Aalto University
School of Science
Master's Programme in Mathematics and Operations Research

Janne Lahti

Model-based learning interventions to enhance decision-making of master fighter controllers

Master's Thesis
Espoo, July 7, 2022

Supervisor: Professor Kai Virtanen
Advisor: Professor Kai Virtanen

Aalto University
School of Science
Master's Programme in Mathematics and Operations Research

ABSTRACT OF MASTER'S THESIS

Author:	Janne Lahti		
Title:	Model-based learning interventions to enhance decision-making of master fighter controllers		
Date:	July 7, 2022	**Pages:**	80
Major:	Systems and Operations Research	**Code:**	SCI3055
Supervisor:	Professor Kai Virtanen		
Advisor:	Professor Kai Virtanen		

A master fighter controller (MFC) is a decision maker in air combat whose responsibility is to control air operations that include multiple flights, i.e., units of four fighters. An MFC faces complex and time-intensive decision making problems, and extensive training is a prerequisite for the ability to perform rapid decisions successfully. This thesis introduces a novel fighter allocation (FA) model to support training and learning of MFCs. The model allows construction of real-life scenarios portraying the current state of the battlefield and enables rapid computations to evaluate MFCs' decisions in these scenarios. The initial situations of the scenarios contain locations of friendly and hostile fighters, and optionally restricted operating zones. The decision of the MFC is set to the FA model as the desired commit locations where the friendly fighters start engaging the enemy. Then, the feasibility of the decision made by the MFC, i.e., if the friendly fighters can reach the commit locations in time without exposure to enemy threat, is evaluated by the FA model. In the evaluation, first, the time optimal routes of the friendly fighters from their initial locations to their commit locations are computed using network optimization. Second, the outcome of the scenario is determined by simulating the flight paths of the friendly fighters along the optimal routes. The outcome of the scenario describes if the decision made by the MFC has been feasible.

In this thesis, an experimental study was also conducted to analyze the benefits of the FA model in training of MFCs. The study was performed during a training course for flight leaders and future MFCs. The use of the model was practiced in a 2-hour tactical training session where real-life scenarios constructed by a subject matter expert were studied. Participants' ability to perceive and understand the state of the battlefield as well as to make feasible decisions in test scenarios was measured before and after the training intervention. Statistically significant improvements in the decision making of the participants were observed due to the intervention. Thus, the study suggests that the use of the FA model - and model-based training and learning practices in general - can enhance MFCs' perception skills, understanding and decision making.

Keywords:	air combat, control of air operations, behavioral operations research, network optimization, simulation
Language:	English

Aalto-yliopisto
Perustieteiden korkeakoulu
Matematiikan ja operaatiotutkimuksen maisteriohjelma

DIPLOMITYÖN TIIVISTELMÄ

Tekijä:	Janne Lahti		
Työn nimi:	Päätaistelujohtajien päätöksenteon kehittäminen mallipohjaisilla käytänteillä		
Päiväys:	7. heinäkuuta 2022	**Sivumäärä:**	80
Pääaine:	Systeemi- ja operaatiotutkimus	**Koodi:**	SCI3055
Valvoja:	Professori Kai Virtanen		
Ohjaaja:	Professori Kai Virtanen		

Päätaistelunjohtaja on ilmataistelussa päätöksentekijä, jonka vastuulla on johtaa erilaisia ilmaoperaatioita. Näitä operaatioita toteuttaa yleensä useita hävittäjäparvia, jotka koostuvat neljästä yksittäisestä hävittäjästä. Päätaistelunjohtajan päätöksenteko-ongelmat ovat monimutkaisia ja nopeaa päätöksentekoa vaativia, joten päätöksenteon harjoittelu edesauttaa onnistuneiden päätösten tekemistä. Tässä työssä esitellään uusi hävittäjien allokointimalli tukemaan päätaistelunjohtajien koulutusta ja oppimista. Malli mahdollistaa realististen ilmataisteluskenaarioiden rakentamisen ja päätaistelunjohtajien päätösten hyvyyden nopean arvioinnin. Skenaarioiden alkuasetelman määrittelyssä malliin syötetään sekä omien että vastustajan hävittäjien sijainnit tarkastelun alkuhetkellä sekä tarvittaessa lentokieltoalueita. Päätaistelunjohtajan tehtävänä on päättää omille hävittäjille sijainnit, joista hävittäjien on tavoiteltavaa aloittaa ilmataistelu vastustajaa vastaan. Mallin avulla määritellään näiden sijaintien ja edelleen päätaistelunjohtajan tekemän päätöksen toteuttamiskelpoisuus, eli ehtivätkö omat hävittäjät annettuihin sijainteihin ajoissa siten, että etäisyys vastustajaan ei muodostu hävittäjien kannalta liian pieneksi. Toteuttamiskelpoisuuden arvioinnissa ratkaistaan omien hävittäjien optimaaliset minimiaikareitit alkusijainneista päätaistelunjohtajan asettamiin sijainteihin verkko-optimoinnilla. Tämän jälkeen skenaarion lopputulos lasketaan simuloimalla omien hävittäjien lento minimiaikareittejä pitkin. Simulaation lopputulos kuvaa päätaistelujohtajan päätöksen toteuttamiskelpoisuuden.

Tässä työssä toteutettiin myös kokeellinen tutkimus, jossa arvioitiin hävittäjien allokointimallin tarjoamaa hyötyä päätaistelunjohtajien koulutuksessa. Mallia käytettiin parvenjohtajien ja tulevien päätaistelunjohtajien päätöksenteon harjoittamiseen realistisissa ilmataisteluskenaariossa. Harjoituksen kesto oli kaksi tuntia. Osallistujien kykyä hahmottaa taistelukenttää ja kykyä tehdä onnistuneita päätöksiä mitattiin ennen ja jälkeen oppitunnin. Osallistujien päätöksenteossa havaittiin tilastollisesti merkittävä kehittyminen oppitunnin seurauksena. Tutkimus antoi lupaavia tuloksia hävittäjien allokointimallin – ja yleisesti mallipohjaisten koulutus- ja oppimiskäytänteiden – käytön hyödyistä päätaistelunjohtajien hahmotuskyvyn ja päätöksenteon kehittämisessä.

Asiasanat:	ilmaoperaatioiden hallinta, ilmataistelu, käyttäytymistutkimuksellinen operaatioanalyysi, simulointi, verkko-optimointi
Kieli:	Englanti

Acknowledgments

First and foremost, I want to thank Professor Kai Virtanen for the possibility to work in his group and conduct this thesis. His extensive guidance, dedication and patience have been nothing but vital for this project. I am also deeply grateful for all the trust with the unique and exciting responsibilities I have had during the years. Furthermore, I really want to thank Adjunct Professor Heikki Mansikka not only for all the knowledge about the subject matter, but also for inspiring discussions and invaluable advice at all stages of this project. Likewise, I want to express my gratitude to the other subject matter experts who contributed in one way or another to this thesis and especially to its experimental study. It has been an absolute privilege to work with you all during these years. Huge respect.

I would also like to thank all my co-workers at Y220 for providing such an inspiring and joyful environment to work. Especially I want to thank Patrik for exceptional peer support and discussions, Heikki for many amazing common endeavors, and Mikko for the help and guidance since day one.

These years at Aalto University have been without a doubt the best of my life. This is mostly due to all the great friends I am privileged to have. Both those I have been lucky to meet during the last seven years and those who have been there for much longer. You have made these years unforgettable and full of memories, and my path through these studies would have been so much rougher without you.

Most importantly, I am deeply grateful to my family for their uncompromising support and understanding during my studies and life. Thank you for always believing in me, without you I would have not been able to get here. And Roosa, thank you for your endless support and patience on my good and bad days when writing this thesis.

Espoo, July 7, 2022

Janne Lahti

Contents

Chapter 1

Introduction

1.1 Background

Air warfare is a fast-paced form of warfare where proactive preparation and quick reactions to enemy threats are inevitable in order to succeed. The time scale of air operations and combat are usually measured in hours or even minutes. Thus, well-prepared missions as well as good reactive decisions from combat units, commanders and other decision makers are required. Decision makers in an aerial warfare face complex problems where decisions need to be made fast and the their consequences are usually significant in the context of entire operations. Thus, the decision makers need to be well-trained to make good decisions in a timely manner. For example, reacting correctly and quickly enough to an enemy air operation is important, as the flight speed of enemy fighter planes can exceed 2000 kilometers per hour, and consequently, the time available for counteracting is limited. This thesis considers improving decision making in such situations using a model-based approach.

A master fighter controller (MFCs) is an important decision maker in air warfare. MFCs are responsible for making near real-time decisions in different air operations and deciding how to utilize aircraft such as fighters. In practice, the MFC controls air operations by deciding general maneuvering of fighters, such that the overall mission effectiveness and the mutual support between fighters is maximized. The MFC has usually several flights, i.e., units typically consisting of four fighters, under his command. A person with similar duties and responsibilities can also be referred to as, e.g., a fighter allocator.

One of the key responsibilities of the MFC is to lead and control defensive counterair (DCA) (see, e.g., Joint Chiefs of Staff, 2017) missions. The objective of the decision made by the MFC in these missions is, e.g., to negate an imminent enemy threat, to defend an important location or other assets (Bennett Jr et al., 2002), or to maintain the control of the air (LeMay Center for Doctrine, 2019). As in warfare in general, the MFC as a decision maker faces uncertainty, time pressure, and high complexity in his responsibilities (Thunholm, 2004). Usually, the information about the enemy threat is not completely perfect and the MFC might not be certain about, e.g., the number of incoming aircraft, their types, flight speeds, weapons, or routes. Thus, the MFC has a challenging decision making problem where he/she needs to assess the utilization of assets with respect to the expected enemy threat and the desired objectives. The ability of the MFC to make good decisions under a significant time pressure is critical for the success of a DCA mission. The MFC has to have a good understanding and a perception of the overall tactical situation as well as the individual factors and variables in the circumstances at hand. With a limited amount of information and time to make the decisions, the proficiency of the MFC or any other decision maker in warfare attained from training is important (Holmes, 1995; Development, Concepts and Doctrine Centre, 2016a). This thesis addresses the need to improve and develop the training of the MFCs in order to strengthen their perception and decision making abilities.

1.2 Objective

An MFC decides the general maneuvering of fighters in a DCA mission. In the maneuvering, it is essential to position the flights to suitable commit locations, i.e., locations where they start engaging the enemy (Joint Chiefs of Staff, 2017). Selecting the commit locations is a critical decision taken by the MFC. The objective of this thesis is to develop a model that can be utilized to improve MFCs' decision making in selecting these commit locations of the flights. In this thesis, a model called "fighter allocation model" (FA model) is constructed and its effect to the decision making of the MFCs are measured. The purpose of the FA model is to allow constructing training scenarios reflecting the real-life situations the MFCs might face and evaluating the decisions the MFCs make in them. The FA model is designed to be utilized in the training of the MFCs to support their learning to make better decisions.

Perceiving and understanding the current state of the battlefield correctly are required in order to improve the decision making regarding the selection of commit locations. In the FA model, an initial state of a battlefield is described by the locations of friendly and enemy fighters. The model also allows defining restricted operation zones (ROZs) to portray restrictions in the airspace for the friendly fighters. Additionally, important assets, areas of interest, or desired engagement frontiers (see, e.g., Harju et al., 2019) can be included to help to illustrate relevant objectives for DCA missions. For the given initial situation, an MFC has to decide how to maneuver friendly fighters or flights. The decision is described in a form of desired destinations, i.e., commit locations, for these fighters or flights. The model then evaluates if it is possible and safe to reach these locations determined by the MFC. The FA model utilizes network optimization (see, e.g., Bertsekas, 1998) and simulation techniques (see, e.g., Law, 2013) to perform these evaluations. Network optimization is employed for determining the optimal routes of the fighters from their initial locations to their commit locations, while the outcomes of the scenarios are provided by continuous-time simulation. These outcomes are presented using visualizations and numerical data.

The FA model allows MFCs to test and develop their ability to correctly perceive and understand the state of the battlefield as well as its geometry. By developing these abilities the MFCs consequently learn also to make better decisions. An MFC faces the problem of assigning fighters to feasible commit locations every time he/she has a decision to make. Evaluating the spatial feasibility of the decision is a core part of the MFC's decision making. With the model, it is quick to determine the outcome of the MFC's decision and to observe how good the decision was and how it could be improved. Utilizing the model in the training of the MFCs can be beneficial and can improve their perception, understanding, and decision making abilities.

In this thesis, an experimental study is also performed where the FA model is employed in a training of future MFCs. The purpose of the experimental study is to evaluate the effects of the use of the FA model in the decision making of MFCs, and consequently, verify the benefits of the model when used in training. A training intervention is performed by applying the FA model in the training of future MFCs and flight leaders. The quality of the decisions made by the participants in test scenarios reflecting real-life decision making situations are measured before and after the training intervention. With these performance measurements, the benefits offered by the model in the participants' decision making are assessed.

1.3 Related literature

The responsibility of an MFC is to maneuver assets, usually flights, to achieve desired outcomes in the warfare. Optimal use of military assets such as weapons or fighters has been a subject of extensive studies. A classic problem regarding the utilization of the military assets is the Weapon Target Assignment (WTA) problem (see, e.g, Ahuja et al., 2007; Huaiping et al., 2006; Kline et al., 2019), and it has been studied from the 1950s (Manne, 1958). The WTA problem is essentially an optimization problem where allocation of defensive weapons to attacking targets is optimized by, e.g., maximizing the probability of destroying the targets or minimizing the probability of destruction of the defending assets. While the WTA problem can be compared to the problem MFCs face, the approach to address the WTA problem is fundamentally different than with the FA model presented in this thesis. The WTA problem is a pure combinatorial problem, and by solving it, an optimal solution for the use of the weapons is found. The FA model, however, is constructed to provide learning opportunities for MFCs and to develop their skills to perform rapid maneuver decisions.

Decision making in air combat is mostly studied and modeled in literature from the perspective of pilots (see, e.g., Virtanen et al., 1999b, 2004, 2006b). From a broader point of view, studies and models usually have focused on, e.g., air combat tactics (see, e.g., Mulgund et al., 1998; Tidhar et al., 1998; Kewley and Embrechts, 2002; Poropudas and Virtanen, 2010) or team performance (see, e.g., Virtanen et al., 2006a; Mansikka et al., 2021a). While decision making from a perspective of military command and control (see, e.g., Joint Chiefs of Staff, 2016) has been studied in the literature (see, e.g., Kaempf et al., 1996), model-based approaches to enhance military decision making are not yet widely explored. Models like the FA model intended for improving decision making via model-based training have not been presented earlier in the open military literature.

A main component of the FA model developed in this thesis consists of finding optimal routes for fighters. Optimization of routes is a widely studied problem in the literature. There exists a variety of established methods for optimizing the routes of the fighters, including but not limited to optimal control (see, e.g., Kirk, 2004) and network optimization (see, e.g., Bertsekas, 1998). In the context of military aircraft routing, approaches using optimal control are presented in, e.g., Zabarankin et al. (2006), Virtanen et al. (1999a), Karelahti et al. (2007), Karelahti et al. (2008), and Miller et al. (2011). With the optimal control, optimization problems and their solutions are continuous in time and based on calculus of variations. These solutions are acquired using various analytical and/or numerical solution techniques. However, solving an optimal control problem can be computationally demanding and require accurate enough models of the fighters. Finding globally optimal solutions may also not be trivial and problems in convergence of numerical solution algorithms can arise. Furthermore, any obstacles for the routes convert the problem into a highly non-convex one, which further complicates the search of the globally optimal solution (Adler et al., 2012).

Network optimization approaches rely on discretizing the airspace into a grid and formulating the problem as a shortest-path problem. In the military context, network optimization approaches are utilized in, e.g., Carlyle et al. (2009), Puustinen (2013), and Gunell (2019). With approaches based on network optimization, the computation times for optimal routes are usually relatively short. Additionally, when using suitable network algorithms, globally optimal solutions can be always acquired and no convergence issues arise, contrary to the optimal control methods (Rippel et al., 2005). On the other hand, due to the discrete nature of the optimization problem formulation, the obtained routes might not describe the freedom of the movement of an aircraft as realistically as with, e.g., with optimal control approaches. In the FA model, acquiring the optimal routes in a negligible amount of time with a reasonably accuracy is desired in order to allow fast repeated computations. As the main purpose of the model is educational, its use consists significantly of studying different what-if scenarios and evaluating decisions with a trial-and-error basis. Consequently, rapid re-computations are desirable for acquiring outcomes of different decisions quickly, preferably in almost real-time. Due to these requirements, network optimization is utilized in the FA model to compute the optimal routes. The routes are then smoothed after the optimization using moving average in order to make them more realistic. Thus, while network optimization does not provide as realistic routes as optimal control approaches, a discrete grid describes the movement of a fighter accurately enough for the training purposes of the FA model. Additionally, adding obstacles such as no-fly zones (NFZs) is straightforward and does not affect the solution time of the problem significantly.

The purpose of the FA model constructed in this thesis is to enhance the decision making abilities of MFCs. The effect of the FA model to the decision making is measured in an experimental study. Studying learning outcomes from simulation models has been a point of interest relatively long in operations research field, but there is only little evidence of learning outcomes and insights generated by such models (Gogi et al., 2016). Generally, the value given by the models is hard to determine as implementations and their benefits are rarely assessed (Fone et al., 2003). However, it has been observed that model users can usually produce better solutions or decisions after interacting with a model, while still poorer than known or optimal answers (O'Keefe, 2016). Furthermore, it is also suggested that appropriate visualizations and especially animations can aid to do better decision making than with just simple graphics or numerical data (O'Keefe, 2016). There are also studies where some evidence have been obtained that, e.g., solving case studies with visual interactive simulations can improve the solutions prior to the usage of the simulation model (Bell and O'Keefe, 1995). Moreover, associations between generating insights and the use of simulation models has been found (Gogi et al., 2016). The learning outcomes of a simulation modeling have also been studied from the perspective of a model building, i.e., if participating in the model building leads to learning (see, e.g., Rouwette et al., 2011; Monks et al., 2014). These studies of behavioral aspects in the use of models in decision making and problem solving as well as experimental research with models belong to the field of behavioral operational research (BOR) (see, e.g., Hämäläinen et al., 2013; O'Keefe, 2016). This thesis contributes to field of BOR by performing an experimental study to evaluate the effect of a model-based learning intervention in training of MFCs. In the military context, calls for such studies have been made (Salas et al., 1998) and procedures to measure effectiveness of the training have been proposed (Bell and Waag, 1998). However, studies where learning outcomes from the use of models have not been widely considered in the unclassified literature.

1.4 Structure

The thesis is structured as follows. Chapter 2 outlines the relevant military context and introduces the responsibilities of MFCs in more detail. The need for a model-based approach in the training of the MFCs is also discussed. The FA model is presented in Chapter 3. In Chapter 4, its practical use is demonstrated and discussed with example scenarios. A real-life experimental study evaluating the capability of the model to improve perception and decision making abilities of MFCs is presented in Chapter 5. The benefits of the FA model to the real-life decision making, its limitations, potential expansions and further research avenues are discussed in Chapter 6, where concluding remarks are also given.

Chapter 2

Decision making in defensive counterair missions

This chapter provides background to the general landscape of air operations where an MFC operates. Air operations are discussed focusing on counterair missions. The concepts of defensive counterair (DCA) and offensive counterair (OCA) missions are introduced. As the FA model constructed in this thesis focuses on the decision making of an MFC in DCA missions, DCA missions and the MFC's role and responsibilities in them are presented in more detail. The MFC's decision making problem in a DCA mission as a whole is discussed, with focus on elements which the FA model addresses.

2.1 Defensive counterair missions

Airspace is a crucial environment for military operations performed in land, sea, and air (Joint Chiefs of Staff, 2014). Control of the air has been considered as a prerequisite to success in modern military operations (Joint Chiefs of Staff, 2019) and one of the first priorities of defense forces (LeMay Center for Doctrine, 2019). It is one of the key roles of utilization of air power in addition to intelligence, surveillance, reconnaissance, air attack, and air mobility (Development, Concepts and Doctrine Centre, 2007). Control of the air describes the level of influence in the airspace relative to an adversary (LeMay Center for Doctrine, 2019). Achieving control of the air allows performing different types of air operations as it facilitates freedom to act in the airspace and denies its use by the adversary (Joint Chiefs of Staff, 2017). In order to challenge, gain and maintain control of the air, counterair missions are performed (LeMay Center for Doctrine, 2019; Development, Concepts and Doctrine Centre, 2016b). The objective of the counterair missions is usually to negate or destroy enemy aircraft or missiles (Joint Chiefs of Staff, 2016).

Counterair missions can be divided further into defensive counterair (DCA) and offensive counterair missions (OCA). DCA missions are conducted to reduce effectiveness of enemy threat in friendly airspace in order to protect friendly forces and assets. The nature of DCA missions is typically reactive, and they are usually located near or inside friendly airspace. DCA missions are generally conducted when an enemy threat is already active, which leads to their reactive nature. OCA missions are instead proactive missions where enemy capabilities and assets such as bases or equipment are targeted in order to prevent launch of aircraft or missiles. Use of OCA missions desirably decreases the need of DCA missions if enemy capabilities and assets are successfully impacted. Additionally, OCA missions force enemy to perform their DCA missions and require rapid decisions and actions from them. Various military capabilities such as fighter aircraft can be used for both DCA and OCA missions, and planning for these missions needs often to be collective and synchronized. Detailed descriptions and discussions for DCA and OCA missions can be found in, e.g., LeMay Center for Doctrine (2019), Joint Chiefs of Staff (2017), Development, Concepts and Doctrine Centre (2007), Development, Concepts and Doctrine Centre (2016b), and Joint Chiefs of Staff (2019). In the scope of this thesis, only DCA missions are considered.

In DCA missions, fighter aircraft typically operate as flights, i.e., as units generally consisting of four fighters. In this thesis, a flight refers to such a unit of four fighters. A flight is lead by a flight leader who is responsible for deciding appropriate tactics, techniques and procedures (TTPs) (see, e.g., Mansikka et al., 2021b,c,d) for his/her flight. Flights are supported by fighter controllers (FCs), who provide, e.g., battle-space awareness as well as command and control services (see, e.g. LeMay Center for Doctrine, 2019; Joint Chiefs of Staff, 2017) to the flights. Due to the highly dynamic and complex nature of DCA missions, usually a single FC is allocated for each flight. The FCs are lead by an MFC, who is responsible for the coordination of these operations (Himanen, 2013).

An air attack is typically performed in a situation where an attacker possesses quantitative, e.g., more fighters or other assets, and/or qualitative, e.g., better equipment or pilots, superiority over a defender (Holmes, 1995). However, even with inferior forces, the defender can reach a local control of the air at least temporarily by utilizing a suitable course of action (COA). A COA is a plan for the coordinated and centralized use of the DCA aircraft and their TTPs. The selection of a COA is based on strategic guidelines which describe the overall objectives of DCA missions and how they should be carried out as a whole. However, they do not necessarily contain detailed instructions to the actual use of aircraft, i.e., consider COAs itself. The selection of a suitable COA to be conducted is made by an MFC. Once a suitable COA has been has identified and selected for the DCA mission, the MFC informs his/her FCs about the selected COA. The FCs then task their respective flights to execute their part of the COA.

2.2 Decision making of master fighter controllers

An MFC faces a complex decision making problem when selecting a COA to be conducted for a DCA mission. The selected COA must fulfill the objectives of the mission in the best possible way, i.e., the COA must be effective. An effective COA assumes that the relative positions between the flights enable mutual support between them and that the actions of the flights are best suited for the DCA mission. Additionally, the selected COA must be feasible. When selecting the COA, the MFC must take into account the tactical situation on the battlefield with its restricting factors. The capabilities and current locations of friendly and enemy aircraft, as well as operational restrictions, essentially limit the possibilities for different COAs to be actually executed. If it is possible to conduct the COA with the current situation on the battlefield, the COA is considered as feasible.

Identifying and selecting an effective and feasible COA is a challenging task for the MFC for several reasons. First, the MFC needs to address the broad-level guidelines and objectives set to the mission and how they affect the selection of the COA. The MFC must also decide the best approach to utilize assets assigned to the mission. Other simultaneous operations such as OCA missions may require same assets as the DCA missions (Joint Chiefs of Staff, 2017), and the assets available for the DCA missions can be limited. As a prerequisite, the MFC naturally has to be aware of his/her assets as well as enemy's capabilities. Additionally, the airspace can be highly congested and restricted during DCA missions. The ability to operate in different areas may also differ significantly due to, e.g., enemy air defense or electronic countermeasures such as radar jamming. Thus, it is not trivial to determine how long it takes for a flight executing the COA to maneuver in the airspace. Furthermore, there can be multiple flights involved in a single COA. With both attacking and defending aircraft maneuvering in three dimensions at high speeds, the MFC's decision making problem is constantly changing. Moreover, the MFC must make the decisions under a significant time pressure. With the time-sensitive nature of DCA missions, streamlined decision making and coordination of the mission are required from the MFC (Joint Chiefs of Staff, 2017, 2019). A certain COA may be effective if launched immediately, but ineffective a few minutes later. Overall, the MFC faces an exceptionally difficult decision making problem when identifying and selecting a COA which is both effective and feasible.

Training for the counterair missions is an important part to set up and maintain an effective system for counterair operations. While training as a whole is time consuming and expensive, it provides "the final cog in the counterair system" (Holmes, 1995). Especially in time-critical situations training and education are essential as decision making in them is usually based on intuition instead of extensive analysis (Development, Concepts and Doctrine Centre, 2016a). Thus, also MFCs receive training for reacting to enemy threat and performing DCA missions. It is essential that the MFC can perceive and understand the current situation of the battlefield well and decide an effective and feasible COA for the DCA mission based on the available information. With training, these decisions become automatized and implementing them in wartime environment becomes natural (Bennett Jr et al., 2002). Providing training that develops understanding is considered as a necessity (Development, Concepts and Doctrine Centre, 2016a), and the FA model developed in this thesis supports this training and learning process of MFCs.

The FA model addresses the selection of a COA from the perspective of feasibility, i.e., the model helps MFCs to identify which COAs are feasible in decision making situations they are facing. A focal point in determining the feasibility of a COA is commit locations given for the flights where they will start engaging the enemy. The model allows an MFC to evaluate if a COA is actually possible to conduct or not with the selected commit locations when considering the limitations of the current state of the battlefield. The model also helps to identify the importance of all relevant factors when evaluating the feasibility of a COA. By offering a new perspective to the training process, the model supports the learning of the MFCs to improve perception and understanding of the geometry of the battlefield and to make better decisions when selecting feasible COAs for DCA missions. The model, however, does not consider the effectiveness of COAs, i.e., if the COA fulfills any strategic guidelines or if it enables the achievement of the objectives set for the DCA mission. The effectiveness of COAs is discussed in Jalovaara (2021), where a multi-attribute decision analysis model for selecting an effective COA is also constructed.

Geographical restrictions play a pivotal role when considering the feasibility of the COA. Restricted operating zones (ROZs) are areas where operation of some or all airspace users is restricted (Joint Chiefs of Staff, 2014). If all the flying activity on the area is restricted, the area can be also referred to as a no-fly zone (NFZ). The reasons for ROZs in the context of a DCA mission may include, e.g., hostile surface-to-air missiles or anti-aircraft artillery, enemy radars, other critical operations in the area, or friendly surface-to-air missiles or anti-aircraft artillery to avoid friendly fire (see, e.g., Puustinen, 2013; Rantala, 2018). When the MFC is selecting a COA, ROZs greatly increase the complexity of the decision making problem. Consequently, they are a crucial factor to the feasibility of the COA and important to take into account in the FA model.

An essential tactical aspect related to the decision making problem of an MFC is also formations of flights. As discussed in Section 2.1, the fighters the MFC commands operate usually in flights. In addition to the desired commit locations of the flights in the DCA missions, also the positions of the fighters in and between the flights are tactically relevant. Flying and performing combat in formations is done to provide mutual support to other fighters and to concentrate fire (Shaw, 1985). The formations differ in both positioning of the fighters as well as in the distances between the fighters, while latter can also be varied within the same formation. Different formations provide various advantages and disadvantages which are discussed in, e.g., Shaw (1985). While the amount of alternative formations is substantial, several of them have became generally utilized (see, e.g., Shaw, 1985; Korean Air Force, 2005; Farrell et al., 1996) and depicted in Figure 2.1. "Wall" formation has all the four fighters flying next to each other in a row of four. "Finger four" formation is similar to a wall, but the row of four fighters is staggered. "Trail" consists of four fighters one after another. "Box" consists of four fighters forming approximately a square, i.e., consisting of two rows of two fighters. "Diamond" formation has a leading fighter followed by a row of two fighters followed by a trailing fighter. While further discussion about the flight formations is not within the scope of this thesis, these formations are relevant for the FA model as they are a significant part of the COAs the MFC has to decide.

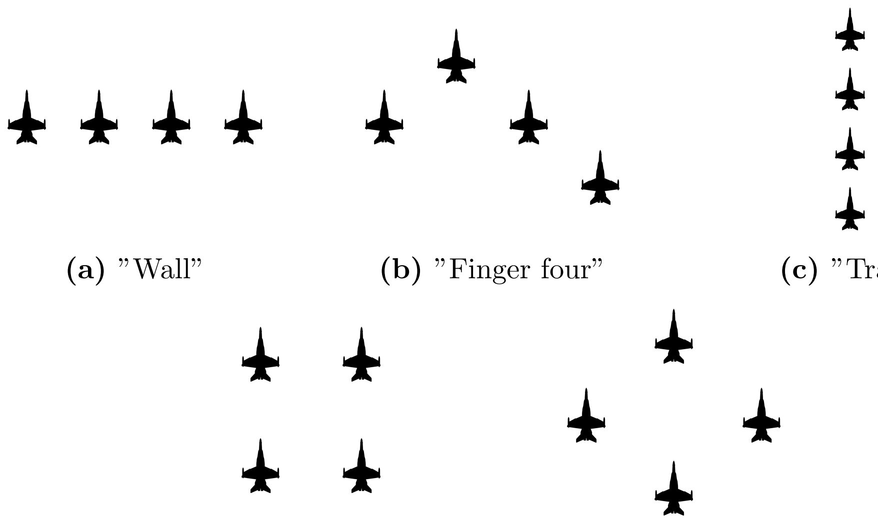

(a) "Wall" **(b)** "Finger four" **(c)** "Trail"

(d) "Box" **(e)** "Diamond"

Figure 2.1: Illustrations of common flight formations. The distances between the fighters are out of scale.

Chapter 3

Fighter allocation model

This chapter introduces a fighter allocation model (FA model) constructed for analyzing the real-life decision making problem of MFCs presented in Section 2.2. The model is utilized to create and solve training scenarios, later simply referred to as scenarios. The scenarios resemble the real-life decision making problems MFCs face, and the model allows answering the question if a decision made by the MFC is good or not. Making a decision for a scenario essentially reflects the selection of a feasible COA including commit locations of fighters or flights. By creating and solving the scenarios, MFCs can train their decision making process by identifying if their decisions, i.e., selection of COAs, are successful or not in different circumstances regarding their feasibility.

3.1 Overview

A problem an MFC essentially faces regarding the feasibility of a COA is how to assign friendly fighters or flights into suitable geographic locations with suitable formations such that they are able to engage successfully against the enemy threat. The decision may involve assigning individual fighters or, e.g., flights containing multiple fighters, where also suitable formations within and between the flights are considered. Creating a scenario in the FA model assumes that the locations of friendly fighters, the enemy threat, and the operational restrictions in a form of ROZs are known. TTPs for each flight are excluded from the model, as they are decided by the flight leaders and FCs, and not by the MFC. Additionally, possible air combat itself is not taken into account.

The model contains three phases. The first phase is referred to as the initialization phase. In this phase, current positions of friendly fighters or

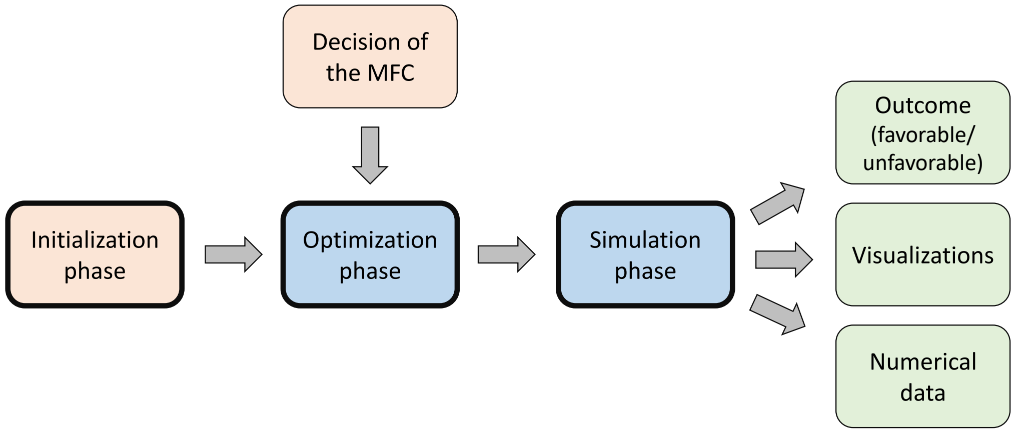

Figure 3.1: Flowchart of the FA model. Steps requiring user interaction are illustrated with red background, computations made by the model with blue background, and outputs provided by the model with green background. The three main phases of the model are highlighted with bold edges.

flights, enemy threat and ROZs are defined. Then, the MFC assigns the friendly fighters or flights to desired commit locations with suitable formations based on his/her choice of a COA. In the optimization phase, the shortest paths of the individual friendly fighters from their initial positions to the destinations, i.e., the commit locations given by the MFC, are computed. The third phase is the simulation phase, where the routes of the friendly and enemy fighters are simulated. The outcome of the scenario is determined with the simulation. The outcome describes if the decision made by the MFC has been successful. Thus, there are two possibilities for the outcome: favorable or unfavorable. In addition to the outcome, the FA model provides visualizations and numerical data as an output. These visualizations and numerical data provide deeper insight on how the outcome is achieved and how the scenario plays out with the decision made by the MFC. The workflow of the FA model is illustrated in Figure 3.1.

The model is implemented with MATLAB (MathWorks, 2015), and a graphical user interface (GUI) is constructed to ease the use of the model. The parameters of the model are set in the GUI and the outputs of the model are presented in it. The GUI provides an option to apply a studied scenario on a map, which improves the construction of scenarios reflecting real-world situations and helps analyzing and perceiving the results of the model.

3.2 Components and assumptions

In the FA model, two opposing forces are considered, referred to as Blue and Red. Red side is considered as hostile, and Blue side as the defensive side performing DCA missions. The FA model describes the situation of the battlefield from the perspective of the MFC of the Blue side. Consequently, it is assumed that the user of the model sees the Blue forces as friendly and Red forces as the enemy. While the model is suitable for any aircraft, in this thesis, fighters are considered. Each fighter of both forces is represented individually. However, the fighters can be also combined into a groups of multiple fighters such as pairs, i.e., 2 fighters, or flights. It is assumed that all the fighters are airborne at the beginning of a scenario.

If the battlefield contains areas where Blue forces cannot operate, ROZs can be applied. The model can also include elements relevant to the defensive intentions of the DCA missions of the Blue side. Important assets, areas of interest, or engagement frontiers are included in the model to describe the battlefield and the MFC's decision making problem in more detail.

The actions of the Red forces are set to the model by the user. These actions are deterministic and no random elements are present. Thus, it is assumed that the threat, i.e., the actions of the Red forces are known and contain no surprise elements or uncertainty.

The airspace is described as a three-dimensional Euclidean space $\mathbb{R}^3$ and the curvature of the Earth is not taken into account. Thus, the geographical locations of all the components in the model are given in Cartesian coordinate system. In addition, the speeds of all fighters are considered as constant ground speeds. In reality, the maximum ground speed for the fighter would be affected by its flying altitude and environmental conditions, but they are here ignored as a simplification. For each fighter, a maximum ground speed is given. In the simulation phase, not all the fighters are necessarily simulated to be flying at their maximum speeds.

3.2.1 Blue forces

Blue forces and their fighters are the ones the MFC is making decisions for. At the initialization phase, the fighters of the Blue forces are set to initial locations. The location of each fighter is described as a point in three-dimensional space, i.e., (x, y, z), where x and y are coordinates of a geographical location of the fighter in a horizontal plane, and z is the current flying altitude above the geographical location. In addition, the fighter is assigned a flying direction at the initial situation in the horizontal plane. This direction describes where the fighter is heading at the initial situation in which the MFC needs to react and make a decision. The direction is given in degrees. Direction of $0°$ implies that the fighter is heading to north and the direction is measured clockwise, e.g., $90°$ corresponds to west and $180°$ to south.

The initial locations, the flying altitudes and the initial directions of the Blue fighters do not necessarily need to be assigned individually. Instead, the fighters can be combined into groups consisting of multiple fighters. Usually the desired group size is four fighters, i.e., a flight. In the model, a flight can be assigned the initial location in xy-plane, the flying altitude and the initial flying direction, which are then inherited to the fighters in the flight. The exact locations of the fighters inside the flight are determined by the formation of the flight. In addition to five formations presented in Figure 2.1, any arbitrary formation can be manually constructed by assigning four fighters initial locations and destinations next to each other such that they form a desired formation.

In addition to the initial location, the fighters are also assigned a desired destination. This destination describes the commit location decided by the MFC. Similarly to the initial location, it is also a point (x, y, z) in Cartesian coordinate system. The final direction of the fighter at the destination needs to be also defined, and it is given similarly as with the initial location. Instead of assigning destinations and final directions for individual fighters, they can be also given for a flight. The destination and final direction are then inherited to the fighters in the flight. The exact destinations of the fighters are again determined by the formation of the flight.

A flight can also have no formation. This corresponds to circumstances where the fighters of the flight are not together at the initial situation and may have arbitrary locations, but they are still assigned to form a flight at the destination, or vice versa.

Fighters with different capabilities can be included in the model. For each fighter, the maximum flying speed as a ground speed is set. Each fighter has also a commit envelope defined as an area in the shape of a spherical cone in front of the fighter. The commit envelope describes the ability of the fighter to start engaging an enemy. When an enemy fighter is under the commit envelope of a fighter, it is assumed that the fighter can start air combat against this enemy fighter. The commit envelope is defined by three parameters: a commit range, and commit aspects in horizontal and vertical dimensions. The commit range determines the length of the envelope. The horizontal and vertical commit aspects define the maximum angles the enemy fighter can be positioned horizontally and vertically with respect to the flying direction of the fighter, respectively. Furthermore, maximum horizontal turn angles, maximum angles of ascent and descent as well as minimum and maximum flying altitudes are defined separately for each fighter.

3.2.2 Red forces

Red forces describe the enemy which poses a threat the MFC has to encounter. In the FA model, it is assumed that the activity of the Red forces, i.e., the initial locations, routes and destinations of the Red fighters are known. A route for each Red fighter is parameterized in the initial phase of the model by setting waypoints in a horizontal plane. At least two waypoints are required, i.e., an initial location and a destination. Intermediate waypoints can be given if needed. It is assumed that the flying altitude of all the Red fighters is fixed, i.e., they fly at a constant altitude. The Red fighters are given a constant ground speed, and the speeds can be assigned for each fighter individually. Commit envelopes are also defined for each Red fighter by specifying their ranges as well as aspects in horizontal and vertical planes.

3.2.3 Restricted operating zones

ROZs are areas where Blue fighters cannot operate. In the FA model, when computing the shortest paths from initial locations to destinations, the paths must avoid all the defined ROZs. Thus, ROZs may significantly increase the distance and time needed to travel between the initial locations and the destinations of the Blue fighters.

ROZs are defined as polyhedra. Their corner points are given as points in Cartesian coordinate system in (x, y, z) space. The number of corners in a ROZ is not limited, and consequently its shape is not constrained. Usually a ROZ is defined to lay in a certain geographical area between certain altitudes. Thus, most commonly a polyhedron describing a ROZ can be simplified as a polygon in a horizontal plane having maximum and minimum altitudes.

3.3 Computation of shortest paths

After the scenario is initialized, shortest paths of Blue fighters are solved. The three-dimensional airspace is discretized in order to compute the shortest paths for the fighters from their initial locations to their destinations. The discretization is a three-dimensional grid of nodes, i.e., a graph. The graph consists of nodes $i \in V$ and edges $(i, j) \in E$ between the nodes, and it is denoted as $G(V, E)$. The nodes represent the points which a fighter can travel between, and the edges describe the possible movements of the fighter between the nodes. Here, the graph is weighted, i.e., each edge $(i, j) \in E$ has an individual cost c_{ij} which represents the distance between nodes i and j. Additionally, the graph is directed, i.e., the cost an edge is dependent on the direction of the movement. This is due to the ability of a fighter to ascend and descend in different angles, i.e., an edge between vertical planes of the graph cannot be necessarily traversed in both directions. When utilizing graphs in models, graphs are often referred also as to networks, and solving shortest path problems or other optimization problems using graphs is referred to as network optimization (Bertsekas, 1998).

3.3.1 Structure of the network

A three-dimensional network $G(V, E)$ is applied to describe a geographical location under consideration. Both the vertical and horizontal dimensions of the network are defined such that the initial locations and destinations of all the fighters in the scenario are covered by the network. Additionally, the vertical dimensions are selected such that they limit the minimum and maximum altitudes according to the possible flying altitudes of the fighter.

The nodes in the network are evenly spaced in a horizontal plane, and the distance d_{hor} between the nodes is defined. Smaller value for d_{hor} can lead to more realistic routes but increases the computational complexity of the optimization problem. The nodes also have an even vertical spacing between the horizontal planes. However, the distance between the nodes in consecutive horizontal planes, i.e., d_{ver}, can differ from the horizontal spacing d_{hor}. When considering the vertical spacing of the nodes, the capabilities of the fighter need to considered. A fighter has physical constraints how steep it can ascend or descend. Thus, the vertical spacing d_{ver} is selected such that it is compatible with the maximum angles of ascent and descent of the fighter. Additionally, the vertical spacing determines the lowest possible ascent or descent angles between two nodes. Thus, using smaller vertical distance d_{ver} than the minimum angles of ascent and descent of a fighter is not reasonable.

The selection of connections, i.e., edges from a node to its neighboring nodes to form the network, is carried out depending on the desired solution time of the problem and the realism of the routes. In a horizontal plane, there are connections to all adjacent nodes as well as to some nodes two and three steps away in the grid. In total, there are 32 horizontal edges from one node. The connections are depicted in Figure 3.2a. This amount of edges is selected to provide enough possibilities of horizontal turn angles that are evenly spread. However, in order to decrease the computational difficulty of the problem, the amount of edges is still kept reasonably low. Between vertical planes, a node is also connected to nodes in three adjacent

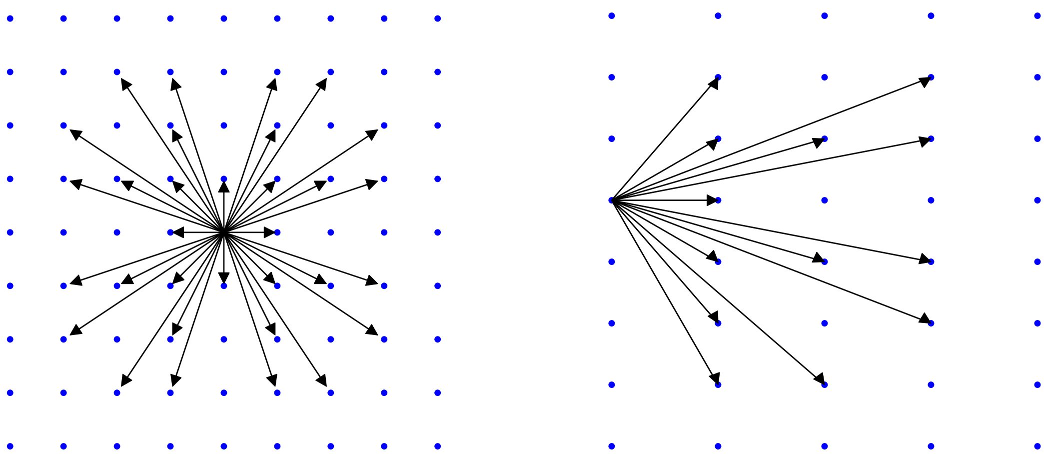

(a) Horizontal connections

(b) An example of potential connections between vertical and horizontal planes

Figure 3.2: Connections between the nodes in the network. There are in total 32 connections from one node in horizontal planes. The vertical planes are connected to three adjacent planes. The horizontal edges are always constructed as presented in Figure a), while the edges between the horizontal and vertical planes seen in Figure b) depend on the maximum ascending and descending angles of a fighter.

vertical planes. The amount of horizontal planes that a node is connected to is selected such that the maximum ascending and descending angles of the fighter are satisfied. Examples of potential edges between vertical and horizontal planes are illustrated in Figure 3.2b.

As some elements of the network are based on capabilities of fighters, separate networks are built for different types of fighters. If there are various types of Blue fighters in one scenario, multiple networks are constructed and the shortest paths are computed in the respective networks for each fighter.

The nodes lying on areas covered by ROZs are denoted by V_{ROZ}. As flying is not permitted within ROZs, using edges connected to nodes in V_{ROZ} are not allowed when solving shortest paths. Therefore, the shortest paths for any of the fighters cannot visit any nodes in V_{ROZ}.

3.3.2 Optimization problem

For each fighter, the shortest path from the source node $s \in V$ (corresponding to the initial location of the fighter) to the target node $t \in V$ (corresponding to the destination of the fighter) in a network $G(V, E)$ is solved. The goal is to find a shortest path P which minimizes the cost of the path from the source node s to the target node t when each edge (i, j) is associated with a cost c_{ij}. The decision variables of the optimization problem are binary variables x_{ij} which describe if the edge (i, j) is included in the shortest path or not, i.e.,

$$x_{ij} = \begin{cases} 1, \text{ if } (i, j) \in P, \\ 0, \text{ if } (i, j) \notin P. \end{cases} \tag{3.1}$$

The total cost of the shortest path P is computed as

$$\sum_{(i,j)\in E} c_{ij} x_{ij}, \tag{3.2}$$

where c_{ij} is the cost of an edge (i, j). Equation (3.2) is the objective function to be minimized. The problem is constrained such that the values of x_{ij} form a path P from the source node s to the target node t. The nodes on the path P are not allowed to be inside ROZs, i.e., $V_{ROZ} \cap P = \emptyset$. Furthermore, consequent edges on the path must respect the maximum horizontal turn angle α_{max} of the fighter.

In order to achieve the desired flying direction at the destination of the fighter, an auxiliary node u is specified in addition to the source and target nodes s and t. The node u is defined as the second-last node of the shortest path. It is selected from the neighbors of the target node t such that the direction of the edge (u, t) is as close to the desired direction of the fighter at the target node as possible. The edge (u, t) is always included in the shortest path. Therefore, the shortest path is solved such that the direction of the fighter at the destination is as desired while the maximum turn angle α_{max} is still always respected when arriving to the target node t.

The shortest path problem is formulated as

$$\min_{x_{ij}} \sum_{(i,j)\in E} c_{ij}x_{ij} \tag{3.3}$$

$$\text{s.t.} \quad \sum_{j|(i,j)\in E} x_{ij} - \sum_{j|(j,i)\in E} x_{ji} = \begin{cases} 1 & \text{if } i = s, \\ -1 & \text{if } i = t, \\ 0 & \text{otherwise.} \end{cases} \tag{3.4}$$

$$x_{ik}x_{kj}\left(\measuredangle((i,k)(k,j)) + \alpha_{max} - 180^\circ\right) \geq 0 \quad \forall\, i,j,k : (i,k),(k,j) \in E, \tag{3.5}$$

$$\sum_{k\in V_{\mathrm{ROZ}}} x_{ik} = 0 \quad \forall\, i \in V, \tag{3.6}$$

$$x_{ut} = 1 \quad , \tag{3.7}$$

$$x_{ij} \in \{0,1\} \quad \forall\, i,j \in V, \tag{3.8}$$

where c_{ij} is the distance of an edge (i,j), s is the source node, t is the target node, α_{max} is the maximum turn angle of the fighter in a horizontal plane, and V_{ROZ} is a set of nodes inside ROZs. Constraint (3.4) defines that the number of incoming and outgoing edges for a node is equal, except that the start node can have only one outgoing edge and the target node can have only one incoming edge. Constraint (3.4) restricts the angle between the consecutive edges to be above $180^\circ - \alpha_{max}$. Consequently, the turns on the shortest path do not exceed the maximum horizontal turn angle α_{max}. Constraint (3.6) prevents any edges leading to nodes in areas covered by ROZs to be included in the shortest path. Constraint (3.7) forces the edge between the auxiliary node u and the target node t to be included in the shortest path. Constraint (3.8) provides the feasible values for the decision variables.

3.3.3 Solution

The optimization problem for each fighter is solved using A* algorithm (Hart et al., 1968), which is an extension of Dijkstra's shortest path algorithm (Dijkstra, 1959). Dijkstra's algorithm is probably the most common and well-known algorithm for solving single-source shortest path problems in a weighted graph with non-negative weights, i.e., costs. Dijkstra's algorithm traverses the graph and updates the smallest cumulative costs $g(i)$ from the source node to the node i of all unvisited nodes until it either has visited all the nodes in the network or reaches the target node. Thus, it finds the shortest path from the source node to all the nodes including also the target node. The algorithm is discussed in more detail in, e.g., Cormen et al. (2009).

A* algorithm extends Dijkstra's algorithm by applying a heuristic function $h(i)$ to guide the algorithm to traverse the graph (see, e.g., Nilsson, 1982). While Dijkstra's algorithm visits all the nodes in the network, the use of the heuristic function reduces the number of visited nodes and can make the graph traversing faster (see, e.g., Pearl, 1984). The heuristic function $h(i)$ is an optimistic estimation of the smallest cost from node i to the target node. Whereas Dijkstra's algorithm selects the next node to be visited based on the current smallest cumulative cost $g(i)$, i.e., selecting node with $\min(g(i))$, A* also considers the heuristic function $h(i)$ and selects the next node to be visited based on

$$\min\left(g(i) + h(i)\right), \tag{3.9}$$

where i is an unvisited node, $g(i)$ is the current smallest cumulative cost from the source node to the node i, and $h(i)$ is an estimate for the remaining path from node i to the target node. Thus, with a reasonable heuristic function, the algorithm avoids to visit nodes that are not likely to be part of the shortest path.

A* algorithm requires the heuristic function $h(i)$ to be admissible to guarantee to return the shortest path with the minimum cost (Hart et al., 1968). The heuristic is admissible if it never overestimates the actual cost from a node i to the target node.

In the FA model, Euclidean distance from the node n to the target node t is employed as the heuristic function $h(n)$. With the nodes n and t, the heuristic function is

$$h(n) = \sqrt{(x_t - x_n)^2 + (y_t - y_t)^2 + (z_t - z_n)^2}. \tag{3.10}$$

If there are ROZs in the network, the heuristic function is defined to also take them into account. When selecting the next node m to be visited and computing the value of the heuristic function $h(m)$, it is also checked if the straight line between the node m and the target node t intersects any ROZ. If it does, the distances of straight lines from the node m via each corner point of the ROZ to the target node t are computed. These distances are estimates for the length of the route between the nodes m and t. Therefore, the value of the heuristic function then estimates better the actual cost of the remaining path, as it cannot go through an ROZ. Furthermore, selecting the lowest distance among these distances via each corner point never overestimates the cost of the actual path, thus keeping the heuristic admissible.

If there are multiple ROZs between the nodes m and t, the minimum distance to travel via one corner point of each ROZ is computed as described above. As the shortest path must avoid all the ROZs, the actual path has a cost at least as large as the largest of these distances. Thus, the value of the heuristic function is selected to correspond to the ROZ which has the maximal smallest distance via any of its corner points to the target node. Therefore, heuristic function in the presence of ROZs is

$$
\begin{aligned}
h(m) = \max_{r \in R} \min_{q_i^r} & \sqrt{\left(x_{q_i^r} - x_m\right)^2 + \left(y_{q_i^r} - y_m\right)^2 + \left(z_{q_i^r} - z_m\right)^2} \\
& + \sqrt{\left(x_t - x_{q_i^r}\right)^2 + \left(y_t - y_{q_i^r}\right)^2 + \left(z_t - z_{q_i^r}\right)^2},
\end{aligned}
\tag{3.11}
$$

where r is the index of a ROZ belonging to the set of ROZs R, q_i^r is the ith corner point of a ROZ r, and (x_m, y_m, z_m), (x_t, y_t, z_t) and $\left(x_{q_i^r}, y_{q_i^r}, z_{q_i^r}\right)$ are coordinates of the next node, the target node and the corner points of the ROZs, respectively. In practice, as the amount of ROZs is usually small, the minimum and maximum values in the heuristic function (3.11) are computed by iterating through all the ROZs and their corner points.

The optimization problem (3.3)–(3.8) is solved using A* algorithm with heuristic functions (3.10) and (3.11) for each Blue fighter. The algorithm returns the shortest path for each fighter which obeys the constraints set by the maximum turn angles of the fighter as well as the geographical limitations set by ROZs.

3.4 Evaluation of routes using simulation

The shortest paths, later referred to as optimal routes, of each Blue fighter computed in the optimization phase are then utilized in the simulation phase of the FA model. In this phase, the whole scenario is evaluated using the routes of Blue and Red fighters. The simulation reveals the moments of potential exposures of the Blue fighters to the Red fighters before the desired commit locations are achieved. This is enabled by taking into account the commit envelopes of the Blue and Red fighters. The simulation determines the outcome of the scenario and if the decision made by the MFC has been successful or not. The outcome of the scenario can be either favorable or unfavorable.

In the simulation, the routes of the Blue and Red fighters are first preprocessed. The routes are divided into samples corresponding to a flight time of one second. Additionally, the routes are smoothed using moving average. This is done in order to modify sharp angles on the routes into smoother turns to reflect real-life turns of fighters. As the routes of Blue fighters are solved with the network optimization, the solution in an even-spaced network might contain turns that do not look realistic. Furthermore, the routes of the Red fighters can be also defined by only few waypoints and contain sharp angles.

The scenario is simulated using the preprocessed routes. In the simulation, the Blue and Red fighters are set to fly their routes with given flying speeds, and the time is incremented with steps of one second. For every second, it is checked if any of the Blue fighters are under the commit envelope of any of the Red fighters, i.e., if the fighter lies inside the spherical cone describing the commit envelope. Additionally, it is checked if any of the destinations of the Blue fighters are under the commit envelope of any Red fighter. The simulation is terminated when the last fighter has arrived to its destination.

The Blue fighters may arrive at their final destinations at different times, as the distances of the optimal routes as well as the flying speeds can vary significantly. As this is not a desired arrangement in a real-life situation, the speeds of the Blue fighters can be also determined such that the Blue fighters arrive to their destinations exactly at the same time. In this case, the fighter having the longest flight time for the optimal route is set to fly the route at its maximum speed and arriving at its destination at time t_f. The flying speeds of all the other Blue fighters are scaled such that they arrive at their final destinations also at time t_f. While this may lead to unrealistically low flying speeds especially at specific altitudes, it is considered a valid approximation in this model. It addresses accurately enough the real-life situation where the fighters need to vary their flying speeds to achieve the desired locations and flight formations. It is also possible to simulate a situation where all the Blue fighters fly at their maximum speeds and may arrive to the final destinations at different times. This allows easier investigation of the margins each fighters has, i.e., how early or late each fighter is when they use their maximum flying speed.

There are four possible end conditions which determine if the outcome of a scenario is favorable or unfavorable:

1. The commit envelope of a Red fighter has reached the final destination of a Blue fighter before the Blue fighter has arrived there, i.e., the desired commit locations cannot be reached in time.
2. One or more Blue fighters are within the commit envelope of a Red fighter before they have reached their destinations, i.e., the Red fighter can engage the Blue fighter before the desired commit locations are reached.
3. One or more Red fighters have reached the ends of their routes before one or more Blue fighter has.
4. The Blue fighters have arrived at their destinations before none of the earlier end conditions are met.

The last end condition 4 is the desired one for the MFC, as it indicates that the commit locations given for the fighters have been feasible and the decision he/she made has been successful. Thus, the fourth end condition leads to a favorable outcome of the scenario. The end conditions 1-3 are not desired, as they all indicate that the Blue fighters are not able to reach their destinations in time to achieve the commit locations decided by the MFC. Thus, when any of the end conditions 1-3 are met, the outcome of the scenario is considered unfavorable. Examples of each end condition are depicted in Figure 3.3.

3.5 Outputs

The FA model produces outputs after the simulation phase is completed. The simplest output is the outcome of the decision: favorable or unfavorable, i.e., if the fighters have arrived to their destinations in time or not. In addition to providing the outcome, the model also allows examining the scenario and its outcome by numerical data and visualizations.

The FA model and its implementation provides a second-by-second visualization of the scenario. The visualization portrays the flights of the Blue and Red fighters from their initial locations to their destinations. The scenario can be visualized from the beginning of the scenario, i.e., t_0, to the moment when the last fighter arrives to its destination, i.e., t_{end}. Any point of time between t_0 and t_{end} can be presented with an accuracy of one second. The visualization can be also viewed as an animation with selected speed. The moments when an end condition is triggered are highlighted in the animation, and these moments can be found instantly in the visualization.

Furthermore, numerical data of the scenario is generated. It contains key numbers regarding each Blue fighter in the scenario. The data includes length and flight times of complete routes. Additionally, the data offers an information if each fighter can arrive their destinations in time, i.e., feasibility of the destination of each fighter. The feasibility is obtained with respect to two perspectives: the whole scenario or the single fighter. Regarding the former case, it is checked if the fighter has made its destination when any of the Blue fighters has first triggered an end condition leading to the unfavorable outcome. This is referred to as the collective feasibility. Only if all the destinations of the fighters in the scenario are collectively feasible, the outcome of the scenario is favorable. Regarding the latter case, only the single fighter is taken into account, i.e., if the fighter could make its destination before it triggers any of the end conditions leading to the unfavorable outcome. This is referred to as the individual feasibility. It describes the outcome if there were no other fighters in the scenario. Even if the destination of a fighter is individually feasible, it might not be collectively feasible. Furthermore, if the destination of a fighter is not individually feasible, it cannot be collectively feasible. If there is only one Blue fighter in the scenario, the collective and individual feasibility are equivalent. If there are multiple Blue fighters, the collective feasibility is significant for the outcome of the whole scenario. For both cases of feasibility, also a margin is provided. The margin is the time difference between the moment when an end condition for the unfavorable outcome is triggered and when the fighter arrives to its destination. If the fighter arrives late, the margin is negative, and otherwise it is positive.

Visualizations and numerical data provide additional insight on why the outcome of the scenario is the obtained one. They allow evaluations of why the decision made by the MFC has been successful or not. Additionally, it can be analyzed how the decision could be adjusted. From numerical data, one can perceive which fighters are clearly too late, which are nearly on time or late, and if some fighters have too conservative destinations. Data also quickly highlights which fighters in the scenario are crucial for a successful decision, i.e., it is critical that they are assigned to sufficient destinations. Visualizations enable a comprehensive observation of the scenario and its evolution at any point of time. Numerical data and visualizations complement each other and both support the learning process of selecting a feasible COA.

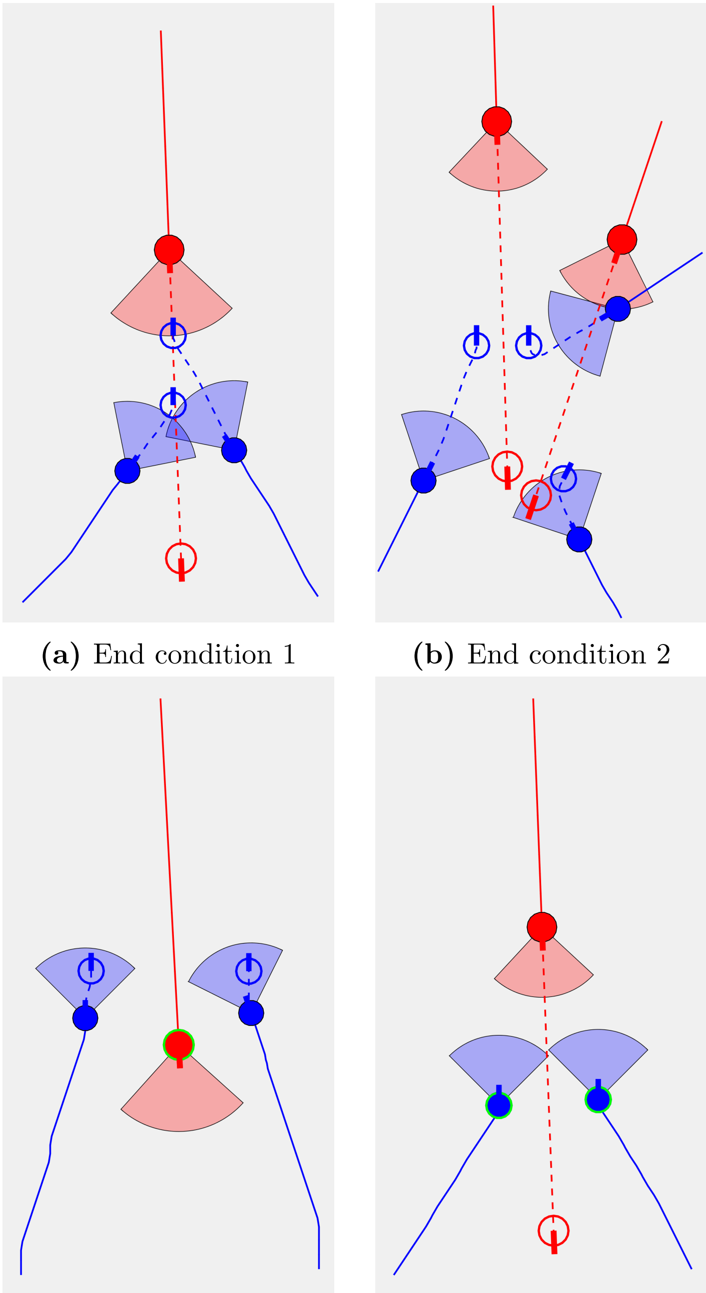

(a) End condition 1 **(b)** End condition 2

(c) End condition 3 **(d)** End condition 4

Figure 3.3: Four end conditions of the simulation presented with visualizations obtained from the FA model. End conditions 1-3 lead to the unfavorable outcome for the scenario, while end condition 4 results in the favorable outcome. The visualization presents the scenario at a certain moment of time. Blue and Red fighters are depicted as filled circles with respective colors, while routes flown are illustrated as solid lines and the rest of the routes as dashed lines. Segments adjacent to the fighters present the commit envelopes, and unfilled circles the destinations of the fighters. Green edges around fighters indicate that they have reached their targets.

Chapter 4

Utilization of the fighter allocation model in training of master fighter controllers

This section introduces utilization practices of the FA model. It is employed to analyze three example scenarios that could be used in the training of MFCs. Examples with both favorable and unfavorable outcomes are given and their solutions are discussed. The first two scenarios involve three Blue fighters and one Red fighter. The third one is a scenario with a larger scale and contains four Blue flights and two Red flights. The outputs of the model, i.e., visualizations and numerical data, are presented and discussed for each scenario. Parameters and the numerical results of the scenarios are given in imperial units, as they are commonly used in aviation. Thus, distances are given in nautical miles (nmi), altitudes in feet, and speeds in knots.

4.1 Scenario with the favorable outcome

The first scenario contains a threat posed by Red forces that is assumed to be one fighter. The MFC has three Blue fighters to perform a DCA mission on the Red fighter. The initial situation of the scenario is depicted in Figure 4.1. The initial situation describes the locations and directions of all the Blue fighters as well as the location of the Red fighter and its route. The MFC has also to consider an important asset which needs to be protected. Flying altitudes, flying speeds as well as the commit ranges and aspects of the fighters are presented in Table 4.1. The commit envelopes are also visualized for the MFC in Figure 4.1. It is assumed that the flying altitudes of all the fighters remain constant, and that the Red fighter is slightly faster than the

Blue fighters. Additionally, all the Blue fighters are flying at their maximum speed and may reach their destinations at different times. The Red fighters are not allowed to reach the important asset within the radius of 55 nautical miles.

The MFC has to assign the three Blue fighters to suitable commit locations to perform a DCA mission. The hypothetical decision of the MFC, i.e., the commit locations of the Blue fighters, are depicted in Figure 4.2.

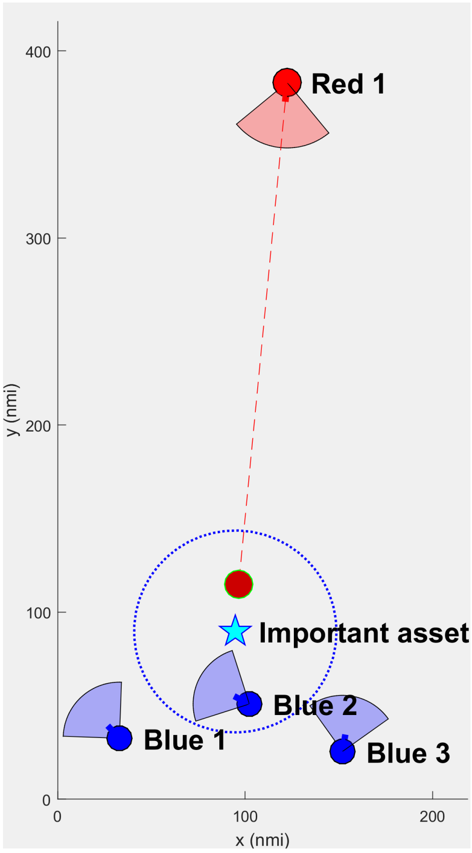

Figure 4.1: Initial situation of the scenario with the favorable outcome describing locations of the Blue fighters and the Red fighter as filled circles with respective colors, the route of the Red fighter, and the important asset. The directions and commit envelopes of all the fighters are also visualized.

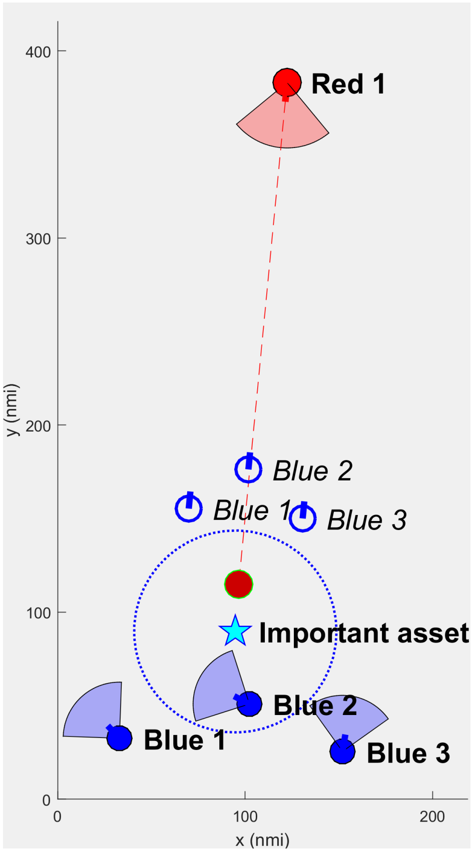

Figure 4.2: Initial situation of the scenario with the favorable outcome with the destinations of the Blue fighters given by the MFC. The destinations are illustrated with unfilled blue circles for each fighter, and the desired directions at the destinations are also depicted.

Table 4.1: Parameters of the Blue and Red fighters in the scenario with the favorable outcome.

Unit	Altitude (feet)	Flying speed (knots)	Commit range (nmi)	Commit aspect (degrees)
Blue 1	32 000	500	30	45
Blue 2	32 000	500	30	45
Blue 3	32 000	500	30	45
Red 1	32 000	570	35	45

After the commit locations decided by the MFC are set to the FA model, the optimization and simulation phases are conducted. First, the optimal routes of the Blue fighters are computed with a network having horizontal spacing of 2.7 nautical miles between the nodes. A turning constraint of 45 degrees is applied, i.e., the route cannot contain turns steeper than 45 degrees. The optimal routes of the fighters are presented in Figure 4.3. The nodes of the network as well as its horizontal dimensions are also illustrated.

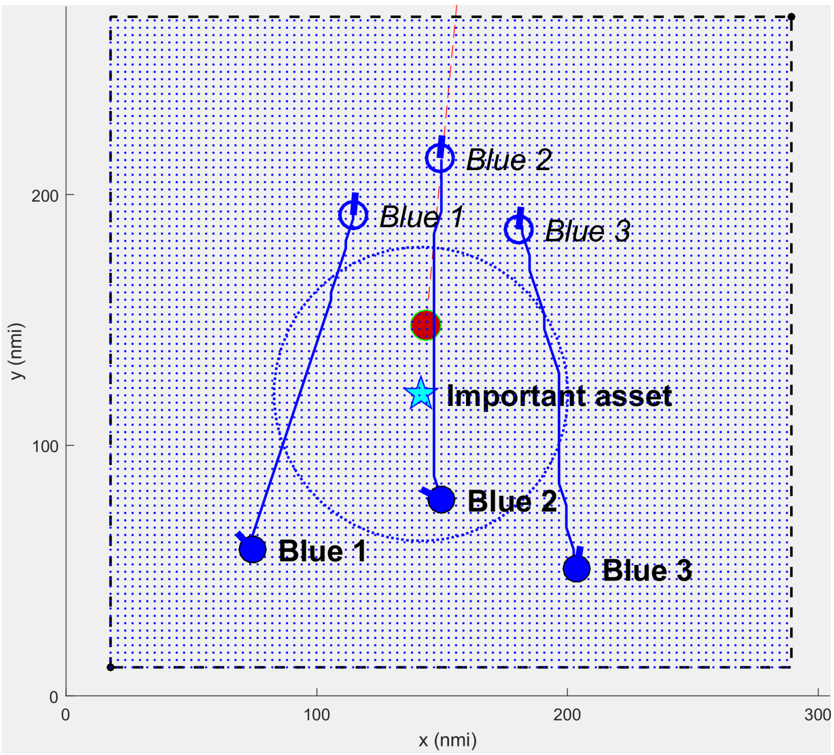

Figure 4.3: Optimization network and the optimal routes computed for the Blue fighters in the scenario with the favorable outcome.

The simulation phase is conducted after the optimization phase. The simulation is performed and the outcome of the scenario is obtained. Visualizations at different points of time are presented in Figure 4.4. Figure 4.4a describes the situation eight minutes after the initial situation. None of the end conditions discussed in Section 3.4 is met and no outcome is yet reached. Figure 4.4b describes the situation when one of the end conditions has been activated. All of the three Blue fighters have reached their destinations, and none of the end conditions leading to the unfavorable outcome, i.e, the end conditions 1-3, is triggered. This indicates that no Blue fighter or destination of a Blue fighter has been in the commit envelope of the Red fighter, and the Red fighter has not reached its target. Thus, the scenario has concluded in the end condition 4, which implies the favorable outcome. Consequently, the destinations given by the MFC have been feasible and the MFC's decision has been successful.

Numerical results are presented in Table 4.2. They contain the distances of the routes and flying times for each Blue fighter. It is also indicated which Blue fighters are able to reach their destinations successfully before any of the end conditions 1-3 are met in the scenario. This corresponds to the collective feasibility of each fighter. Additionally, it is stated if each fighter can make to its destination without triggering any of the end conditions, i.e., if it is individually feasible. Furthermore, a margin for each fighter regarding both collective and individual feasibility is given.

The length of the routes and flying times are similar for each fighter. The difference in flying times between the first and the last Blue fighter to arrive to their destinations is only 34 seconds. As all the fighters reached their destinations in time, all of them are both collectively and individually feasible and have positive margin in both cases. This implies that all the fighters reach their destinations successfully, which can be seen also with the visualization presented in Figure 4.4b. The margins for collective feasibility describe that Blue 1 is two minutes and 39 seconds earlier at its destination than when an end condition for the unfavorable outcome would occur. Based on the visualizations of the scenario in Figures 4.2 and 4.4b, it would occur when the commit envelope of the Red fighter would hit the destination of Blue 2. Similarly, Blue 2 is three minutes and 17 seconds early from that moment, and Blue 3 is two minutes and 43 seconds early. Thus, each fighter would have had a few minutes to spare during their routes.

The margins for individual feasibility of the fighters are 12 minutes and 45 seconds for Blue 1, three minutes and 17 seconds for Blue 2, and 12 minutes and 49 minutes for Blue 3. For Blue 2, the individual margin is identical to the collective margin. This occurs because the same event leading for an end condition, i.e., when the commit envelope of the Red fighter reaches the destination of Blue 2, defines both the margins. For Blue 1 and Blue 3, the individual margins are determined by the end condition 3, i.e., when the Red fighter arrives to its destinations. This is due to the commit envelope of the Red fighter reaching neither of these fighters nor their destinations. Consequently, the individual margins are large as Blue 2 and Blue 3 are at their destinations significantly earlier than the Red fighter at its destination.

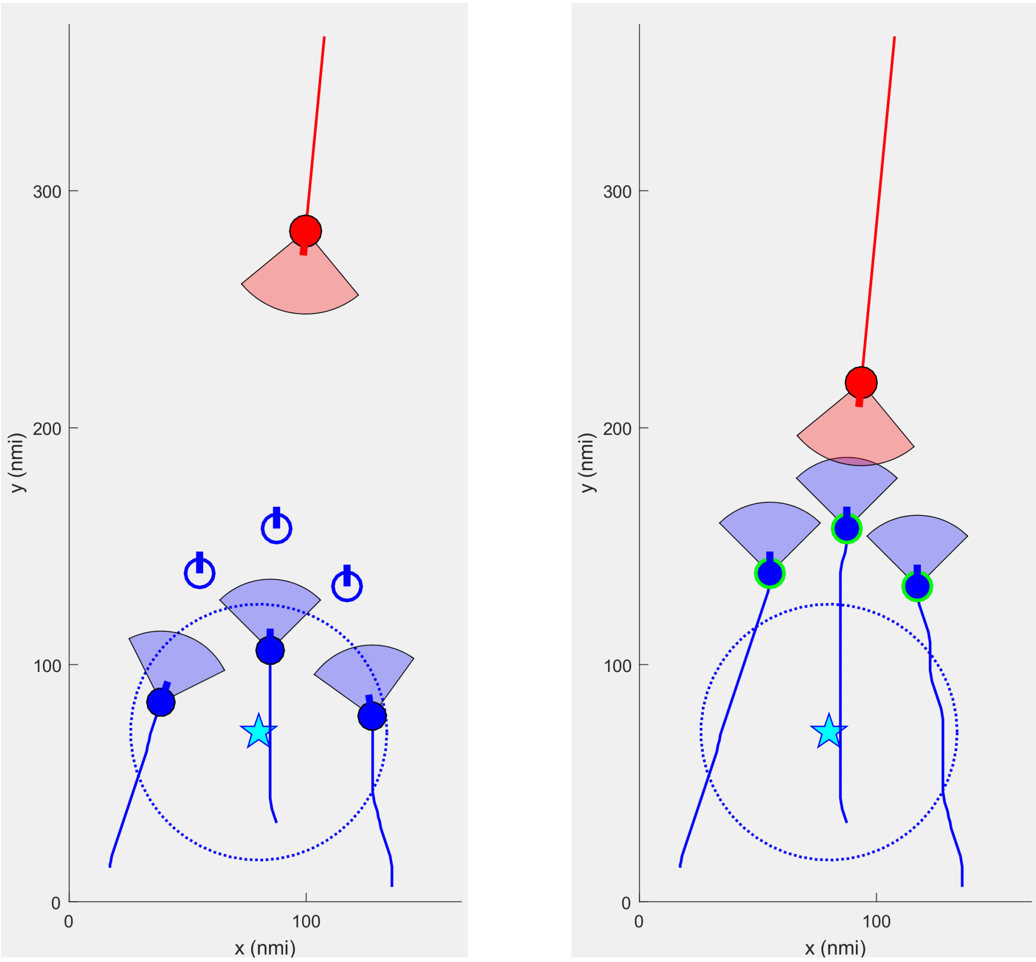

(a) **(b)**

Figure 4.4: Locations of the fighters and routes they have flown in the scenario with the favorable outcome a) after 8 minutes from the initial situation, when none of the end conditions of the simulation is met, and b) all the Blue fighters have achieved their destinations and an end condition of the simulation is met. The decision of the MFC is successful. Green edges around the fighters indicate that they have reached their destinations.

Table 4.2: Numerical results of the simulation in the scenario with the favorable outcome.

Unit	Flying speed (knots)	Distance of the route (nmi)	Flying time (minutes)	Feasible? (collective)	Collective margin (minutes)	Feasible? (individual)	Individual margin (minutes)
Blue 1	500	130	15:39	Yes	+2:39	Yes	+12:45
Blue 2	500	125	15:01	Yes	+3:17	Yes	+3:17
Blue 3	500	130	15:35	Yes	+2:43	Yes	+12:49

4.2 Scenario with the unfavorable outcome

In the second scenario, a threat posed by Red forces is again in a form of one fighter. The MFC also has three Blue fighters to perform a DCA mission on the Red fighter. The initial situation of the scenario is presented in Figure 4.5. Flying altitudes, flying speeds, and ranges and aspects for the commit envelopes of the fighters are similar to the ones in the first scenario and given in Table 4.1. It is assumed that the flying altitudes of all the fighters remain constant, and that the Red fighter is slightly faster than the Blue fighters. All the Blue fighters are flying at their maximum speed and may reach their destinations at different times. The Red fighters are not allowed to reach the important asset within the radius of 55 nautical miles.

The MFC has to assign the three Blue fighters to suitable locations to perform a DCA mission. Figure 4.6 presents the visualization of the hypothetical decision of the MFC, i.e., the destinations of the three Blue fighters.

The optimal routes of the Blue fighters are again computed with a network having horizontal spacing of 2.7 nautical miles between the nodes and with a turning constraint of 45 degrees. The optimal routes and the utilized network are visualized in Figure 4.7.

The outcome of the scenario is obtained after the simulation phase. Visualizations at two different moments are presented in Figure 4.8. Figure 4.8a describes the situation 14 minutes and 21 seconds after the initial situation. At that moment, the commit envelope of the Red fighter reaches the destination of the fighter Blue 2 before the fighters Blue 2 and Blue 3 have arrived to their destinations. Thus, the end condition 1 is met, and the outcome of the scenario is unfavorable. The decision made by the MFC is unsuccessful. The situation when all the Blue fighters have arrived to their destinations is depicted in Figure 4.8b. This is achieved 23 minutes and 3 seconds after the initial situation. It is revealed that the Red fighter has passed all the destinations of the Blue fighters at that moment. The commit locations given by the MFC for the Blue fighters are clearly too far to be reached in time, and

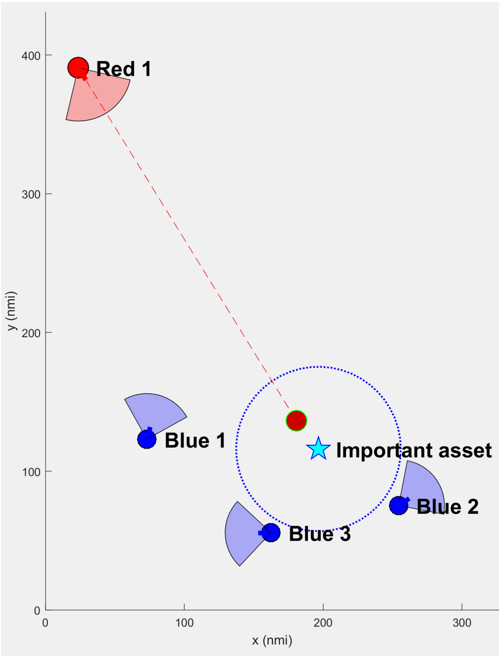

Figure 4.5: Initial situation of the scenario with the unfavorable outcome describing locations of the Blue fighters and the Red fighter as filled circles with respective colors, the route of the Red fighter, and the important asset. The directions and commit envelopes of all the fighters are also visualized.

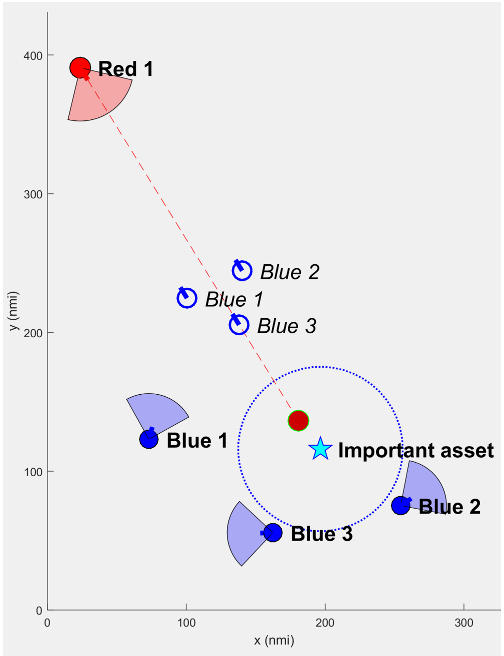

Figure 4.6: Initial situation of the scenario with the unfavorable outcome with the destinations of the Blue fighters given by the MFC. The destinations are illustrated with unfilled blue circles for each fighter, and the desired directions at the destinations are also depicted.

some of the Blue fighters are significantly too late.

Numerical results of the scenario are given in Table 4.3. The fighter Blue 1 is the only one that can make its destination in time, which is also seen in the visualizations. This is because its route and therefore its flying time are significantly shorter, 99 nautical miles and slightly under 12 minutes, than those of Blue 2 (192 nautical miles and over 23 minutes) and Blue 3 (141 nautical miles and slightly under 17 minutes). While Blue 1 is two minutes and 30 seconds early at its destination, Blue 2 is eight minutes and 42 seconds late and Blue 3 two minutes and 36 seconds late for achieving the favorable outcome. Consequently, the margin for the collective feasibility is positive for Blue 1, but negative for Blue 2 and Blue 3. Blue 1 has identical margins for both collective and individual cases, which implies that Blue 1 activates an end condition for the scenario. The destination given for Blue 2 is clearly not reachable, as its individual margin is eight minutes and 22 seconds negative. With a negative individual margin, it would not make to its destination in time without triggering an end condition for the unfavorable outcome itself. While the destination of Blue 3 is not collectively feasible, it has a positive margin for the individual feasibility. Thus, it can make its destination before the Red fighter would reach it, while being late when considering the feasibility of the whole scenario.

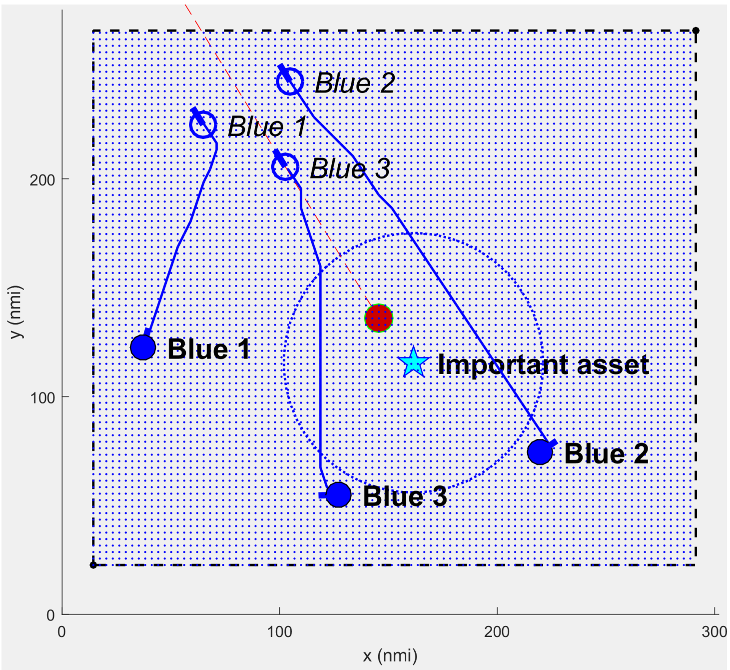

Figure 4.7: Optimization network and the optimal routes computed for the Blue fighters in the scenario with the unfavorable outcome.

Table 4.3: Numerical results of the simulation in the scenario with the unfavorable outcome.

Unit	Flying speed (knots)	Distance of the route (nmi)	Flying time (minutes)	Feasible? (collective)	Collective margin (minutes)	Feasible? (individual)	Individual margin (minutes)
Blue 1	500	99	11:51	Yes	+2:30	Yes	+2:30
Blue 2	500	192	23:03	No	-8:42	No	-8:22
Blue 3	500	141	16:57	No	-2:36	Yes	+0:25

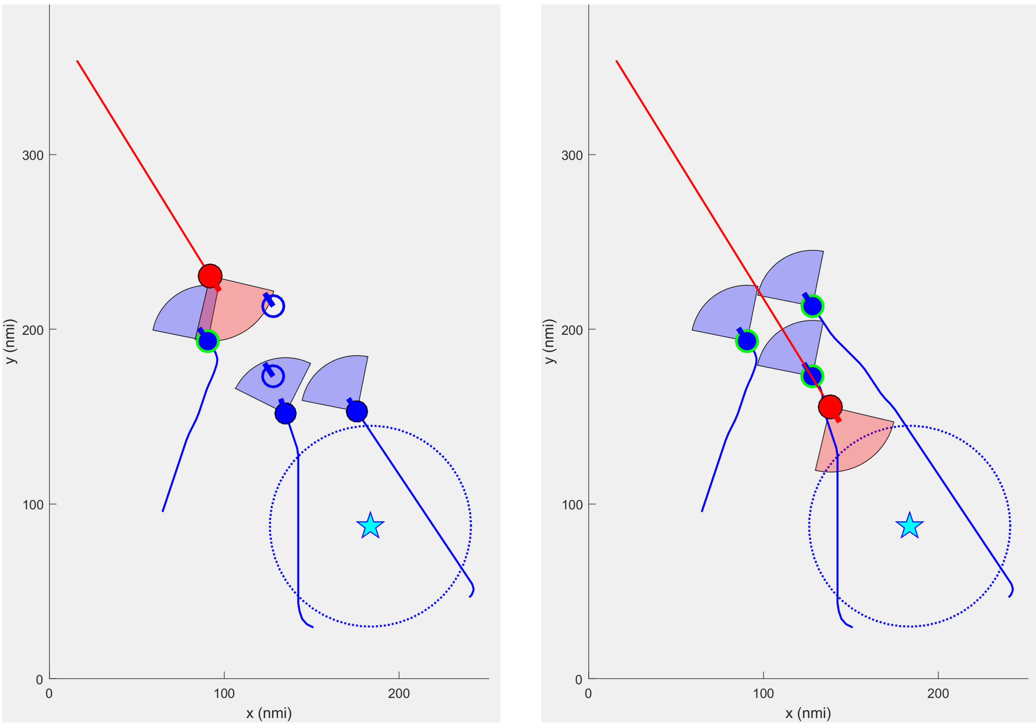

(a) (b)

Figure 4.8: Locations of the fighters and routes they have flown in the scenario with the unfavorable outcome a) 14 minutes and 21 seconds after the initial situation, when the end condition 1 is met, and b) 23 minutes and 3 seconds after the initial situation, when last Blue fighters have arrived to their destinations too late, and the Red fighter has already passed all the destinations of the Blue fighters. Green edges around the fighters indicate that they have reached their destinations.

4.3 Large scale scenario

The third scenario contains two threats posed by Red forces which are here assumed to be two flights of Red forces. The MFC has four Blue flights, i.e., 16 fighters, to perform a DCA mission on the two Red flights. Now, instead of individual fighters, flights are investigated. The initial situation of the scenario is presented in Figure 4.9. The initial locations for the Blue fighters are given as formations of the flights at their initial locations. In flights Blue F1 and Blue F2, the initial locations are determined by the “box” formation. Similarly, flights Blue F3 and Blue F4 are given the ”diamond” formation at their initial locations. The positions of Red flights are described by the

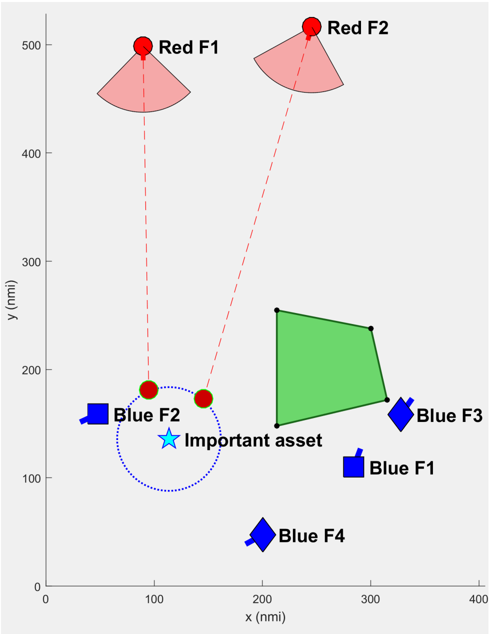

Figure 4.9: Initial situation of the large scale scenario describing locations and directions of the Blue flights and the Red flights, the routes of the Red flights, the important asset, and a ROZ illustrated as a green polygon. The shapes describing the Blue flights are determined by their initial formations.

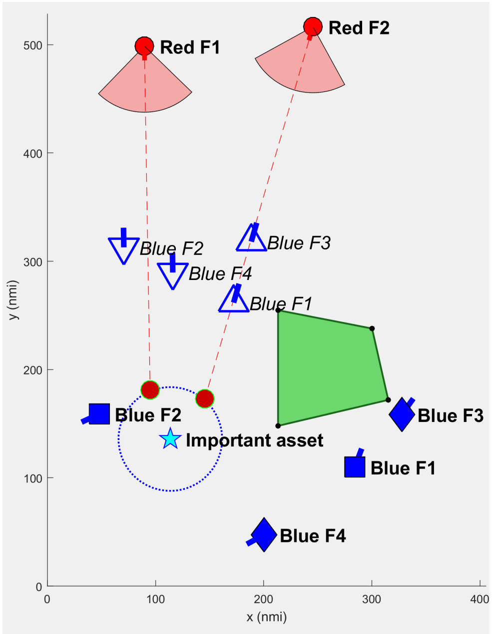

Figure 4.10: Initial situation of the large scale scenario with the destinations of the Blue flights given by the MFC. The destinations are illustrated with unfilled blue shapes for each flight, and the direction at the destination is also visualized. The shapes describing the destinations of the flights are determined by the desired formation of the flights.

first fighters in the flight. Flying altitudes, flying speeds as well as commit ranges and aspects for the flights are given in Table 4.4. These parameters are inherited to the individual fighters in the flights. Again, it is assumed that the flying altitudes of all the flights remain constant, and that the Red flights are slightly faster than the Blue flights. Furthermore, all the Blue fighters are flying at their maximum speed and can reach their destinations at different times. The Red fighters are not allowed to reach the important asset within a radius of 55 nautical miles. Furthermore, a ROZ is defined, which constraints the possible routes of the Blue flights.

The MFC has to assign the four Blue flights to suitable locations to perform a DCA mission. The destinations of the Blue flights corresponding to the hypothetical commit locations given by the MFC are illustrated in

Table 4.4: Parameters of the Blue and Red flights in the large scale scenario. The parameters are inherited for the individual fighters in the flights.

Unit	Altitude (feet)	Flying speed (knots)	Commit range (nmi)	Commit aspect (degrees)
Blue F1	32 000	600	30	45
Blue F2	32 000	600	30	45
Blue F3	32 000	600	30	45
Blue F4	32 000	600	30	45
Red F1	32 000	690	55	45
Red F2	32 000	690	55	45

Figure 4.10. Now, the destinations of the individual fighters are based on the destinations and formations given for the flights. Flights Blue F1 and Blue F3 are given the "trail" formation at their destinations, while flights Blue F2 and Blue F4 are given the "wall" formation. The individual initial locations and the destinations of each Blue fighter are presented in Figure 4.11.

A network with 2.7 nmi horizontal spacing between the nodes and 45 degrees as the turning constraint are again applied to compute the optimal

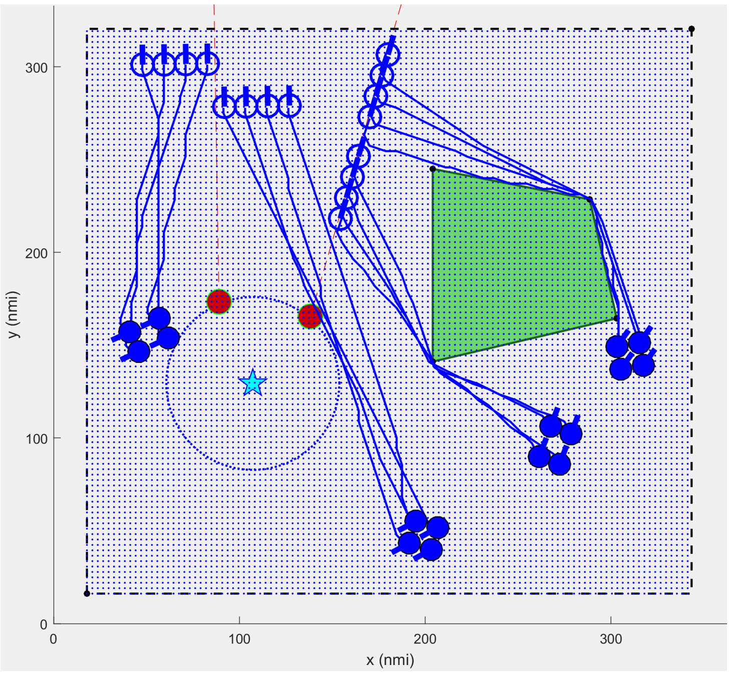

Figure 4.11: Optimization network and the optimal routes computed for the Blue fighters in the large scale scenario.

routes of the Blue flights. The optimal routes are determined individually for each fighter in the flights. The network and the optimal routes are visualized in Figure 4.11.

Visualizations provided by the simulation at two moments of time are presented in Figure 4.12. The state of the battlefield after 9 minutes and 50 seconds since the initial situation is depicted in Figure 4.12a. At that time, the commit envelope of a Red fighter reaches the destination of a Blue fighter before all the Blue fighters have arrived to their destinations. The end condition 1 is met and the outcome of the scenario is unfavorable. Thus, the decision made by the MFC is unsuccessful. Figure 4.12b visualizes the situation 23 minutes and 51 seconds after the initial situation. At that moment, all the Blue fighters have arrived to their destinations. Figure 4.12b shows that one of the Red flights has reached the terminal point of its route and the another one is also almost at its destination. The commit locations given by the MFC for the Blue flights are clearly too far to be reached in time, and all the Blue flights are significantly too late.

Numerical data is presented in Table 4.5. Fighters in the flight Blue F2 have the shortest routes to fly, between 134 and 149 nautical miles, while Blue F3 and Blue F4 have the longest routes, over 200 nautical miles for each fighter. None of the flights can make their destinations before an end condition for the unfavorable outcome realizes. Consequently, none of the fighters are collectively feasible and all have negative collective margins. The flight Blue F2 has the least negative collective margins, between three and five minutes. Due to the longest routes, the flights Blue F3 and Blue F4 have the most negative collective margins, over ten minutes. When considering the individual margins, none of the flights are entirely individually feasible. Three of the fighters in flight Blue F1 could just make their destinations before a Red fighter, while the fourth one is late. Thus, none of the flights could make their destinations without triggering an end condition for the unfavorable outcome themselves.

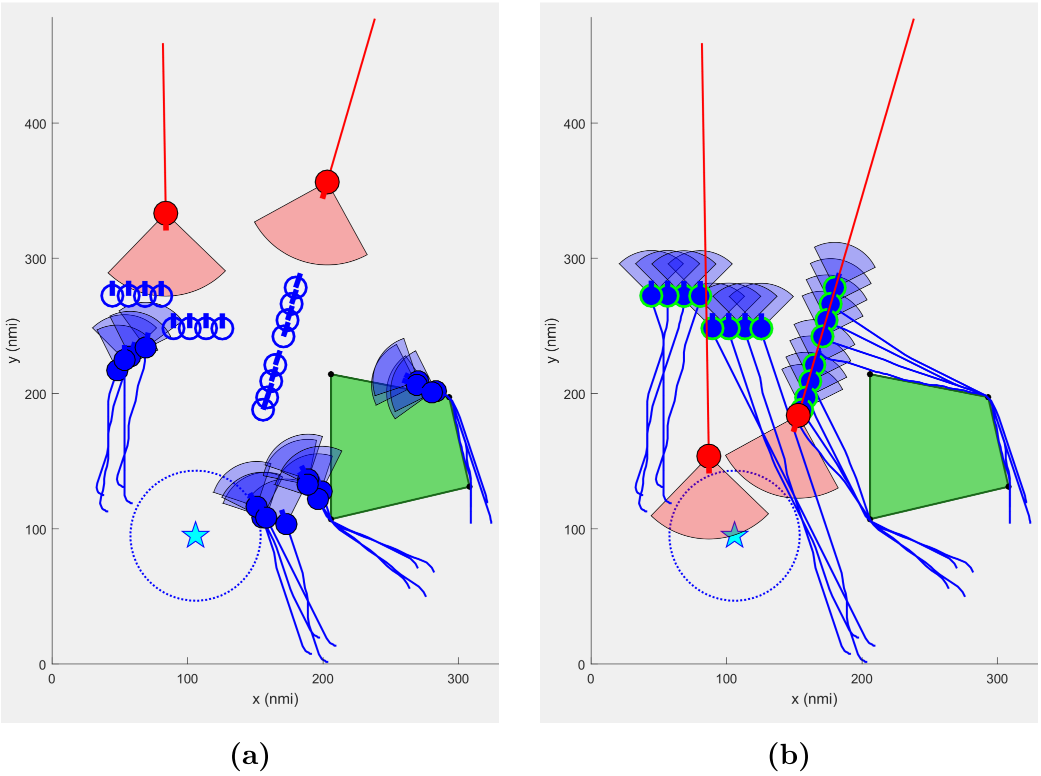

Figure 4.12: Locations of the fighters and routes they have flown in the large scale scenario a) after 9 minutes and 50 seconds since the initial situation, when the end condition 1 is met, and b) after 23 minutes and 51 seconds since the initial situation, when the last blue fighters have arrived to their destinations too late. The Red fighters have passed all the destinations of the Blue fighters and have almost reached the important asset. Green edges around the fighters illustrate that they have reached their destinations.

Table 4.5: Numerical results of the simulation in the large scale scenario.

Unit	Flying speed (knots)	Distance of the route (nmi)	Flying time (minutes)	Feasible? (individual)	Individual margin (minutes)	Feasible? (collective)	Collective margin (minutes)
Blue 1 (F1)	600	189	18:58	No	-9:08	No	-2:56
Blue 2 (F1)	600	170	16:59	No	-7:09	Yes	+0:01
Blue 3 (F1)	600	176	17:37	No	-7:47	Yes	+0:21
Blue 4 (F1)	600	159	15:57	No	-6:07	Yes	+2:45
Blue 5 (F2)	600	140	14:00	No	-4:10	No	-4:01
Blue 6 (F2)	600	149	14:54	No	-5:04	No	-4:34
Blue 7 (F2)	600	134	13:25	No	-3:35	No	-3:35
Blue 8 (F2)	600	142	14:13	No	-4:23	No	-3:15
Blue 9 (F3)	600	207	20:44	No	-10:54	No	-9:19
Blue 10 (F3)	600	218	21:47	No	-11:57	No	-9:24
Blue 11 (F3)	600	203	20:17	No	-10:27	No	-6:56
Blue 12 (F3)	600	214	21:23	No	-11:33	No	-7:04
Blue 13 (F4)	600	229	22:58	No	-13:08	No	-10:41
Blue 14 (F4)	600	225	22:31	No	-12:41	No	-10:37
Blue 15 (F4)	600	235	23:30	No	-13:40	No	-10:33
Blue 16 (F4)	600	238	23:51	No	-14:01	No	-12:08

Chapter 5

Experimental study on a training intervention with the fighter allocation model

This chapter presents an experimental study performed to evaluate benefits of the FA model. The study was carried out with a training intervention conducted by the FA model during the training of MFCs. The purpose of the experiment was to analyze whether the learning of the MFCs can benefit from the use of the model. In the experiment, test scenarios were employed and the performance of participants were measured before and after the use of the model in their training.

5.1 Arrangement of the experiment

5.1.1 Participants

The participants of the experiment were current fighter controllers and flight leaders. The experiment was carried out during a formal qualification course of MFCs. In total ten participants took part in the experiment. Three of the participants were experienced fighter controllers who were in training to obtain the qualification of an MFC. The fighter controllers had an age between 35 and 39. Seven of the participants were flight leaders aged 30 or 31. The average flying experience of the flight leaders was 445 hours, while the minimum was 400 hours and maximum 500 hours. All participants were male. A written informed consent was collected from all the participants prior to the experiment.

5.1.2 Training intervention

A training intervention was performed during morning of a course day as an approximately two-hour session where the FA model was utilized. The model was introduced and experimented during the session to demonstrate real-life scenarios as well as the most common biases and shortcomings when perceiving decision making situations of an MFC. The participants also had the possibility to interactively take part in the experimentation and contribute to the scenarios which were evaluated during the session. The session was led by an expert who had previous experience of the model. Due to the time limits of the course, participants had no possibility to individually use the model.

5.1.3 Test scenarios

The participants' abilities to perceive and understand the current state of the battlefield were measured with test scenarios. In a test scenario, the participants were given a situation on a map. Each scenario consisted of three Blue fighters and one Red fighter on the battlefield. For the Blue fighters, their initial locations and destinations, i.e., the commit locations, were shown. For the Red fighters, the complete routes from the initial locations to the destinations were presented. However, the participants were not shown if the Blue flights actually could reach their destinations in desired time, i.e., before any of the end conditions leading to unfavorable outcomes of the scenarios would occur.

Between the test scenarios, the locations and destinations of the fighters were varied. Flying speeds and commit envelopes of both Blue and Red fighters were same in all the test scenarios and reflected values in real world. Furthermore, the altitudes of the initial locations and the destinations of the Blue and Red fighters were identical in all the test scenarios.

In each scenario, the participants had to assess if the three Blue fighters could reach their destinations before the Red fighter triggers any of the end conditions leading to the unfavorable outcome. If a participant assessed that a Blue fighter cannot reach its destination before an end condition for the unfavorable outcome takes place, he was additionally required to estimate where the Blue fighter would have been located when the end condition would occur. As the scenarios were given on a map, only the location on the horizontal plane was to be assessed and the assessment of the flying altitudes was excluded.

The scenarios were analyzed beforehand with the model and the real feasibility of each route according to the model, i.e., can or cannot reach in time, was determined. The real outcomes of the scenarios were not presented for the participants. The routes were classified as ”in time” or ”late” according to their outcomes, as the two types of routes were analyzed separately later.

40 test scenarios were prepared and constructed by an expert and used in the experiment. Consequently, each participant had to assess a total of 120 different routes. The difficulties of the scenarios were varied, and some of the scenarios were more difficult to assess than others. As an example, two of the test scenarios are presented in Figure 5.1. In order to distinguish the individual Blue fighters and their destinations, one of the Blue fighters was illustrated with blue color, while for the two others, green and black colors were used. However, these three fighters to be assessed are still later referred to as the Blue fighters, despite them being depicted with different colors in visualizations.

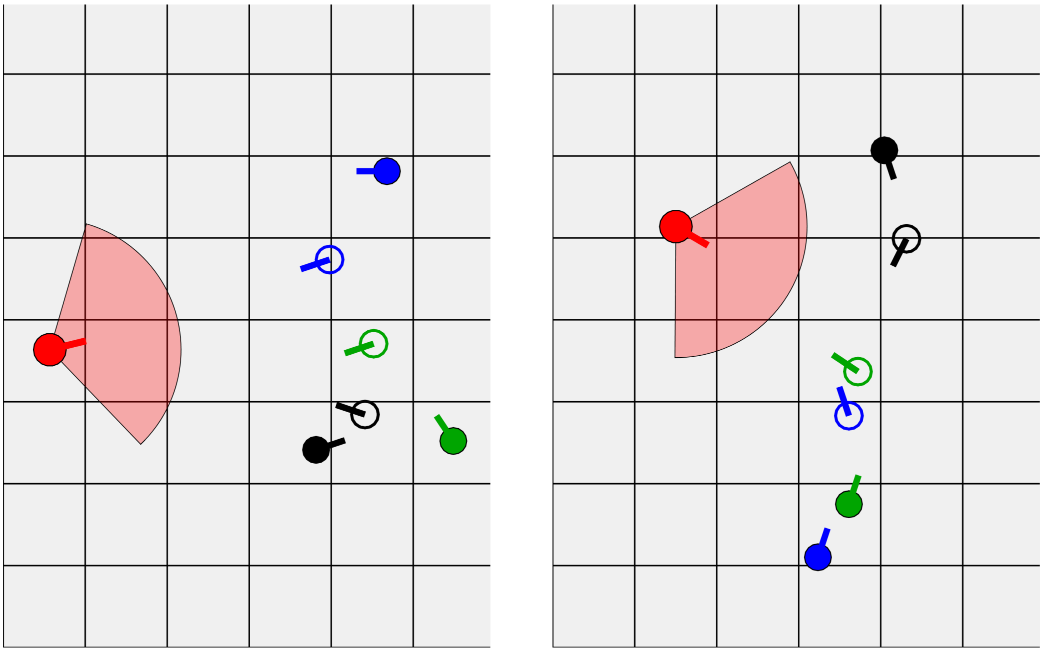

(a) **(b)**

Figure 5.1: Two of the 40 test scenarios assessed by the participants in the experiment. The scenarios are illustrated as given to the participants on a paper, except a real map was presented on the background. The squares of the grid are sized 50 nmi x 50 nmi. The three Blue fighters are colored here as blue, green, and black to distinguish them. In Figure a), a scenario where the correct assessment would have been that all the three fighters can arrive to their destinations in time. In Figure b), a scenario where only the black fighter can make its destination in time.

5.1.4 Test event

The 40 test scenarios were divided into two sets of 20 test scenarios. One of the sets was conducted before the training intervention utilizing the FA model, and the other set was carried out after the intervention. These sets were respectively labeled as "before" and "after". The sets were designed such that they were as homogeneous as possible, i.e., both sets contained the same amount of scenarios with similar difficulties. Additionally, the scenarios were arranged to the sets such that routes they included had a balanced amount of routes labeled as "in time" and "late".

The participants first assessed the first 20 scenarios, i.e., the "before" set of the scenarios. This was followed by the training intervention described in Section 5.1.2. After the intervention, the participants assessed the last 20 scenarios, i.e., the "after" set of the scenarios.

The scenarios were printed on a paper, where the initial locations, destinations and flying speeds of the Blue and Red fighters were presented on a map. Additionally, the complete route of the Red fighter was shown. A grid with squares sized 50 x 50 nautical miles was applied on the map to help perceiving the distances. Each scenario was printed on a single paper sheet, and the participants were assigned to mark their assessments on the papers. The appearance of the papers was as depicted in Figure 5.1. However, a real map was presented on the background of the scenario and the fighters were tied to real-life locations, as opposed to Figure 5.1. The participants were briefed how to fill the papers before the test was conducted during the course. All the scenarios were presented in the same order for each participant. The papers were collected after each set was finished. The assessments were entered to the implementation of the model in order to numerically analyze them.

5.1.5 Assessment errors

When the assessments were done, errors in them were computed. Two different metrics for the assessment errors were defined. The purposes of the error metrics were to quantify and measure how good an assessment was and enable numerical analysis and comparison of the assessments. The errors were measured in kilometers.

In both error metrics, the actual location of a Blue fighter at the end of the scenario, i.e., when an condition occurs, and the assessed location of the Blue fighter given by a participant are compared. The actual location of the Blue fighter at the time when an end condition is met is referred to as a ground truth. In scenarios with the favorable outcome, the ground truth for each

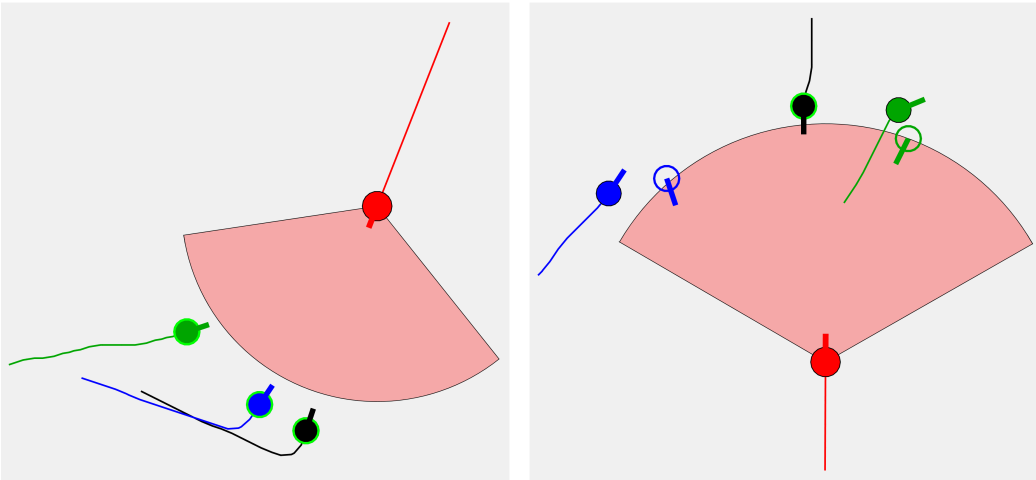

(a) (b)

Figure 5.2: Blue fighters at their ground truth locations. When the outcome of the scenario is favorable (a), the ground truths of all the Blue fighters are at their destinations. When the outcome of the scenario is unfavorable (b), the ground truth of a Blue fighter is its location at the time when an end condition leading to the unfavorable outcome is triggered. In Figure b), for the fighter visualized in black, the ground truth is its destination, while for the blue and green fighters, it is somewhere along their routes.

Blue fighter is its destination. In scenarios with the unfavorable outcome, the ground truths are defined to be the locations of the Blue fighters at their optimal routes when any of the end conditions for the unfavorable outcome occurs. The definitions of the ground truths for both cases are illustrated in Figure 5.2.

The first metric for assessment errors was defined as a straight-line distance, i.e., Euclidean distance between the assessed location of a Blue fighter and its ground truth. The errors computed using this metric are later referred to as straight-line errors. The computation of the straight-line error is illustrated in Figure 5.3 for both ”in time” and ”late” routes. In Figures 5.3a and 5.3b, the error is visualized for ”in time” routes. When a participant correctly assessed that a Blue fighter can reach its destination before an end condition for the unfavorable outcome occurs, the error was zero (Figure 5.3a). If a participant incorrectly assessed that the Blue fighter cannot reach its destination in time, the error was the straight-line distance between the assessment and the ground truth, which was the destination of the fighter (Figure 5.3b). Figures 5.3c and 5.3d depict the error when the correct assessment was that the Blue fighter will not reach its destination

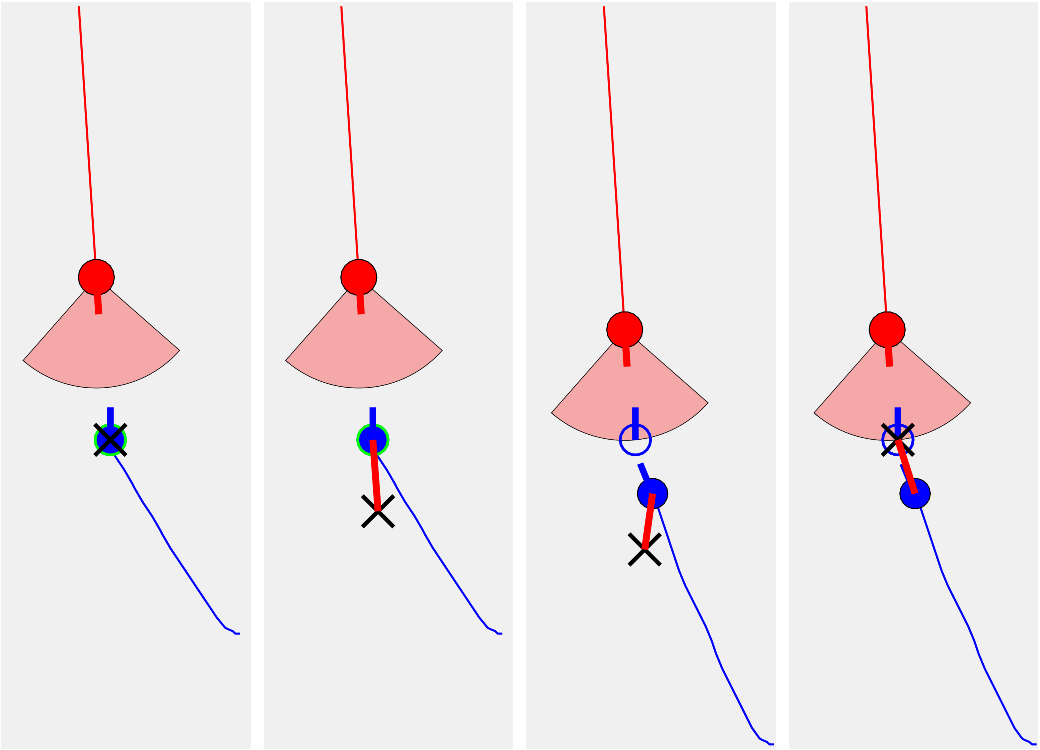

(a) "In time", correct **(b)** "In time", incorrect **(c)** "Late", correct **(d)** "Late", incorrect

Figure 5.3: Definitions of the assessment errors based on the straight-line error for both "in time" and "late" routes. The assessment is visualized as a black cross, and the error is visualized as a red bar between the assessment and the Blue fighter at its ground truth location. With a correct assessment for "in time" routes, the error is zero (a). With an incorrect assessment for "in time" routes (b), and both correct (c) and incorrect (d) assessments for "late" routes, the error is nonzero.

in time. When a participant correctly assessed that the Blue fighter cannot make its destination in time and assessed a location where the Blue fighter is when an end condition occurs, the error was the distance between the assessed location and the ground truth (Figure 5.3c). If a participant incorrectly assessed that the Blue fighter would be in time to its destination, the error is the distance between the assessment, i.e., the destination, and the ground truth (Figure 5.3d).

The second metric for assessment errors was determined to be relative to the actual route of the Blue fighter. It was defined to be sum of two distances. First, a straight-line distance between the assessed location of

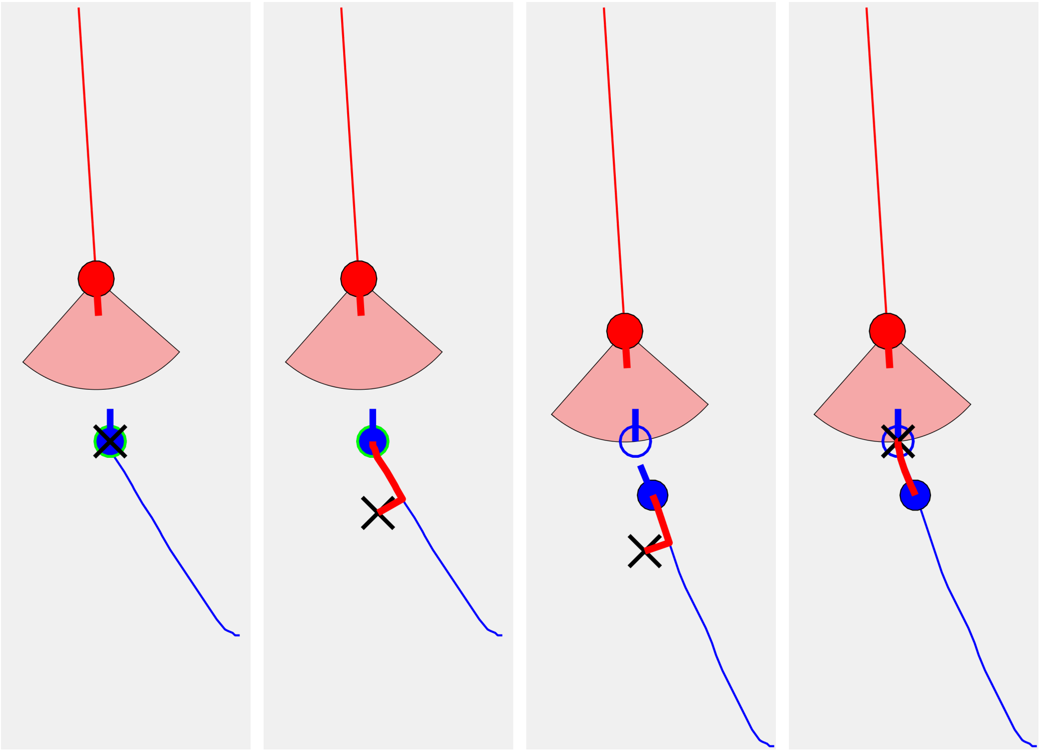

(a) "In time", correct (b) "In time", incorrect (c) "Late", correct (d) "Late", incorrect

Figure 5.4: Definitions of the assessment errors based on the route-relative error for both "in time" and "late" routes. The assessment is visualized as a black cross, and the error is visualized as a red bar between the assessment and the Blue fighter. With a correct assessment for "in time" routes, the error is zero (a). With an incorrect assessment for "in time" routes (b), and both correct (c) and incorrect (d) assessments for "late" routes, the error is non-zero.

the Blue fighter and the nearest point of the route of the Blue fighter was computed. Then, the length of the route between the nearest point of the route and the ground truth was calculated. The value of the error was then the sum of these distances. As this error metric was relative to the actual route of the Blue fighters, the error is later referred to as a route-relative error. The definitions of the route-relative error are visualized in Figure 5.4. The errors in "in time" routes are illustrated in Figures 5.4a and 5.4b. As in the straight-line error, if a participant correctly assessed that a Blue fighter arrives in time, the error was zero (Figure 5.4a). When an incorrect assessment was made for an "in time" route, the error was the distance between the assessment and the nearest point of the route added to the length of the rest of the route (Figure 5.4b). The errors for "late" routes are presented in Figures 5.4c and 5.4d. The error for a correct assessment in a "late" route, i.e., that the Blue fighter cannot reach its destination in time, was computed between the assessment and the ground truth (Figure 5.4c) as described earlier. If a participant incorrectly assessed that the Blue fighter arrives to its destination in time while it really does not, the error was then the length of the remaining route (Figure 5.4d).

While the straight-line error was a simple and mostly effective metric to determine the deviation between the assessments and the ground truth, it could fall short if the shape of a route was curved. When a heavily curved route was present, the location of a Blue fighter could have been nearer to the ground truth earlier on its route and then temporarily moved away from it before again approaching it. Consequently, it was possible that a location that was further away from the ground truth was actually more correct assessment, while the value of the straight-line error was similar or even smaller. This weakness of the straight-line error was addressed when using the route-relative error. This is demonstrated in Figure 5.5, where Assessment 1 was clearly better than Assessment 2, as Assessment 2 overestimated the location of the fighter more. However, the error was relatively similar for both assessments with the straight-line error, as illustrated in Figure 5.5a. When using the route-relative error, the error was significantly larger for Assessment 2, which described the actual error in the perception better. This is pointed out in Figure 5.5b.

The route-relative error had also shortcomings in certain situations. It did not necessarily reflect the differences in errors reasonably in, e.g., situations where the ground truth was located in curves of the route. A problem rose especially when an assessment was made perpendicular to the direction of the movement at the ground truth. In that case, the route-relative error essentially reduced to the straight-line error. Thus, assessing the curve of the route significantly wrong was not necessarily reflected in the value of the route-related error. An example is presented in Figure 5.6. Assessment 2 clearly did not address the turn of the Blue fighter at all. While Assessment 1 underestimated the progress of the Blue fighter, it had taken into account the turn better. However, the value of the route-relative error was almost twice as large for Assessment 1, while it was questionable which one of the assessments was actually better. The values of the straight-line errors were almost equal, as pointed out in Figure 5.6b, which reflected the actual errors in the spatial perception more accurately. Additionally, if an assessment was made in the center of a 180 degree turn, it would have not been unambiguous to determine which was the nearest point of the route. Thus, two assessments

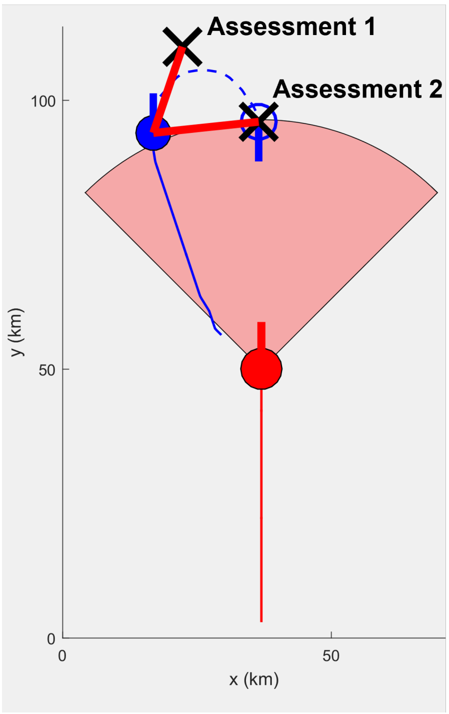

(a)

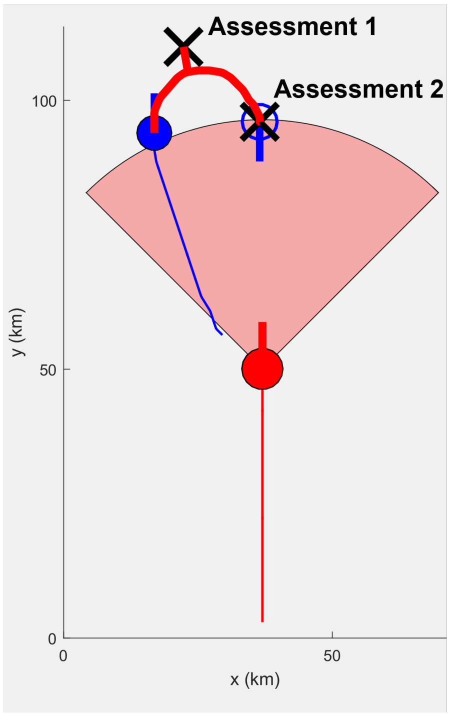

(b)

Figure 5.5: Example route where the drawback of the straight-line assessment error is demonstrated. The location of the Blue fighter is the ground truth, and the rest of the route is visualized as a blue dotted line. Assessment 1 is better than Assessment 2, as it is closer to the ground truth with respect to the route. However, with the straight-line error (a), the errors are almost equal (15 km and 17 km). With the route-relative error (b), the error is 17 km for Assessment 1 and 29 km for Assessment 2.

in the center of a 180 degree turn but marginally on the different sides of it would have had significantly different magnitudes of route-related errors, as the nearest points of the route would have been on the different sides of the turn.

Straight-line and route-relative errors were both sufficient metrics for measuring errors of the assessments, but both also had their shortcomings. Neither metric was able to capture the characteristics of the errors in the actual perception of the routes unambiguously better. Therefore, both error metrics were utilized in further analysis of the experiment.

Assessment errors were computed separately using both error metrics for each route in all the test scenarios and individually for each participant. Thus, for each route, ten error values were acquired with each error metric,

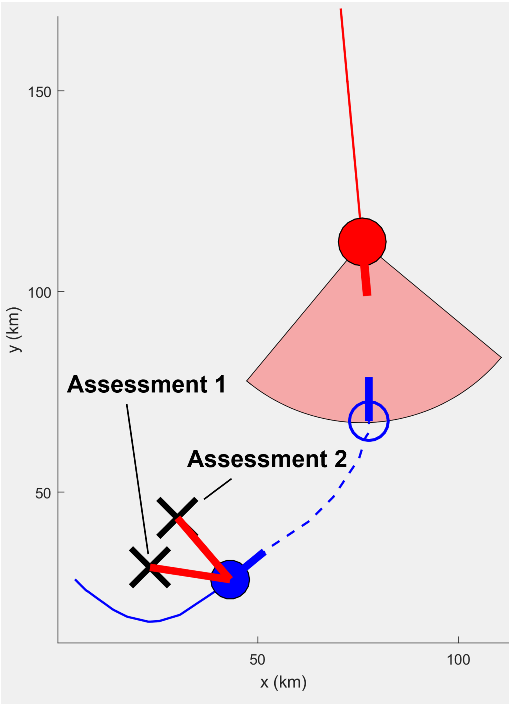

(a)

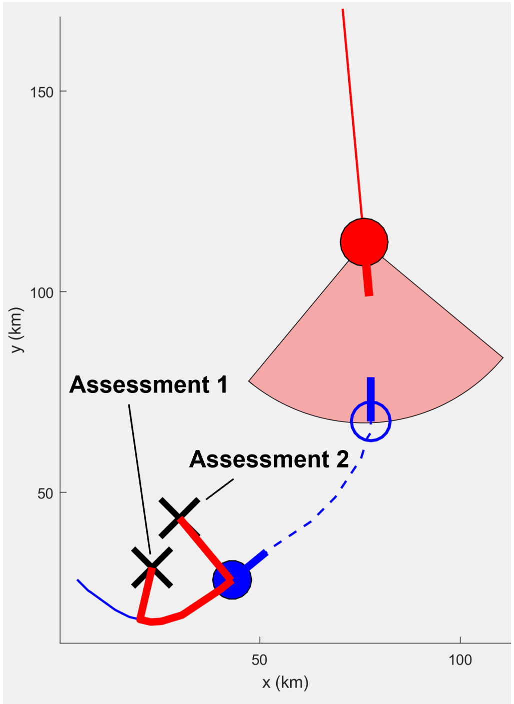

(b)

Figure 5.6: Example route where the shortcoming of the route-relative error with respect to the straight-line error is demonstrated. The location of the Blue fighter is the ground truth, and the rest of the route is illustrated as a blue dotted line. It is not unambiguous which assessment is better, as Assessment 1 underestimates the progress of the Blue fighter, while Assessment 2 fails to predict the curve correctly. Values of the straight-line error are equal for both assessments, i.e., 18 km. However, the route-relative error is 35 km for Assessment 1, while only 18 km for Assessment 2.

one for each participant. The computation of the errors was done automatically with the implementation of the model.

The error metrics allowed studying the differences between the participants' assessments numerically. Example illustrations of assessments made by the participants in two test scenarios are presented in Figure 5.7 after the assessments had been imported to the implementation of the model. The assessments for each fighter are presented as crosses with corresponding colors, and a cross represents an assessment of one participant. The scenarios are presented at the time when an end condition is triggered, whereas the initial situations are shown in Figure 5.1. The locations of the fighters in the Figure 5.7 thus correspond to the ground truths of the fighters.

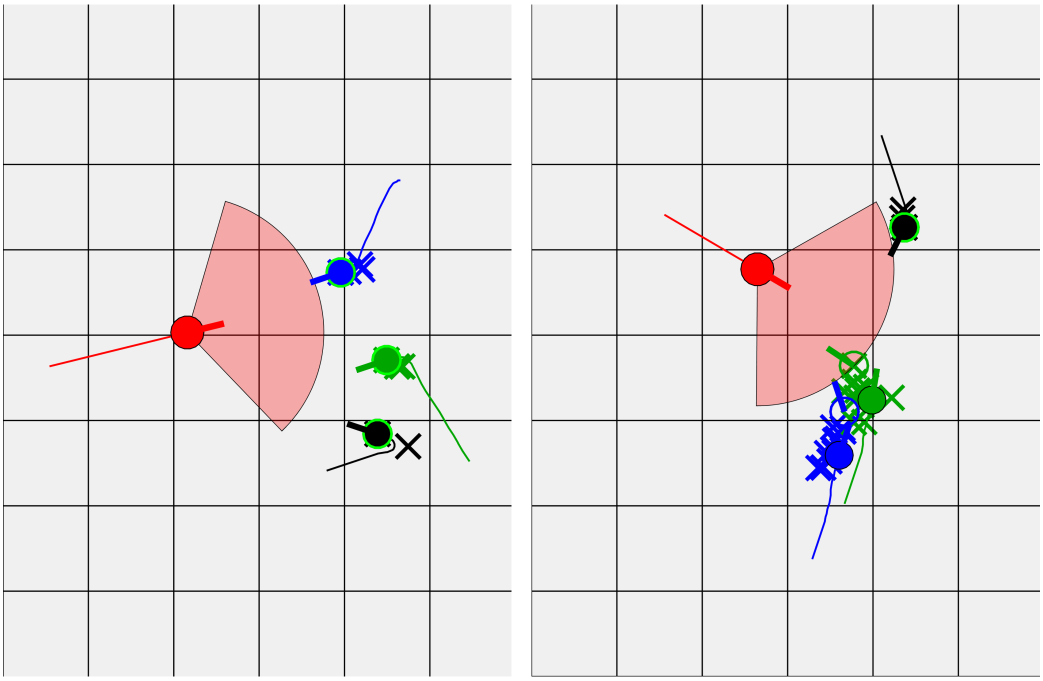

(a) (b)

Figure 5.7: Assessments of the participants in two test scenarios. The fighters are presented as filled circles and their routes as lines adjacent to the circles. The assessments are illustrated as crosses with a color corresponding to the assessed fighter. The scenarios are visualized at the time when an end condition is triggered, i.e., revealing the ground truths of the locations of the fighters which the participants tried to assess. The grid, 50 nmi x 50 nmi, is presented identically as on the papers to which the participants gave the assessments.

In the test scenario in Figure 5.7a, all the three Blue fighters could make their destinations in time and their ground truths were equal to their destinations. There were also assessments that were not at the final destinations of the fighters. Thus, several participants had assessed that those fighters would not make their destinations before the commit envelope of the Red fighter would reach the destination of a Blue fighter. Consequently, they had given assessments for the locations of the Blue fighters that differed from their destinations. In the scenario in Figure 5.7b, the outcome was unfavorable and not all the fighters made their destinations in time. Only the black fighter was able to arrive to its destination before the commit envelope of the Red fighter reached the destination of the green fighter. Most of the participants had correctly assessed that the green and blue fighters would not be at their destinations when the end condition was triggered, while there was clear dispersion in the assessments. The location of the black fighter was also assessed fairly well, and there were only two participants who had not assessed the black fighter to be at its destination.

5.2 Results

The data obtained from the test event was analyzed in order to study the effects of the training intervention. The 40 test scenarios contained 120 routes of Blue fighters assessed by the ten participants. Thus, 1200 assessments in total were obtained. The differences in the assessment errors between the 60 routes assessed before the intervention and the 60 routes after the intervention were studied. The before-after comparison was carried out by first finding pairs of routes from the "before" and "after" sets of the routes that had same characteristics and were thus equally difficult to assess. Then, for each pair of routes, Wilcoxon signed-rank test (see, e.g., Conover, 1999) for paired data was performed to test if the assessment errors of the ten participants had a statistically significant difference before and after the intervention. Since the samples are dependent as they are assessments of same persons before and after, Wilcoxon signed-rank test was selected. This non-parametric test was employed instead of paired t-test as the number of samples is low ($n = 10$) and the normality of the data could not be reliably verified. As the assessment errors were computed with two different error metrics, the test was performed separately using both assessment errors. First, the test was conducted using the straight-line errors and then using the route-relative errors. The tests were performed also with error values relative to the length of the route, i.e., the errors were scaled with the total length of the corresponding route. However, the results were not different when these relative errors were studied. Thus, only the results with the absolute errors are presented.

The study was conducted using single routes instead of test scenarios consisting of three routes in order to enable the numerical analysis of the assessment errors. The test scenarios as a whole were complex and consisted of multiple variables to consider when making the assessments. As there were no identical scenarios before and after the training intervention, finding justifiable pairs between the scenarios would had been difficult. Identical scenarios before and after the intervention were not used because it was intended to show the correct answers of the "before" scenarios during the intervention. However, due to practical reasons, this was not eventually possible. Thus, studying and comparing the test scenarios would have not been reasonable, and instead, single routes were analyzed.

5.2.1 Pairing the routes

To perform the pairwise statistical analysis for route pairs, each single route assessed before the intervention was associated with a route assessed after the intervention such that each route had an unique pair. First, the pairing was restricted based on if the route was labeled as "in time" or "late" according to their correct outcomes. This is due to the differently behaving values of the errors with these two route types (see Section 5.1.5). For instance, correct answers in "in time" routes would always yield an error of 0, while in "late" routes a correct answer would require an exactly correct assessment for the ground truth within a magnitude of one meter. Given the practical implementation of the experiment and the difficulty of the assessment procedure, such a precise assessment was not expected and probable. Thus, for the errors to be reasonably comparable, "in time" routes were only paired with "in time" routes, and "late" routes were only paired with "late" routes.

The pairing of the routes was based on following criteria which characterized each route:

1. The length of the route.
2. The distance between the Red fighter's initial location and the nearest destination of any Blue fighter in the same scenario including the destination of the fighter whose route is studied.

An illustration of the criteria is given in Figure 5.8. Both criteria are measured in kilometers.

These criteria were judged most relevant for describing the relative difficulty of the assessment of a route. While the assessment was conducted similarly for all lengths of routes, the possibility for larger errors grew if routes were longer. Additionally, the distance between the Red fighter's initial location and the nearest destination of any Blue fighter in the same scenario determined heavily when an end condition for the scenario was triggered. With longer distances, the Blue fighters could fly further which made the assessment more difficult, while if the Red fighter was near to a destination of a Blue fighter, an end condition was triggered significantly earlier and the assessment was easier.

The criteria were computed for each of the 120 routes. As the values of the two criteria had slightly different magnitudes, they were normalized to a range from 0 to 1 in order to give equal importance for both criteria. It was experimented that the normalization had a negligible effect on the results of the pairing, but it was judged to be reasonable to be performed. The pairwise differences of the routes based on the normalized values of the

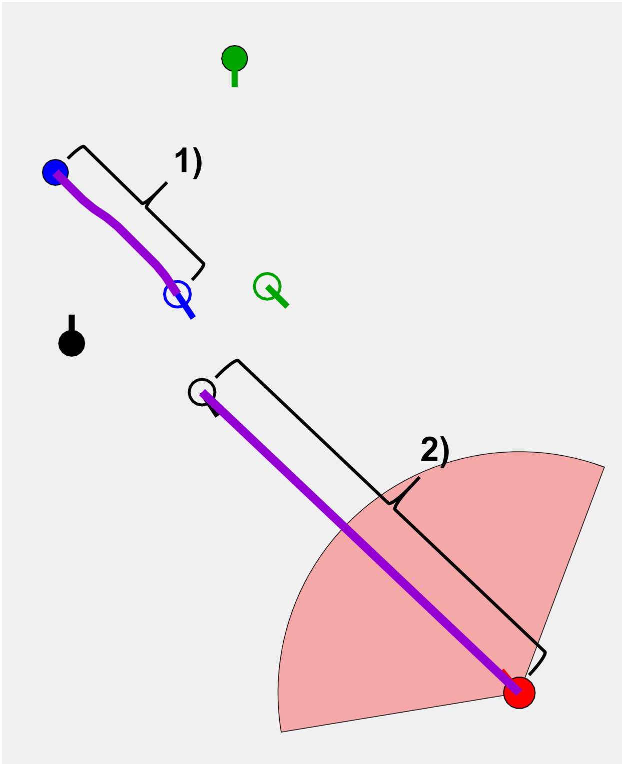

Figure 5.8: Two criteria used for pairing the routes exemplified for the fighter colored as blue. The criteria are visualized as purple lines. The first criterion (1) is the length of the route of the fighter. The second criterion (2) is the distance between the Red fighter's initial location and the nearest destination of any Blue fighter in the scenario.

criteria were computed. This pairwise difference measure was used for the pairing, both with the "in time" and "late" routes.

Based on the pairwise differences, the routes labeled "in time" were first paired. As some of the routes assessed before the intervention had already an error of 0 for all the participants, i.e., all the participants had assessed them correctly, they were removed from the study. This was done since the goal of the study was to provide insight if the use of the FA model could improve the decision making of the participants. When there were no errors in any of the assessments for a route already before the intervention, no improvements could have been discovered. After these routes were removed, the differences between all "before" and "after" routes were compared, and the best possible pairs were selected. There were no sufficient pairs for all the routes. Eventually, 17 pairs of "in time" routes were found for the statistical analysis.

The routes labeled ”late” were paired using the same difference measure. In order to keep the analysis consistent with the ”in time” routes, routes with the smallest errors in ”before” assessments were ignored. Again, when the error was already low, which corresponds to the error of 0 in the case of ”in time” routes, before the intervention, no improvement could have been discovered. Again, 17 pairs of ”late” routes were found for the statistical analysis.

5.2.2 ”In time” routes

Wilcoxon signed-rank test for paired data with the significance level of $\alpha = 0.05$ was performed for each of the 17 ”in time” route pairs. As there were two different metrics for the assessment errors, the test was performed first using the straight-line errors as the paired data. The results are given in Table 5.1. Then, the test was conducted using the route-relative errors as the paired data. These results are presented in Table 5.2.

When using the straight-line error, the means improved in almost all the route pairs after the intervention. However, from the medians, no clear inferences could be drawn, as the medians were mostly zero already before and also after the intervention. No statistically significant differences were found in any of the pairs.

With the route-relative error, the results were essentially similar. The means mainly improved after the intervention, while the medians were mostly zero even before the intervention and remained zero after it. However, in one route pair, a statistically significant difference was found. The difference was contrary to desired, as the median and the mean clearly increased after the intervention.

5.2.3 ”Late” routes

Wilcoxon signed-rank test was also performed for the ”late” routes. 17 route pairs were tested and the significance level of $\alpha = 0.05$ was applied. The test was conducted separately for straight-line and route-relative errors. The results using the straight-line errors are presented in Table 5.3 and using the route-relative errors in Table 5.4.

With the straight-line error the difference was found to be statistically significant in five route pairs. In all these route pairs, the mean and the median error were smaller after the training intervention. In majority of the pairs, mean and median errors reduced after the training intervention. In some of the pairs, the error also increased after the training intervention, but the difference was not statistically significant in these cases.

Table 5.1: Mean and median errors before and after the training intervention, values of the test statistic, i.e., signed-rank sum, and p-values for each "in time" route pair. The assessment errors were the straight-line errors.

Route pair #	Mean error before	Mean error after	Median error before	Median error after	Signed-rank sum	p-value
1	7.52	0.00	0.00	0.00	10	0.125
2	1.81	0.00	0.00	0.00	1	1.000
3	2.95	0.00	0.00	0.00	3	0.500
4	0.93	0.00	0.00	0.00	1	1.000
5	2.93	0.00	0.00	0.00	3	0.500
6	1.80	1.67	0.00	0.00	2	1.000
7	1.03	0.79	0.00	0.00	2	1.000
8	6.67	0.00	0.00	0.00	10	0.125
9	10.94	0.00	4.00	0.00	15	0.063
10	9.67	0.63	0.00	0.00	10	0.125
11	5.82	11.96	0.00	7.90	7	0.297
12	7.94	25.64	3.44	25.96	3	0.078
13	1.34	0.00	0.00	0.00	1	1.000
14	1.91	1.82	0.00	0.00	2	1.000
15	11.59	0.00	0.00	0.00	10	0.125
16	2.61	7.50	0.00	0.00	2	0.750
17	1.57	0.46	0.00	0.00	2	1.000

When the route-relative error was employed, a statistically significant difference was found in seven route pairs. As with the straight-line error, in all these pairs, the mean and the median errors reduced after the training. The differences of the five pairs that had a statistically significant difference using the straight-line error were all statistically significant with the route-relative error. Again, in majority of the pairs, the mean and the median errors were smaller after the intervention, while the difference was statistically significant in none of the opposite cases.

5.3 Discussion

The results of the experimental study show a promise of the usefulness of the FA model. With the "late" routes, statistically significant differences in the route pairs were found. In all these pairs, the mean and the median errors were smaller after the intervention. With the straight-line error, five pairs

Table 5.2: Mean and median errors before and after the training intervention, values of the test statistic, i.e., signed-rank sum, and p-values for each "in time" route pair. The assessment errors were the route-relative errors.

Route pair #	Mean error before	Mean error after	Median error before	Median error after	Signed-rank sum	p-value
1	9.14	0.00	0.00	0.00	10	0.125
2	2.27	0.00	0.00	0.00	1	1.000
3	3.40	0.00	0.00	0.00	3	0.500
4	1.06	0.00	0.00	0.00	1	1.000
5	3.20	0.00	0.00	0.00	3	0.500
6	2.28	2.03	0.00	0.00	2	1.000
7	1.18	1.13	0.00	0.00	2	1.000
8	7.79	0.00	0.00	0.00	10	0.125
9	16.14	0.00	4.60	0.00	15	0.063
10	12.03	0.71	0.00	0.00	10	0.125
11	7.36	15.36	0.00	10.04	8	0.375
12	10.20	47.09	4.35	54.36	1	**0.031**
13	1.90	0.00	0.00	0.00	1	1.000
14	2.46	2.32	0.00	0.00	2	1.000
15	15.09	0.00	0.00	0.00	10	0.125
16	3.07	10.86	0.00	0.00	2	0.750
17	1.76	0.54	0.00	0.00	2	1.000

with statistically significant difference out of 17 pairs were found, and seven pairs out of 17 with the route-relative error. This implies that the use of the model has improved the participants' perception and understanding of the state of the battlefield in certain cases as desired. However, the improvement was not statistically significant in all the route pairs. With the statistically insignificant pairs, there were also pairs where the mean and the median increased after the training intervention. Thus, there were also cases where the use of the model did not have the desired effect on the assessments of the participants or they could not be observed, and the benefits of the model to the learning process could not be revealed.

In the case of "in time" routes, no statistically significant differences in the errors before and after the intervention could be perceived using the straight-line error. With the route-relative error, only one statistically significant difference in a route pair was found out of 17 route pairs. Interestingly, the difference was opposite to the presumed, as the mean and the median errors

Table 5.3: Mean and median errors before and after the training intervention, values of the test statistic, i.e., signed-rank sum, and p-values for each "late" route pair. The assessment errors were the straight-line errors.

Route pair #	Mean error before	Mean error after	Median error before	Median error after	Signed-rank sum	p-value
1	28.25	14.41	22.30	10.60	54	**0.004**
2	11.79	14.87	7.95	9.34	21	0.557
3	13.85	14.06	9.58	10.15	32	0.695
4	23.14	13.52	16.34	8.80	52	**0.010**
5	18.96	12.47	20.68	6.64	42	0.160
6	12.95	11.63	12.86	5.41	42	0.160
7	26.83	26.80	19.95	26.83	26	0.922
8	17.31	6.04	17.09	1.31	47	**0.047**
9	17.75	22.68	18.22	23.36	18	0.375
10	19.63	16.89	19.90	13.14	35	0.492
11	25.95	13.93	24.53	13.18	51	**0.014**
12	14.20	18.77	9.63	17.07	17	0.322
13	17.93	12.46	7.21	11.61	29	0.922
14	16.71	12.34	9.47	9.91	34	0.557
15	19.97	14.43	18.60	15.18	43	0.131
16	24.54	20.64	22.27	18.89	40	0.232
17	33.99	16.16	32.23	15.03	55	**0.002**

were significantly larger after the training intervention. With a more detailed investigation to the routes, this was due to the route after the intervention in this pair being very problematic to assess. In the scenario containing the route, one of the destinations of the Blue fighters was located such that it was difficult to assess if the commit envelope of the Red fighter actually hit the destination or passed by nearby. Most of the participants estimated that the destination of the Blue fighter would have ended up under the commit envelope of the Red fighter, contrary to the true outcome of the scenario. Consequently, most of the participants then assessed that the end condition for the scenario occurred earlier than it really did, while the real outcome of the scenario was actually favorable. In addition, the route contained two curves, which complicated the assessment even more. Thus, the assessment errors for the route were very large. The scenario is depicted in Figure 5.9, where the studied route is illustrated in blue. As there were no other scenarios with such a difficult setting that would have led to similar errors, the after

Table 5.4: Mean and median errors before and after the training intervention, values of the test statistic, i.e., signed-rank sum, and p-values for each "late" route pair. The assessment errors were the route-relative errors.

Route pair #	Mean error before	Mean error after	Median error before	Median error after	Signed-rank sum	p-value
1	31.01	17.55	27.18	14.39	47	**0.049**
2	14.37	18.11	10.97	13.29	23	0.695
3	15.98	15.96	11.99	12.24	32	0.695
4	30.50	16.05	24.28	10.59	52	**0.010**
5	35.19	14.54	44.47	8.51	55	**0.002**
6	15.33	16.10	15.05	6.77	35	0.492
7	32.44	39.05	27.40	37.24	19	0.432
8	20.99	7.84	21.78	1.47	47	**0.047**
9	31.31	28.70	39.57	30.16	34	0.557
10	25.45	22.02	26.69	18.18	32	0.695
11	35.93	17.94	35.52	19.18	51	**0.014**
12	17.22	22.27	13.68	20.70	19	0.432
13	21.52	16.43	9.03	15.21	28	1.000
14	22.21	15.59	12.35	13.55	36	0.432
15	29.78	17.66	27.01	17.72	50	**0.020**
16	31.10	26.29	31.69	23.19	41	0.193
17	41.21	18.63	38.15	15.84	55	**0.002**

route of this scenario could be seen as an outlier.

There were no other statistically significant differences in the "in time" routes than the one previously discussed. Therefore, no real inferences can be drawn from the "in time" routes regarding the impact of the usage of the model. The statistical inference using Wilcoxon signed-rank test was complicated by the amount of zero values for the errors. As pointed out by the medians in Tables 5.1 and 5.2, many of the values of the errors were zero before or after the intervention. This was due to the definitions of the errors in "in time" routes (see Section 5.1.5, and Figures 5.3 and 5.4), i.e., that the errors were zero when a correct assessment was made. This led to a skewed distribution for the assessment errors, as a correct assessment always yielded zero whereas an incorrect assessment led to a nonzero error. The behavior was different than with "late" routes, as with them, some sort of nonzero error was always obtained even with relatively correct assessments. The skewed distribution with many zeros before and after the intervention led to more easily having assessments where there were no change, i.e., the values of the errors were zero before and after the intervention. Having a zero difference reduced the already small sample size even more.

When a participant assessed that a Blue fighter makes its route in time, auxiliary assessment would have been required to make reasonable differences between the "in time" assessments. For example, it could have been assessed that how far the Blue fighter would continue in its terminal direction until an end condition for the scenario is reached. Then, the distribution of the "in time" assessments would have been more even and, most importantly, not being prone to be heavily skewed to zero. Consequently, it could have been possible to make more inferences also using the "in time" assessments.

To conclude, there were encouraging results when studying the "late" routes. However, in the case of "in time" routes, no findings could be made. Due to the differences in evaluating the errors in these two route types, the "late" routes were more useful in studying the results of the training intervention. Furthermore, the sample size of the study ($n = 10$) was small, which also significantly accounts for not perceiving more differences and having more statistically significant results. Overall, it can be stated that even with a short training intervention, the FA model affected positively on the perception abilities of the participants in the test scenarios. Thus, the FA model has promise to support the training of MFCs and improve their decision making.

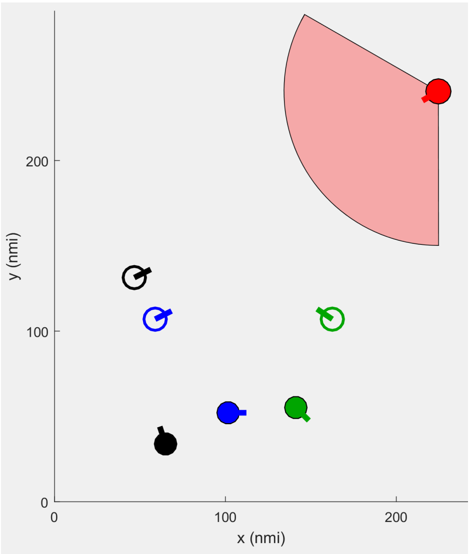

(a)

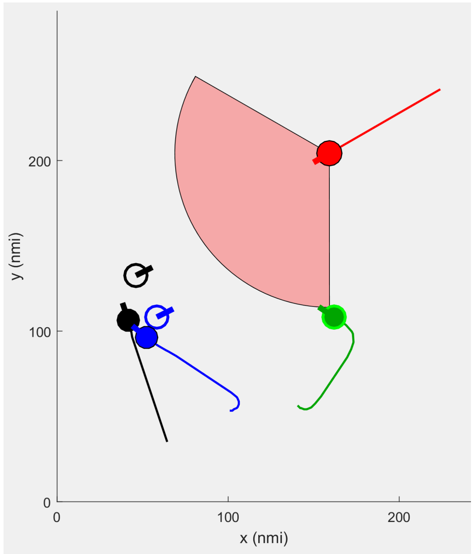

(b)

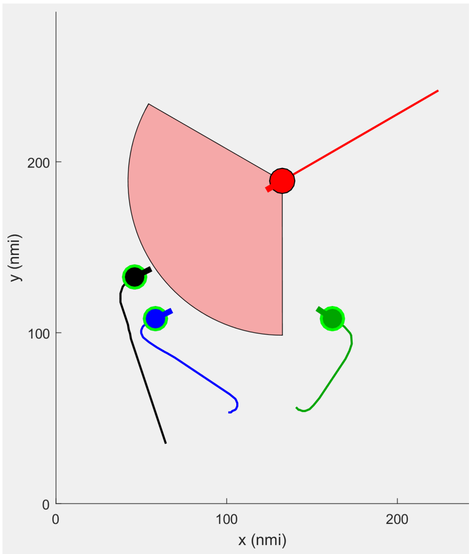

(c)

Figure 5.9: Test scenario that contained a route that caused a statistically significant difference in "in time" pairs. As it was difficult assess if the commit envelope of the Red fighter reached the destination of the green fighter, there were notable errors in the assessments of the blue fighter. Two curves on the route of the blue fighter also complicated the estimation of the location. The scenario is presented at a) its initial moment, b) after 7 minutes, and c) after 9 minutes and 52 seconds when all the Blue fighters have arrived to their destinations.

Chapter 6

Conclusions

6.1 Fighter allocation model

This thesis introduced a fighter allocation (FA) model constructed to support the training and learning of master fighter controllers (MFCs). An MFC is responsible for controlling and coordinating defensive counterair (DCA) missions in order to, e.g., negate threat from enemy fighters or missiles, to protect important assets or to maintain the control of the air. DCA missions are usually performed with multiple flights for which the MFC has to decide a suitable course of action (COA).

The selection of the COA is a complex problem and requires rapid decision making, as the DCA missions are usually reactive in nature. One of the important considerations the MFC has to make is if potential COAs are feasible, i.e., if a COA can be actually conducted under the current geographical situation in the battlefield. In particular, crucial factors to determine the feasibility of a COA are commit locations of fighters or flights. The MFC assigns the commit locations for each fighter or flight, and at these locations, they start engaging the enemy. In order to perform good decisions quickly, it is important that the MFC correctly perceives and understand the current and near future geometry of the battlefield. Therefore, training regarding the perception and understanding of the MFCs is required. The FA model provides a method for the MFCs to learn and train these abilities and improve their decision making.

The FA model consists of three phases: initialization, optimization and simulation. In the initialization phase, a scenario including an initial situation the MFC needs to react to is created and parameterized. The initial situation reflects the current state of the battlefield and includes the locations of the friendly fighters as well as enemy fighters and their routes from initial locations to destinations. Restricted operating zones (ROZs) can be defined to restrict the flyable area. Furthermore, important assets, areas of interest, or desired engagement frontiers can be optionally included in the initial situation as relevant factors for the DCA mission. At the end of the initialization phase, the MFC makes the decision in order to achieve the objectives of the DCA mission by assigning fighters to the most suitable commit locations to engage the enemy. After the decision is set in the model in a form of destinations, i.e., commit locations, for all the friendly flights or single fighters, network optimization is applied in the second phase. The airspace is described as a network, and the fastest routes of all the friendly fighters from their initial locations to their destinations are solved using A* algorithm. The third phase is simulation, where a continuous-time simulation is performed. The routes of the fighters are evaluated over time by checking encounters of friendly and enemy fighters until they have reached their targets. The outcome of the scenario is determined by the simulation. The outcome is favorable if the friendly fighters are able to reach the destinations set by MFC before the enemy fighters have reached them, before they encounter the enemy fighters, or before enemy fighters have reached their targets. Otherwise, the outcome is considered unfavorable. The favorable outcome indicates that the decision made by the MFC has been successful, while the unfavorable outcome corresponds to an unsuccessful decision. The outcomes of the scenarios are presented as visualizations and numerical data, which allow studying the scenario and the decision made by the MFC in more detail.

The main purpose of the FA model is to provide a new way to train decision making of MFCs. The FA model allows the MFCs to easily and quickly construct and evaluate different complex scenarios corresponding to real-life DCA missions they may encounter in their duties. By using the FA model, the scenarios can be repeated and modified in order to determine the shortcomings in the perception and understanding of an MFC and to develop those abilities. No such model with similar purpose of use has been introduced earlier in the existing unclassified literature.

The FA model can also be utilized in the training and learning of other relevant personnel involved in DCA missions, such as fighter controllers (FCs) and fighter pilots, especially flight leaders. As the FCs are responsible for providing their flights battle-space awareness, improving their perception and understanding on their flights' spatial capabilities is beneficial. Thus, the FA model could provide additional value also in training of the FCs. Additionally, the FA model has potential to be employed also outside of training use. As its use is relatively simple and computations are fast, it is flexible enough to be used in other purposes than training and learning. By offering the evaluation of the routes and timings of the fighters, it could be utilized, e.g., in operational planning of counterair missions. Furthermore, development processes of COAs as well as tactics, techniques and procedures (TTPs) can be supported by the model.

While the FA model benefits the training of MFCs in its current form, it has potential to be developed further. Currently, the approach to solve optimal routes in a network is relatively simplified. With constant ground speeds, the effect of the flying altitude in ground speed is ignored. In reality, the maximum ground speed of a fighter differs as a function of altitude. Consequently, using, e.g., a fixed Mach number as the flying speed would lead to routes that use varying altitudes to complete the routes faster. The optimization network could be also constructed differently, e.g., using a non-grid network based on kinematic capabilities of fighters (see, e.g., Babel, 2013) to obtain more realistic routes.

There exists also avenues to expand the FA model. In the simulation phase of the model, additional relevant factors in the battlefield could be taken into account. For example, hostile surface-to-air missiles could be included (see, e.g., Puustinen, 2013; Gunell, 2019). Additionally, the fuel consumption of fighters could be taken into account to provide limits for feasible routes. For instance, to succeed in a DCA mission, it may be required that fighters have certain amount of fuel left when reaching their destinations to be able to complete the tasked COA. Consequently, in some cases, the fighters would need to take routes where the consumption is less than on the fastest route by flying at higher altitudes. The implementation of the fuel consumption as a constraint or an objective of the network optimization problem would be straightforward. The aforementioned additions could be advantageous especially when considering the potential use of the FA model in an operational environment.

6.2 Experimental study

An experimental study was conducted to investigate the value provided by the FA model in the training of MFCs. In the study, a training intervention was performed with the model during a training course of future MFCs and flight leaders. Ten participants took part in the experiment. The participants had to assess outcomes of test scenarios and locations of fighters in them. In each scenario, three fighters and their locations at the end of the scenario were assessed. The experiment was conducted during a training session where the participants first assessed 20 test scenarios, which was followed by a training intervention in a form of a 2-hour session where the FA model was utilized. After the intervention, 20 other test scenarios were assessed by the participants.

The real outcomes and the locations of the fighters in the test scenarios were determined by the FA model. Two different error metrics were defined to measure the errors made in the assessments. The errors before and after the experiment were compared and differences in them were analyzed using Wilcoxon signed-rank test. The experimental study provided promising results, as statistically significant improvements in some of the assessments were observed. Thus, the potential of the FA model to improve perception and decision making was illustrated. With the experimental study, this thesis also contributed to the field of behavioral operations research (BOR) by evaluating the effects of the FA model in the learning of the participants. No such a study regarding this type of model-based training intervention has been presented earlier in the unclassified literature.

Although the performed study revealed positive results, the benefits of the FA model should be studied further. The sample size of the current study was small, and the study would have benefited from more participants. The length and style of the training intervention were also limited. Thus, the findings of its impact to the perception, understanding and decision making of the participants are mostly preliminary. In order to observe the effects of the FA model to the learning in more detail, the training intervention would preferably be longer and include extensive hands-on use of the model for all the participants. By allowing the participants to use the model themselves and to evaluate the scenarios and their own decisions in them, the educative effect of the model would be presumably stronger. Additionally, presenting the correct results of the assessments performed before the intervention and discussing the participants' errors could enhance the educative impact of the intervention.

Many aspects of the setting of the study performed could be elaborated. In order to have unambiguous before-after pairs for routes, identical scenarios should be presented before and after the intervention. With the current implementation and pairing process for routes, some of the assessment data was lost as no sufficient pairs were found for all the routes. However, as it would be beneficial to present and discuss correct answers and point individual errors in the "before" scenarios during the intervention, the identical scenarios could not be presented as such. In order to mitigate the possibility that a participant would remember the correct outcome of the scenario, the scenarios could be, e.g., mirrored and rotated on the map as well as presented in different order after the intervention.

Better metrics for the assessment errors could be also considered to provide better measures for the quality of the assessments. Furthermore, the assessment task could be set up differently. For example, the participants could be presented a scenario with a route of a Red fighter and posed a question "where would you maneuver the fighters (or flights) in this situation described in this scenario?". Then, with the FA model, the scenario with locations given by the participant would be evaluated. However, different measures for assessment errors should be determined, as this type of phrasing of a question would favor more conservative assessments that could be reached more easily. The assessments could be also performed with and without time pressure, and the benefits of the model could be investigated for both cases.

Measuring and analyzing situational awareness (see, e.g., Mansikka et al., 2021e) and mental workload of participants during assessments could reveal interesting insights into the decision making process of MFCs. These attributes have been measured and studied in simulation settings with aircraft pilots (Mansikka et al., 2016, 2019; Virtanen et al., 2021). Similar studies could also be extended to MFCs. The FA model would allow generating suitable decision making situations for such studies.

Bibliography

Adler, A., Bar-Gill, A., and Shimkin, N. Optimal flight paths for engine-out emergency landing. In *2012 24th Chinese Control and Decision Conference*, pp. 2908–2915, 2012.

Ahuja, R. K., Kumar, A., Jha, K. C., and Orlin, J. B. Exact and Heuristic Algorithms for the Weapon-Target Assignment Problem. *Operations Research*, 55(6):1136–1146, 2007.

Babel, L. Three-dimensional Route Planning for Unmanned Aerial Vehicles in a Risk Environment. *Journal of Intelligent & Robotic Systems*, 71(2): 255–269, 2013.

Bell, H. H. and Waag, W. L. Evaluating the Effectiveness of Flight Simulators for Training Combat Skills: A Review. *The International Journal of Aviation Psychology*, 8(3):223–242, 1998.

Bell, P. C. and O'Keefe, R. M. An Experimental Investigation into the Efficacy of Visual Interactive Simulation. *Management Science*, 41(6): 1018–1038, 1995.

Bennett Jr, W., Schreiber, B. T., and Andrews, D. H. Developing competency-based methods for near-real-time air combat problem solving assessment. *Computers in Human Behavior*, 18(6):773–782, 2002.

Bertsekas, D. P. *Network optimization: continuous and discrete methods*. Athena Scientific, Belmont, MA, 1998.

Carlyle, W. M., Royset, J. O., and Wood, R. K. Routing Military Aircraft With A Constrained Shortest-Path Algorithm. *Military Operations Research*, 14(3):31–52, 2009.

Conover, W. J. *Practical Nonparametric Statistics*. John Wiley & Sons, New York, NY, third edition, 1999.

Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. *Introduction to Algorithms*. MIT Press, Cambridge, MA, third edition, 2009.

Development, Concepts and Doctrine Centre. Joint Doctrine Publication 0-30: UK Air and Space Power. Second edition. UK Ministry of Defence, 2007.

Development, Concepts and Doctrine Centre. Joint Doctrine Publication 04: Understanding and Decision-making. Second edition. UK Ministry of Defence, 2016.

Development, Concepts and Doctrine Centre. Allied Joint Doctrine for Air and Space Operations (AJP-3.3(B)). Edition B, Version 1. NATO Standardization Office, 2016.

Dijkstra, E. W. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.

Farrell, B., Drendel, L., Williamson, K., Harrison, J., Cherwin, B., Ellis, J., and Nogle, J. *Formation Flight Manual*. T-34 Association, Inc., Cameron Park, CA, fourth edition, 1996.

Fone, D., Hollinghurst, S., Temple, M., Round, A., Lester, N., Weightman, A., Roberts, K., Coyle, E., Bevan, G., and Palmer, S. Systematic review of the use and value of computer simulation modelling in population health and health care delivery. *Journal of Public Health*, 25(4):325–335, 2003.

Gogi, A., Tako, A. A., and Robinson, S. An experimental investigation into the role of simulation models in generating insights. *European Journal of Operational Research*, 249(3):931–944, 2016.

Gunell, N. Optimal Route of Military Aircraft in Multi-Stage Surveillance, Reconnaissance and Target Acquisition Missions. Master's thesis, Department of Mathematics and Systems Analysis, Aalto University, Finland, Espoo, 2019.

Harju, M., Liesiö, J., and Virtanen, K. Spatial multi-attribute decision analysis: Axiomatic foundations and incomplete preference information. *European Journal of Operational Research*, 275(1):167–181, 2019.

Hart, P. E., Nilsson, N. J., and Raphael, B. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

Himanen, P. Ilmavoimien taistelunjohtajan tehtäväanalyysi: tilannetietoisen päätöksenteon tarkastelu critical decision methodin avulla. Master's thesis, National Defence University, Helsinki, Finland, 2013.

Holmes, J. M. *The Counterair Companion: Short Guide to Air Superiority for Joint Force Commanders*. Air University Press, Maxwell Air Force Base, AL, 1995.

Huaiping, C., Jingxu, L., Yingwu, C., and Hao, W. Survey of the research on dynamic weapon-target assignment problem. *Journal of Systems Engineering and Electronics*, 17(3):559–565, 2006.

Hämäläinen, R. P., Luoma, J., and Saarinen, E. On the importance of behavioral operational research: The case of understanding and communicating about dynamic systems. *European Journal of Operational Research*, 228 (3):623–634, 2013.

Jalovaara, P. Selection of Air Combat Tactics using a Multi-Attribute Decision Analysis Model with Incomplete Preference Information. Bachelor's thesis, Department of Mathematics and Systems Analysis, Aalto University, Espoo, 2021.

Joint Chiefs of Staff. Joint Publication 3-52: Joint Airspace Control. Joint Chiefs of Staff, 2014.

Joint Chiefs of Staff. Joint Publication 1-02: Department of Defence Dictionary of Military and Associated Terms. Joint Chiefs of Staff, 2016.

Joint Chiefs of Staff. Joint Publication 3-01: Countering Air and Missile Threats. US Joint Chiefs of Staff, 2017.

Joint Chiefs of Staff. Joint Publication 3-30: Joint Air Operations. Joint Chiefs of Staff, 2019.

Kaempf, G. L., Klein, G., Thordsen, M. L., and Wolf, S. Decision Making in Complex Naval Command-and-Control Environments. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 38(2):220–231, 1996.

Karelahti, J., Virtanen, K., and Raivio, T. Near-Optimal Missile Avoidance Trajectories via Receding Horizon Control. *Journal of Guidance, Control, and Dynamics*, 30(5):1287–1298, 2007.

Karelahti, J., Virtanen, K., and Öström, J. Automated Generation of Realistic Near-Optimal Aircraft Trajectories. *Journal of Guidance, Control, and Dynamics*, 31(3):674–688, 2008.

Kewley, R. H. and Embrechts, M. J. Computational military tactical planning system. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 32(2):161–171, 2002.

Kirk, D. E. *Optimal Control Theory: An Introduction*. Dover Publications, Mineola, NY, 2004.

Kline, A., Ahner, D., and Hill, R. The Weapon-Target Assignment Problem. *Computers & Operations Research*, 105:226–236, 2019.

Korean Air Force. Basic Employement Manual - F-16C. *Korean Air Force Tactics, Techniques and Procedures 3-3*, volume 5, 2005.

Law, A. M. *Simulation modeling and analysis*. McGraw-Hill Education, New York, NY, fifth edition, 2013.

LeMay Center for Doctrine. Air Force Doctrine Publication 3-01: Counterair Operations. U.S. Air Force, 2019.

Manne, A. S. A Target-Assignment Problem. *Operations Research*, 6(3): 346–351, 1958.

Mansikka, H., Virtanen, K., Harris, D., and Simola, P. Fighter pilots' heart rate, heart rate variation and performance during an instrument flight rules proficiency test. *Applied Ergonomics*, 56:213–219, 2016.

Mansikka, H., Virtanen, K., and Harris, D. Dissociation Between Mental Workload, Performance, and Task Awareness in Pilots of High Performance Aircraft. *IEEE Transactions on Human-Machine Systems*, 49(1):1–9, 2019.

Mansikka, H., Virtanen, K., Harris, D., and Jalava, M. Measurement of team performance in air combat – have we been underperforming? *Theoretical Issues in Ergonomics Science*, 22(3):338–359, 2021.

Mansikka, H., Virtanen, K., Harris, D., and Salomäki, J. Live–virtual–constructive simulation for testing and evaluation of air combat tactics, techniques, and procedures, Part 1: assessment framework. *The Journal of Defense Modeling and Simulation*, 18(4): 285–293, 2021.

Mansikka, H., Virtanen, K., Harris, D., and Salomäki, J. Live–virtual–constructive simulation for testing and evaluation of air combat tactics, techniques, and procedures, Part 2: demonstration of the framework. *The Journal of Defense Modeling and Simulation*, 18(4): 295–308, 2021.

Mansikka, H., Virtanen, K., Mäkinen, L., and Harris, D. Normative Performance Measurement in Simulated Air Combat. *Aerospace Medicine and Human Performance*, 92(11):908–912, 2021.

Mansikka, H., Virtanen, K., Uggeldahl, V., and Harris, D. Team situation awareness accuracy measurement technique for simulated air combat - Curvilinear relationship between awareness and performance. *Applied Ergonomics*, 96:103473, 2021.

MathWorks. MATLAB, version 8.6.0.267246 (R2015b). The MathWorks Inc., Natick, MA, 2015.

Miller, B., Stepanyan, K., Miller, A., and Andreev, M. 3D path planning in a threat environment. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 6864–6869, 2011.

Monks, T., Robinson, S., and Kotiadis, K. Learning from discrete-event simulation: Exploring the high involvement hypothesis. *European Journal of Operational Research*, 235(1):195–205, 2014.

Mulgund, S., Harper, K., Krishnakumar, K., and Zacharias, G. Air combat tactics optimization using stochastic genetic algorithms. In *SMC'98 Conference Proceedings. 1998 IEEE International Conference on Systems, Man, and Cybernetics*, volume 4, pp. 3136–3141, 1998.

Nilsson, N. J. *Principles of artificial intelligence*. Springer, Berlin, 1982.

O'Keefe, R. M. Experimental behavioural research in operational research: What we know and what we might come to know. *European Journal of Operational Research*, 249(3):899–907, 2016.

Pearl, J. *Heuristics: intelligent search strategies for computer problem solving*. Addison-Wesley, Reading, MA, 1984.

Poropudas, J. and Virtanen, K. Game-Theoretic Validation and Analysis of Air Combat Simulation Models. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 40(5):1057–1070, 2010.

Puustinen, H. Military aircraft routing with multi-objective network optimization and simulation. Master's thesis, Department of Mathematics and Systems Analysis, Aalto University, Finland, Espoo, 2013.

Rantala, E. Analyzing defense capability of a fleet of military aircraft through simulation. Master's thesis, Department of Mathematics and Systems Analysis, Aalto University, Finland, Espoo, 2018.

Rippel, E., Bar-Gill, A., and Shimkin, N. Fast Graph-Search Algorithms for General-Aviation Flight Trajectory Generation. *Journal of Guidance, Control, and Dynamics*, 28(4):801–811, 2005.

Rouwette, E. A. J. A., Korzilius, H., Vennix, J. A. M., and Jacobs, E. Modeling as persuasion: the impact of group model building on attitudes and behavior. *System Dynamics Review*, 27(1):1–21, 2011.

Salas, E., Bowers, C. A., and Rhodenizer, L. It Is Not How Much You Have but How You Use It: Toward a Rational Use of Simulation to Support Aviation Training. *The International Journal of Aviation Psychology*, 8 (3):197–208, 1998.

Shaw, R. L. *Fighter combat: tactics and maneuvering.* Naval Institute Press, Annapolis, MD, 1985.

Thunholm, P. Planning Under Time Pressure: An Attempt Toward a Prescriptive Model of Military Tactical Decision Making. In Montgomery, H., Lipshitz, R., and Brehmer, B., editors, *How Professionals Make Decisions*, pp. 53–66. CRC Press, Mahwah, NJ, 2004.

Tidhar, G., Heinze, C., and Selvestrel, M. Flying Together: Modelling Air Mission Teams. *Applied Intelligence*, 8(3):195–218, 1998.

Virtanen, K., Ehtamo, H., Raivio, T., and Hämäläinen, R. VIATO-visual interactive aircraft trajectory optimization. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 29(3):409–421, 1999.

Virtanen, K., Hämäläinen, R., and Mattila, V. Team optimal signaling strategies in air combat. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 36(4):643–660, 2006.

Virtanen, K., Raivio, T., and Hämäläinen, R. P. Decision Theoretical Approach to Pilot Simulation. *Journal of Aircraft*, 36(4):632–641, 1999.

Virtanen, K., Raivio, T., and Hämäläinen, R. P. Modeling Pilot's Sequential Maneuvering Decisions by a Multistage Influence Diagram. *Journal of Guidance, Control, and Dynamics*, 27(4):665–677, 2004.

Virtanen, K., Karelahti, J., and Raivio, T. Modeling Air Combat by a Moving Horizon Influence Diagram Game. *Journal of Guidance, Control, and Dynamics*, 29(5):1080–1091, 2006.

Virtanen, K., Mansikka, H., Kontio, H., and Harris, D. Weight watchers: NASA-TLX weights revisited. *Theoretical Issues in Ergonomics Science*, accepted for publication, 2021.

Zabarankin, M., Uryasev, S., and Murphey, R. Aircraft routing under the risk of detection. *Naval Research Logistics*, 53(8):728–747, 2006.