 3.10: An illustration of the intra-controller consistency in merging waypoint selection, based on the flight paths of maneuvered aircraft for conflict at waypoint MUMSO. To obtain a visual comparison of the overall resolution patterns for ATCOs, only the maneuvered aircraft's trajectory is plotted for each conflict scenario.

though the ATCO *B* had varying preferences for the maneuver direction, the merging point for all the flights was consistently KILOT, although MABAL was a viable option in numerous scenarios. This is indicative of some visual markers that ATCOs use to reduce the airspace complexity in conflict situations with higher surrounding traffic and maintain consistency of flight route patterns in the concerned sector. ATCOs *C* and *E* were also consistent in their merging waypoint selection for the majority of conflict scenarios. ATCO *D* merged the aircraft to any one of the available options of KILOT or PEKLA. The trajectory change point (TCP) location influences the choice of merging waypoint for both of these ATCOs. According to ATCOs' feedback, it is acceptable in operations to return the maneuvered aircraft to its original flight path outside the ATCO's sector if the ATCO of the subsequent sector is agreeable (as demonstrated by ATCOs *C*, *D*, and *E*).

Aircraft Separation achieved for resolved conflicts

Understanding the ATCOs' preferences for desired safe separations is critical. ATCO C achieved a minimum separation of 7.7 NM on average. It is evident that the distributions for ATCO A and B are skewed towards the higher values as compared to other ATCOs, with separation values of approximately 8 NM or less for 50% of the data for both ATCOs. Furthermore, because of some extreme outliers, the average values for these ATCOs were higher. For the other ATCOs, the data distribution is relatively consistent and similar. Overall, it becomes apparent that all ATCOs preferred maintaining safety buffers beyond the 5 NM lateral safety separation standard.

3.4.2 ATCOs' Conflict Resolution Strategies

Based on the findings in the previous subsection, we extract the strategies used by the ATCOs for resolving air traffic conflicts. Consistent with the information outlined in Section 3.4.1, these strategies are discussed in light of the elements comprising the aircraft's end-to-end conflict resolution maneuver. Except for ATCO B, ATCOs prefer farther aircraft to maneuver. This is similar to conflict resolution patterns identified by other researchers [27]. With average values of 8.7 and 8.1 minutes respectively, ATCOs B and C seem to prefer early resolutions when we compare the mean values for maneuver initiation time, while the other ATCOs prefer to maneuver the aircraft later. The direction of maneuver is consistent (i.e. within the region bound by the two aircraft trajectories) for all the conflict points except MUMSO, where the ATCOs tend to maneuver the aircraft in the other direction in some scenarios. This is to avoid getting close to the sector boundary. This behavior is consistent for the group of ATCOs unfamiliar with the airspace. Thus, we conclude that the location of the conflict and the timing of the maneuver have an impact on the ATCOs' subsequent decisions, such as maneuver direction. In terms of deviations given to the aircraft, there is a clear strategy distinction, where the group unfamiliar with the airspace prefers smaller headings in comparison to ATCOs A and B (Figure 3.8). On similar lines, CTD values show that ATCO B has relatively lower consistency in the deviation provided when compared to other ATCOs. Overall, the safe separations achieved by all the ATCOs are comparable. The group of ATCOs unfamiliar with the airspace have a comparatively lower spread of the data as compared to ATCOs A and B.

ATCO A has a consistent strategy which is as follows: identification of the farther aircraft from the conflict pair and its selection for the maneuver, assigning relatively higher heading angle (average of 35.8°) and directing the aircraft toward the tail of the other aircraft, with average cross-track deviations ranging from 6.4 to 15.4 NM. The preference of merging waypoint is always within the sector.

ATCO $B's$ strategy is not as consistent as ATCO A. For instance, ATCO B chooses to maneuver the farther aircraft in about 43% cases. ATCO B is consistent in the direction of resolution maneuver for all conflict locations except MUMSO (as seen in Figure 3.10) where the direction of maneuver changes based on the required magnitude of cross-track deviations. Thus, ATCO $B's$ strategies based on the identified consistencies are as follows: maneuver the selected aircraft relatively early (average of 8.7 minutes before L.O.S), with higher heading (average of 31.5°) and cross-track deviation ranging from 11.1 to 23.1 NM. Since the cross-track deviation is higher, the location of the conflict affects the direction of the maneuver. For ATCO B also, the merging waypoint is always inside the sector.

For the group of ATCOs unfamiliar with the airspace, the conflict resolution strategies are similar which is demonstrated by the data distribution, especially for the heading angle, cross-track deviation, and the safe separations achieved are similar. All the ATCOs have absolute heading change values lower than 20°and average cross-track deviation values for all are below 12 NM. It is also visible that even for conflict waypoint MUMSO, these ATCOs maintain a relatively greater distance from the sector boundary, as compared to ATCOs A and B.

3.5 Summary

To summarize, various factors and their combinations influence the ATCOs' conflict resolution strategies. These include factors such as the leading and the farther aircraft in the conflict pair, the location and proximity of the conflict point to the sector boundary, aircraft that enters the sector earlier, the potential visual markers in the sector where aircraft are directed to consistently, and list of available merging waypoints. As such, the identified factors that influence the conflict resolution strategies were used as input features for the machine learning models. Conflict resolution strategies are also subject to the individual preferences of the ATCOs in terms of safety, orderliness, and efficiency of the traffic flow in the sector. While safety is of the highest priority in air traffic control, different ATCOs prefer different safety buffers in their decisions, which proportionally affects the 'efficiency' of operations. In other words, there is a trade-off between safe and efficient traffic operations which different ATCOs manage differently, based on their individual preferences. Conflict resolution preferences, according to ATCO post-experiment feedback, are also affected by the preceding and following sectors, aircraft hand-over/take-over practices, and traffic flows.

Apart from the individual preferences of the ATCOs, the differences in the strategies of ATCOs A and B when compared to the other ATCOs can also be attributed to these two ATCOs' familiarity with the airspace as a whole. Thus, the strategies of ATCOs A and B may represent the influence of these external factors such as traffic patterns as well. However, since the other ATCOs (C to H) are unfamiliar with the traffic patterns and dynamics of the airspace, the resolution strategies might be a result of only the presented conflict scenarios. This is an interesting hypothesis and experiments to test it must be designed in the future. Because the data from these ATCOs has a similar distribution, we can test whether a group conformal machine learning model would be suitable for these. Thus, we combine them into a pseudo-ATCO, ATCO X, to test the machine learning model's performance on this aggregated data, which is discussed in the next chapter. This compiled dataset includes 398 conflict scenarios.

Chapter 4

Air Traffic Conflict Resolution Strategies: Learning and Prediction

The conflict resolution data collection experiments with the ATCOs' enabled the identification of their conflict resolution strategies and also the identification of factors that affect these strategies. The second challenge in achieving conformity of automation methods in performing experts' tasks such as air traffic conflict resolution is the automation method's ability to learn and recommend solutions that incorporate these strategies or propose solutions that are conformal to the ATCOs' conflict resolution strategies.

This chapter discusses two learning-based approaches, supervised learning and reinforcement learning, which utilize the ATCOs' data to generate ATCO conformal conflict resolution advisories. Further, the operational implications and challenges associated with these methods are discussed.

4.1 Learning ATCOs' strategies: Supervised Learning Approach

4.1.1 Overview

The data collected from the 8 ATCOs was used to train different machine learning models that generate ATCO conformal conflict resolution predictions. As discussed in the previous chapter, ATCO A and ATCO B have differing conflict resolution strategies. ATCOs C to H demonstrated similar strategies. Thus, individual sensitive models were developed for ATCOs A and B, and a group conformal model was developed for the other ATCOs. This required using the data to extract relevant features from the conflict scenarios which were used as input data for the prediction models. The performance of the trained models was evaluated in terms of their prediction accuracy and the level of conformance with the ATCOs' data using Frechet distance as the metric. Furthermore, features influencing the prediction models were identified and the robustness of the models was evaluated through a sensitivity analysis.

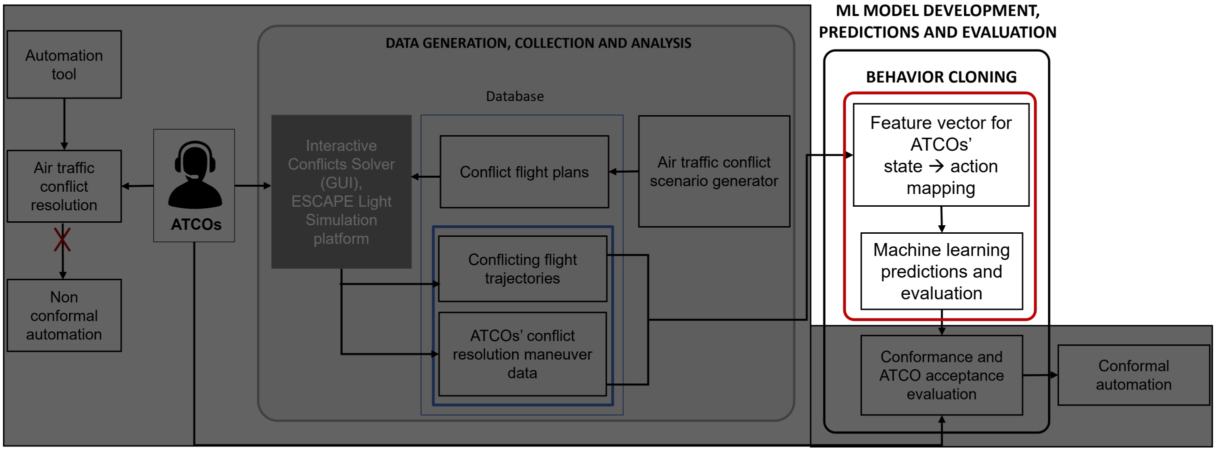

Figure 4.1: The overall concept diagram highlights the machine learning model development for ATCO conformal predictions, using the ATCOs' data collected through human-in-loop experiments, as discussed in Chapter 3.

The subsequent sections discuss formulating the ATCO conformal resolution prediction as a machine learning problem, the learning architecture, the details on feature engineering, model selection, and model evaluation, in detail.

4.1.2 Problem Formulation

The proposed learning framework is articulated in four modules with their internal architecture in Figure 4.2. Module 1 consists of generating conflict simulation data with specifications of Sector 6 of the Singapore Flight Information Region (FIR). The conflict resolution maneuvers from the ATCOs are obtained for these conflict scenarios. Module 2 consists of analytics, visualization, and comparison of the resolved trajectories obtained from Module 1. This enables the identification of the ATCOs' conflict resolution strategies. Module 3 comprises the feature extraction process. This uses the sector information and comparison of initial conflicting flight trajectories and the resolved flight trajectories at each time step to extract the feature for the machine learning model. Module 4 uses the proposed machine learning architecture to predict the components of the conflict resolution maneuver, and enables the generation of an end-to-end conflict resolution maneuver. These predicted trajectories are analyzed for conformity with the original maneuvered trajectories obtained from module 2.

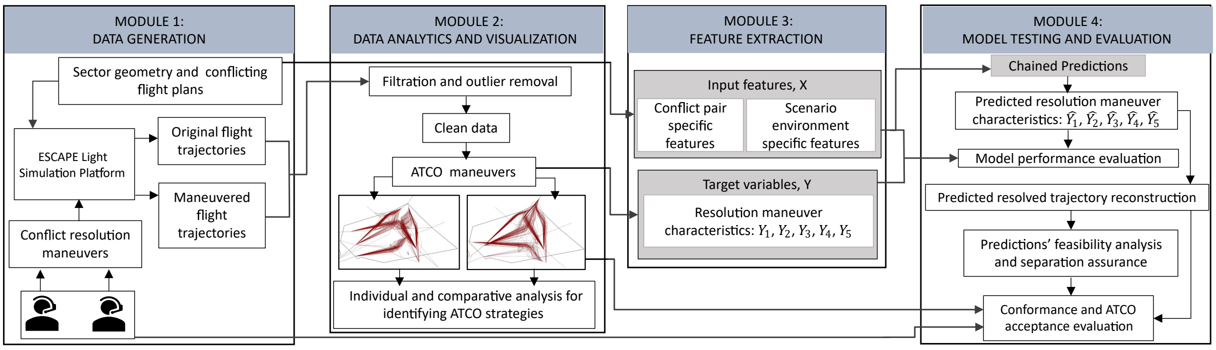

Figure 4.2: A visual representation of the Machine Learning (ML) framework comprising of four key modules. Data generation included creating conflicting flight plans and the ESCAPE Light Simulation platform to record ATCO conflict resolution strategies. The data analytics and visualization module consisted of the procedure to filter data and extract the ATCOs' strategies. Next, the filtered data was used to extract the input features and target variables in the Feature extraction module. Finally, these features were used for the ML predictions and generation of an end-to-end conflict resolution maneuver of the aircraft.

4.1.3 Conflict Resolution Dataset

The data used in this work was obtained from the human-in-loop experiments, as discussed in Chapter 3. In summary, the data was collected from 8 experienced ATCOs. Out of these, 2 ATCOs were familiar with the airspace and the remaining 6 ATCOs were unfamiliar with the airspace. Please refer to Chapter 3, Section 3.3.4 for details. ATCO A and ATCO B demonstrated different conflict resolution strategies for the same conflict scenarios presented to them during the data collection experiments. Thus, personalized ML models were developed for these ATCOs. The corresponding dataset consisted of 612 and 572 conflict scenarios for *ATCO A* and *ATCO B* respectively. Since ATCOs C to H demonstrated similar conflict resolution strategies, a group conformal model was developed for these ATCOs. The model was trained on the combined dataset of these ATCOs which was 398 conflict scenarios.

4.1.4 Learning Architecture

Behavior cloning for air traffic conflict resolution involves predicting a sequence of actions performed by air traffic controllers (ATCOs). For a given conflict scenario in the current work, an ATCO's conflict resolution maneuver is a path stretch [100] maneuver, that can be decomposed into the following action space. Based on the characteristics or features of the conflict scenario the first action is selecting the aircraft which should be maneuvered. This is a binary classification problem. Then, a decision should be taken regarding when the selected aircraft should be maneuvered, which is a regression problem. This action can be recorded as either time (minutes) or distance (NM) prior to a loss of separation. Subsequently, the direction and magnitude of the maneuver must be decided. In this regard, the heading angle must be decided by the ATCO accordingly, which is transformed into a regression problem. This is followed by a decision on the magnitude of cross-track deviation. This is again a regression task. Finally, the maneuvered aircraft must be merged into its original flight path, based on the available reporting waypoints. This is a multi-class classification problem. Based on these regression and classification tasks, the performance of the individual models is evaluated by mean absolute error (MAE) for regression, and accuracy (%) for classification. Thus, predicting an end-to-end conflict resolution maneuver, which is derived from the ATCO's conflict resolution preferences can be formulated as a supervised machine learning problem in the following manner.

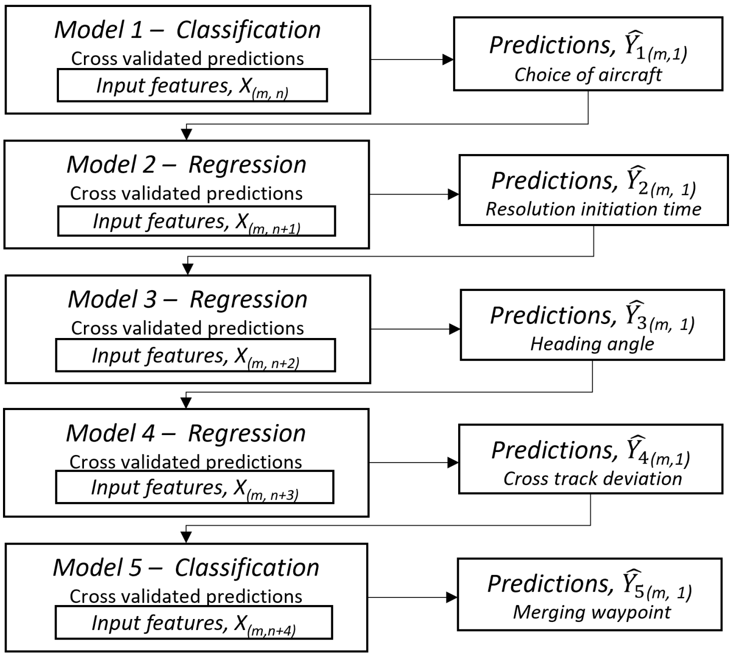

Figure 4.3: Multiple prediction models representing a chained prediction sequence to predict an end-to-end conflict resolution maneuver for the aircraft conflicts. The predictions of the previous models were used as inputs to the subsequent models for the predictions. Here *'m'* represents the number of rows and *'n'* represents the number of columns of the datasets.

Given a set of $\boldsymbol{M}$ training examples of the form (X_1, Y_1), (X_2, Y_2)..., (X_m, Y_m), such that X_i is the feature vector of the i^{th} training example of shape *(*1,n*)* where n represents the numbers of features and Y_i is the target variable, the learning algorithm seeks a function f: $\boldsymbol{X} \rightarrow \boldsymbol{Y}$ which best maps $\boldsymbol{X}$ to $\boldsymbol{Y}$, where $\boldsymbol{X}$ is the input space and $\boldsymbol{Y}$ is the output space. In other words, $\boldsymbol{X}$ represents the input features for the model and $\boldsymbol{Y}$ is the prediction of the model. In this research, a chained prediction architecture is utilized. This architecture incorporates predictions from previous models as inputs along with other features in subsequent models. For instance, given that the first model predicts the choice of aircraft to maneuver using the input features $X_{(m,n)}$, this prediction $\hat{Y}_{(m,1)}$ is added as an input feature to the next model, which predicts the maneuver initiation time. This is repeated till the last model of the chained prediction sequence. The details of the input features used by each machine learning model are in Table 4.1 and Table 4.3. Thus, for the p^{th} model in a chained prediction sequence, the dimensions of the initial dataset $X_{m,n}$ change to $X_{m,(n+p-1)}$ as the predictions from the previous model are added to the input features for subsequent predictions. The architecture of the chained predictions with the type of prediction model is shown in Figure 4.3.

4.1.5 Feature Engineering

Features pertaining to the conflicting flight pair (Cp), and the surrounding environment features (Sf) were extracted through the information of the sector dimensions, airways information, and the locations of flights. For the target variables i.e. the resolution maneuver features (Rf), the maneuvered flight trajectories were compared with the flight trajectories following the original flight plans. Table 4.1 shows the list of features extracted from the data and used for the learning algorithm. The input feature list was created after discussions with the ATCOs on the factors that they considered important while resolving air traffic conflict and also the features documented in the literature [45,49,80]. The abbreviations CPA and L.O.S refer to the closest point of approach, and loss of separation, respectively. Trajectory change point (TCP) refers to the point of maximum cross-track deviation (CTD) from where the aircraft is merged into its original flight path.

Table 4.1: Description of the features extracted from ATCOs' data and their abbreviations used in the figures.

S.no	Feature description	Abbreviation	Type
1	Initial heading: resolved flight	initialhead_resolved	Cp
2	Initial heading: unresolved flight	initialhead_unresolved	
3	Conflict angle	conflict_angle	
4	Distance from conflict point to maneuver start	mantoconflict_dres	
5	Distance from conflict point to the location of unresolved flight when maneuver starts	mantoconflict_dunres	
6	Direction of TCP wrt maneuvered flight	tcp_direction	
7	Direction of unresolved flight wrt TCP	unresolvedflight_dirn	
8	Distance of sector boundary from maneuver start	d_bound_manstart	Sf
9	Distance of CPA of resolved flight from sector boundary	d_bound_CPA	
10	Number of neighbouring aircraft	neighbouring_ac	
11	Mean distance of neighbouring aircraft from resolved flight at maneuver initiation	meandistance	
12	Time prior to L.O.S when resolution maneuver is initiated	timetoresolution	Rf
13	Distance prior to L.O.S when resolution maneuver is initiated	d_resstart_LOS	
14	Heading angle	headingAngle	
15	Max. cross-track deviation	crosstrack_dist_max	
16	Merging waypoint name	mergewp_name	
17	Distance between merging waypoint and TCP	mergpt_TCPdist	

Cp: Conflict pair feature; Sf: Surrounding environment feature
Rf: Resolution feature

4.1.6 Model Selection And Performance Evaluation

We utilized tree-based models, namely random forest and XGBoost, for our work due to their superior performance on tabular datasets [101]. To select the most suitable model for our approach, we compared their performance to that of a kernel-based model, specifically support vector machines (SVM). To ensure accurate model selection and performance evaluation, we employed a combination of hyper-parameter tuning and cross-validation techniques. For this purpose, we conducted 5-fold cross-validation and performed a grid search on key parameters to optimize the tree-based models. The tree depth varied from 3 to 20, while the number of trees ranged from 100 to 600 at intervals of 50. Other parameters such as *bootstrap* and *max_features* were kept at their default values as documented in the Scikit-Learn library [102]. The parameter settings are shown in Table 4.2.

Table 4.2: Range of hyperparameters used in grid search to find the optimal values.

Hyperparameter	Interpretation	Range / Value
n_ estimators	Number of trees in the model	(100, 150, 200, ..., 550, 600)
max_ depth	Maximum depth of each tree	(3, 4, 5, 6, ..., 19, 20)
max_features	Number of features to consider when looking for the best split	"sqrt" (default)
bootstrap	Whether bootstrap samples are used when building trees	"True" (default)

After hyper-parameter tuning the values of *n_ estimators* and *max_ depth* was set to 500 and 10 respectively. Five models were used to predict an end-to-end conflict resolution maneuver for a conflict, which incorporates the following components- choice of the aircraft to maneuver, conflict resolution initiation time, assigned heading, cross-track deviation, and the choice of merging waypoint. The predictions for classification or regression made by the preceding model were used as input features to the subsequent models. This methodology is also known as chained prediction [103]. Classification accuracy and mean absolute error (MAE) were used to evaluate the model's performance. Figure 4.3 (Section 4.1.4) demonstrates the flow between the various model instances used for prediction.

Prediction Performance

The performance of random forest, XGBoost, and SVM on the ATCOs' datasets was evaluated using 5 metrics ($P1$ to $P5$). Since the sequence of the individual prediction models is a classification model, followed by three regression models and again a multi-class classification model, the performance of the individual models is evaluated by mean absolute error (MAE) for regression ($P2$, $P3$, and $P4$) and accuracy (%) for classification ($P1$ and $P5$).

To identify the most suitable model out of random forest, XGBoost, and Support vector machines (SVM), these three models were trained and tested on the three ATCOs' datasets (ATCO A, B, and X). A 5-fold cross-validation setup was used to evaluate the performance of the three models (Table 4.3). In this process, the dataset is randomly split into 5 equal parts. Model training is performed on four out of the 5 parts and tested on the remaining part of the dataset. This entire process is repeated 5 times, for different test datasets, and an average test set value is obtained. For the 5 prediction metrics in Table 4.3 ($P1$ to $P5$), this value was obtained for the three ATCOs and the mean was compared for the three models. For instance, the classification accuracy of 93.6% is the average value of R.F for ATCOs A, B, and X, and so on. The averaged values are shown in Table 4.3. Since random forest has the best overall performance metrics, it was used as the preferred model for further analysis.

Table 4.3: Performance of random forests (R.F), XGBoost, and SVM on the dataset of ATCOs A, B and X averaged for comparison. P1 to P5 are the 5 prediction models used to predict an end-to-end conflict resolution maneuver of the aircraft. The best performance is highlighted in bold.

Prediction Model	Metric	Features Used	Avg. Model Performance		
			R.F	XGB	SVM
P1	Accuracy, % (Choice of aircraft)	1-5	**93.6**	93.4	92.4
P2	MAE, NM (Maneuver initiation)	1-3, 7, 9-11	**4.4**	4.8	4.4
P3	MAE, ° (Heading angle)	1-5, 7-12	**5.3**	5.4	6.8
P4	MAE, NM (CTD)	1-12	**1.5**	1.8	1.6
P5	Accuracy, % (Merging waypoint)	1-12	**90.1**	87.8	88.7

Individual Sensitive Conformal Models

Due to differing conflict resolution strategies for ATCOs A and B, we develop personalized machine learning models for conflict resolution for their datasets. For each prediction metric, the averaged value of the 5 cross-validation folds with standard deviation is shown in Table 4.4. Since the metrics for the train and test set for the ATCOs have relatively similar average values, it implies that the extent of overfitting is very small. For ATCOs A and B, the majority of the misclassified cases (in predicting the choice of aircraft to maneuver) involved scenarios in which both aircraft in the conflicting pair were roughly the same distance from the conflict point. This indicates that the performance drops in situations where either of the aircraft is equally likely to maneuver. Similarly, the classification accuracy for ATCO As's dataset in predicting the merging waypoint was lower potentially due to the usage of both MABAL and KILOT as merging waypoints, as long as the overall pattern of the resolution maneuvers is preserved.

Table 4.4: Average values of performance metrics on train and test sets of the ATCOs A and B with their standard deviation, using random forest. P1 to P5 are the 5 predictions made by the machine learning model.

	ATCO A		ATCO B	
	Train	Test	Train	Test
P1: Accuracy, % (Choice of aircraft)	95.1±0.7	94.6±1.4	95.5±0.7	92.7±0.6
P2: MAE, NM (maneuver initiation)	3.5±0.2	4.3±0.8	2.8±0.3	3.5±1.4
P3: MAE, (Heading angle)	4.2±0.2	5.5±0.5	3.7±0.1	5.2±0.3
P4: MAE, NM (CTD)	0.9±0.1	1.2±0.1	1.2±0.1	1.6±0.1
P5: Accuracy, % (merging waypoint)	94.6±0.5	93.9±3.1	99.3±0.1	99.2±0.7

Using the predicted values of the five components of the conflict resolution maneuver, P, predicted trajectories can be generated (Figure 4.4).

These trajectories were generated for the scenarios where the first prediction i.e. the choice of aircraft is correct. It is visible that the predicted trajectories closely conform to the actual maneuvered trajectories. In other words, the machine learning architecture is able to successfully clone the behavior of the ATCOs with the chained predictions. For example, for the end-to-end conflict resolution maneuver,

$P = (SQ160,\ 6.26\ minutes,\ Left,\ 14.82\ Nm,\ WP071)$, obtained from ATCO A, the

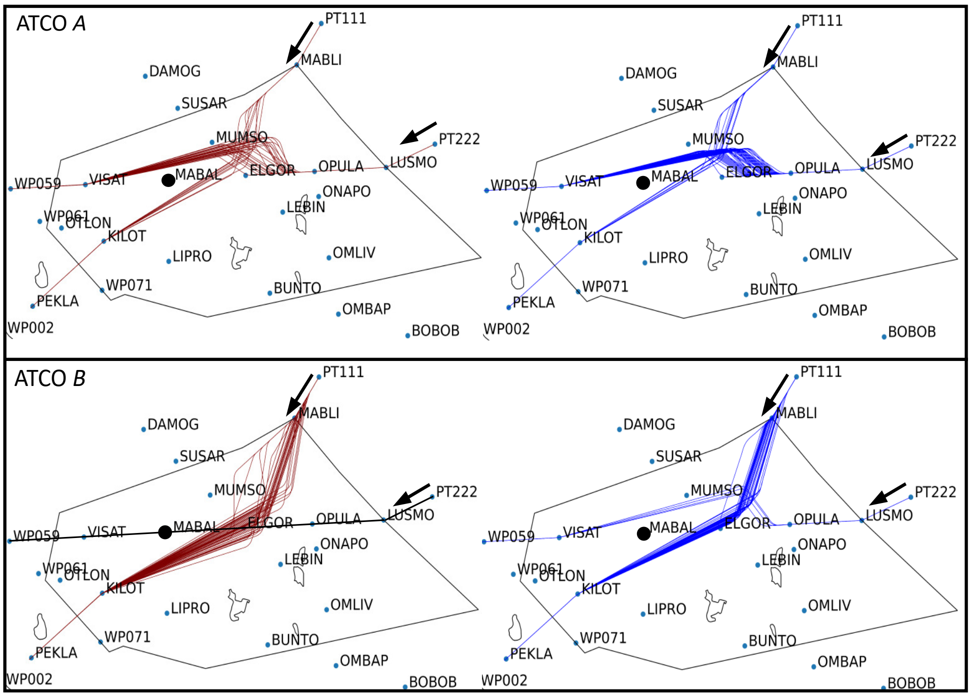

Figure 4.4: A comparison of the original resolutions (left) and the trajectories predicted through chained predictions (right) for ATCO *A* (top) and ATCO *B* (bottom). Close conformance of the predicted trajectories with the original maneuvered trajectories is evident from these images.

predicted maneuver is:

$$\hat{P} = (SQ160,\ 6.19\ minutes,\ Left,\ 12.32\ Nm,\ WP071).$$

Group Conformal Models

As discussed, the ATCOs unfamiliar with the airspace have similar maneuver preferences and hence, we tested a group conformal model for their dataset. ATCO $X's$ dataset represents the combined dataset of all the ATCOs in this group. The model was trained and tested on the combined dataset of the ATCOs. This was also based on a 5-fold cross-validation setup. The performance is documented in Table 4.5. This model achieves high prediction performance on the dataset, except for predicting the merging waypoint. The classification accuracy for this prediction is lower due to the small sample size, imbalance in the sample size from each ATCO, and their varying preferences in the selection of the merging waypoint. For instance, Figure 3.10 shows that ATCOs *C* and *E* consistently prefer PEKLA as the merging waypoint, whereas ATCO *D* chooses KILOT or PEKLA depending on where the merging is initiated from the maximum CTD point. As previously discussed, increasing the dataset size and decreasing the imbalance in the dataset size of each ATCO may improve model performance for ATCO X.

Table 4.5: Average values of performance metrics on train and test sets of the ATCO X with their standard deviations, using random forests. P1 to P5 are the 5 predictions made by the machine learning model.

	ATCO X	
	Train	Test
P1: Accuracy, % (choice of aircraft)	97.6±0.6	93.5 ±1.8
P2: MAE, NM (maneuver initiation)	4.2±0.1	5.3±0.4
P3: MAE, (Heading angle)	3.8±0.1	5.3±0.5
P4: MAE, NM (CTD)	1.4±0.1	1.8±0.1
P5: Accuracy, % (merging waypoint)	85.9±1.5	77.1±5.9

Furthermore, to ensure the feasibility of the personalized and group conformal predictions, a simple heuristic was implemented. In case the predicted trajectory of the maneuvered aircraft witnessed L.O.S, the CTD point was extended by 0.5 NM in the direction of the given heading deviation. The rationale behind this value is that the simulation environment updates the trajectory data after every 5 seconds, which corresponds to approximately 0.5 NM. A new trajectory was generated based on this and safe separation was evaluated again. Iterating over this process ensured that the outputs were always safely separated. The pseudo-code for the algorithm is in Algorithm 1.

Algorithm 1 Ensure Safe Separation

```
procedure EnsureSafeSeparation(P : ($C$, $T$, $D$, $D_c$, $M_{wp}$))
    $current\_heading \leftarrow D$
    $predicted\_traj \leftarrow$ reconstruct($P$)
    $CPA \leftarrow$ calculateCPA($predicted\_traj, traj\_intruder$)
    while $CPA <$ safetybuffer do
        $new\_D_c \leftarrow$ update_CTD($D_c$, step = 0.5 NM)
        $D_c \leftarrow new\_D_c$
        $predicted\_traj \leftarrow$ reconstruct($P$)
        $CPA \leftarrow$ calculateCPA($predicted\_traj, traj\_intruder$)
    end while
end procedure
```

Trajectory Conformance analysis

Comparison of the three machine learning models with the existing studies [2,68,72,74,75] is quite challenging due to the multiple reasons, which include differences in simulation environments, the methodologies used (deep learning [74] and reinforcement learning [2,75]), the differences in data collection methods, and participants in the experiments as well.

Furthermore, given the dynamic and stochastic nature of air traffic control, the ATCOs have the freedom to maneuver the aircraft using heading vectors, speed changes, altitude changes, or a combination of these, based on their maneuver preferences, traffic complexity, and a multitude of other external factors. Thus, there are no standards to establish how a particular conflict should be resolved, as long as safety is ensured. Other standards such as low aircraft deviations or efficient conflict resolution maneuvers are contextual and depend on factors such as ATCOs' experience and judgment, and their maneuver preferences. A standard practice in machine learning research is to compare different models on the researchers' own dataset or on the established benchmark datasets, as in [104, 105]. Due to the unavailability of a benchmark dataset in our research, we have compared tree-based machine learning models (Random forest and XGBoost) and a kernel-based model (Support Vector Machines) on our datasets to identify the model with the best performance (refer to Table 4.3 and section 4.1.6). The best performance for each metric has been highlighted in bold.

To quantitatively measure the similarity or conformance of the predictions with the ATCOs' maneuver preferences, we use Frechet distance (Figure 4.5), which is a commonly used metric to measure the similarity of curves [106] and perform trajectory clustering analysis [107]. By definition, this is a measure of similarity between curves, which take into consideration the ordering or orientation of the points on the curves [108]. Frechet distance is based on the following idea: "Let us consider two ‘walkers’ who are walking on two different curves and are connected by a string. The speed of the walkers may be the same or different (which represents the step size). Given that moving back is not allowed, what is the shortest distance of the string sufficient for these walkers to traverse their curve [109].

We calculated the Frechet distance between the following aircraft trajectory pairs: the maneuvered trajectory obtained from the ATCO and the predicted trajectory provided by the machine learning model, and the maneuvered trajectory obtained from the ATCO and optimal conflict resolution maneuver trajectories based on a grid-search method. We then compared the similarity between these pairs. To obtain an optimal resolution for a conflict,

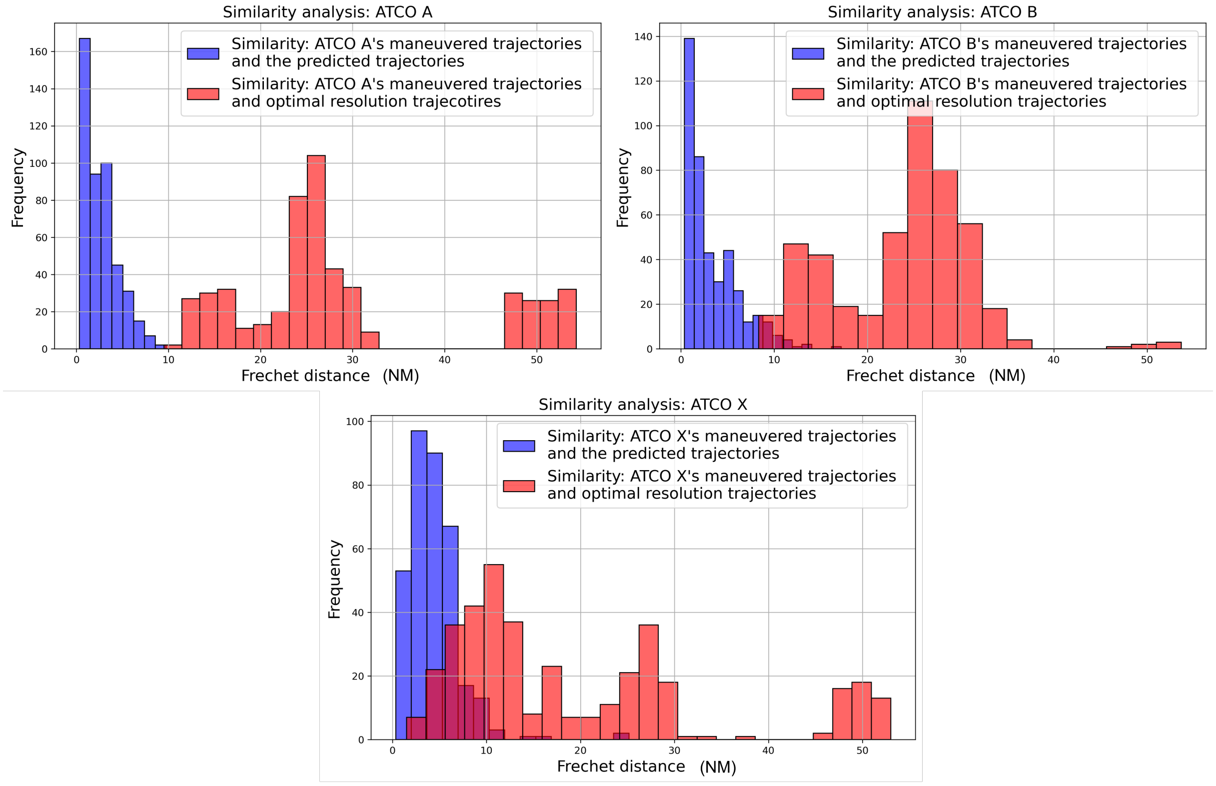

Figure 4.5: Estimation of the degree of conformance through measurement of similarity between the aircraft trajectories using Frechet distance.

we identified the heading and the cross-track deviation that leads to minimum maneuver time and used it as the solution for the conflict. Although the trajectories based on optimal resolution maneuvers and the trajectories based on the prediction models provide safe separation, we can observe from Figure 4.5 that the predicted trajectories are more similar to the ATCOs' maneuvers, or are more conformal to the ATCOs' preferences. It is inferred from the lower Frechet distance distribution between the ATCOs' trajectories and the predicted trajectories (in blue), as compared to the ATCOs' trajectories and the optimal trajectories (red) for the three ATCOs. Here, optimal trajectories are based on a resolution maneuver that provides safe separation and minimizes the maneuver time as well.

Features Influencing Models' Predictions

To investigate how the input features affect the models' prediction performance, we performed a permutation feature importance analysis. The overall idea is to randomly shuffle the values of a feature, keeping the values of other features in the feature list constant, to see the change in the model score. The process is repeated for all the input features of the model. This model inspection technique is called permutation feature importance and is defined as the decrease in the model score when a single feature value is randomly shuffled. The drop in the model score indicates how much the model depends on a particular feature. This also provides insights into what features the respective ATCOs consider while making decisions. Figure 4.6 shows the top 5 features for each prediction.

Figure 4.6 can be interpreted as follows: while predicting the maneuver initiation distance (Figure 4.6, ii), conflict angle is the most important feature for all the ATCOs since changing its values affects the prediction the most. In the case of ATCO $A's$ prediction model (red) the baseline mean absolute error is 4.3 NM. The range of the boxplot whiskers is 0.15 to approximately 0.30. Thus, the range of MAE varies from 4.45 NM to 4.6NM on changing the feature (the error increases). From the figure, we can identify the features that are considered important by all ATCOs while making maneuver decisions. Furthermore, the different magnitudes of the importance of each of these features indicate varying preferences within the ATCOs. For example, while considering the prediction of maneuver initiation distance, although conflict angle is important for all the ATCOs, the magnitude of importance is different, indicating varying preferences.

Models' Sensitivity Analysis

Conflict resolution datasets obtained from the ATCOs are small and may be not representative of the uncertainties such as environment (weather), variations in the ATCOs' decisions, and the time difference between when the ATCOs' take a decision and when the action is implemented. Therefore, we up-sampled the data to test the model's robustness as well as evaluate the model's generalization capability. For the training data, up-sampling was performed by adding Gaussian noise with a mean equal to the value of each feature element and the standard deviation set to 5% of the feature's average. We observed that the model trained on this dataset (up-sampling with 5% noise) provided the most stable performance over the range of noisy test sets.

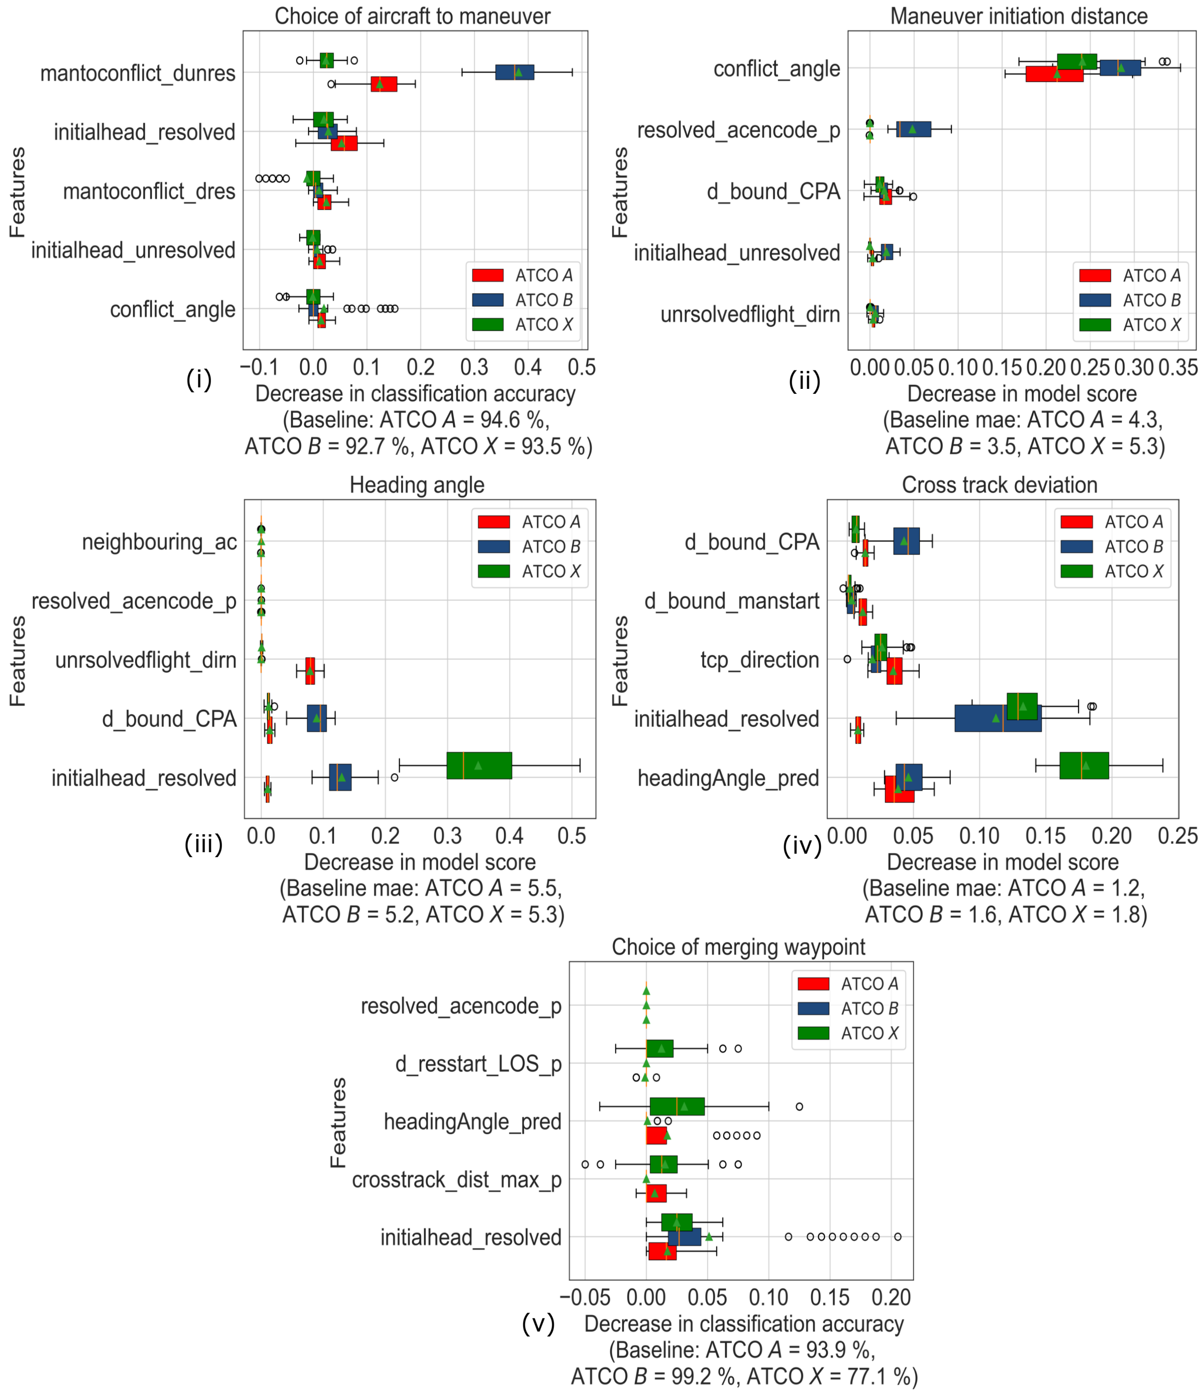

Figure 4.6: Permutation feature importance analysis for the random forest models developed for ATCOs *A*, ATCO *B* and ATCO *X*.‘*resolved_acencode_p*’, ‘*d_resstart_LOS_p*’, ‘*headingAngle_pred*’ and ‘*crosstrack_dist_max_p*’ represent the predicted values. The baseline performance of the models for each ATCO’s dataset is mentioned below each sub-figure.

For the test set, the noise varied from 0% to 10% at intervals of 2.5%. The rationale for selecting this noise range is our understanding that the ATCOs’ resolution preferences for ‘similar’ conflict scenarios should not vary significantly. Adding excessive noise to the data may generate conflict scenarios for which different conflict resolution maneuvers might be preferred, or the prediction labels might not remain suitable. For instance, on adding 10% noise to features such as a mean conflict angle of 57.03°, MIT(avg.= 6.62 minutes for ATCO *A*) and CTD (avg. = 10.98 NM for ATCO *A*), the new ranges are 51.3° to 62.73°, 6.02 to 7.22 minutes and 9.88 NM to 12.08 NM, respectively. Thus, the majority of the uncertainties related to the environment and ATCOs' decisions should be captured within this noise range. Figure 4.7 shows that the models are robust in the noisy test sets with up to 7.5% noise since the decrease in the performance metrics is small. On adding more noise (10%) there is a visible drop in the performance. This is especially the case for ATCOs *B* and X in heading angle predictions.

For ATCO $B's$ prediction model, the MAE in heading angle prediction is 4.39°at 0% noise and 4.77°at 7.5% noise. This implies a percentage error increase of 8.6% when 7.5% Gaussian noise is added to the data. On similar lines, the MAE in heading angle prediction for ATCO *B* is 5.39°at 10% noise in the data. Thus, from the baseline of 4.39°(at 0% noise level), this is an increase of 22.7%. A similar trend is observable for ATCO *X* (Figure 4.7 (iii)). Thus, the proposed models are robust to data uncertainties and noise within a noise range of 7.5%.

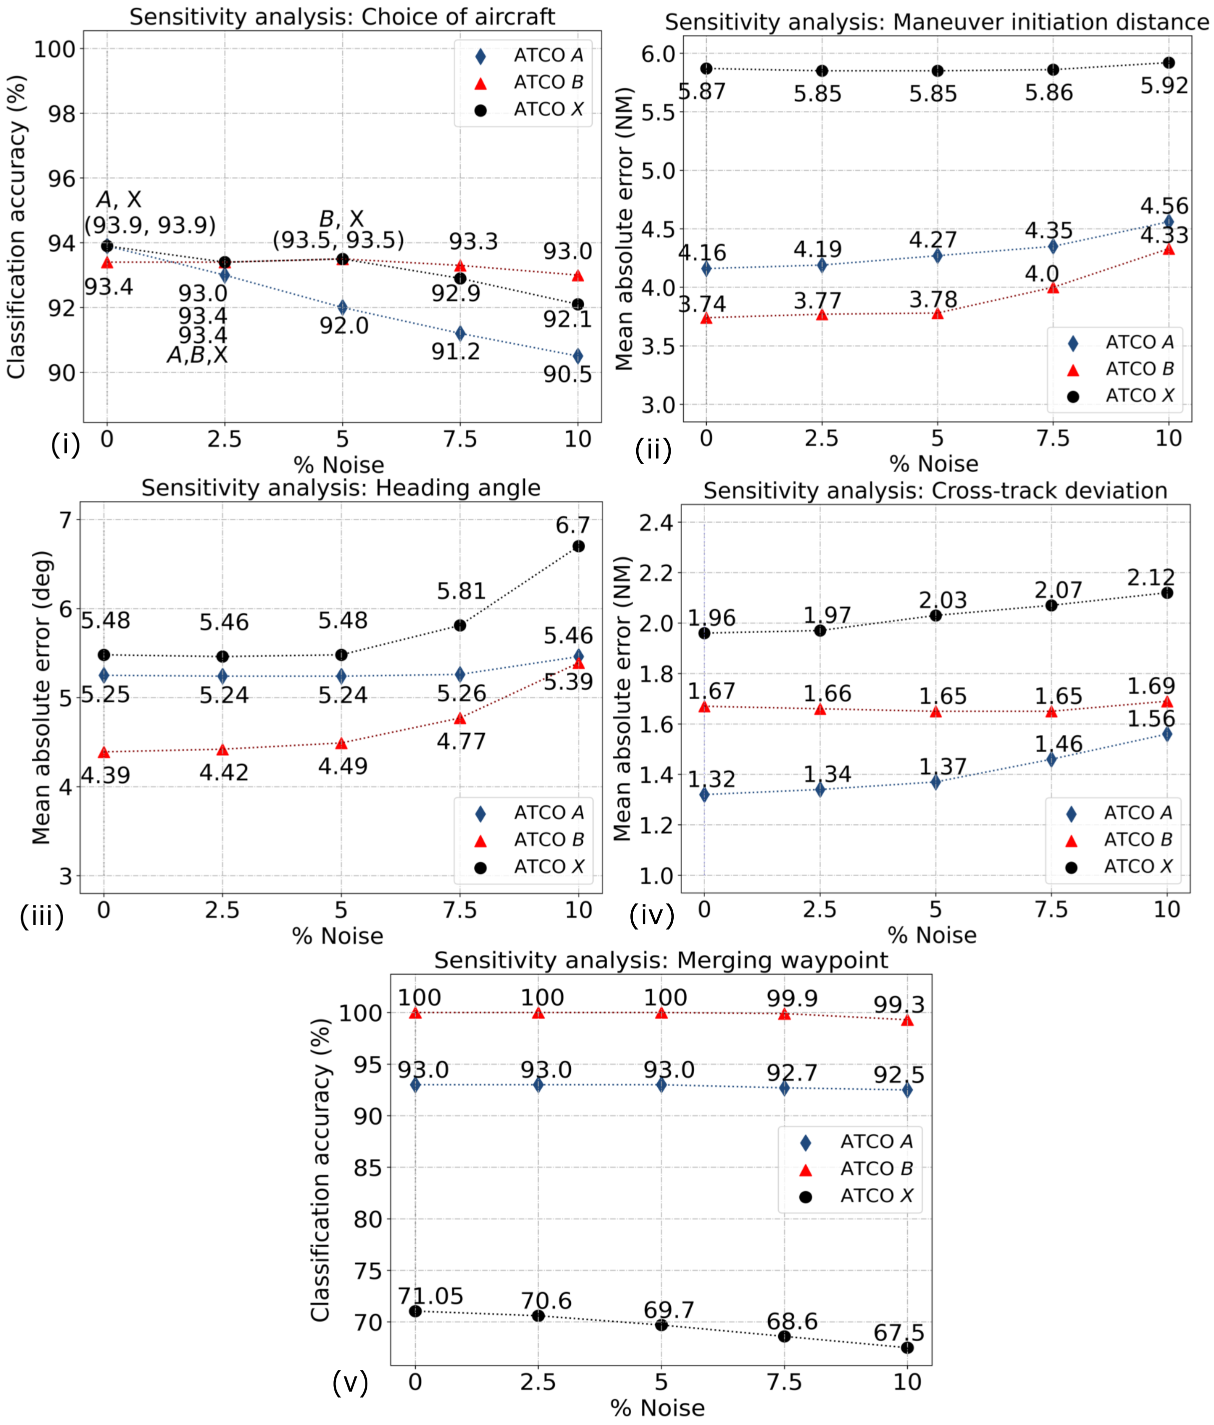

Figure 4.7: Sensitivity analysis for the ML models developed for ATCO A, ATCO B and ATCO X, over increasing test set noise.

4.2 Learning ATCOs' Strategies: Deep Reinforcement Learning Approach

4.2.1 Overview

Reinforcement learning has been used in multiple research areas including autonomous driving and robotics [110,111], to learn to perform a task directly from an expert's demonstrations. On similar lines, an agent can be trained to take actions that mimic an ATCO's conflict resolution pattern based on the information provided to the agent and the design of the reward function. We use the conflict resolution data collected from an ATCO and use the resolution maneuver trajectory as a reference for the agent to resolve the corresponding conflict. The following sections discuss the problem formulation, the Markov Decision Process framework, the experimental setting, and results and discussions.

4.2.2 Problem Formulation

This work proposes a deep reinforcement learning (DRL) approach for air traffic conflict resolution using behavior cloning, wherein, the conflict resolution data collected from an experienced ATCO is used as demonstrations to train the reinforcement learning agent, using a model-free DRL policy. Figure 4.8 highlights the overall idea of the research. The problem of air traffic conflict resolution is formulated as a sequential decision-making task, requiring a series of decisions to be made over time to resolve the conflicts. In operations, in case of a potential conflict between the aircraft, an ATCO performs a series of actions to achieve safe separation. These actions are based on sequential decisions such as the selection of the aircraft to maneuver, the maneuver initiation time, magnitude and extent of deviation to be provided, and subsequently, merging the aircraft into its original flight path. This constitutes an end-to-end conflict resolution maneuver. Conflict scenarios are randomly sampled from the data collected from the ATCO and a vector representation of each scenario, along with the information on the ATCO's resolution, is used by the agent to propose a series of actions to generate a similar end-to-end conflict resolution maneuver. One step in the simulation translates to 2 minutes in real time. Thus, a decision is taken by the agent every 2 minutes. The formal framework of Markov Decision Processes (MDPs) is utilized to model the sequential decision-making problem, incorporating states, actions, and rewards.

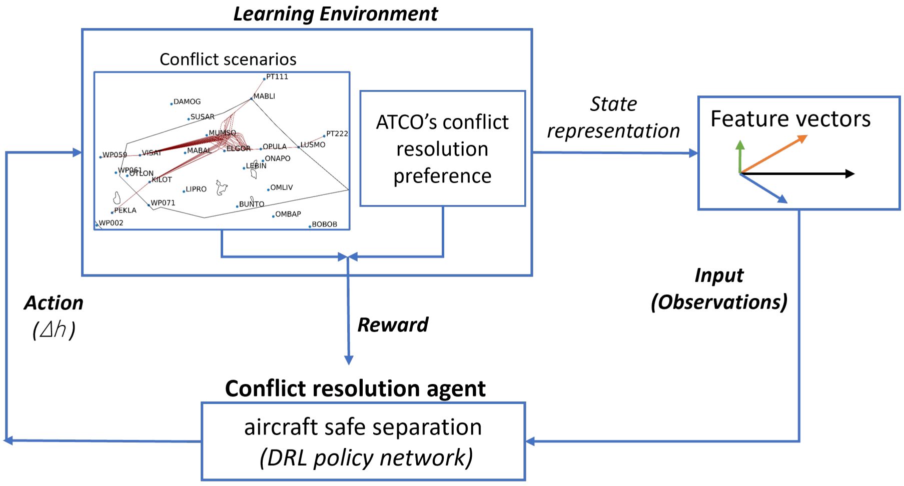

Figure 4.8: Overall idea of the research where the expert's data, along with the observation extracted from the conflict scenario are used to train the agent.

4.2.3 Conflict Resolution Dataset

The data used in this work is obtained from the human-in-loop experiments discussed in Chapter 3. We use ATCO A's dataset which consists of 612 conflict scenarios. During the training process, these conflict scenarios along with the conflict resolution provided by the ATCO, are randomly sampled to be input to the environment in the form of flight trajectories. In terms of ATCO A's maneuvers, the distribution of the closest point of approach between the two aircraft and the cross-track deviation of the maneuvered aircraft are shown in Figure 4.9.

4.2.4 Learning Mechanism

The learning mechanism has several key components, which include providing the agent with sufficient information to support decision-making, evaluating the agent's actions, and providing feedback to the agent in the form of a reward. To facilitate the learning process for resolving conflicts, we have developed a scenario generator that produces

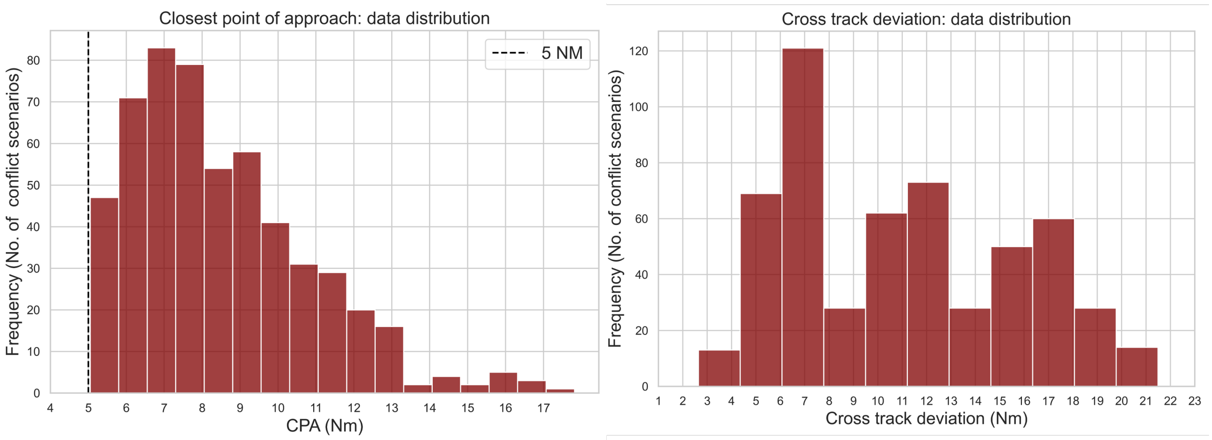

Figure 4.9: Distribution of the closest point of approach (CPA, left) and the cross-track deviation (CTD, right) of the maneuvering aircraft from its intended flight path obtained from the ATCO's conflict resolution data

conflict scenarios and represents them in a way that can be perceived by the agent. The agent's actions are defined and the mapping from these actions to the maneuvers taken by the agent is established. A reward function has been designed to assess the effectiveness of the maneuvers suggested by the agent. These components of the learning mechanism, which define the MDP, are discussed in detail as follows:

Observation space

The observation space contains all the information that the agent can receive at each step in order to decide the next action. In this work, the observation space includes information on the location of the two aircraft, the location of the destination for the maneuvering aircraft, the number of actions taken, the distance between the two aircraft, the distances of the two aircraft from their respective destinations, the cross-track deviations of the maneuvering aircraft (i.e the distance of the aircraft from its flight plan), the deviation of the maneuvering aircraft from the ATCO's conflict resolution trajectory, and the heading of the maneuvering aircraft. Overall, this is a vector of 13 elements.

Action space

In the current work, the action space involves heading maneuvers to achieve safe separation. This is based on the data collected from the ATCO during conflict resolution data collection experiments with the ATCO in the en-route flight phase. The action space consists of discrete actions in intervals of 15°, with a range of -45° to +45° permissible change. In other words, the agent can select an action from the options of -45°, -30°, . . . +30°, +45° at each time step.

Reward mechanism

The reward mechanism is primarily based on two criteria: conformance with the ATCO's maneuver preference and efficiency of the maneuver. Efficiency here implies minimizing the cross-track deviation of the maneuvering aircraft. In this regard, it is possible to tweak the reward function from a fully conformal to an efficient maneuver, by changing the weights of the reward (Equation 4.1). A loss of separation in the simulations is a situation where the separation between the aircraft is less than 5 nautical miles (NM). At each step, the agent receives a reward $R_s(s', a)$, which depends on the action a and the resultant state vector s', based on the following reward structure:

$$R_s(s', a) = \begin{cases} -\alpha * n_{actions} \\ -\beta * dev_{(maneuver)} - \gamma * dev_{(CTD)} \\ -\delta * n_{simulation\_steps} \end{cases} \tag{4.1}$$

In equation 4.1, $dev_{(maneuver)}$ refers to the deviation of the agent from the ATCO's maneuvered aircraft trajectory and $dev_{(CTD)}$ refers to the deviation of the agent's trajectory from the original flight path. The agent receives a reward at episode termination based on the criteria in Equation 4.2. Here, d_{min} refers to the distance between the two aircraft, and d_{sep} is the minimum safety separation. An invalid state would involve a situation where the aircraft turns back after a sequence of simulation steps. d is the distance between the maneuvering aircraft and the destination.

$$R_t = \begin{cases} -\Lambda & \text{if } d_{min} < d_{sep} \\ -\lambda & \text{if actions lead to an invalid state} \\ \Gamma & \text{if destination is reached} \\ \frac{d}{\Delta} & \text{otherwise} \end{cases} \tag{4.2}$$

The total reward after each episode is represented by equation 4.3, where N is the total steps in the episode.

$$R_T = \sum_{n=1}^{N} R_s + R_t \tag{4.3}$$

Learning algorithm

Training the agent was based on Proximal Policy Optimization (PPO) [112]. PPO is a policy gradient method that was developed to be robust and relatively easy to implement and tune. The algorithm parameterizes the policy as a neural network and employs a clipped surrogate objective function to control the policy updates, ensuring training stability (refer Equation 4.4). In reinforcement learning, the objective is to find a policy that maximizes the expected sum of rewards over time. The naive approach would be to directly maximize this expected cumulative reward. However, in practice, this can lead to large policy changes, making the learning process unstable and prone to divergence. PPO addresses this issue by using a surrogate objective that incorporates a constraint on the policy update. The surrogate objective is constructed in such a way that it provides a good approximation to the true objective while preventing overly large policy changes. This constraint helps in maintaining stability during the learning process and avoiding situations where the policy diverges. It involves clipping the ratio of the current policy and the old policy, $\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$, between $1-\epsilon$ and $1+\epsilon$. Here, ϵ is a hyperparameter. The minimum operation ensures that the policy update is not too aggressive, and the clipping function bounds the policy ratio to a certain range, preventing it from deviating too far from the old policy.

$$L^{CLIP}(\theta) = \mathbb{E}_t \left[\min \left(\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)} A_t, \text{clip} \left(\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}, 1-\epsilon, 1+\epsilon \right) A_t \right) \right] \tag{4.4}$$

Here $\pi_\theta(a_t|s_t)$ is the probability of taking an action a_t given a state s_t, under a policy π_θ. The advantage function A_t, represents how much better the taken action is, as compared to the expected value. This can be represented by equation 4.5, where the first term on the right hand side of the equation is the cumulative reward from time step t onwards, and $V(s_t)$ is the value function at state s.

$$A_t \approx \sum_{k=0}^{T-t} \gamma^k r_{t+k} - V(s_t) \tag{4.5}$$

Experimental setting

For the experiments, the values of the parameters $\alpha, \beta, \gamma, \delta, \lambda, \Lambda, \Gamma$ and Δ are fixed at 0.01, 0.01, 0.01, 0.01, 10, -20, 10 and 10, respectively. These values are obtained after multiple iterations and model performance evaluations. The Proximal policy optimization algorithm [112] used in this work is adapted from the Stable-Baselines 3 [113]. The parameters for the PPO algorithm are in Table 4.6. Training is performed on the Intel(R) Core(TM) i9-9900X CPU which takes about 4 hours.

Table 4.6: Parameters for PPO training.

Parameters	Value
Training timesteps	3e+6
Learning rate	3e-4
Discount factor	0.99
Clipping coefficient	0.2
ANN architecture	Multi-layer Perceptron
Optimizer	Adam
Hidden layers	64 X 64
Activation function	Tanh

4.2.5 Results and Discussions

Figure 4.10 shows the model's convergence after 1,000,000 scenarios. During training, the model's performance is measured by the average reward achieved and the number of steps taken to complete each episode. There is observable variation in the convergence plot for the average number of episodes because the length of the conflict scenarios is not the same. On similar lines, in an ideal case with no actions and no conflicts, the theoretical maximum reward is 10. Since actions taken, and deviations from the flight plan and the ATCO's data incur negative rewards, the average reward achieved stabilizes to around 5. To evaluate the trained model's performance, results are documented for 612 test set scenarios. This is done to obtain a dataset of the same size as the ATCO's conflict resolution data.

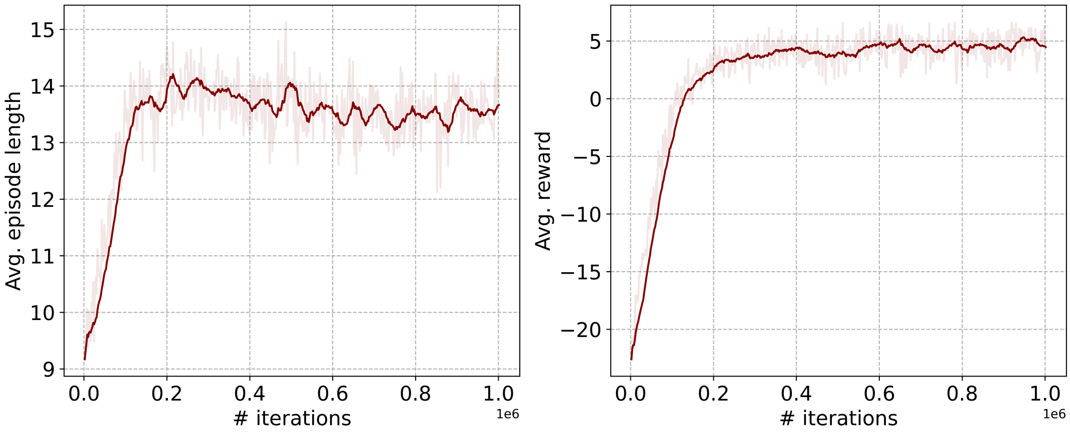

Figure 4.10: Model convergence: episode length (left) and the average reward (right).

In terms of safe separation, the agent achieves safe separation for 99.8% of the scenarios, i.e. a loss of separation occurs in one scenario only. The CPA distribution from the DRL policy is within the limits for the majority of scenarios with a few cases of outliers. The results for the cross-track deviation from the DRL policy are comparable to the ATCO's preferences.

On comparing the data distributions of Figure 4.11 and Figure 4.12, it is observable that although the ATCO has slight variations in the resolution outputs (leading to a relatively continuous distribution), the DRL policy provides consistent and identical resolutions with safe separation, which can be inferred from the distribution. The variation

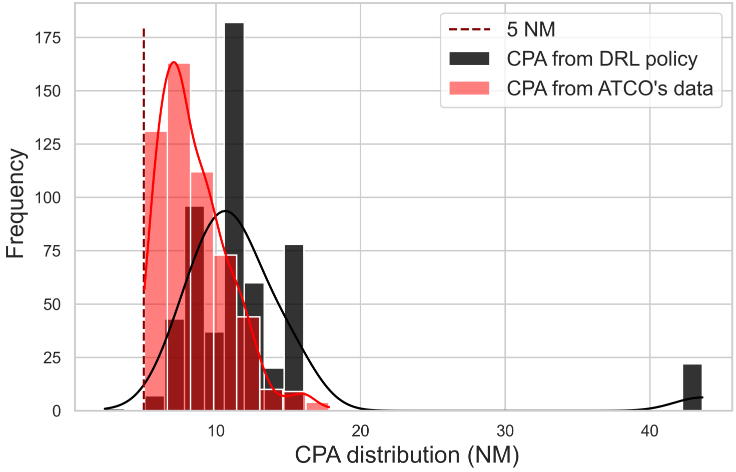

Figure 4.11: Comparison of the closest point of approach (CPA) distribution of the ATCO and the DRL policy.

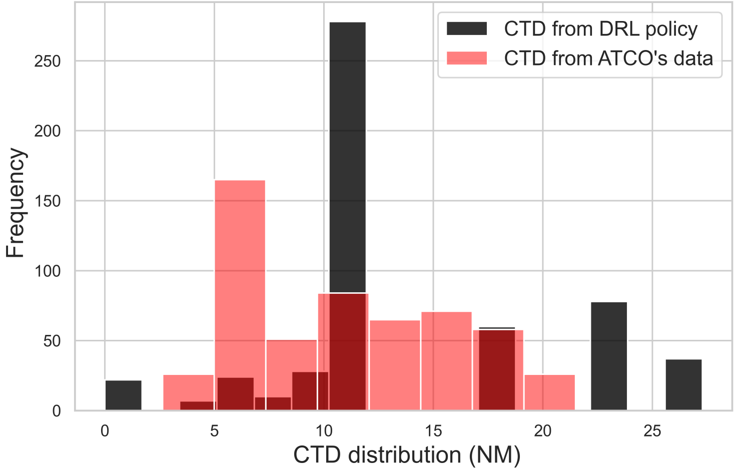

Figure 4.12: Comparison of the cross-track deviation (CTD) distribution of the ATCO and the DRL policy.

of the ATCO can thus be combined into a consistent maneuver by the agent, which might be an interesting characteristic, which must be investigated further. For certain conflict points in the sector such as MUMSO (Figure 4.13), it is difficult for the agent to follow the ATCO's maneuvered trajectory while ensuring safe separation. The ATCO prefers to deviate the aircraft towards the region bound by the two aircraft trajectories (cone of conflict). On the contrary, the agent opts for the other direction for the maneuver.

This behavior is demonstrated possibly because the conflict appears significantly early as compared to other conflict locations and the design of the reward function guides the

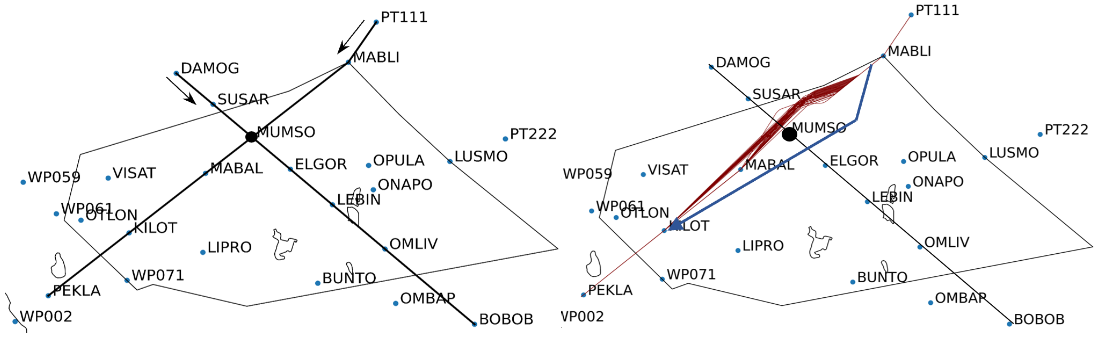

Figure 4.13: ATCO's preferred conflict resolution direction for a conflict at waypoint MUMSO is shown in maroon color. On the contrary, the DRL agent's maneuver direction is opposite to the ATCO (represented in blue color).

agent to opt for a safe maneuver immediately. Furthermore, the action space involves steps of 15°, which should be reduced to enable the agent to mimic the ATCO's conflict resolution maneuver more precisely. For the same reason, the step size of the simulation should be reduced. Overall, more parameter tuning is required for the agent to demonstrate ATCO conformal conflict resolution strategies.

4.3 Operational Implications

When integrated into an operational setting, automation tools of this nature should consistently serve as recommender systems, offering advisories that air traffic controllers (ATCOs) can choose to accept or reject. In all situations, ATCOs are the final decision-makers. Nonetheless, such tools that are trained on an expert's behavior have significant operational implications, one of which is a decision aid to reduce ATCO's workload. They can expedite the training process for novice ATCOs by reducing their learning curve, and serve as a means of implicit knowledge transfer from experienced ATCOs to their less-experienced counterparts. This does not imply that novice ATCOs must strictly adhere to the expert's solutions but rather have a reference that has demonstrated success in the past. Similar applications of using expert knowledge to train novice operators have also been proposed in other domains [114].

There are also several challenges such as safety assurance, when integrating such machine learning-based advanced cyber-physical-human systems in operations. Along with extensive experiments including stress tests to ensure that the systems produce acceptable if not optimal results in case of outlier events, human oversight, and feedback to the system are important. Validation and certification of such methodologies in more operational environments (involving weather and ATCOs' decision uncertainties) by the safety agencies is the biggest milestone to be achieved.

In this regard, the European Union Aviation Safety Agency (EASA) has provided a first set of technical objectives and organization provisions necessary for the approval of Level 1 AI applications (assistance to humans) [115]. EASA has also discussed the safety assessment and guidance for safety-related ML applications in terms of the initial safety assessment (during the design phase) and continuous safety assessment (based on operational data and in-service events) [116]. Furthermore, to facilitate a smoother integration of these systems into operations, a paradigm shift in ATCO training is required to accommodate the presence of AI/ML-based tools.

4.4 Summary

This chapter presents two machine learning-based methods for ATCO conformal air traffic conflict resolution. The supervised learning-based chained-prediction framework can learn and predict the conflict resolution strategies of ATCOs with high accuracy for the classification tasks and with low mean absolute error (MAE) for the regression task (for instance, the classification accuracy of above 92.7% for predicting the maneuvering aircraft, MAE for maneuver initiation distance < 5.3 NM, MAE for predicting the heading angle $< 5.3°$). The trajectory conformance analysis shows that the predicted aircraft trajectories are more conformal to the ATCOs' preferences as compared to optimal resolution trajectories. Furthermore, it can generate an end-to-end conflict resolution maneuver of the maneuvered aircraft for a given conflict scenario. The sensitivity analysis established that the machine learning models used in the work are robust up to a 7.5% added Gaussian noise to the ATCOs' conflict resolution data. Nonetheless, access to more data for each ATCO would reveal if the performance can be increased further. For all the ATCOs except ATCO B, the trailing aircraft has been the consistent choice to perform a maneuver. Further, ATCOs' feedback also revealed that factors such as the location of the conflict point affect the direction of the maneuver and even the aircraft selected to perform the conflict resolution maneuver. Similarly, strategies such as consistency and magnitudes of headings given, preferred CTD, and consistency and preferences of merging waypoint selection were identified for the ATCOs.

From a wider perspective, we recognize that capturing ATCOs' strategies solely through these experiments is not possible. This is because ATCOs' behavior can vary due to a variety of other external factors (related or unrelated to the job), which may influence their preferences directly or indirectly. These findings suggest the following potential research directions for our future work. First, adding an optimization layer to tweak the existing model predictions in order to provide 'efficient' ATCO conformal conflict resolution is a critical research extension. Here, efficiency would imply maneuvers with smaller heading and cross-track deviations but 'similar' to the ATCOs' conflict resolution preferences. It is important to evaluate ATCOs' acceptance of the machine learning models' optimized solutions. This may help to determine how flexible the ATCOs' preferences are. Scaling up these experiments to include factors such as the structure of neighboring sectors, the flight path of the aircraft outside the concerned sector, weather-related uncertainties, and complex traffic scenarios is also necessary in order to test the model's performance in such conflict scenarios and to further evaluate the model's robustness. Furthermore, this may provide insight into whether other conflict resolution strategies exist.

An interesting consideration is whether a single machine learning framework can be proposed, which is trained on multiple ATCOs' data to improve the application efficiency of this approach. Developing a single machine-learning model that is trained on multiple ATCOs may not be feasible in situations where the preferences of the ATCOs differ significantly. In such situations, the problem may transform into a single input to a multi-output mapping (i.e. the different ATCOs may have contrasting conflict resolution preferences for the same conflict scenario). Nevertheless, developing a baseline or a generic prediction model, which is initially trained on the data of all available ATCOs, and can further be fine-tuned based on individual preferences, is operationally more feasible. Such an approach draws parallels from the domain of recommender systems, wherein a baseline system is fine-tuned according to individual users' preferences. This research idea warrants an extension of the current work with more data collection and prediction model validation experiments involving the ATCOs.

The second proposed approach involves the use of deep reinforcement learning to train a DRL agent to generate ATCO conformal conflict resolution maneuvers by using an ATCO's data as a reference for the agent. To align the model resolutions with ATCO preferences, the reward function is designed to include the cross-track deviation from the ATCO resolution maneuver. As a result, the trained deep reinforcement learning (DRL) model achieves safe separation for 99.8% of the scenarios, with a comparable minimum distance distribution at the Closest Point of Approach (CPA) and cross-track deviations to ATCO behaviors. Additionally, unlike the slight variation in human resolution maneuvers, the proposed DRL tends to provide consistent and identical resolutions for similar conflict resolution scenarios. This characteristic should be further investigated to develop a reliable conflict resolution advisory tool for ATCOs. The discussion of the results also highlights that improving the simulation steps and the action space may potentially improve the agent's performance.

Chapter 5

Air Traffic Conflict Resolution Strategies: Conformance and Evaluation

Numerous methods have been proposed to assist ATCOs in air traffic conflict resolution. While the results of these methods are promising, they suffer from a major limitation in terms of operational acceptance by the ATCOs. It has been documented that the automation assistance methods receive a lower acceptance/are disused because they do not incorporate the ATCOs' conflict resolution strategies in the proposed solutions. On these lines, Chapter 3 discussed a methodology to identify the ATCOs' conflict resolution strategies and the factors affecting these strategies. Further, Chapter 4 discussed a supervised learning-based chained-prediction model that can learn these strategies and generate ATCO conformal conflict resolution maneuvers. Individual sensitive and group conformal machine learning models were proposed in this regard. Chapter 4 also discussed how a DRL agent can learn these strategies. This chapter provides details on two human-in-loop experiments conducted to evaluate the degree of acceptance of the individual sensitive models' predictions by the ATCOs, the consistency of the ATCOs' in their strategy selection, and further, evaluation of the ATCOs' conflict resolution maneuver selection from the given options of a conformal, optimized and a balanced resolution advisory.

5.1 ATCO Validation and Conformance Evaluation

5.1.1 Experiment Objectives

As discussed in Chapter 3 and Chapter 4, personalized chained-prediction models were developed for ATCOs *A* and *B* due to their distinct conflict resolution strategies. To further investigate the consistency of their strategies and the extent of acceptance of the machine learning model predictions by these ATCOs, we performed additional experiments based on the Institutional Review Board Approval (reference number: IRB-2020-11-030). The objective of these experiments was to evaluate how well the ATCOs' responses in the exercise aligned with the strategies identified in Chapter 3. The experiments also sought to identify instances in which ATCO responses differed from previously identified strategies, as well as the extent to which ATCOs accepted the machine learning models' predictions. ATCOs *A* and *B* participated in the experiment. In subsequent discussions, the term 'original strategy' refers to the corresponding ATCO's own strategy.

5.1.2 Experiment Design and Interface

100 conflict scenarios were randomly selected from the datasets of both ATCOs. To each ATCO, the original conflict resolution strategy, the model's predicted maneuver strategy from the ATCO's dataset, and the other ATCO's maneuver strategy were shown. The sequence of these three resolution options was randomized for each scenario in the exercise. Figure 5.1 shows an example from the validation exercise. For each scenario, the ATCOs were provided with the following information: the call signs of the aircraft in conflict, their headings, and the three resolution options. For the resolution options, the heading of the maneuvered aircraft was provided. The ATCOs were asked to choose their preferred conflict resolution strategy from the given options (any one, more than one, or none of the options). Nonetheless, ATCO *B* provided only one resolution option for the presented conflicts.

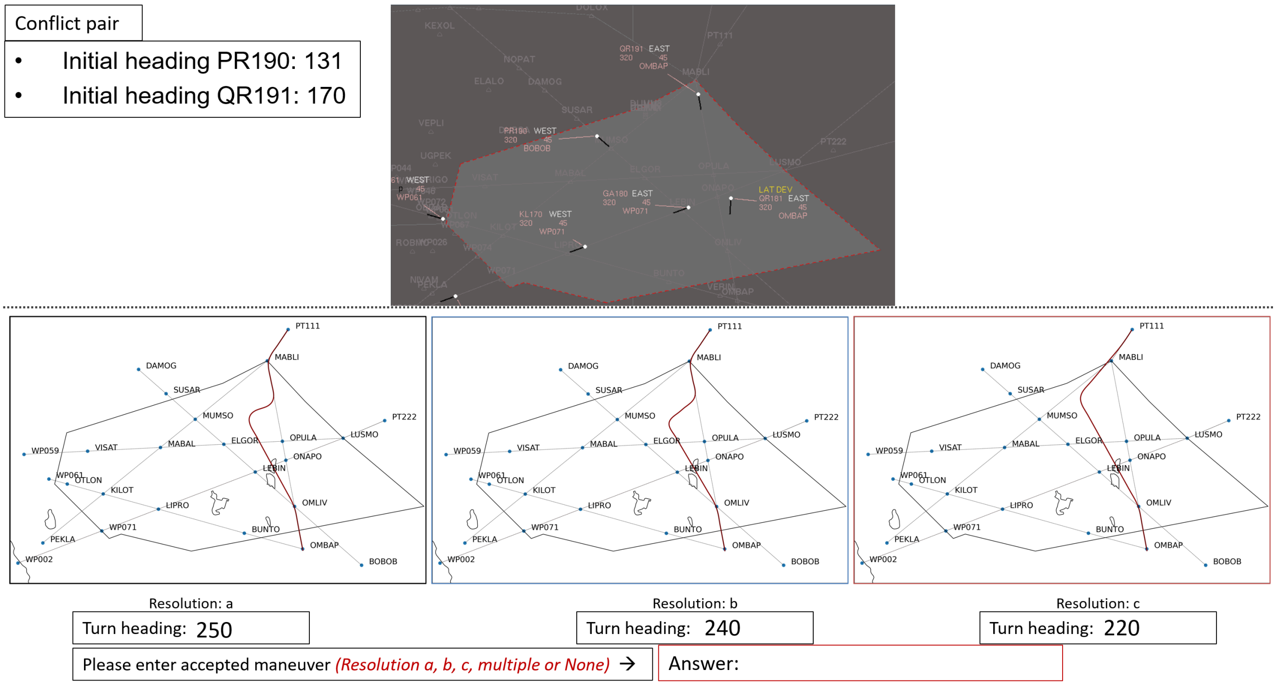

Conflict pair

- Initial heading PR190: 131
- Initial heading QR191: 170

Resolution: a

Turn heading: 250

Resolution: b

Turn heading: 240

Resolution: c

Turn heading: 220

Please enter accepted maneuver *(Resolution a, b, c, multiple or None)* →

Answer:

Figure 5.1: A sample case from the ATCO acceptance evaluation exercise. The information regarding the conflicting flight pair and their initial headings were provided to the ATCOs as shown in the figure. Further, the heading values of the maneuvered aircraft were also presented. The resolution strategy options included the ATCO's original strategy (own strategy), the strategy predicted by the machine learning model on the ATCO's dataset, and the other ATCO's strategy. The ATCOs were asked to record their conflict resolution preferences from the available options.

5.1.3 Results and Discussions

Table 5.1 shows the frequency of selection of the available options for both ATCOs. The subsequent sections provide detailed discussions for both, ATCO A's and ATCO B's data.

Table 5.1: Conflict resolution maneuver selection percentage for ATCO A and ATCO B.

Strategy	Strategy selection percentage	
	ATCO A	ATCO B
Original strategy	16%	24%
Predicted strategy	37%	44%
Other ATCO's strategy	3%	32%
Predicted and Original strategy	32%	–
Original and other ATCO's strategy	2%	–
Predicted and other ATCO's strategy	5%	–
All three strategies	5%	–

Analysis for ATCO A

ATCO A demonstrated high conformance in the selection of strategies similar to those identified in Chapter 3. Overall, ATCO A chose the original conflict resolution strategy or the predicted strategy as at least one of the options in 97% of the scenarios. For 85% of the scenarios, either the predicted or original strategy, or both, were selected for conflict resolution. Further, the predicted strategy was selected as the option or at least one of the options in 79% of the scenarios. Since the predicted and the original conflict resolution strategies have high conformance, this highlights that ATCO A's conflict resolution strategy selection in this exercise was highly consistent and similar to the earlier identified strategies. For 12% of the scenarios, ATCO A accepted ATCO B's strategy as one of the resolution options. These were the cases where ATCO $B's$ strategies were visually identical to ATCO A's original strategy or the predicted strategy. For only 3% of the scenarios, ATCO A opted for ATCO B's resolution only. These were the cases where ATCO B's strategies were similar to ATCO A's strategies in terms of the resolution maneuver direction and MIT, but less aggressive in merging the maneuvered aircraft to the original flight path.

These findings lead to the following conclusion. ATCO A had consistent strategies for air traffic conflict resolution, which can be observed through the repeated selection of ATCO A's original or the predicted conflict resolution strategies in the exercise. In other words, ATCO A demonstrated a strong bias towards his original conflict resolution strategies. This can be attributed to ATCO A's extensive experience in radar operations through which ATCO A has developed an optimal representation of conflicts and their resolutions in the sector. Thus, ATCO A does not prefer to deviate from these even when presented with an alternative resolution strategy. On the whole, ATCO A has well-developed, consistent strategies and prefers to exploit these strategies for conflict resolution. Further, since the model-predicted strategies show high strategic conformance with the original strategies, ATCO A demonstrates a high degree of acceptance of the machine learning model predictions.

Analysis for ATCO B

ATCO *B* opted for the original conflict resolution strategies for 24% of the scenarios and the predicted strategies for 44% of the scenarios. Since the machine learning model's predicted strategies and ATCO B's original strategies conform significantly, it can be safely argued that ATCO *B* had a preference for resolution patterns depicting his original strategies, for 68% of the scenarios. It was interesting to note that ATCO *A*'s strategies were selected for 32% of the scenarios. This was significantly high when compared to the validation results from ATCO *A*. Further analysis of the exercise results revealed mixed preferences for ATCO B in terms of choice of aircraft to maneuver. Also, ATCO *B* was consistent in selecting the strategies that included relatively early resolutions in terms of MIT. These selections were similar to the identified strategies documented in Section 3.4.2 in Chapter 3.

For the scenarios where ATCO *A*'s strategies were selected, both the ATCOs' strategies were similar in terms of maneuver direction but ATCO *A*'s strategies were more efficient in terms of the cross-track deviation, and the heading assigned to the maneuvering aircraft. These results are important as they highlight the scenarios where ATCO *B* is consistent in the selection of conflict resolution strategies. Parallelly, though ATCO *B* has inherent strategies for resolving air traffic conflicts, there are numerous situations where ATCO *B* prefers ATCO *A*'s strategies and the machine learning models' predicted strategies. This demonstrates that ATCO *B* accepts the model's conformal predictions for conflict resolution. Moreover, when provided with relatively efficient alternatives for conflict resolution, ATCO *B* adopts ATCO *A*'s reasonably efficient (involving lower aircraft deviations) strategies. This allows ATCO *B* to refine his original techniques and make better-informed decisions regarding conflict resolution. One potential reason for this behavior is that ATCO *B*, having less experience in radar operations, is more inclined to explore and accept other better alternatives, without a significant bias against his original resolution strategies. This raises an interesting question on the relation between the ATCO experience and the relevance of conformance, and whether conformance becomes more relevant only for experienced ATCOs who are consistent with their strategies.

5.2 Conflict Resolution Preference Evaluation

5.2.1 Experiment Objectives

Another set of experiments was performed with the ATCOs after seeking approval from the Institutional Review Board (IRB) with the approval reference number: IRB-2023-046. 6 different ATCOs (other than ATCO A and B) with over two years of operational experience were invited for the experiments. The objective of this set of experiments was to identify the type of conflict resolution maneuver that they would prefer for a given conflict scenario, out of the given resolution maneuver options. The experiments also sought to investigate the ATCOs' response to questions such as the usefulness of such an automation assistance tool, and whether it improved their decision-making, to name a few, through a post-experiment questionnaire.

5.2.2 Experiment Design and Interactive Interface

The ATCOs were presented with air traffic conflicts in Sector 6 of the Singapore FIR, and for each conflict scenario, three resolution options were provided. Each experimental run consisted of 20 conflict scenarios that were presented to each ATCOs during the experiments. These scenarios were the same for all the participants. Before the start of the experiments, the participants were briefed about the purpose of the experiments and were given time to familiarize themselves with the interface and its functions, to avoid any effects of the interface unfamiliarity on the collected data. After each run, the ATCOs were requested to take a questionnaire, that consisted of Modified Bedford Scales for Workload and China Lakes Scales for situational Awareness. Such scales are commonly used as standards to collect ATCOs' feedback in similar experiments [117]. The questionnaire is attached in Appendix A.

Conflict resolution options

For each conflict scenario, the following three resolution options were provided to the ATCOs:

- **Conformal resolution:** This conflict resolution trajectory was obtained from the chained-prediction model, which is discussed in Chapter 4.
- **Optimal resolution:** This conflict resolution trajectory refers to a trajectory with the least travel time in the sector. In other words, a conflict resolution trajectory with minimum deviations.
- **Balanced resolution:** This conflict resolution trajectory was better than the conformal advisory in terms of the travel time i.e. the maneuver deviations, but had relatively greater deviations when compared to the optimal resolution.

The balanced and optimal resolution options were based on the conformal conflict resolution obtained from the ML model. The following sequence of steps was carried out to obtain these resolutions. For each conflict, the conformal conflict resolution, P', obtained from the ML model was considered as the baseline. A range of values in its vicinity were considered for each component of P'. For instance, for a predicted heading of 270°, a range from -70°to +70°at intervals of 10°was used. Then, each combined resolution was checked and filtered using a conflict detection algorithm to remove all infeasible conflict resolution maneuvers (maneuvers that do not resolve the conflict). Subsequently, using different values for the components of P', a database **D** of feasible conflict resolution was generated.

Each resolution in the database was scored based on two criteria: its optimality score (OS) and its conformance score (CS). OS was based on the maneuvered aircraft's travel time in the sector. CS was evaluated in terms of the difference in the travel time of the conformal resolution and the resolution from the database under consideration. The total

Algorithm 2 Pseudo-code to extract optimal and balanced conflict resolutions

```
$\mathbf{P} \leftarrow (\mathcal{C}, \mathcal{T}, \mathcal{D}, \mathcal{D}_c, \mathcal{M}_{wp})$        ▷ Baseline conflict resolution
Create a conflict resolution database using a range of parameters for each component of $\mathbf{P}$
Initialize $\alpha$ values: $\alpha_1 \leftarrow 0$, $\alpha_2 \leftarrow 0.5$
for each $\alpha$ in $[\alpha_1, \alpha_2]$ do
    Calculate total score $TS$ for each resolution in the database:
        $TS \leftarrow (1-\alpha) \times OS + \alpha \times CS$
    Sort resolutions by $TS$ to find the resolution with the minimum score
    Check resolution feasibility (safe separation)
    Extract the resolution with the minimum score and feasible solution
end for
return Optimal and balanced resolutions
```

score (TS) of the concerned resolution is thus obtained from the equation:

$$T(m) = T_{Exit}(m) - T_{Entry}(m) \tag{5.1}$$

$$OS(m) = T(m) \tag{5.2}$$

$$CS(m) = T(m) - T(P') \tag{5.3}$$

$$TS(m) = (1-\alpha) * OS(m) + \alpha * CS(m) \tag{5.4}$$

Here, the parameter α varied from 0 to 1. ‘m’ refers to the maneuvers, and $T(m)$ is the travel time in the sector. For a given value of α total score of the entire database was calculated and the resolution with the lowest score was selected. Thus, $\alpha = 0$ provided an optimal conflict resolution, and $\alpha = 0.5$ provided a balanced conflict resolution.

The interface visualized the Singapore Flight Information Region (FIR) and focused on Sector 6 for conflict resolution. Due to the high density of airway intersections in Sector 6, the conflict scenarios have higher occurrence probability and complexity compared with other sectors. Additional information on airspace including waypoints, sectors, traveling trajectories, and information on aircraft are displayed accordingly at each timestep. The interface for the experiments is shown in Figure 5.2. The interface is implemented with Python 3 with main libraries including Pygame and Cartopy for interactions and mapping traffic information.

Interface Components

The main components of the interface include:

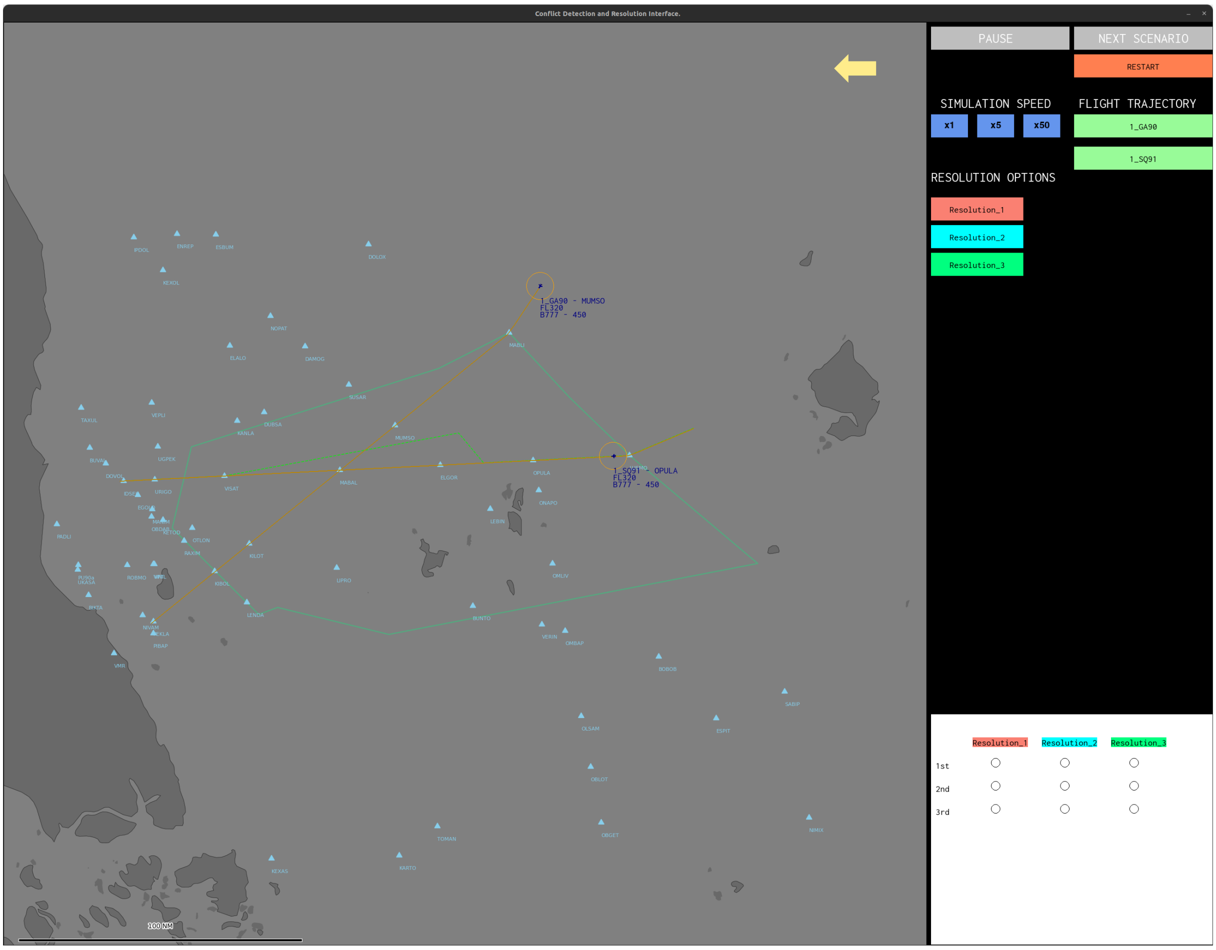

Figure 5.2: Conflict resolution interactive interface with radio buttons for ATCO ranking responses at the bottom right corner, along with other functionalities on the top right.

- **Main visualization window:** The main visualization window displays the traffic scenario of aircraft traveling within the Singapore FIR. A circle with a radius of 5NM surrounded each aircraft icon depicting the required minimum safe separation. The sector's boundaries and waypoints were also plotted in the background for reference. The area of Sector 6 features an interactive button for zooming in, while a yellow backward arrow button allows for zooming out of Sector 6.

- **Upper right button menu:** On the right-hand side of the interface, there are three buttons: the Start button, the Next Scenario button, and the Restart button. Upon clicking, the Start button initiates the scenario and changes into a Pause/Resume button, which is used to pause or resume the scenario when clicked. The Next Scenario button concludes the current scenario and loads the next traffic scenario from the dataset into the main visualization window. The Restart button reloads the data and resets the traffic scenario to its initial state.

- **Speed control buttons:** Below the "Simulation Speed" label on the interface, there are three buttons labeled "x1," "x5," and "x50." Clicking on each button runs the traffic scenario at the corresponding speed.

- **Flight trajectory buttons:** Below the "Flight Trajectory" label on the interface, there is a set of buttons. The number and labels of these buttons depend on the number of aircraft in the current traffic scenario. Clicking on each button displays the corresponding trajectory of the matching call sign in the main visualization window.

- **Resolution option buttons:** There are three buttons below the speed control buttons, which are used to visualize the available conflict resolution options. "Resolution_1" (red color) refers to the optimal resolution, "Resolution_2" (blue color) refers to the balanced resolution, and "Resolution_3" (green color) refers to the conformal resolution.

- **Recording ATCOs' interactions:** For each conflict scenario, there are nine radio buttons for each ATCO to rank the recommended resolution (bottom right of Figure 5.2). The response is recorded into a log file for further analysis, as shown in Figure 5.2.

5.2.3 Results and Discussions

Table 5.2 summarizes the overall conflict resolution preferences across ATCOs. In total, 120 conflicts were shown to all ATCOs (20 each). Due to errors in recording ATCO F's data, the total data collected was less than 120. Cumulatively, it is evident that conformal conflict resolution was the most favorable choice for 70 scenarios (61.4% selection). Overall, the optimal conflict resolution was the most favorable for only 22.8% of the scenarios, with conformal and balanced resolutions being selected for 77.19% of the scenarios. On similar lines, the balanced conflict resolution was the second choice for 62 scenarios (54.3% selection), with a combined selection of balanced and conformal resolutions for 75.4% of the scenarios. The optimal conflict resolution was the second most favorable option for 24.5% of the scenarios. These results are interesting in the sense that although ATCOs mainly preferred conformal resolutions to resolve conflicts, there is an inclination towards balanced conflict resolutions, which are an improvement over conformal resolutions in terms of maneuver efficiency. This also justifies our assumptions of group conformal behavior of the participants, since they demonstrate group-conformal characteristics in terms of their conflict resolution strategies.

Table 5.2: Overall Summary: Distribution of ATCOs' preferences for the most favorable, second favorable, and least favorable conflict resolution options across all the scenarios.

	Optimal	**Balanced**	**Conformal**
Most favorable resolution	26	18	**70**
Second favorable resolution	28	**62**	24
Least favorable resolution	**61**	34	10

Figure 5.3 shows the ATCOs' selections for the most favorable, the second favorable, and the least favorable conflict resolution choice for the presented conflict scenarios. Except for ATCOs A and F, the other ATCOs prefer the conformal prediction as the most favorable option for the majority of scenarios. The data distribution in Figure 5.3 (a) suggests that while A strongly prefers the optimal resolution (red), other ATCOs lean more towards the conformance model (green). In the case of ATCO B and F, the optimal option and the balanced option are not preferred in any scenario, respectively.

Figures 5.3 (b) and 5.3 (c) show the second and least preferred options for recommended resolutions, respectively. Except for ATCO A, ATCOs generally preferred balanced recommendation options over the optimal ones. This preference could arise from the benefits of balanced options in maintaining sufficient safety separation, as opposed to focusing solely on minimizing cross-track distance, which is the goal of the optimal resolution options. Based on feedback from ATCOs, conformal prediction is the most preferred option for conflict resolution, followed by the balanced option. The optimal option is not favored by ATCOs due to the narrow safety separation achieved and limited room for accommodating uncertainty in the environment. After conducting the experiments, the ATCOs provided feedback on the Conflict Resolution system through a post-experiment

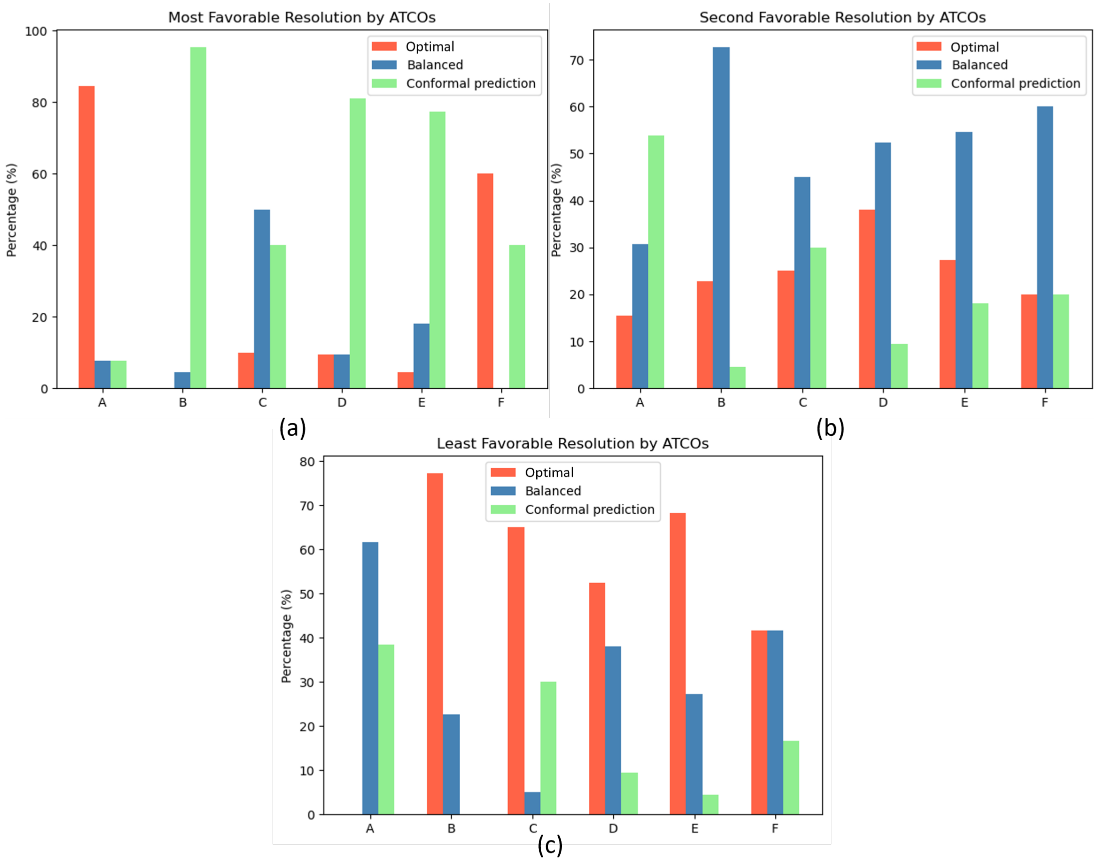

Figure 5.3: Conflict resolution preferences of the ATCOs. To each ATCO, three conflict resolution options were presented: optimal, balanced, and conformal prediction.

questionnaire. The features that influenced ATCOs' decisions on resolution options included cross-track deviation of the maneuvering aircraft and the achieved safe separation in the resolution. Notably, ATCOs stated that their choice of resolution options would also depend on the current weather conditions and the traffic status in both the current sector and the adjacent sectors, as well as the available maneuvering space.

Additionally, 2x2 and 3x3 tests of independence were performed to check the statistical significance of the collected results. First, a contingency test was performed to analyze if there is a dependency between the resolution options (optimal, balanced, and conformal) and the favorability (most favorable, second most favorable, and least favorable). ($\chi^2(4)$ $= 102.51$, $p < 0.001$) affirms that the results are significant and that the favorability levels and the choice of resolution are dependent. A second test was performed to analyze the dependence of the resolutions (balanced and conformal) on the resolution option (most favorable and second most favorable). ($\chi^2(1) = 44.63$, $p < 0.001$) affirms that the results are significant.

Table 5.3: Post-experiment feedback from ATCOs on the conflict resolution experiments and the proposed decision support system.

ATCO	**Usefulness**	**Realistic Scenarios**	**Improved decision making**	**Reasonable Conflict Resolution strategies**	**Enough Resolution strategies**
A	7	7	7	6	6
B	6	6	6	6	6
C	5	2	3	4	4
D	5	6	6	6	5
E	5	4	5	4	5
F	4	4	4	3	3

Table 5.3 shows the feedback from ATCOs on 5 different criteria, where the highest score (7) is associated with the strongest agreement and the lowest score (1) is associated with strong disagreement. All ATCOs found the system useful, except ATCO F, whose feedback was neutral (neither agree nor disagree). Since the scenarios consisted of two aircraft conflicts, the ATCOs mentioned that the scenarios could be made more realistic, with the presence of surrounding traffic and incorporating weather effects. All the ATCOs reported high levels of situational awareness and low workload, due to the nature of the conflict scenarios. Further, most of the ATCOs agreed that such a decision support system would improve their decision-making. All ATCOs except C and F agreed that the provided conflict resolutions were reasonable in terms of the maneuver execution and that three options were sufficient for the ATCOs to take a decision. As such, the ATCOs' feedback included increasing the complexity of the scenarios in terms of the sector traffic and including other maneuvers such as flight level change and speed control for future experiments.

5.3 Summary

This chapter discussed two human-in-loop experiments involving experienced ATCOs. The first experiment aimed to investigate the consistency of ATCOs' strategies and the extent of acceptance of the machine learning model predictions by these ATCOs. The second set of experiments was performed to identify the type of conflict resolution maneuver (optimal, balanced, or conformal) that the ATCOs would prefer when presented with conflict scenarios.

The experiment to test ATCOs' acceptance highlighted important insights into ATCO strategies based on their experience levels. ATCO A demonstrated consistent strategies for conflict resolution, which can be attributed to ATCO A's significantly greater experience in radar operations. ATCO B showed consistency in strategy selection but for significant cases, opted for ATCO A's relatively efficient strategies. Overall, ATCO A selected the original strategy as one of the resolution preferences for 97% of the scenarios and the predicted strategy as one of the options for 78% of the scenarios. ATCO B selected the conflict resolution strategies depicting ATCO B's original strategies for 68% of the scenarios. The results from this exercise also indicate that conformance can increase the acceptance of automation tools in air traffic control operations and are in line with the previous work in literature [68]. Thus, such an approach has the potential for use in assisting ATCOs in air traffic conflict resolution. From a wider perspective, we recognize that capturing ATCOs' strategies solely through these experiments is not possible. This is because ATCOs' behavior can vary due to a variety of other external factors (related or unrelated to the job), which may influence their preferences directly or indirectly.

The experiments to identify the preference of the type of conflict resolution highlighted that the majority of the ATCOs prefer conformal prediction and balanced resolution as the top two preferred choices. Two ATCOs selected the optimal resolution option as their first choice of conflict resolution maneuver. The feedback from the ATCOs also revealed that such a decision support tool may be useful for the ATCOs and assist them in improved decision-making. Nonetheless, additional experiments with increased complexity and more ATCOs must be performed to better estimate the metrics in Table 5.3.

Chapter 6

Agent-Based Conflict Resolution Strategies: Flow-Centric Paradigms

The air traffic control paradigm is shifting from sector-based operations to flow-centric approaches to address the scalability limitations of geographically-bound air traffic sectors and to meet the growing demands of air traffic. These future concepts of operations differ from traditional air traffic operations, especially in maintaining safe separation between flights. Flow-centric operations are characterized by maintaining safe separation between traffic flows (both inter-flow as well as intra-flow), in contrast to current standards of maintaining separation between pairs of aircraft.

This chapter presents a novel approach for resolving air traffic conflicts in flow-centric en-route airspace by employing a combination of a model-free Deep Reinforcement Learning policy and a self-stabilizing graph structure. The problem is formulated as a sequential decision-making task in a large action space, requiring a series of decisions to be made over time to resolve potential conflicts at both the inter-flow and intra-flow levels, while adhering to the flow plans and subsequently reaching the destination.

6.1 Flow-Centric Operations: Introduction

The scalability constraints of the current sector-based air traffic control hinder meeting future demands and ensuring sustainable growth [97]. As discussed earlier, splitting the sectors to accommodate increasing traffic becomes infeasible beyond a certain threshold, rendering the sector size impractical for operations. As a result, alternative concepts, such as flight-centric and flow-centric operations (FCOs) are being explored. In flow-centric operations, a traffic flow refers to a group of aircraft operating on multiple airways with distinct geographical characteristics, including flight trajectory orientation, proximity to their current geographic area, flights originating from and heading towards the same region, and flights proceeding to destinations within the same area/region. [17]. FCOs aim to optimize traffic at a regional level, improve the workload distribution across under and over-loaded ATC units, and overcome the scalability constraints of current geographic sectors. This approach shifts the ATCOs' responsibility from managing all traffic within a given sector to overseeing a specified number of aircraft throughout their flight segment within an airspace.

Air traffic conflict resolution is a crucial aspect of air traffic control, wherein ATCOs bear the responsibility of ensuring the safety and efficiency of flight operations. In the current operations, each sector is controlled by a separate ATCO, who is in charge of ensuring safe and efficient flight operations in that sector. ATCOs possess knowledge of individual aircraft's movements and intentions, maintain situational awareness regarding potential safety violations, and are trained to respond effectively to such scenarios. Communication between ATCOs from different sectors typically occurs during aircraft handover/takeover or in situations where aircraft deviate from their designated reporting waypoints due to weather or other factors. In FCOs, the responsibilities of ATCOs remain unchanged, but the characteristics significantly differ from existing structured airspace operations. ATCOs handling a flow must ensure safe separations between the flows as well as the aircraft of the same flow. This makes it a two-level conflict resolution problem. Thus, greater situational awareness is required to monitor and ensure safety at both levels.

Furthermore, a collaborative approach between the ATCOs managing different flows is required at all times. For conflict resolution, factors such as the type of maneuvers, the maneuver frequency, and information sharing from both, intra-flow and inter-flow aspects also become crucial.
With the fundamental differences in sector-based and flow-centric air traffic operations and the increasing air traffic demand, the existing air traffic conflict resolution advisory tools based on structured airspace are not applicable in a flow-centric setting. To address this challenge, we propose a two-level model-free and learning-based approach to resolve air traffic conflicts using deep reinforcement learning (DRL). The proposed approach formulates the conflict resolution task as a decision-making problem within a large and complex action space. It employs a DRL-based policy to ensure safe separation in inter-flow conflicts and a self-stabilizing graph structure to ensure safe separation in intra-flow conflicts. By combining speed and heading maneuvers, our approach ensures safe separations between aircraft and facilitates conflict resolution in FCOs.

According to the existing definitions, flow-centric operations (FCOs) involve groups of aircraft operating on multiple airways, guided by their geographical characteristics [17]. Based on the insights and reasoning derived from the limited literature on FCOs, we formulate the FCOs as illustrated in Figure 6.1. The flow routes are determined based on future demands, and the flow size (width and aircraft number) may vary depending on the traffic volume. Figure 1 illustrates the flow routes within a specific airspace region, where the width of each flow route represents the level of traffic demand on that particular route. Furthermore, in order to accommodate increased airspace requirements, a flow route may be extended to additional flight levels. To demonstrate the potential conflicts that may arise, Figure 6.1(b) presents an enlarged view of two flows, with arrows indicating the direction of flow for each route.

With multi-level dimensions of each flow, the most preferred resolution would involve either flow speed change or heading change, or a combination of both. Furthermore, from the operational aspect, there are significant challenges that must be addressed. For instance, the cruise speed of different aircraft in a flow may vary, resulting in dynamic

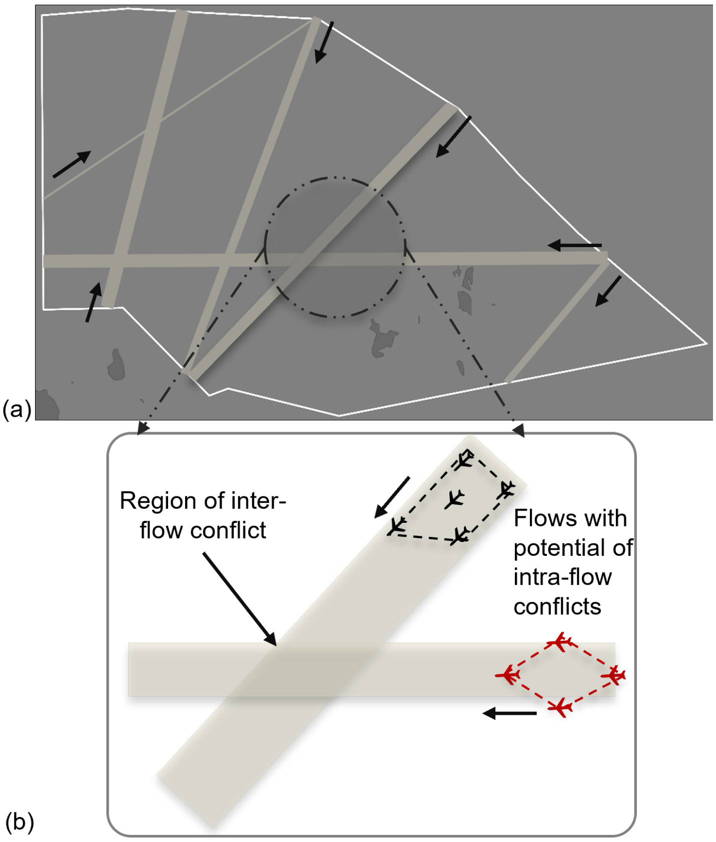

Figure 6.1: An illustration of flow-centric operations (FCOs). (a) Multiple flow routes with directions in a given airspace: the width of each flow depicts the corresponding traffic demand passing through it. (b) Magnified view of two crossing flows with the potential of inter and intra-flow conflicts.

flow structures. In other words, the flow sizes may increase or decrease depending on the aircraft speeds. This may lead to situations involving loss of separation within the flows. Thus, along with monitoring inter-flow separations, a mechanism to further ensure safe separation within the flows is also required. Decisions should also be made whether, in case of a detected conflict, the maneuver should be implemented by the entire flow/both flows or just the affected aircraft. This will govern the interactions between the ATCOs and the aircraft, necessitating mechanisms for splitting and merging flow for conflict resolution. It also involves detecting and resolving potential secondary conflicts. Overall, at first glance, FCOs may appear similar to conventional structured airspace operations. However, upon deeper analysis, the underlying complexities of FCOs become evident.

6.2 Flow Centric Conflict Resolution as a Reinforcement Learning Problem

This study introduces a novel approach for resolving air traffic conflicts in flow-centric airspace, employing a combination of a model-free Deep Reinforcement Learning (DRL) policy and a self-stabilizing graph structure. The problem of air traffic conflict resolution is formulated as a sequential decision-making task, requiring a series of decisions to be made over time to resolve potential conflicts at both the inter-flow and intra-flow levels, while adhering to the flow plans and subsequently reaching the destination. This is a centralized coordination scheme where an agent takes actions to ensure safe separation between the flows.

In this context, a flow plan is defined similarly to a flight plan, where each flow (representing a group of aircraft) is expected to reach the subsequent reporting waypoint at a specified time, ensuring minimal chances of secondary flow conflicts in the airspace. The formal framework of Markov Decision Processes (MDPs) is utilized to model the sequential decision-making problem, incorporating states, actions, and rewards. The selected actions not only impact immediate rewards but also influence subsequent states and, consequently, future rewards. In the proposed approach, the system's state, subsequent actions, and rewards are influenced not only by the learned policy of the agent but also by the output of the self-stabilizing graph structure, which guarantees intra-flow safe separation. This integration ensures the effective resolution of air traffic conflicts in FCOs. Figure 6.2 shows the overall concept diagram for the proposed approach.

The present study focuses on the specific scope of addressing conflict resolution in FCOs involving two crossing flows with varying numbers of aircraft and speeds within each flow. The dynamic nature of these flows, which evolve over time steps due to different aircraft cruise speeds, is taken into account. The action space includes options such as heading change, speed change, or a combination of both, applicable to one or both flows. These actions are based on the nominal aircraft performance parameters in the cruise flight phase. Furthermore, since the research focuses on en-route flight phase with sufficient

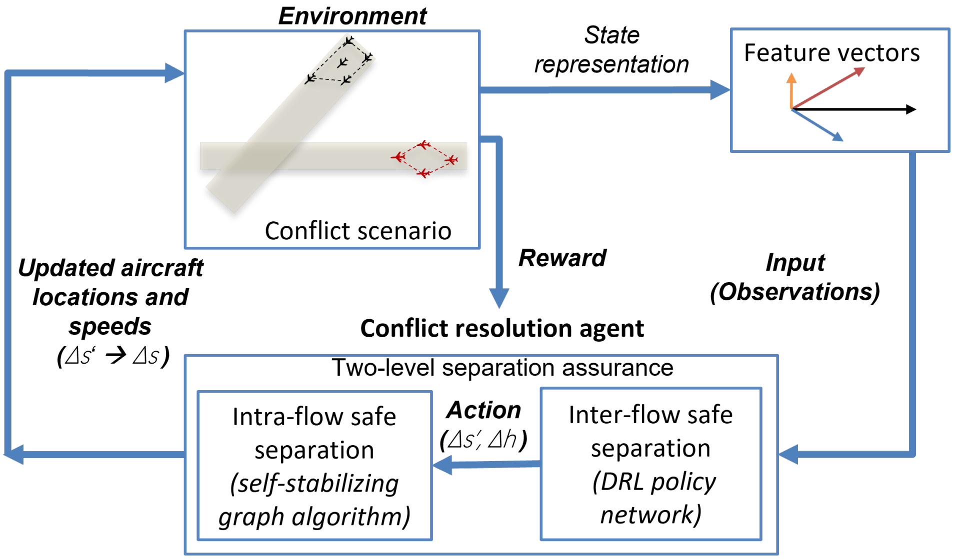

Figure 6.2: A concept diagram for the interaction between the agent and the learning environment. Scenarios involving interflow and intra-flow conflict are generated and the vector representation of the extracted features is used by the agent to propose an action based on the learned policy, thereby reaching a new state and receiving a certain reward. The updated actions pass into a self-stabilizing algorithm which ensures intra-flow safe separation and outputs the updated location and speed of the aircraft in both flows.

inter-flow and intra-flow separation, effect of factors such as aircraft wake turbulence is not considered. In this approach, the selected action by the agent is implemented on the entire flow, affecting each aircraft within the flow. The allowed range of aircraft speeds is determined based on Automatic Dependent Surveillance-Broadcast (ADS-B) data (Figure 6.3). By training a policy that maximizes the cumulative reward over time, the proposed approach aims to effectively resolve air traffic conflicts while minimizing deviations from the flow plan. The model-free nature of the approach enables it to handle the added uncertainties, and achieve the global objectives of safety, efficiency, and adherence to flow plans.

6.3 Methodology

6.3.1 Conflict Scnearios

In this study, we use one day's ADS-B data of flights in the South East Asian region to obtain the inter-flow and intra-flow speeds for flow movements. The cruise speed distribution obtained from the data is depicted in Figure 6.3. For each conflict scenario, inter-flow velocities are randomly selected from this distribution. Thus, the inter-flow velocities range between 360 knots (kts) and 540 kts. Additionally, the velocity of each aircraft within a flow, known as the intra-flow velocity, is sampled from a normal distribution. The mean of this distribution is set to the flow speed, while the standard deviation is set to 30 kts. This approach ensures that there is variation in aircraft velocities within a flow, reflecting real-world operational conditions. In accordance with the definition provided by the International Civil Aviation Organization (ICAO), a crossing conflict is characterized by conflicts occurring within the range of 45° to 135° on both, the left and the right of the reference aircraft [5]. Therefore, the scenarios considered in this study involve crossing conflicts at three specific angles: 45°, 90°, and 135°, in an abstract airspace. The time interval preceding the occurrence of a conflict is approximately 19-23 minutes, allowing for a sufficient look-ahead time. To introduce variability, the starting points of each flow are randomized by applying certain offsets, ensuring that the loss of separation occurs at different locations in each scenario. Additionally, the number of aircraft within each flow ranges from 3 to 10 in each scenario. This range of is selected to ensure sufficient scenario variability while maintaining a reasonable flow size. These aircraft are randomly generated within a square region of side 20 NM at the start of the scenario. These aforementioned variations are included to enhance the generalization capabilities

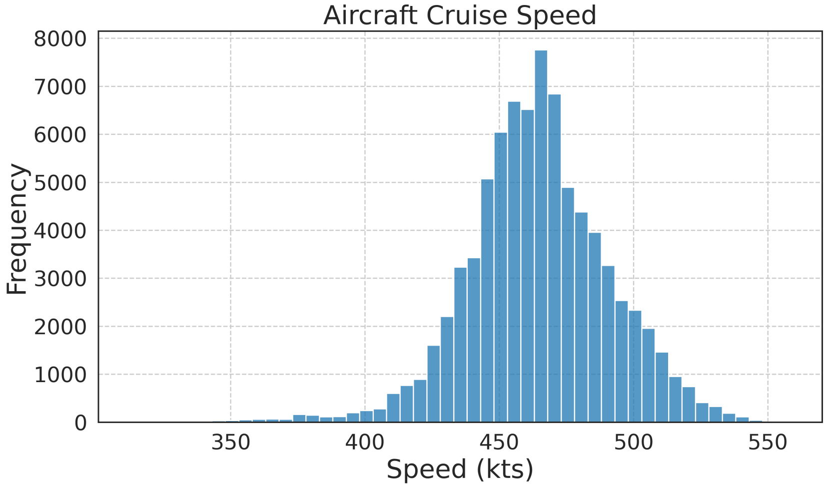

Figure 6.3: Distribution of aircraft cruise speed from the ADS-B data. For each scenario, inter-flow and intra-flow speeds are sampled from this distribution.

of the proposed model. For the remainder of this work, the terms 'conflict scenario' and 'episode' are used interchangeably. Figure 6.4 shows a representative conflict scenario with two flows. Here, D_{min} is the minimum distance between the flows, and h_{i1} and h_{i2} are the initial headings of the flows. If no maneuver is performed D_{min} is bound to be zero. In case of a maneuver performed by the agent on any of the flows,D_{min} gives the closest point of approach (the minimum distance between the flows). These flows must reach their destinations (reporting points) based on the flow plans.

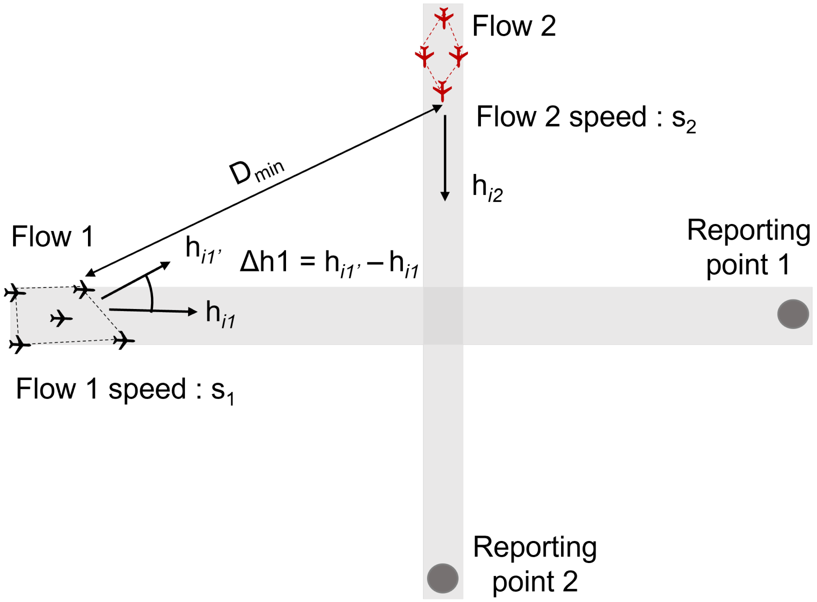

Figure 6.4: An illustration of the conflict scenarios with two flows in a 90° crossing conflict. $\triangle$h1 represents the difference between the original heading of the flow (h_{i1}) and the current heading ($h_{i1'}$) based on the selected action.

6.3.2 Markov Decision Process Framework

The Markov Decision Process (MDP) framework has several key components, which include providing the agent with sufficient information to support decision-making, evaluating the agent's actions, and providing feedback to the agent in the form of a reward. To facilitate the learning process for resolving conflicts, we have developed a scenario generator that produces conflict scenarios (discussed in the previous subsection) and represents them in a way that can be perceived by the agent. The agent's actions are defined and the mapping from these actions to the maneuvers taken by the agent is established. A reward function has been designed to assess the effectiveness of the maneuvers suggested by the agent. These components are discussed in detail as follows:

Observation space

The observation space contains all information that the agent can receive at each step in order to decide the next action. In this work, the observation space includes information like the distance between the two flows, the number of aircraft in each flow, the initial (at simulation start) and current headings and speeds of the two flows, the cross-track deviations of the flows (i.e the distance of each flow from its original path) and the distance of the flows from the target locations. Furthermore, to indicate the current direction of the flow with respect to the original flow path, the change in the heading (positive or negative) is also added to the observation space. Overall, this is a vector of 17 elements.

Action space

In the current work, the agent takes two actions at each time step i.e. one for each flow. One time step implies one minute in the scenarios. The actions include speed change, heading change, or a combination of the two. Given the initial flow speed, s, and the initial heading of the flow, α, and the intra-flow speeds, the agent has the following available action space for each flow:

$$A_S = [\delta s_-, 0, \delta s_+]$$

$$A_h = [\delta \alpha_-, 0, \delta \alpha_+]$$

At each step, a speed change of ± 30 kts or 0 kts, a heading change of ± 30° or 0°, or a combination of these can be made. Thus, a total of 9 actions are available for each flow, scaling the total number of available actions for both flows to 81. The action space is encoded in Table 6.1. Due to such a large action space, DRL is a suitable approach for conflict resolution.

Table 6.1: Encoding of the agent's action space.

Encoding	Action(δs, δh)	Encoding	Action(δs, δh)
0	-30 kts, -30°	5	0 kts, 30°
1	-30 kts, 0°	6	30 kts, -30°
2	-30 kts, 30°	7	30 kts, 0°
3	0 kts, -30°	8	30 kts, 30°
4	0 kts, 0°		

Reward structure

The reward mechanism is based on two primary criteria: safety and efficiency. Thus, a maneuver that successfully separates the two flows is given a positive reward and any maneuver that causes loss of separation is given a heavy negative reward. On similar lines, maneuvers that reduce the induced delay and allow the flows to successfully reach the destination are promoted with positive rewards. An inter-flow loss of separation in the simulations is a situation where the separation between flows is less than 5 nautical miles (NM). Efficiency is measured in the context of the deviation from the flow plan, the cross-track deviations, and closest point of approach between the flows. The agent is also penalized for taking excessive actions while resolving the conflicts. At each step, the agent receives a reward $R_s(s', a)$, which is based on the action a and the resultant state vector s'. This reward is based on the number of actions ($n_{actions}$) and the deviation from the flow plan ($dev_{(flowplan)}$) which is estimated by the change in the number of steps required by both flows to reach the destination, as compared to the initial number of steps (based on the flow velocities at the simulation start).

$$R_s(s', a) = \begin{cases} -\alpha * n_{actions} \\ \\ -\beta * dev_{(flowplan)} \end{cases} \tag{6.1}$$

The agent receives rewards at episode termination based on the criteria in equation 6.2. Here, d_{min} refers to the distance between the two flows and d_{sep} is the minimum safety separation. Invalid maneuvers involve aircraft speeds and headings outside the specified range. $d_{(f1+f2)}$ is the sum of the distances of the flows from their respective destination. An episode finished when both the flows reach their destinations. The total reward after each episode is represented by equation 6.3, where N is the total steps in the episode.

$$R_t(s', a) = \begin{cases} -\gamma & \text{if } d_{min} < d_{sep} \\ -\delta & \text{if } a \text{ is an invalid maneuver} \\ \lambda & \text{if both flows reach the destination} \\ \frac{\Gamma}{d_{(f1+f2)}} & \text{otherwise} \end{cases} \tag{6.2}$$

$$R_T = \sum_{n=1}^{N} R_s + R_t \tag{6.3}$$

Here, the values of the parameters $\alpha, \beta, \gamma, \delta, \lambda$ and Γ are 0.002, 0.0002, 10, 5, 10, and 5, respectively. These values were obtained after multiple iterations and model performance evaluations.

6.3.3 Learning Algorithm

Inter-flow conflict resolution

FCOs require inter-flow and intra-flow level air traffic conflict resolution. Therefore, in this work, the learning algorithm also consists of two stages: inter-flow conflict resolution using a reinforcement learning algorithm and intra-flow safety separation using a self-stabilizing graph algorithm. Proximal Policy Optimization (PPO) [112] has been used to train the agent for inter-flow air traffic conflict resolution, due to its faster convergence as compared to off-policy methods, for our specific research problem. PPO is a deep reinforcement learning algorithm that is widely used in various domains of intelligent transportation systems, and has achieved state-of-the-art performance in several benchmarks. It is a model-free policy optimization algorithm that operates by updating the policy with a clipped surrogate objective function to ensure stable training and avoid large policy updates. PPO has several advantages, including the ability to handle large, continuous, or discrete action spaces, providing a guarantee of monotonic improvements in the objective, and good sample efficiency. Please refer to section 4.2.4 for details of the algorithm.

Algorithm 3 Intra-flow Separation Algorithm

```
procedure INTRA-FLOW SEPARATION($G, (x, y)$)
    $sorted\_nodes \leftarrow$ sortNodes($G.V, (x, y)$)
    for $i$ in $sorted\_nodes$ do
        $basenode \leftarrow i$
        for j in $sorted\_nodes$ do
            if index$j >$ index$i$ then
                $dist \leftarrow$ calculateDistance($basenode, j$)
                while $dist <$ safetybuffer do
                    $delta \leftarrow$ safetybuffer $- distance$
                    $j' \leftarrow$ moveNode($j, delta, \alpha$)
                    $dist \leftarrow$ calculateDistance($basenode, j'$)
                end while
            end if
            $sorted\_nodes \leftarrow$ sortNodes($G.V, (x, y)$)
        end for
    end for
end procedure
```

Intra-flow conflict resolution

To ensure intra-flow safe separation, we represent each flow as a self-stabilizing graph structure. Self-stabilizing algorithms are a fundamental branch of fault-tolerant computing and were first introduced by Dijkstra [118]. Based on pre-defined criteria, two states for the system can be defined, which are (i) the legitimate state, and (ii) the illegitimate state. Self-stabilizing algorithms are resilient to transient faults i.e. if a perturbation brings the system to an illegitimate state, then the system must be able to again reach a legitimate state after a finite number of moves without any external intervention. These have been used to solve synchronization problems [119], constructing breadth-first trees [120] and graph explorations [121]. From the context of air traffic conflict resolution, an illegitimate state implies a situation where aircraft witness a loss of safe separation. A state where all the aircraft in a flow are safely separated is considered a legitimate state.

The graph spacing algorithm to ensure intra-flow separation is depicted in Algorithm 3. Let $G = (V, E)$ represent a uni-directed, complete, weighted graph for each flow, where V represents the vertices/nodes (the aircraft) and E represents the edges, with the distance between the nodes implying the weights. At the inter-flow level, an action is taken to change the state of the environment at each step of the episode based on the current observations received by the agent. In other words, this action updates the locations of the aircraft in the flow. At intra-flow level, since the aircraft's speeds differ, this may cause loss of separation if the trailing aircraft in a flow are flying faster, as the episode progresses. Thus, before the state is updated after each step, the self-stabilization algorithm checks if the distance between any two aircraft in a flow is below a specified threshold. If yes, the algorithm provides updated speed changes and consequently, the updated locations of the aircraft, which are then used to transition to a new state. For this, the farthest aircraft with respect to the destination is taken as a reference for sorting the nodes (Algorithm 3). Each episode always starts with a legitimate state. The threshold is currently set to 7 NM. This accommodates the operational uncertainties associated with aircraft location and speed.

6.4 Experiments and Results

6.4.1 Experimental Setting

The PPO algorithm used in this work is adapted from the Stable-Baselines 3 [113]. The parameters for the PPO algorithm are in Table 6.2. The training process consists of 3,000,000 scenarios. Each scenario lasts for a maximum of 60 time steps. Training is performed on the Intel(R) Core(TM) i9-9900X CPU which takes about 4 hours.

Table 6.2: Parameters for PPO training.

Parameters	Value
Training timesteps	3e+6
Learning rate	3e-4
Discount factor	0.99
Clipping coefficient	0.2
ANN architecture	Multi-layer Perceptron
Optimizer	Adam
Hidden layers	64 X 64
Activation function	Tanh

Figure 6.5 shows the model's convergence after 3,000,000 scenarios. During training, the model's performance is measured by the average reward achieved and the number of steps taken to complete each episode. One episode should take approximately 45 steps to

finish. Initially, the agent takes over 50 steps to reach the destination but as the iterations increase, it identifies optimal paths for both flows to reach their destinations. The number of steps for the episodes to terminate eventually stabilizes at around 45. This is because additional steps imply a deviation from the flow plan, incurring a negative reward for every additional step. On similar lines, in an ideal case with no actions and no conflicts, the theoretical maximum reward is 10. Since actions and deviations from the flow plan incur negative rewards, the rewards stabilize around 8.5. Even though the model can achieve high average rewards after 1,000,000 iterations, its performance is unstable. Therefore, we run the simulations to continue training till the model converges.

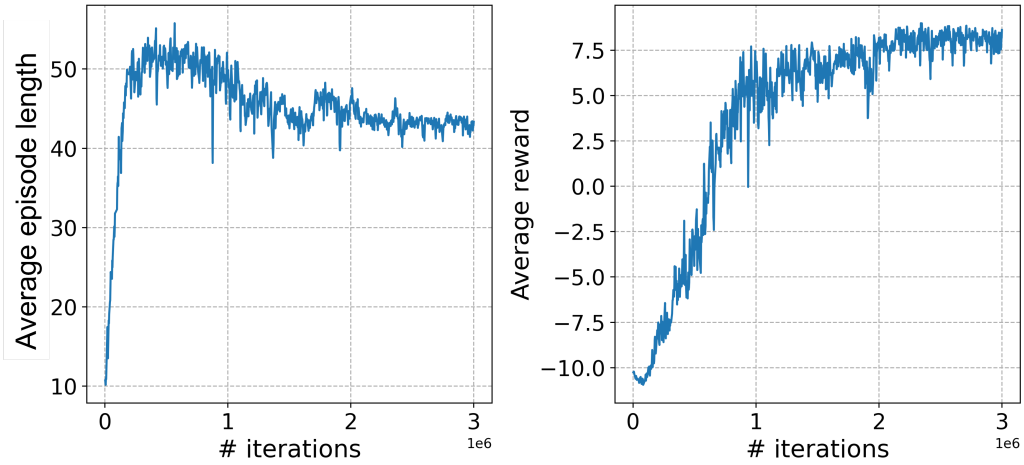

Figure 6.5: Model convergence: episode length (left) and the average reward (right).

To evaluate the trained model's performance, results are documented for 1000 test set scenarios in the following subsections. In the discussions, Flow 1 refers to the flow moving in the horizontal direction. Flow 2 refers to the second flow which intersects the path of Flow 1.

6.4.2 Safe Separation Analysis

Inter-flow safe separation

The Air transportation system is a safety-critical system. Thus, model performance in terms of safe separations achieved is important. The model achieves a safe separation for 100% of scenarios. This implies that the policy learned by the agent is successful in resolving the conflicts despite the associated uncertainties in the speed of the flows, and their dynamic shape and size.

Intra-flow safe separation

Figure 6.6 shows the number of actions (speed adjustments) taken by the intra-flow safe separation algorithm at each time step. The figure shows the results for 10 aircraft in each flow. Due to the varying speed of the aircraft in a flow, there are instances where the intra-flow separation reduces below the specified threshold as the scenario evolves. To ensure safe separation, the algorithm updates the aircraft speed (and hence the location) before the state of the environment is changed. Furthermore, Figure 6.10, right, shows the minimum distance between the aircraft in Flow 1, the lowest value of which is 7 NM. Thus, the intra-flow safe separation algorithm successfully ensures that the aircraft within a flow are separated by the defined threshold at each time step.

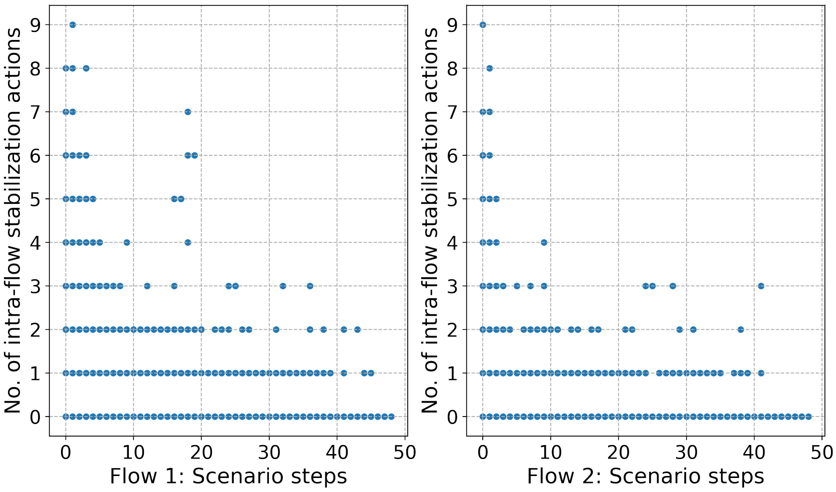

Figure 6.6: Number of actions taken by the intra-flow safe separation algorithm at each time step before the state of the environment is updated.

6.4.3 Maneuver Efficiency Analysis

We analyze maneuver efficiency in terms of adhering to the flow plans, the cross-track deviations (CTD) during each episode, and the closest points of approach (CPA) between the two flows. Figure 6.7 highlights the delay distribution of the two flows. The negative values imply reaching the destination early, which is not preferred since this might also

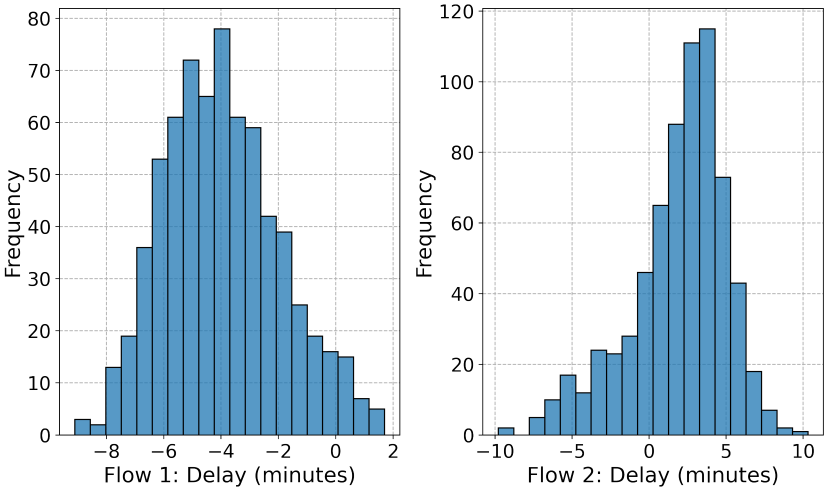

Figure 6.7: Delay distribution in terms of the difference between the original flow plan and the actual flow plan.

influence the events (secondary conflicts) further down the line. In terms of the absolute delays, the average values for Flow 1 and Flow 2 are 2.53 minutes and 9.49 minutes.

The maximum CTDs of the two flows from their original paths are shown in Figure 6.8. There are two potential reasons for the higher CTDs of Flow 1 in some scenarios. First, with the increase in the number of aircraft, the flow structure itself increases, given the intra-flow separation must be 7 NM. Second, the shape and size of the flows are dynamic. As each conflict scenario evolves, these attributes of the flows change depending on the velocities of the aircraft. Thus, the size of a flow changes based on the slowest and the fastest aircraft as its member.

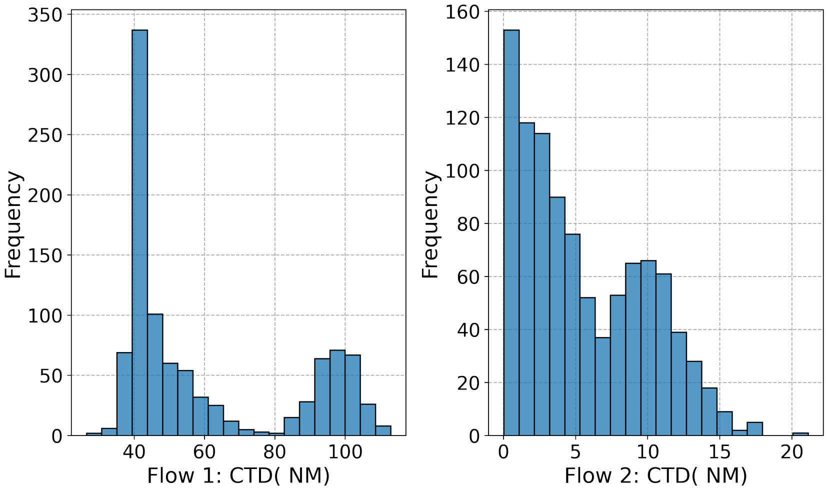

Figure 6.8: The maximum cross-track deviation distribution (NM) for Flow 1 and Flow 2.

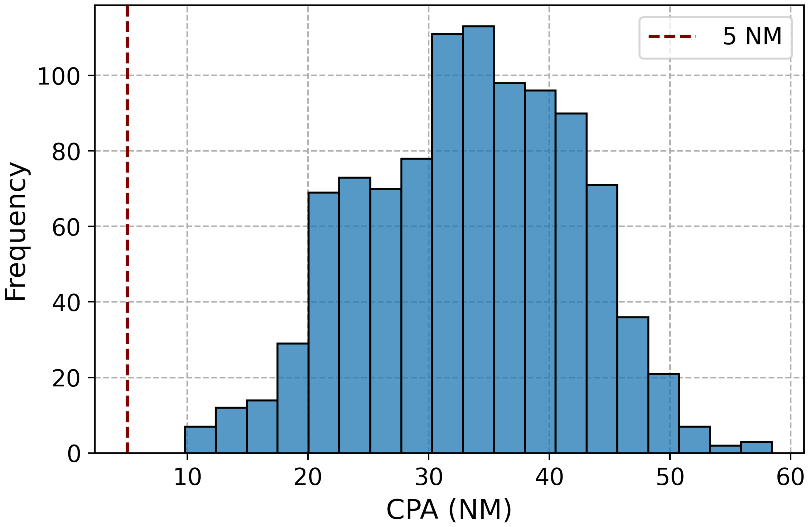

Figure 6.9: CPA distribution between Flow 1 and Flow 2. The CPA values are always above 5 NM which indicates 100% safe separation assurance.

Similarly, the CPA between the two flows is relatively higher. This is measured as the smallest distance between any two aircraft in the two flows (Figure 6.9). As discussed, the higher values of the efficiency metrics in some cases are due to the dynamic nature of the flows. Figure 6.10 shows the change in the flow size for 100 conflict scenarios with the evolution of each scenario (time steps) for Flow 1. The size is measured in terms of the maximum distance between an aircraft pair in the flow and the minimum distance between an aircraft pair in that flow.

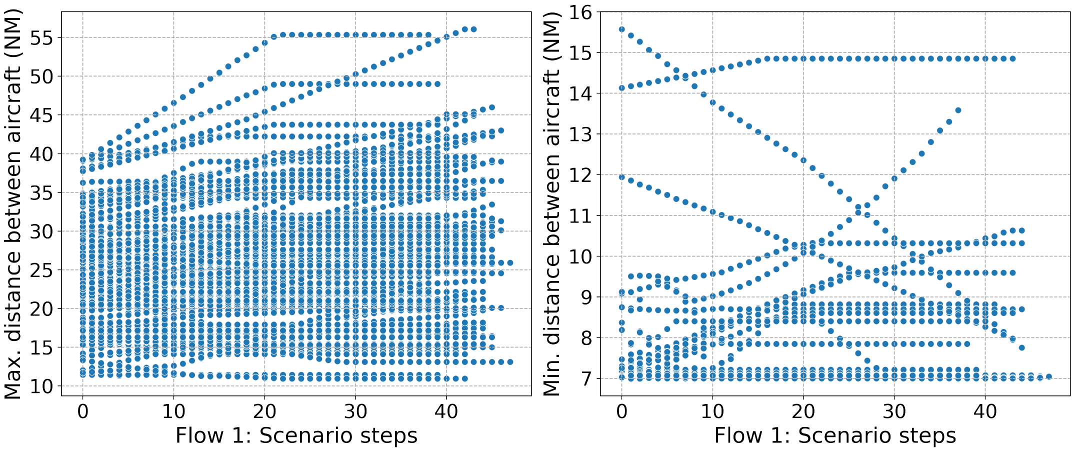

Figure 6.10: Flow1 size and its variation with scenario evolution, for 100 conflict scenarios. Variation of the maximum distance (left) and the minimum distance between aircraft (right) with respect to the episode length.

It can be seen that the maximum distance between aircraft (Figure 6.10 left), varies from approximately 12 NM to 39 NM depending on the number of aircraft in the flow. Further, this size also varies as the scenarios evolve, due to different speeds of the aircraft. With such dynamics associated with both flows, it is significantly complex to achieve better efficiency metrics in terms of lower CTD, CPA, and delay. In terms of the conflict scenarios and the current action space, the following have been observed. Since the acute conflict scenarios (45°) with larger flow sizes witness a loss of separation sooner, they require higher deviations to ensure safe separations. Thus, a larger look-ahead time would allow for smaller and relatively more efficient maneuvers. Furthermore, the current action space allows for a change of 30 kts in one minute, which is very gradual. Higher magnitudes of speed change may also allow for conflict to be resolved by speed change only, leading to significantly lower CTD and improving other metrics as well.

6.5 Summary

The flow-centric concepts of operations differ significantly from the current sector-based operations, especially in maintaining safe separation between flights. The existing conflict resolution approaches (mathematical and learning-based) are not suitable for applications in a flow-centric setting. Thus, in this work, we have proposed a novel approach for air traffic conflict resolution for FCOs which involves a policy-based learning for inter-flow conflicts and a self-stabilizing graph approach for intra-flow conflicts. The methodology ensures 100% safe separations despite the uncertainties present in the learning environment. The absolute values of delays for the flow are 2.53 and 9.49 minutes respectively, which are within acceptable limits, given the uncertainties and dynamics associated with the flows' size, speed, and evolution over time. Results and discussion are also presented regarding the performance of the learned policy in terms of maneuver efficiency and potential improvements in the action space and scenario design to improve the model performance. Nonetheless, further experiments and analysis are still required to improve the model performance in terms of efficiency metrics.

Future work in this direction can be extended in various directions. Since the current research focuses on all aircraft at the same flight level, with no flight level changes in the maneuvers, the research can be extended to flows in 3-dimensional volumes. Further, an investigation on the concept of operations in the cases of merging of the flows or splitting of a flow into two or more flows is required. Since FCOs involve different flows being regulated by separate ATCOs, a multi-agent framework could facilitate the investigation of the interaction between the agents managing different flows, under the supervision of the ATCOs. Other research directions involve the development of algorithms to generate stable topological graph structures to represent the flows, increasing the traffic complexity and extensive simulations and analysis of the challenges such as splitting and merging of flow, that the FCOs bring to the air traffic management systems.

Chapter 7

Conclusions and Future Work

This chapter concludes the thesis. First, a brief summary of the research performed in this thesis is provided. Then, the research questions that were identified in Chapter 2 are revisited and the corresponding findings are summarized. Several challenges associated with the integration of ML-based conflict resolution models in air traffic control are discussed, along with the contributions of this thesis towards addressing some of these challenges. Finally, the opportunities identified for the extensions of the current work and potential future research direction are discussed.

7.1 Thesis Summary

Present-day air traffic control is reaching its operational limits and its effects are manifesting as increased flight delays and its environmental impact due to more fuel burn, passenger inconvenience, and reduced safety and efficiency due to increased ATCO workload. Numerous automation methods have been proposed to assist the ATCOs and alleviate their workload. New concepts of operations such as flow-centric paradigms are also being explored to manage the projected future air traffic. The existing automation methods to assist the ATCOs in their tasks, especially air traffic conflict resolution face a common challenge of limited acceptance, because these methods do not incorporate the ATCOs' conflict resolution strategies in the proposed solutions. Developing methods to address this challenge involves the identification of the strategies, developing a representation of these strategies that can be learned by a machine learning model, and developing machine learning models for air traffic conflict resolution that can generate ATCO-conformal conflict resolutions. Furthermore, advanced air traffic management concepts such as flow-centric operations (FCOs) introduce significant challenges in adapting the present-day air traffic conflict resolution approaches for such novel paradigms.

On these lines, this thesis makes novel contributions in the following three aspects. First, the thesis proposes a methodology to identify the ATCOs' conflict resolution strategies and identify the factors influencing these strategies. This is achieved through novel human-in-loop experiments for resolving air traffic conflict resolution in a high-fidelity simulation environment involving 8 experienced ATCOs. Second, the thesis develops a representation for the identified strategies and proposes supervised learning and reinforcement learning-based conflict resolution models to generate ATCO-conformal prediction. The thesis evaluates the consistency of the ATCOs' strategies and the acceptance of conformal machine learning predictions by the ATCOs. Through human-in-loop experiments on unseen conflict scenarios in a purpose-built simulation environment, the thesis also evaluates the preferences of the ATCOs in selecting optimal, balanced, or conformal conflict resolution predictions for the presented conflicts. Third, to address the future flow-centric air traffic paradigm, the thesis proposes a two-stage conflict resolution approach for flow-centric operations. Conflict resolution in FCOs requires ensuring safe separation at the inter-flow level (between the flows) and intra-flow level (within each flow). Here, inter-flow conflict resolution is modeled as a sequential decision-making problem, and intra-flow separation is addressed using the concept of self-stabilization. The simulated conflict scenarios involved crossing conflicts at 45°, 90°, and 135°angles with aircraft numbers in each flow varying between 3 to 10 and differing aircraft speeds in each conflict scenario. With a constraint that the action selected by the agent is applied to all the aircraft in a flow, the proposed methodology achieves 100% safe separation for both inter-flow and intra-flow conflicts.

7.2 Revisiting the Research Questions

Research Question 1:

What is a suitable representation of the ATCOs' conflict resolution strategies? How can we identify the ATCOs' strategies based on this representation and analyze the factors affecting the strategies?

Findings:

A strategy is a generalization of the actions taken by the ATCOs during a conflict resolution process. These actions are a result of sequential decisions taken by the ATCOs and constitute an end-to-end conflict resolution maneuver. Such a maneuver is represented as a sequenced tuple, $P = (C, T, D, D_c, M_{wp})$.

ATCOs utilize various strategies while resolving air traffic conflicts. The conflict resolution data was collected from 8 experienced ATCOs, out of which 2 were experienced in the airspace simulated for the experiments. The conflict scenarios consisted of crossing conflicts with various conflict angles. From the data collected during the experiments, distinct strategies were identified, with varying consistency demonstrated by the ATCOs.

For example, ATCO A preferred to maneuver the trailing aircraft with a relatively higher heading (average of 35.8°) and cross-track deviations. The direction of the maneuver was always toward the tail of the leading aircraft (i.e. within the cone of conflict) and the merging point was always within the sector. ATCO B exhibited less consistent strategy selection than ATCO A in terms of the direction of maneuver. ATCO B preferred significantly higher cross-track deviations (ranging from 11.1 to 23.1 NM) which also affected the maneuver direction. The remaining ATCOs, who were unfamiliar with the airspace, had similar strategies for conflict resolution, especially for the heading and cross-track deviation. All the ATCOs preferred certain buffers over the 5NM safe-separation requirement.

The ATCOs' strategies were found to be influenced by a multitude of factors, which were identified from the data analysis and feedback from the ATCOs. For instance, the location of the conflict and the time maneuver initiation have an impact on the ATCOs' subsequent decisions, such as maneuver direction. Other factors include the leading and the trailing aircraft in the conflict pair, the location and proximity of the conflict point to the sector boundary, aircraft that enters the sector earlier, the potential visual markers in the sector where aircraft are directed consistently, and the available merging waypoints. individual preferences of the ATCOs in terms of safety, orderliness, and efficiency of the traffic flow in the sector also affect the conflict resolution strategies. Furthermore, different ATCOs prefer different safety buffers in their decisions, which proportionally affects the 'efficiency' of operations. From the data, the average values of the safe separations achieved by the ATCOs ranged between 7.7 to 9.6 NM. Feedback from the ATCOs revealed that the conflict resolution strategies are also affected by the preceding and following sectors, aircraft hand-over/takeover practices, weather-related effects, and traffic flows.

Research Question 2:

What type of machine learning models are suitable for generating ATCO-conformal conflict resolution advisories?

Findings:

This thesis proposes two machine learning-based methods for generating ATCO-conformal air traffic conflict resolution predictions through behavior cloning. The supervised learning approach uses a chained prediction architecture to develop individual sensitive prediction models for ATCO A and ATCO B, and a group conformal model for the ATCOs' who demonstrated similar conflict resolution strategies. The proposed models can learn and predict the conflict resolution strategies of ATCOs with high accuracy for the classification tasks and with low mean absolute error (MAE) for the regression task (for instance, the classification accuracy of above 92.7% for predicting the maneuvering aircraft, MAE for maneuver initiation distance < 5.3 NM, MAE for predicting the heading angle $< 5.3°$). Further, the models demonstrate robust performance up to 7.5% Gaussian noise in the data. The human-in-loop experiments further validate that the conformal predictions are accepted by the ATCOs for the majority of the conflict scenarios.

The second approach models air traffic conflict resolution as a sequential decision-making problem and uses DRL for air traffic conflict resolution, wherein the ATCO's conflict resolution data is used as demonstrations to train an RL agent to perform ATCO-conformal conflict resolution. The trained agent achieves safe separation for 99.8% of the scenarios. Further, the distribution of the closest point of approach and the cross-track deviation for the resolutions by the agent is similar to the from the ATCO's data. Nonetheless, the DRL-based approach requires further tuning of the reward function and the action space to improve the agent's performance.

Research Question 3:

How can air traffic conflicts be resolved in a flow-centric paradigm where traffic is modeled as intersecting flows?

Findings:

This thesis proposes a novel approach for air traffic conflict resolution in Flow-centric operations (FCOs) that utilizes a model-free deep reinforcement learning policy and self-stabilizing graph structure to ensure inter-flow and intra-flow safe separation. The conflict scenarios consisted of crossing conflicts with varying numbers of aircraft and different aircraft speeds for every scenario. The performance of the trained model was evaluated for safety and efficiency (adhering to flow plans, cross-track deviation (CTD), and the closest point of approach (CPA)) of the maneuvers. The trained model was successful in resolving both, inter-flow conflicts and intra-flow conflicts with 100% accuracy. In terms of the absolute delays in the flow plans, the average values for the two flows are 2.53 minutes and 9.49 minutes respectively. Due to the large number of aircraft in the flows in certain scenarios (maximum of 10 per flow), and the intra-flow safe separation requirement of 7 NM, the values of cross-track deviation and CPA were higher. Also, conflicts at acute angles (45°) necessitated such larger deviations in CTD for achieving safe separation.

7.3 Challenges in Integration of ML Based Conflict Resolution Models in Air Traffic Control

The transformative power of machine learning (ML) is reshaping industries across the globe. The key enablers for this transformation are the advancements in data availability, improved computational power, the development of faster and more efficient algorithms, and the industry recognition of the potential of these technologies. With the emergence of ML models as promising tools to optimize decision-making processes, enhance safety measures, and improve overall efficiency, their integration into air traffic control has become a topic of paramount importance and keen interest. This has also been recognized by international organizations such as the International Civil Aviation Organisation (ICAO), the Federal Aviation Administration (FAA), and the European Union Aviation Safety Agency (EASA). However, the integration of such cutting-edge technologies into the highly regulated and safety-critical domain of air traffic control poses substantial challenges for multiple stakeholders, especially from the aspect of air traffic conflict resolution. This includes scientists and researchers engaged in the development of ML models, certification agencies responsible for ensuring adherence to safety standards, and the end-users i.e. the ATCOs, who play a central role in implementing and relying on these technological advancements.

Several challenges have been identified in relation to safety critical systems [122, 123] that may hinder the adoption and integration of these technologies in operations. Recently, EASA published the first set of technical objectives and organization provisions necessary for the approval of Level 1 AI applications (assistance to humans) [115]. EASA has also discussed the safety assessment and guidance for safety-related ML applications in terms of the initial safety assessment (during the design phase) and continuous safety assessment (based on operational data and in-service events) [116]. Researchers have also proposed methodologies to develop AI competencies in aviation education [124] to upskill the workforce to leverage AI safely, effectively, and responsibly.
Some notable challenges from the perspective of the stakeholders in terms of the integration of ML in air transportation are:

7.3.1 Challenges in Conflict Resolution Model Development

These challenges are related to the scientific community that develops the ML models for air traffic conflict resolution. The challenges include:

Robust Decision Making Capability

Air traffic control is characterized by inherent uncertainties due to changing environmental conditions, unforeseen events, and outlier events or edge cases due to the evolving system dynamics. the safety-critical nature of this domain increases the complexity further. The scope of the conflict resolution problems addressed using ML is very specific, with model training and testing in controlled environments. Thus, the current models struggle to handle such variabilities which poses a significant operational challenge. Robust models should be able to provide at least feasible, if not optimal solutions in case of data noise so that certain non-negotiable characteristics such as safety are always assured. Furthermore, with the advancements in ML methods, the threat of adversarial attacks on such safety-critical systems has also increased [125]. Researchers have discussed the susceptibility of ML methods (especially deep learning-based classifiers) against adversaries [126, 127]. Such adversaries are capable of manipulating ML system vulnerabilities and cause them to make mistakes [128], which in the case of safety-critical systems may prove to be catastrophic.

For our research, we have ensured the feasibility of solutions (safe separations) at all times for the chained prediction framework. The performance of this framework has also been tested after the addition of Gaussian noise to the dataset, to ascertain model robustness in the presence of noise in the input dataset.

Monitoring Model Performance

In safety-critical systems like air traffic conflict resolution, monitoring the ML model is a critical challenge that extends beyond the final performance metrics of ML models. Rather, it involves tracking the adaptation of the models to new data, addressing issues like concept drift and model degradation in real-world dynamic environments [129], the identification of potential biases in the models which might be introduced due to the data collection source (eg. the ATCOs and the simulation environments), and safeguarding the system from adversarial attacks as discussed before. Another aspect of system monitoring within the context of conflict resolution is knowing the capabilities of the developed models and hidden functionalities that might surface in case of outlier events or noisy data input, and avoid unintended behaviors. This becomes a challenge as the size and complexity of the models increase, and it becomes difficult to track how different components of the model contribute to the final result.

Trustworthy Models

Trust is a critical factor that must be addressed to ensure safe and efficient use of the developed conflict resolution models, and promote successful ATCO-ML teaming. In this regard, human trust in novel ML-based algorithmic systems is a major challenge. A proposed model must demonstrate characteristics such as conformance to the users' behavior, transparency, and explainability to be trustworthy. While conformance has been demonstrated to improve the acceptance of such model by the end users, factors like transparency and explainability can help the user understand the functioning of the model, and why if at all, the model commits the sort of errors that it does [130]. On these lines, the current research has proposed machine learning-based models for ATCO-conformal conflict resolution advisories, which may assist in the development of trustworthy models that are acceptable to the ATCOs in an operational environment.

Technology Readiness Level

The existing ML-based automation methods, especially in air traffic conflict resolution, are focused on Lower technological readiness levels (TRL 1 and TRL 2). While these methods demonstrate promising results in controlled and isolated environments, the end goal is to integrate these models with larger sub-systems in an operational environment. Increasing the TRL also poses challenges from the aspect of data availability, simulation environments that are sufficiently complex to incorporate operational stochasticity and uncertainty of the air traffic control environment, and the involvement of operational subject matter experts from air traffic control along with the ML scientists for model improvement. Furthermore, flexibility and modularity should be an integral part of the initial model development, to ease the process of subsequent operational integration.

Data Management and Privacy

Developing ATCO-conformal conflict resolution models requires ensuring the collected data quality in terms of being representative of the ATCOs' strategies without errors, and being representative of the possible conflict scenarios that might arise in the sector. Concerns about ensuring that ATCOs' identities remain anonymous during data collection and ensuring privacy of data arise due to the sensitivity of air traffic data. Further, how the data is stored and who has access to the collected data must be identified and adhered to. This also involves the role of the certification agencies in the validation of data that is being used to develop the ML models for safety-critical operations such as air traffic conflict resolution.

7.3.2 Challenges in Validation and Certification of the ML Based Conflict Resolution Models

These challenges are associated with international aviation regulatory organizations such as FAA, ICAO, and EASA, who are responsible for the validation and safety certification of new automation methods. With the advent of increasingly complex and capable ML models, the following challenges are foreseen:

Developing Certification Processes

The existing automation tools are not 'learning-based' i.e. the software and their parameters do not change or evolve with time. In contrast, ML methods, especially those pertaining to online learning or continual learning change as they receive new data. In other words, the performance of the ML-based models is data-dependent. Thus, developing a certification process for the models must also involve certification of the data that is used in model training and testing. This is completely incompatible with the current certification process and would require significant changes in the current regulations and guidance to develop a certification process for the ML models. Furthermore, certification of such models for air traffic conflict resolution is more intricate as compared to other domains, due to the safety-critical and dynamics characteristics of air traffic conflict resolution.

Developing Staff Competency

Since the core functions of the regulatory organization are certification, and the development of rules and standards for the integration of new methodologies and tools into operations, it is of utmost importance that the staff receive the right level of ML expertise to carry out their tasks. Although the staff is not directly involved in the development of ML models, they must be exposed to the foundational concepts and techniques to better understand how a particular model will behave and perform. Further, this is a resource and time-intensive task, and bridging the knowledge gap may pose practical challenges for the organizations.

In order to assist in developing staff competency, we delivered a 3-day workshop to the EASA staff from 11^{th} September to 13^{th} September 2023 at the EASA headquarters in Cologne, Germany. The workshop focused on the utilization of reinforcement learning in ATM/ANS (Air Traffic Management/Air Navigation Services) applications, which was attended by a total of approximately 20 participants from EASA. The workshop was a mixture of theoretical concepts of reinforcement learning and hands-on exercises on its applications in air traffic conflict resolution. The workshop was very well received by the participants who rated it at 4.4/5 for an overall satisfaction score.

Deciding Scope of Models and Operational Authority

Deciding the scope of the ML models is crucial, especially in situations of human-AI teaming (HAT) such as air traffic conflict resolution, where the ML model acts as an advisory assistant or functions alongside an ATCO as a team member. It has been identified earlier that there is a tipping point or a stage after which the acceptance of automation by the user decreases. This is a stage where the role of automation shifts from an advisory to a decision maker [66]. Decisions on the scope and operational authority of the models will also set clear demarcation of responsibility for the decisions taken. A critical question that still needs to be addressed in HAT is who takes responsibility in case of a mistake/error by the ML/AI tool.

7.3.3 Challenges in Adoption of ML Models by ATCOs

These challenges are associated with the end-users i.e the ATCOs who will use such ML models in day-to-day operations.

ATCO Competency

The introduction of advanced technologies demands a paradigm shift in the skill set and knowledge base of the ATCOs. Addressing the challenge of user competency involves establishing effective communication channels between ML model developers and the ATCOs, wherein the developers must understand the requirements of the ATCOs and the ATCOs must develop their competencies with the advancing ML technologies. Bridging the knowledge gap of the end users can help enhance user confidence in adopting ML technologies in operations.

ATCO Trust

ATCO trust presents a formidable challenge in the integration of ML-based conflict resolution models into safety-critical systems like air traffic control. The inherent complexity of such ML models often results in black-box systems, where the decision-making processes are not readily interpretable by the ATCOs. This lack of transparency and explainability in how the ML models generate results can lead to skepticism and reluctance among the ATCOs to fully embrace ML recommendations. For instance, the ATCOs may not be confident in accepting the ML-based conflict resolution advisories because of the lack of trust and skepticism in how the models propose the solutions. A dilemma here is that with the increase in mathematical complexity of the models, including such characteristics of transparency and explainability becomes difficult. The question of trust becomes more critical as we increase the reliance on, and delegate more responsibility to ML-based systems [131].

Automation-Induced Complacency

Automation-induced complacency and its adverse effects have been the research focus for many decades [132]. Researchers have discussed that automation bias and automation-induced complacency continue to be factors contributing to aviation-related incidents and accidents [133]. From the aspect of air traffic control, the reliance on automated systems for air traffic conflict resolution can lead to a diminished level of vigilance and attentiveness among the ATCOs, creating a potential risk for complacency. As the ML models take on more responsibilities in decision-making processes, there is a risk that ATCOs may become overly reliant on the technology, potentially diminishing their responsiveness and critical decision-making skills required for conflict resolution. This over-reliance may also be detrimental in subsequent situations when support from the ML models is not available due to unforeseen reasons. This challenge directly links to the the operational authority and scope of the ML models which must be decided by the regulatory agencies.

7.4 Opportunities for Future Research

This thesis proposes methods for enhancing air traffic conflict resolution through machine learning, conformal automation, and flow-centric paradigms. There are several opportunities for improvement and extension of the work presented, which are discussed in the following sections.

7.4.1 Models for Increased Conflict Scenario Complexity

The conflict scenarios in this thesis were primarily two aircraft conflicts, in the presence of surrounding aircraft. Air traffic conflict resolution is affected by a multitude of factors such as the neighbouring sector geometry, no-go-zones, and the surrounding traffic. Further, the action space available to the ATCOs should be increased to include speed change and flight level change for conflict resolution. Human-in-loop experiments should be performed with the ATCOs in scenarios with increased complexity, to identify whether they employ different strategies to resolve such conflicts. The machine learning models trained on

ATCOs' data in such scenarios shall be closer to the operational environment.

On similar lines, there are significant opportunities for research extension in conflict resolution for flow-centric operations. First, the action space of the conflict resolution maneuvers can be increased to include flight level change. Nonetheless, this brings additional complexities and requirements for analysis in intra-flow conflict resolution such as issues of level-bust during climb or descent maneuvers. Further, the complexity of FCOs will further increase during the merging or splitting of flows because it will require the organization of the aircraft in each flow to assist such operations. Since ATCOs will still be in the loop to ensure safety, it is important to include their feedback in these initial conceptualizations, to facilitate the TRL increase of FCOs.

7.4.2 Investigation of Generalizability and Model Robustness

Generalizability is crucial in machine learning to ensure that the model can effectively handle varying scenarios and is not overly tailored to the specific characteristics of the training data. From the aspect of generalizability in air traffic conflict resolution, we should focus on the ability of a model to perform well in conflict scenarios presented in sectors that are different from the one where the model was trained. Though tests were performed to analyze the robustness of the models in our work, testing the models in different sectors with new conflict scenarios would aid in a detailed performance evaluation. This analysis shall also indicate how robust the developed models are, to the changes in the features of the conflicts and environment. Transfer learning is another aspect of machine learning wherein a pre-trained model is used as a starting point for a new model in a new environment. The advantage of transfer learning is faster learning through the transfer of already acquired knowledge. This can be specifically used within the scope of reinforcement learning for ATCO-conformal advisories in different sectors.

7.4.3 Development of Methods for Adaptive Conformal Conflict Resolution Advisories

The presented work has demonstrated that conformal conflict resolution advisories are more acceptable to the ATCOs. It was also observed that in some scenarios ATCOs would also prefer an improvement over the conformal advisory from the machine learning model, in terms of the cross-track deviation to the aircraft, and selected a balanced advisory for conflict resolution. There is a potential benefit in devising a method that enables the model to dynamically adjust its balance towards or away from the optimal resolution or machine learning prediction based on the ATCO's feedback over time. This would enhance the model's responsiveness to ATCO preferences and evolving situations.

7.4.4 Integration With Other Air Traffic Control Subsystems

The proposed conflict resolution methods have been developed solely in a conflict resolution-specific environment. This has been a common trend in the existing work discussed in the literature. We are aware that in operations, conflict resolution is a sub-process that is related to other processes such as air traffic flow management at a global level, and processes such as conflict detection and ATCO-pilot communication for action execution at the local level. Thus, it is valuable to develop a seamless pipeline integrating these operations and involving experts to test the performance of ATFM operations, conflict detection, its resolution, and execution, similar to the operational settings. This will require analysis of scenarios wherein a delay in communication and execution renders the current proposed solution infeasible, and a new, updated solution is required by the resolution method. Research in this direction has already taken shape in exploratory projects such as the Hypersolver under SESAR joint undertaking [134].

7.5 Concluding Remarks

Air traffic controllers are a critical part of the air traffic control system. With the increasing air traffic and forecasted future demand, the workload of the ATCOs will be one of the limiting factors in ensuring sustained growth. While numerous research efforts have been proposed to assist the ATCOs in the conflict resolution task, methods to generate ATCO-conformal conflict resolution advisories hold the most significant potential of acceptance in operations. This is because air traffic conflict resolution is a human-centric domain that requires ATCOs' inputs to ensure safe separation. Further, while we acknowledge the advantages of conformal automation, we must realize that the current concepts of operation are incapable of meeting long-term future demands and that novel paradigms such as flow-centric operations may offer the desired solutions. Thus, it is of great value to pursue further research in these directions.

Bibliography

[1] D.-T. Pham, N. P. Tran, S. Alam, V. Duong, and D. Delahaye, “A machine learning approach for conflict resolution in dense traffic scenarios with uncertainties,” in ATM 2019, 13th USA/Europe Air Traffic Management Research and Development Seminar, 2019.

[2] P. N. Tran, D.-T. Pham, S. K. Goh, S. Alam, and V. Duong, “An interactive conflict solver for learning air traffic conflict resolutions,” Journal of Aerospace Information Systems, vol. 17, no. 6, pp. 271–277, 2020.

[3] C. Ma, S. Alam, Q. Cai, and D. Delahaye, “Dynamic air traffic flow coordination for flow-centric airspace management,” in USA-Europe ATM Seminar, 2023.

[4] G. Dowek, C. Munoz, and A. Geser, “Tactical conflict detection and resolution in a 3-d airspace,” Institute For Computer Applications In Science And Engineering Hampton VA, Tech. Rep., 2001.

[5] ICAO, “4444–procedures for air navigation services–air traffic management,” Montreal, QC, Canada: The International Civil Aviation Organization (ICAO), 2016.

[6] S. Ontanón, G. Synnaeve, A. Uriarte, F. Richoux, D. Churchill, and M. Preuss, “A survey of real-time strategy game ai research and competition in starcraft,” IEEE Transactions on Computational Intelligence and AI in games, vol. 5, no. 4, pp. 293–311, 2013.

[7] S. Loft, P. Sanderson, A. Neal, and M. Mooij, "Modeling and predicting mental workload in en route air traffic control: Critical review and broader implications," Human factors, vol. 49, no. 3, pp. 376–399, 2007.

[8] M. Schuver-van Blanken and J. Van Merriënboer, "Air traffic controller strategies in operational disturbances-an exploratory study in radar control," in Proceedings of the 30th EAAP Conference, vol. 8, 2012.

[9] International Civil Aviation Organisation, "ICAO Long Term Traffic Forecasts," Tech. Rep., 2007. [Online]. Available: https://www.icao.int/sustainability/documents/ltf_charts-results_2018edition.pdf

[10] EUROCONTROL, "Eurocontrol aviation outlook 2050," Tech. Rep., 2022. [Online]. Available: https://www.eurocontrol.int/publication/eurocontrol-aviation-outlook-2050

[11] Statista, "Annual growth in global air traffic passenger demand from 2006 to 2022, with forecasts until 2023," Tech. Rep., 2023. [Online]. Available: https://www.statista.com/statistics/193533/growth-of-global-air-traffic-passenger-demand/

[12] J. Hoekstra, R. Ruigrok, and R. Van Gent, "Free flight in a crowded airspace?" Progress in Astronautics and Aeronautics, vol. 193, pp. 533–546, 2001.

[13] S. M. Galster, J. A. Duley, A. J. Masalonis, and R. Parasuraman, "Air traffic controller performance and workload under mature free flight: Conflict detection and resolution of aircraft self-separation," The international journal of aviation psychology, vol. 11, no. 1, pp. 71–93, 2001.

[14] A. Majumdar and J. Polak, "Estimating capacity of europe's airspace using a simulation model of air traffic controller workload," Transportation research record, vol. 1744, no. 1, pp. 30–43, 2001.

[15] J. d'Engelbronner, C. Borst, J. Ellerbroek, M. Van Paassen, and M. Mulder, "Solution-space–based analysis of dynamic air traffic controller workload," Journal of Aircraft, vol. 52, no. 4, pp. 1146–1160, 2015.

[16] EUROCONTROL. Free route airspace. [Online]. Available: https://www.eurocontrol.int/concept/free-route-airspace

[17] SESAR, "A proposal for the future architecture of the european airspace," The Single European Sky ATM Research (SESAR), 2019.

[18] H. Ritchie. (2024) What share of global co emissions come from aviation? [Online]. Available: https://ourworldindata.org/global-aviation-emissions

[19] ICAO. (2024) Environmental trends in aviation. [Online]. Available: https://www.icao.int/environmental-protection/Documents/EnvironmentalReports/2022/ENVReport2022_Art7.pdf

[20] N. Durand, J.-M. Alliot, and O. Chansou, "Optimal resolution of en route conflicts," vol. 1, no. 5, 1997, p. 2.

[21] N. Durand and J.-M. Alliot, "Ant colony optimization for air traffic conflict resolution," in 8th USA/Europe Air Traffic Management R&D Seminar, 2009.

[22] C. Westin, C. Borst, and B. Hilburn, "Strategic conformance: Overcoming acceptance issues of decision aiding automation?" IEEE Transactions on Human-Machine Systems, vol. 46, no. 1, pp. 41–52, 2015.

[23] S. J. Undertaking, "European atm master plan," Tech. Rep., 2020. [Online]. Available: https://www.sesarju.eu/masterplan2020

[24] J. Planning and D. Development Office, Washington, "Concept of operations for the next generation air transportation system," Tech. Rep., 2011. [Online]. Available: https://apps.dtic.mil/sti/pdfs/ADA535795.pdf

[25] J. Krozel and M. Peters, "Strategic conflict detection and resolution for free flight," in Proceedings of the 36th IEEE Conference on Decision and Control, vol. 2. IEEE, 1997, pp. 1822–1828.

[26] Z.-H. Mao, E. Feron, and K. Bilimoria, "Stability and performance of intersecting aircraft flows under decentralized conflict avoidance rules," IEEE Transactions on Intelligent Transportation Systems, vol. 2, no. 2, pp. 101–109, 2001.

[27] S. Fothergill and A. Neal, "Conflict-resolution heuristics for en route air traffic management," in Proceedings of the Human Factors and Ergonomics Society Annual Meeting, vol. 57, no. 1. SAGE Publications Sage CA: Los Angeles, CA, 2013, pp. 71–75.

[28] C. A. Manning et al., "Measuring air traffic controller performance in a high-fidelity simulation," United States. Department of Transportation. Federal Aviation Administration., Tech. Rep., 2000.

[29] N. Durand, J.-M. Alliot, and J. Noailles, "Automatic aircraft conflict resolution using genetic algorithms," in Proceedings of the 1996 ACM symposium on Applied Computing, 1996, pp. 289–298.

[30] K. Bilimoria, B. Sridhar, and G. Chatterji, "Effects of conflict resolution maneuvers and traffic density on free flight," in Guidance, navigation, and control conference, 1996, p. 3767.

[31] J. Hoekstra, R. Van Gent, and R. Ruigrok, "Conceptual design of free flight with airborne separation assurance," in Guidance, Navigation, and Control Conference and Exhibit, 1998, p. 4239.

[32] J. Kosecka, C. Tomlin, G. Pappas, and S. Sastry, "Generation of conflict resolution manoeuvres for air traffic management," in Proceedings of the 1997 IEEE/RSJ International Conference on Intelligent Robot and Systems. Innovative Robotics for Real-World Applications. IROS'97, vol. 3. IEEE, 1997, pp. 1598–1603.

[33] K. Zaghal, "Towards the logic of an airborne collision avoidance system which ensures coordination with multiple cooperative intruders," in ICAS Proceedings, vol. 19. American Institute of Aeronautics and Astronautics, 1994, pp. 2208–2208.

[34] M. Christodoulou and C. Costoulakis, “Nonlinear mixed integer programming for aircraft collision avoidance in free flight,” in Proceedings of the 12th IEEE Mediterranean Electrotechnical Conference (IEEE Cat. No. 04CH37521), vol. 1. IEEE, 2004, pp. 327–330.

[35] D. Delahaye, C. Peyronne, M. Mongeau, and S. Puechmorel, “Aircraft conflict resolution by genetic algorithm and b-spline approximation,” in EIWAC 2010, 2nd ENRI International Workshop on ATM/CNS, 2010, pp. pp–71.

[36] E. Feron, A. Bicchi, and L. Pallottino, “Mixed integer programming for aircraft conflict resolution,” in AIAA guidance, navigation, and control conference and exhibit, 2013, p. 4295.

[37] L. Pallottino, E. M. Feron, and A. Bicchi, “Conflict resolution problems for air traffic management systems solved with mixed integer programming,” IEEE transactions on intelligent transportation systems, vol. 3, no. 1, pp. 3–11, 2002.

[38] A. Alonso-Ayuso, L. F. Escudero, F. J. Martín-Campo, and N. Mladenović, “A vns metaheuristic for solving the aircraft conflict detection and resolution problem by performing turn changes,” Journal of Global Optimization, vol. 63, pp. 583–596, 2015.

[39] C. Allignol, N. Barnier, N. Durand, A. Gondran, and R. Wang, “Large scale 3d en-route conflict resolution,” in 12^{th} USA/Europe Air Traffic Management R&D Seminar, 2017.

[40] C. Allignol, N. Barnier, P. Flener, and J. Pearson, “Constraint programming for air traffic management: a survey 1: In memory of pascal brisset,” The Knowledge Engineering Review, vol. 27, no. 3, pp. 361–392, 2012.

[41] T. Stollenwerk, B. O’Gorman, D. Venturelli, S. Mandra, O. Rodionova, H. Ng, B. Sridhar, E. G. Rieffel, and R. Biswas, “Quantum annealing applied to deconflicting optimal trajectories for air traffic management,” IEEE transactions on intelligent transportation systems, vol. 21, no. 1, pp. 285–297, 2019.

[42] N. Durand and N. Barnier, “Does atm need centralized coordination? autonomous conflict resolution analysis in a constrained speed environment,” in 11th USA/Europe Air Traffic Management R&D Seminar, 2015.

[43] E. M. Rantanen and C. D. Wickens, “Conflict resolution maneuvers in air traffic control: Investigation of operational data,” The International Journal of Aviation Psychology, vol. 22, no. 3, pp. 266–281, 2012.

[44] S. Shin, K. Kim, and I. Hwang, “Identification of aircraft conflict resolution types from recorded flight data,” in 15th AIAA Aviation Technology, Integration, and Operations Conference, 2015, p. 2733.

[45] K. Kim, I. Hwang, and B.-J. Yang, “Classification of conflict resolution methods using data-mining techniques,” in 16th AIAA Aviation Technology, Integration, and Operations Conference, 2016, p. 4075.

[46] X. Olive, J. Grignard, T. Dubot, and J. Saint-Lot, “Detecting controllers’ actions in past mode s data by autoencoder-based anomaly detection,” in 8th SESAR Innovation Days, 2018.

[47] S. Malakis, P. Psaros, T. Kontogiannis, and C. Malaki, “Classification of air traffic control scenarios using decision trees: insights from a field study in terminal approach radar environment,” Cognition, Technology & Work, vol. 22, no. 1, pp. 159–179, 2020.

[48] P. H. Dang, P. N. Tran, S. Alam, and V. N. Duong, “A machine learning-based framework for aircraft maneuver detection and classification,” in Proceedings of the 14th USA/Europe Air Traffic Management R&D Seminar, 2021.

[49] M. Liang, “Case-based reasoning for flight conflicts resolution in en-route short term conflict alert situation,” The Open Automation and Control Systems Journal, vol. 7, no. 1, 2015.

[50] D.-T. Pham, "Machine learning-based flight trajectories prediction and air traffic conflict resolution advisory," Ph.D. dissertation, Université Paris sciences et lettres, 2019.

[51] M. Brittain and P. Wei, "Autonomous separation assurance in an high-density en route sector: A deep multi-agent reinforcement learning approach," in 2019 IEEE Intelligent Transportation Systems Conference (ITSC). IEEE, 2019, pp. 3256–3262.

[52] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural computation, vol. 9, no. 8, pp. 1735–1780, 1997.

[53] M. Brittain, X. Yang, and P. Wei, "A deep multi-agent reinforcement learning approach to autonomous separation assurance," arXiv preprint arXiv:2003.08353, 2020.

[54] M. W. Brittain and P. Wei, "One to any: Distributed conflict resolution with deep multi-agent reinforcement learning and long short-term memory," in AIAA Scitech 2021 Forum, 2021, p. 1952.

[55] J. Mollinga and H. van Hoof, "An autonomous free airspace en-route controller using deep reinforcement learning techniques," 9^{th} International Conference on Research in Air Transportation, ICRAT, 2020.

[56] M. Ribeiro, J. Ellerbroek, and J. Hoekstra, "Improvement of conflict detection and resolution at high densities through reinforcement learning," Proceedings of the ICRAT, 2020.

[57] P. Zhao and Y. Liu, "Physics informed deep reinforcement learning for aircraft conflict resolution," IEEE Transactions on Intelligent Transportation Systems, vol. 23, no. 7, pp. 8288–8301, 2021.

[58] S. Ghosh, S. Laguna, S. H. Lim, L. Wynter, and H. Poonawala, "A deep ensemble multi-agent reinforcement learning approach for air traffic control," arXiv preprint arXiv:2004.01387, 2020.

[59] A. Mukherjee, Y. Guleria, and S. Alam, "Deep reinforcement learning for air traffic conflict resolution under traffic uncertainties," in 2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC), 2022, pp. 1–8.

[60] D. Karikawa, H. Aoyama, M. Takahashi, K. Furuta, A. Ishibashi, and M. Kitamura, "Analysis of the performance characteristics of controllers' strategies in en route air traffic control tasks," Cognition, technology & work, vol. 16, no. 3, pp. 389–403, 2014.

[61] E. Xiaotian and J. Zhang, "Holistic thinking and air traffic controllers' decision making in conflict resolution," Transportation research part F: traffic psychology and behaviour, vol. 45, pp. 110–121, 2017.

[62] B. Kirwan, M. Flynn, and F. Bretigny, "Identification of air traffic controller conflict resolution strategies for the cora (conflict resolution assistant) project," in Proceedings of the 4th USA/Europe Air Traffic Management R&D Seminar, 2001, pp. 3–7.

[63] A. Nunes and R. H. Mogford, "Identifying controller strategies that support the 'picture'," in Proceedings of the Human Factors and Ergonomics Society Annual Meeting, vol. 47, no. 1. SAGE Publications Sage CA: Los Angeles, CA, 2003, pp. 71–75.

[64] S. Corver and G. Grote, "Uncertainty management in enroute air traffic control: a field study exploring controller strategies and requirements for automation," Cognition, Technology & Work, vol. 18, no. 3, pp. 541–565, 2016.

[65] M. Bekier, B. R. Molesworth, and A. Williamson, "Defining the drivers for accepting decision making automation in air traffic management," Ergonomics, vol. 54, no. 4, pp. 347–356, 2011.

[66] ——, "Tipping point: The narrow path between automation acceptance and rejection in air traffic management," Safety science, vol. 50, no. 2, pp. 259–265, 2012.

[67] K. Lee, K. Kerns, R. Bone, and M. Nickelson, “Development and validation of the controller acceptance rating scale (cars): Results of empirical research,” in 4^{th} USA/Europe Air Traffic Management R&D Seminar, 2001.

[68] C. Westin, B. Hilburn, C. Borst, and D. Schaefer, “The effect of strategic conformance on acceptance of automated advice: concluding the mufasa project,” 3^{rd} SESAR Innovation Days, vol. 3, 2013.

[69] G. Klien, D. D. Woods, J. M. Bradshaw, R. R. Hoffman, and P. J. Feltovich, “Ten challenges for making automation a" team player" in joint human-agent activity,” IEEE Intelligent Systems, vol. 19, no. 6, pp. 91–95, 2004.

[70] Å. Svensson, J. Lundberg, C. Forsell, and N. Rönnberg, “Automation, teamwork, and the feared loss of safety: Air traffic controllers’ experiences and expectations on current and future atm systems,” in European Conference on Cognitive Ergonomics 2021, 2021, pp. 1–8.

[71] C. Westin, “Strategic conformance: Exploring acceptance of individual-sensitive automation for air traffic control,” Ph.D. dissertation, 2017. [Online]. Available: https://doi.org/10.4233/uuid:49c6fe9d-2d29-420a-91a2-a97e2049e15e

[72] C. Westin, C. Borst, E. Kampen, T. M. Nunes, S. Boonsong, B. Hilburn, M. Cocchioni, and S. Bonelli, “Personalized and transparent ai support for atc conflict detection and resolution: an empirical study,” 12^{th} SESAR Innovation Days, Budapest, Hungary, pp. 5–8, 2022.

[73] C. Westin, C. Borst, and B. Hilburn, “An empirical investigation into three underlying factors affecting automation acceptance,” 5^{th} SESAR Innovation Days, pp. 1–9, 2015.

[74] S. van Rooijen, J. Ellerbroek, C. Borst, and E. van Kampen, “Toward individual-sensitive automation for air traffic control using convolutional neural networks,” Journal of Air Transportation, vol. 28, no. 3, pp. 105–113, 2020.

[75] R. Regtuit, C. Borst, and E.-J. Van Kampen, "Building strategic conformal automation for air traffic control using machine learning," in 2018 AIAA Information Systems-AIAA Infotech@ Aerospace, 2018, p. 0074.

[76] S. K. Goh, N. P. Tran, D.-T. Pham, S. Alam, K. Izzetoglu, and V. Duong, "Construction of air traffic controller's decision network using error-related potential," in Augmented Cognition: 13th International Conference, AC 2019, Held as Part of the 21st HCI International Conference, HCII 2019, Orlando, FL, USA, July 26–31, 2019, Proceedings 21. Springer, 2019, pp. 384–393.

[77] N. Durand, J.-B. Gotteland, N. Matton, L. Bortolotti, and M. Sandt, "Understanding and overcoming horizontal separation complexity in air traffic control: an expert/novice comparison," Cognition, Technology & Work, vol. 23, no. 3, pp. 481–496, 2021.

[78] D. Delahaye, C. Ma, S. Alam, and Q. Cai, "Air traffic flow representation and prediction using transformer in flow-centric airspace," in 12^{th} SESAR Innovation Days, 2022.

[79] J. K. Kuchar and L. C. Yang, "A review of conflict detection and resolution modeling methods," IEEE Transactions on intelligent transportation systems, vol. 1, no. 4, pp. 179–189, 2000.

[80] D.-T. Pham, S. Alam, and V. Duong, "An air traffic controller action extraction-prediction model using machine learning approach," Complexity, vol. 2020, 2020.

[81] R. S. Sutton and A. G. Barto, Reinforcement learning: An introduction. MIT press, 2018.

[82] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," Advances in neural information processing systems, NIPS, vol. 30, 2017.

[83] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," Advances in neural information processing systems, NIPS, vol. 1, 1988.

[84] S. Ross and D. Bagnell, “Efficient reductions for imitation learning,” in Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings, 2010, pp. 661–668.

[85] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, “Playing atari with deep reinforcement learning,” arXiv preprint arXiv:1312.5602, 2013.

[86] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton et al., “Mastering the game of go without human knowledge,” Nature, vol. 550, no. 7676, pp. 354–359, 2017.

[87] S. Lehmann, S. Bolland, R. Remington, M. S. Humphreys, and A. Neal, “Using a* graph traversal to model conflict resolution in air traffic control,” in Proceedings of the 10th International Conference on Cognitive Modeling, ICCM 2010. Drexel University, 2010, pp. 139–144.

[88] M. Özgür and A. Cavcar, “A knowledge-based conflict resolution tool for en-route air traffic controllers,” Aircraft Engineering and Aerospace Technology, 2008.

[89] K. W. Kallus, D. V. Damme, and A. Dittmann, “Integrated task and job analysis of air traffic controllers - phase 2: Task analysis of en-route controllers, hum.et1.st01.1000-rep-04,,” EUROCONTROL, Tech. Rep., 1999.

[90] Y. Liu, Y. Lee, and A. N. Chen, “Evaluating the effects of task–individual–technology fit in multi-dss models context: A two-phase view,” Decision Support Systems, vol. 51, no. 3, pp. 688–700, 2011.

[91] E. F. Morales and C. Sammut, “Learning to fly by combining reinforcement learning with behavioural cloning,” in Proceedings of the twenty-first international conference on Machine learning, 2004, p. 76.

[92] C. Sammut, S. Hurst, D. Kedzier, and D. Michie, “Learning to fly,” in Machine Learning Proceedings 1992. Elsevier, 1992, pp. 385–393.

[93] J. Saunders, C. L. Nehaniv, and K. Dautenhahn, "Teaching robots by moulding behavior and scaffolding the environment," in Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction, 2006, pp. 118–125.

[94] S. Chernova and M. Veloso, "Confidence-based policy learning from demonstration using gaussian mixture models," in Proceedings of the 6th International Joint Conference on Autonomous agents and Multiagent Systems, 2007, pp. 1–8.

[95] T. Inamura, "Acquisition of probabilistic decision model based on interactive teaching method," in Proc. 9th Int. Conf. on Advanced Robotics, Tokyo, 1999, 1999, pp. 523–528.

[96] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," in Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI, 2018, pp. 4950–4957.

[97] E. Zeki, P. L. de Matos, K. Purves, and M. Gibellini, "Business models for flight-centric air traffic control," Competition and Regulation in Network Industries, vol. 20, no. 4, pp. 319–332, 2019.

[98] D. Gianazza, "Learning air traffic controller workload from past sector operations," in 12^{th} USA/Europe Air Traffic Management R&D Seminar, 2017.

[99] EUROCONTROL. Eurocontrol simulation capabilities and platform for experimentation (escape). [Online]. Available: https://www.eurocontrol.int/simulator/escape

[100] H. Erzberger, T. Lauderdale, and Y. Chu, "Automated conflict resolution, arrival management, and weather avoidance for air traffic management," Proceedings of the Institution of Mechanical Engineers, Part G: Journal of aerospace engineering, vol. 226, no. 8, pp. 930–949, 2012.

[101] L. Grinsztajn, E. Oyallon, and G. Varoquaux, "Why do tree-based models still outperform deep learning on typical tabular data?" in Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track, vol. 35, 2022, pp. 507–520.

[102] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

[103] G. Gkioxari, A. Toshev, and N. Jaitly, "Chained predictions using convolutional neural networks," in Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14. Springer, 2016, pp. 728–743.

[104] Z. Guan, L. Chen, W. Zhao, Y. Zheng, S. Tan, and D. Cai, "Weakly-supervised deep learning for customer review sentiment classification." in International Joint Conference on Artificial Intelligence, IJCAI, 2016, pp. 3719–3725.

[105] Z. Fang, D. Wu, L. Pan et al., "When transfer learning meets cross-city urban flow prediction: spatio-temporal adaptation matters," International Joint Conference on Artificial Intelligence, IJCAI, pp. 2030–2036, 2022.

[106] A. Bombelli, L. Soler, E. Trumbauer, and K. D. Mease, "Strategic air traffic planning with fréchet distance aggregation and rerouting," Journal of Guidance, Control, and Dynamics, vol. 40, no. 5, pp. 1117–1129, 2017.

[107] P. C. Besse, B. Guillouet, J.-M. Loubes, and F. Royer, "Review and perspective for distance-based clustering of vehicle trajectories," IEEE Transactions on Intelligent Transportation Systems, vol. 17, no. 11, pp. 3306–3317, 2016.

[108] H. Alt and M. Godau, "Computing the frechet distance between two polygonal curves," International Journal of Computational Geometry & Applications, vol. 5, no. 01n02, pp. 75–91, 1995.

[109] T. Eiter and H. Mannila, "Computing discrete fréchet distance," Technical Report CD-TR 94/64, Christian Doppler Laboratory., 1994.

[110] H. Liu, Z. Huang, J. Wu, and C. Lv, "Improved deep reinforcement learning with expert demonstrations for urban autonomous driving," in 2022 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2022, pp. 921–928.

[111] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," Annual review of control, robotics, and autonomous systems, vol. 3, pp. 297–330, 2020.

[112] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," arXiv preprint arXiv:1707.06347, 2017.

[113] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," Journal of Machine Learning Research, vol. 22, no. 268, pp. 1–8, 2021. [Online]. Available: http://jmlr.org/papers/v22/20-1364.html

[114] C. Chauvin and S. Lardjane, "Decision making and strategies in an interaction situation: Collision avoidance at sea," Transportation Research Part F: Traffic Psychology and Behaviour, vol. 11, no. 4, pp. 259–269, 2008.

[115] EASA. First usable guidance for level 1 machine learning applications. [Online]. Available: https://www.easa.europa.eu/en/newsroom-and-events/news/easa-releases-its-concept-paper-first-usable-guidance-level-1-machine-0

[116] ——. EASA Artificial Intelligence Roadmap. [Online]. Available: https://www.easa.europa.eu/en/newsroom-and-events/news/easa-artificial-intelligence-roadmap-20-published

[117] K. Pushparaj, P. Reddy, D. Vu-Tran, K. Izzetoglu, and S. Alam, "A multi-modal approach to measuring the effect of xai on air traffic controller trust during off-nominal runway exits," in IEEE Conference on Systems, Man, and Cybernetics, 2023.

[118] E. W. Dijkstra, "Self-stabilizing systems in spite of distributed control," Communications of the ACM, vol. 17, no. 11, pp. 643–644, 1974.

[119] C. Boulinier, F. Petit, and V. Villain, “When graph theory helps self-stabilization,” in Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing, 2004, pp. 150–159.

[120] N.-S. Chen, H.-P. Yu, and S.-T. Huang, “A self-stabilizing algorithm for constructing spanning trees,” Information Processing Letters, vol. 39, no. 3, pp. 147–151, 1991.

[121] Y. Sudo, F. Ooshita, and S. Kamei, “Self-stabilizing graph exploration by a single agent,” arXiv preprint arXiv:2010.08929, 2020.

[122] I. Stoica, D. Song, R. A. Popa, D. Patterson, M. W. Mahoney, R. Katz, A. D. Joseph, M. Jordan, J. M. Hellerstein, J. E. Gonzalez et al., “A berkeley view of systems challenges for ai,” arXiv preprint arXiv:1712.05855, 2017.

[123] D. Hendrycks, N. Carlini, J. Schulman, and J. Steinhardt, “Unsolved problems in ml safety,” arXiv preprint arXiv:2109.13916, 2021.

[124] I. Kabashkin, B. Misnevs, and O. Zervina, “Artificial intelligence in aviation: New professionals for new technologies,” Applied Sciences, vol. 13, no. 21, p. 11660, 2023.

[125] C. Laidlaw, S. Singla, and S. Feizi, “Perceptual adversarial robustness: Defense against unseen threat models,” in International Conference on Learning Representations (ICLR), 2021.

[126] M. Kaufmann, D. Kang, Y. Sun, S. Basart, X. Yin, M. Mazeika, A. Arora, A. Dziedzic, F. Boenisch, T. Brown et al., “Testing robustness against unforeseen adversaries,” arXiv preprint arXiv:1908.08016, 2019.

[127] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” arXiv preprint arXiv:1412.6572, 2014.

[128] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli, “Evasion attacks against machine learning at test time,” in Machine Learning and Knowledge Discovery in Databases: European Conference, ECML

PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III 13. Springer, 2013, pp. 387–402.

[129] F. Bayram, B. S. Ahmed, and A. Kassler, “From concept drift to model degradation: An overview on performance-aware drift detectors,” Knowledge-Based Systems, vol. 245, p. 108632, 2022.

[130] J. Zerilli, U. Bhatt, and A. Weller, “How transparency modulates trust in artificial intelligence,” Patterns, 2022.

[131] W. J. von Eschenbach, “Transparency and the black box problem: Why we do not trust ai,” Philosophy & Technology, vol. 34, no. 4, pp. 1607–1622, 2021.

[132] R. Parasuraman, R. Molloy, and I. L. Singh, “Performance consequences of automation-induced’complacency’,” The International Journal of Aviation Psychology, vol. 3, no. 1, pp. 1–23, 1993.

[133] J. P. Brown, “The effect of automation on human factors in aviation,” The Journal of Instrumentation, Automation and Systems, vol. 3, no. 2, pp. 31–46, 2016.

[134] SESAR. HYPERSOLVER - Artificial Intelligence controller able to manage Air Traffic Control and Air Traffic Flow Management within a single framework. [Online]. Available: https://www.sesarju.eu/projects/hypersolver

Appendix A

Post Experiment Questionnaire

A.1 Workload

Please rate the **AVERAGE** (usual level throughout the session) workload that you experienced during this session.

Task abandoned. I was unable to supply sufficient effort.	10	○
Extremely high workload, no spare capacity. **Serious doubts** as to the ability to maintain level of service.	9	○
Very high workload with almost no spare capacity. **Difficulty** in maintaining level of work	8	○
Very high workload with almost no spare capacity but **no impact** to the primary ATM task.	7	○
Little spare capacity. Level of effort allows **little attention** to additional or other tasks.	6	○
Reduced spare capacity. Additional or other tasks cannot be given the **desired amount of attention**.	5	○
Insufficient spare capacity for **early attention** to additional tasks.	4	○
Enough spare capacity for all desirable additional tasks.	3	○
Workload low.	2	○
Workload insignificant.	1	○

A.2 Situational Awareness

Please read the statements below and rate the overall level of Situational Awareness (SA) that you experienced during this session.

My SA with respect to the task was far too low. I could not perform the task because **I did not possess the necessary information**.	1	○
The SA with respect to my task was very low. I was unaware of **almost all of the information** required to perform the task effectively.	2	○
My SA with respect to the task was low. I was unaware of **most of the information** required to perform the task effectively.	3	○
My SA with respect to the task was low. I was unaware of about **half of the information** required to perform the task effectively.	4	○
My SA with respect to the task was reduced. I was unaware of **some of the important information** required to perform the task effectively.	5	○
My SA with respect to the task was insufficient. I was not aware of **all the information** required to perform the task effectively.	6	○
My SA with respect to the task was not complete. I was able to perform the task, but **not satisfactorily**.	7	○
My SA with respect to the task was good. I was able to perform the task **well most of the time**.	8	○
My SA with respect to the task was very good. I was able to perform the task **well all of the time**.	9	○
My SA with respect to the task was excellent. I was able to perform my task **extremely well all of the time**.	10	○

A.3 Trust

	Not at all	2	3	Neutral	5	6	Extremely
The system was useful	○	○	○	○	○	○	○
The system presented realistic traffic scenarios	○	○	○	○	○	○	○
The system can help me improve my decision making	○	○	○	○	○	○	○
The system provided me with reasonable conflict resolution strategies	○	○	○	○	○	○	○
The system provided me with enough resolution strategies	○	○	○	○	○	○	○

A.4 Additional Questions

1. In the event of a detected conflict, how early or late (in minutes or NM) would you prefer to initiate a conflict resolution maneuver?

2. In the experiments, what was the general basis for your resolution maneuver preference selection for the conflict scenarios?

3. What additional information regarding the conflict resolution maneuver would be useful when accepting/rejecting a proposed resolution?

4. In operations, what features regarding the conflicting aircraft pair, the sector, and the surroundings do you consider before making a decision for conflict resolution?

5. Currently, what type of automation assistance tools are available to assist the ATCOs

in en-route air traffic conflict resolution?

6. What features in a future conflict resolution assistance tool would you prefer to have?

7. On a scale of 1 to 5, how likely are you to accept a conflict resolution advisory given by such an automation assistance tool (1 – lowest, 5 highly likely) in operations? What factors would govern the acceptance or rejection of the advisory?

Not at all	1	○
Less likely (under rare circumstances)	2	○
Neutral	3	○
Likely	4	○
Highly likely	5	○

What factors govern your selection?