Selection of Air Combat Tactics using a Multi-Attribute Decision Analysis Model with Incomplete Preference Information

Perttu Jalovaara

School of Science

Bachelor's thesis
Espoo 22.10.2021

Supervisor

Prof. Kai Virtanen

Advisor

Prof. Kai Virtanen

Aalto University
School of Science

Aalto University
School of Science

Aalto University, P.O. BOX 11000, 00076 AALTO
www.aalto.fi
Abstract of the bachelor's thesis

Author Perttu Jalovaara

Title Selection of Air Combat Tactics using a Multi-Attribute Decision Analysis Model with Incomplete Preference Information

Degree programme Engineering Physics and Mathematics

Major Mathematics and Systems Sciences **Code of major** SCI3029

Teacher in charge Prof. Kai Virtanen

Advisor Prof. Kai Virtanen

Date 22.10.2021 **Number of pages** 31 **Language** English

Abstract

This thesis focuses on the development of a multi-attribute decision analysis model for the selection of air combat tactics. The model is used to support the process of converting the air component commander's intent into a course of action (COA). Such a COA consists of the air combat tactics of four flights, each of which comprises four fighter aircraft. The primary objective of the thesis is to develop a model capable of reducing the number of COAs in consideration to a feasible amount.

Selecting air combat tactics efficiently is critical to the success of an air operation. The difficulty of this task lies in the complexity of evaluating and comparing the COAs. Ideally, the selected COA should represent the commander's intent accurately. In practice, however, the commander's intent is interpreted at different levels of the chain of command and can thus suffer unintentional modifications before being realized as a COA.

In this thesis, air combat tactics are evaluated with respect to three attributes: the probability of killing the target, the probability of surviving the encounter and the efficiency of missiles launched. The commander's intent is interpreted as preference information concerning the attribute weights, which describe the relative importance of the attributes. The COAs are compared using the concept of dominance, which allows for identification of ineffective and effective alternatives with respect to the interpreted preference information.

Two experiments are performed to study the effect of varying forms of preference information on the COA recommendations provided by the model. Based on the results of these experiments, we can conclude that an ordinal ranking of attributes suffices to reduce the number of COAs in consideration from thousands to a handful. In addition, the model developed in this thesis identifies effective COAs that are generally sensible with regard to the given preference information.

Keywords air combat, decision analysis, additive value function, incomplete preference information

Aalto University
School of Science

Aalto-yliopisto, PL 11000, 00076 AALTO
www.aalto.fi
Tekniikan kandidaatintyön tiivistelmä

Tekijä Perttu Jalovaara

Työn nimi Monitavoitteinen päätösanalyysimalli ilmataistelutaktiikoiden vertailuun epätäydellisellä preferenssi-informaatiolla

Koulutusohjelma Teknillinen fysiikka ja matematiikka

Pääaine Matematiikka ja systeemitieteet **Pääaineen koodi** SCI3029

Vastuuopettaja Prof. Kai Virtanen

Työn ohjaaja Prof. Kai Virtanen

Päivämäärä 22.10.2021 **Sivumäärä** 31 **Kieli** Englanti

Tiivistelmä

Tässä kandidaatintyössä kehitetään monitavoitteinen päätösanalyysimalli ilmataistelutaktiikoiden vertailuun. Mallilla tuetaan ilmavoimien komentajan tahdon muuntamista toimintavaihtoehdoksi (engl. course of action, COA). COA määrittää ilmataistelutaktiikat neljälle parvelle, joista jokainen koostuu neljästä hävittäjästä. Työn päätavoitteena on kehittää malli, joka kykenee tunnistamaan parhaat komentajan tahdon mukaiset COA:t kaikista vaihtoehdoista.

Ilmataistelutaktiikoiden tehokas valitseminen on tärkeää ilmaoperaation menestykselle. Tämän valintatehtävän haastavuus johtuu COA:iden arvioinnin ja vertailun monimutkaisuudesta. Valitun COA:n tulisi edustaa komentajan tahtoa täsmällisesti. Todellisuudessa komentajan tahto tulkitaan monella komentoketjun tasolla, mikä altistaa sen tahattomille muutoksille ja saattaa johtaa tahdon vastaisen COA:n valintaan.

Tässä työssä ilmataistelutaktiikoita arvioidaan kolmen kriteerin perusteella: vastustajan torjunnan todennäköisyys, selviytymisen todennäköisyys ja ammuttujen ohjusten tehokkuus. Komentajan tahto tulkitaan preferenssi-informaationa kriteeripainoista, jotka kuvaavat kriteerien suhteellista tärkeyttä. COA:iden vertailuun käytetään dominanssia, joka mahdollistaa tehottomien ja tehokkaiden vaihtoehtojen tunnistamisen annetun preferenssi-informaation suhteen.

Tässä työssä tarkastellaan kokeellisesti erilaisen preferenssi-informaation vaikutusta mallin tuottamiin COA-suosituksiin. Kokeiden tulosten perusteella havaitaan, että kriiterien asettaminen ordinaaliseen tärkeysjärjestykseen riittää potentiaalisesti parhaiden, komentajan tahtoa noudattavien COA:iden rajaamiseen tuhansista kymmeniin. Lisäksi koetulokset osoittavat, että työssä kehitetyn mallin tuottamat COA-suositukset ovat johdonmukaisia annetun preferenssi-informaation kanssa.

Avainsanat ilmataistelu, päätösanalyysi, additiivinen arvofunktio, epätäydellinen preferenssi-informaatio

Contents

Abstract **3**

Abstract (in Finnish) **4**

Contents **5**

1 Introduction **6**

2 Selection of Air Combat Tactics **7**

3 Multi-Attribute Decision Analysis (MADA) Model **10**
3.1 Attributes 10
3.2 Additive Value Function 11
3.3 Incomplete Preference Information 12
3.4 Comparing Overall Value Intervals 14
3.5 Core Index 18
3.6 Implementation 19

4 Experiments **20**
4.1 Experiment I: Preference Order and Attribute Scoring 20
4.1.1 Effect on the number of non-dominated COAs 20
4.1.2 Alternative scoring functions 23
4.2 Experiment II: Convergence of Weight Intervals 24
4.2.1 Effect on the number of non-dominated COAs 24
4.2.2 Effect on the core indices of individual TTPs 26

5 Conclusion **29**

1 Introduction

The primary role of a fighter aircraft is to gain and maintain control of the air. Establishing control of the air is seen as a prerequisite for all air operations (JP 1-02). The establishment of control of the air is dependent on many factors ranging from the capabilities of the aircraft to the skill of the pilots and the tactics used by them. The ability to evaluate tactical alternatives efficiently is critical to the success of air operations (AJP-3.3). The development of practices and tools for evaluating and selecting air combat tactics is therefore warranted.

The combination of tactics for multiple flights, i.e., units of four fighter aircraft, is referred to as a course of action (COA). The COA is commonly decided by the Air Operations Center (AOC), which is responsible for actualizing the intent of the air component commander (JP 3-30). Fighter aircraft can move at extreme speeds in a three-dimensional domain, which results in a wide range of possible engagement alternatives. Selecting an effective COA is a difficult task not only due to the sheer number of alternatives but also due to the complexity of measuring the value of a COA. Instead of attempting to evaluate each COA in an objective sense, we can evaluate them with respect to some set of preferences, preferably that of the air component commander.

The value of a COA is easier to determine if it is evaluated with respect to multiple attributes, instead of trying to holistically rate it as a whole. The purpose and the desired end state of a military operation is communicated across the chain of command through the commander's intent, which is clear and concise in form (JP 1-02). The commander's intent can be incorporated into a multi-attribute approach by interpreting it as preferences concerning the attributes. These preferences can be represented with a value function, assuming that they satisfy specific axioms characterizing rational decision making (see, e.g., Salo and Hämäläinen, 2010). The value function, which is a prevalent decision analysis tool for comparing alternatives, can then be used to calculate the overall values of COAs (see, e.g., Mustajoki et al., 2005). In this thesis, we use an additive value function (Keeney et al., 1993) to evaluate the overall value of each COA as a weighted sum of its scores, i.e., partial values with respect to each attribute.

This thesis focuses on the development of a multi-attribute decision analysis (MADA) model for COA selection. The model is used to support the process of converting the commander's intent into an effective COA. In the sequel, this process is referred to as the decision making process. The primary objective of this thesis is to decrease the number of COAs in consideration from thousands to a handful, thus enabling the AOC to choose the most suitable alternative efficiently. The effect of varying attribute weights on the effective COAs identified by the MADA model is studied experimentally. Interpreting the commander's intent into exact attribute weights is often difficult, so we allow incomplete information, i.e., ordinal or otherwise imprecisely defined preferences, in the model (Weber, 1987). The modeling of incomplete preference information is motivated by its successful use in multi-attribute value models as well as other areas, including spatial decision analysis (Harju et al., 2019), simulation based decision making (Mattila and Virtanen, 2015), game theoretical settings (Kokkala et al., 2019), data envelopment analysis (Salo and Punkka, 2011) and project portfolio selection (Liesiö et al., 2007, 2008; Tervonen et al., 2017). The use of incomplete preference information leads to an interval representation of overall values of the COAs. The overall value intervals can then be analyzed to identify the efficient COAs, whose potentials for usage in the prevailing real-world scenario are then evaluated by the AOC.

The thesis is structured as follows. Section 2 briefly introduces the necessary concepts of air combat and the current practices in COA selection. Section 3 describes the specifications and the implementation of the MADA model in addition to outlining the decision making process. Section 4 details two numerical experiments and presents their results. Lastly, Section 5 concludes the thesis.

2 Selection of Air Combat Tactics

In order to limit the scope of this thesis and to simplify the problem of selecting an effective COA, we must define what constitutes a COA. In general, any planned sequence of activities developed to accomplish a mission can be regarded as a COA (JP 1-02). The missions considered in this thesis fall under the category of defensive counter air operations, in which friendly fighter aircraft (blue) are used against threat fighter aircraft (red). Thus, we use a specific definition of the COA as a coordinated employment of the air combat tactics, techniques and procedures (TTPs) of four blue flights, each consisting of four fighter aircraft. A TTP can generally include a variety of tactical principles and schemes that the pilots of a flight follow to coordinate their efforts towards mission success. However, the concept of TTPs can be simplified to include only two essential tactical contracts: 1) the minimum acceptable range between blue and red aircraft and 2) the missile launch ranges of blue aircraft.

Before proceeding further, it is necessary to specify the assumptions and restrictions of the MADA model developed in this thesis. In addition to using a simplified definition of TTPs, the model relies on two significant assumptions. First, it is assumed that there are no geographical restrictions on COA selection. In reality, however, such restrictions can be imposed by a multitude of factors, ranging from ground-based air defense of the red to the command and control capability of the blue (JP 3-30). Because of the disregard for any geographical restrictions, the model can recommend COAs that are theoretically effective but infeasible in practice. Second, we are not incorporating any information on the red aircraft into the model. For example, in reality, intercepting a flight of multi-role fighters requires a different tactical approach than targeting a single bomber. These assumptions need to be taken into account when considering the recommendations provided by the model.

Let us consider three possible alternatives for the minimum acceptable range between blue and red aircraft (CNATRA P-825): maximum range (R_{max}), mid-range (R_{mid}) and minimum range (R_{min}). We refer to these three ranges as geometry alternatives. There are important distinctions in the risks and opportunities of flights using TTPs with different geometry alternatives. Assuming that the geometry management of the flight is successful, using R_{max} allows mission completion while minimizing the risk of red missile launches against the flight. A flight using R_{mid} actively moves towards the target while maintaining the ability to kinematically defeat, i.e., dodge, any incoming missiles. Lastly, a flight employing R_{min} seeks to get close to the target and compromises the ability to kinematically defeat incoming missiles.

In a similar fashion, we recognize three possible ranges at which the blue flight can employ missiles against the enemy (CNATRA P-825): long-range launch (Fox_{long}), medium range launch (Fox_{med}) and close-range launch (Fox_{close}). These three alternatives are referred to as launch range alternatives. Fox is a brevity word in air combat terminology that is used to indicate the launch of an air-to-air missile. Fox_{long} is the safest launch range alternative, as it maximizes the range between the blue and the red aircraft at missile launch. Missiles launched at this range are the easiest for the target to kinematically defeat. Fox_{med} is a balanced launch range alternative that moderately exposes the blue flight to red missiles and increases the effectiveness of blue missiles compared to Fox_{long}. Finally, Fox_{close} maximizes the effectiveness of the launched missiles but it also exposes the blue flight to missile launches of the red. Overall, the closer the flight moves to the target, the higher is the probability of killing the target. At the same time, moving closer to the target reduces the probability of survival of the blue flight. Thus, the selection of geometry and launch range alternatives revolves around balancing these probabilities.

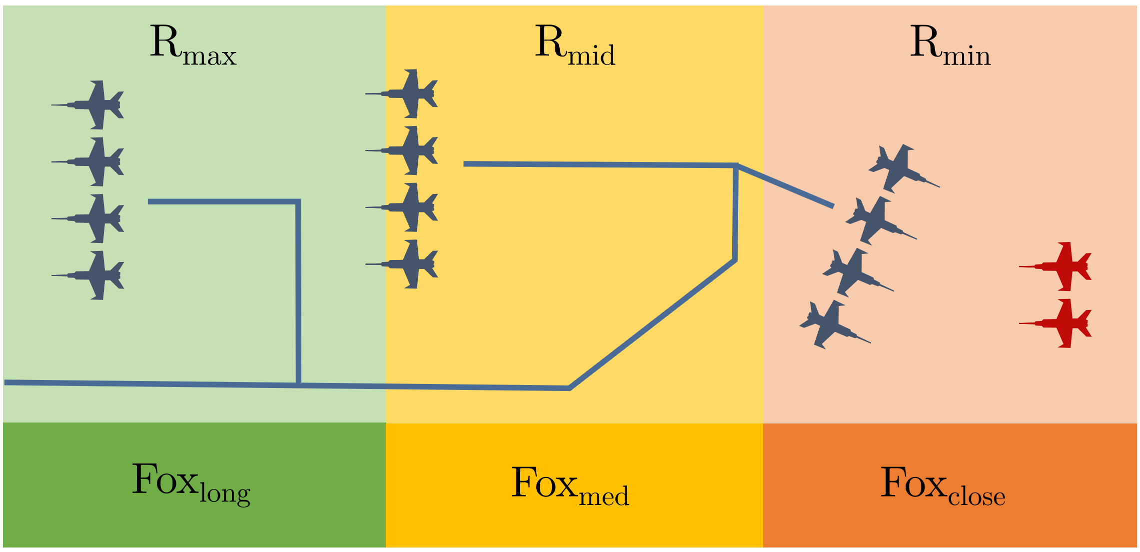

Figure 1: Geometry ($R_{max/mid/min}$) and launch range alternatives ($Fox_{long/med/close}$).

The geometry alternatives and the launch range alternatives of a flight are linked in the following manner. A flight can only employ Fox_{long} at R_{max}, Fox_{med} at R_{mid} and $\text{Fox}_{\text{close}}$ at R_{min}. This relationship is visualized in Figure 1. It is natural to assume that a blue flight will move closer to the red aircraft only if it intends to perform missiles launches from a closer range. Thus, a flight that uses R_{max} geometry must also use Fox_{long} as its only launch range. Moreover, the usage of R_{mid} geometry leads to two possible combinations of launch range alternatives, namely Fox_{med} and a combination of Fox_{med} and Fox_{long}. Let us now consider the possible launch range combinations of a flight that follows R_{min} geometry. Each of the combinations must include $\text{Fox}_{\text{close}}$ by assumption. Therefore, the combinations vary only in the use of Fox_{long} and Fox_{med}: The flight can use one of the two, both, or neither. Thus, the number of possible launch range combinations for R_{min} is four. Enumerating the launch range combinations of all three geometry alternatives results in seven unique combinations. The seven combinations are listed as the rows of Table 1.

Table 1: TTPs as combinations of geometry and launch range alternatives. The letters L, M and C refer to Fox_{long}, Fox_{med} and $\text{Fox}_{\text{close}}$, respectively. Cells with an 'x' indicate the use of the columnar launch range with the geometry of that row.

TTP	Geometry alternative	Launch range alternatives		
		Fox_{long}	Fox_{med}	$\text{Fox}_{\text{close}}$
L	R_{max}	x		
LM	R_{mid}	x	x	
M	R_{mid}		x	
LMC	R_{min}	x	x	x
LC	R_{min}	x		x
MC	R_{min}		x	x
C	R_{min}			x

The same geometry alternatives and launch range combinations are used by the whole flight, i.e., the four aircraft will always follow a common plan. Furthermore, each of the seven combinations listed in Table 1 defines a unique TTP. The TTPs are labeled with a combination of the letters L, M and C according to the launch range alternatives, where L refers to Fox_{long}, M refers to Fox_{med} and C refers to $\text{Fox}_{\text{close}}$. The geometry alternative of a TPP can be deduced from the shortest launch range used. A COA consists of the TTPs of a formation of four flights. For most formations, we can identify flights as leading or trailing based on their relative positions. As each flight can choose from seven different TTPs, there are a total of $7^4 = 2401$ unique COAs. Considering the many simplifications we made while defining air combat tactics, the number of COAs is still rather large.

Before manifesting as a selected COA, the commander's intent is passed through the chain of command (AJP-3.3). The AOC is responsible for selecting an effective COA that also conforms to the commander's intent. This task is clearly demanding. Ideally, the selected COA should represent the air component commander's intent accurately. In practice, however, the commander's intent is interpreted at different levels of the chain of command before it is realized as the TTPs of blue flights. The commander's intent is typically concise and provides guidance only at the operational level (JP 3-30). Thus, an interpretation is necessary in order for the AOC to select an effective COA. Verbal interpretations are naturally susceptible to unintentional modifications. A more robust approach would be to interpret the commander's intent as preference statements concerning the relative importance of some well-chosen attributes of COAs.

In addition to providing robustness throughout the chain of command, the MADA model developed in this thesis enables evaluation and ranking of COAs. During the planning phase of an operation, the AOC has enough time to evaluate a vast range of COAs with the decision support provided by the model (AJP-3.3). Therefore, it is probable that the first COA can be selected so that it represents the commander's intent accurately. However, the situation might change rapidly after the first contact with the red and thus the COA may have to be dynamically modified. Using the model in a fast evolving combat scenario might prove to be challenging because there is no constantly updating source of preference information. Therefore, the strength of the model lies in the planning phase of an operation, when preference information is reliably obtainable and time is not scarce.

3 Multi-Attribute Decision Analysis (MADA) Model

3.1 Attributes

Having identified 2401 unique COAs in Section 2, the next step is to evaluate the overall value of each COA according to the commander's intent. Using standard decision analysis terms, the COAs are decision alternatives and the air component commander is the decision maker (DM). The value of a decision alternative is measured with respect to attributes. A multi-attribute value function representation of the DM's preferences can be used to assess the overall value of the decision alternatives (Keeney et al., 1993). For this multi-attribute approach, we need to define attributes that are sensible for measuring the value of a COA.

In air combat, the ratio between red aircraft shot down and blue aircraft lost is a traditional measure of performance output (Mansikka et al., 2021). The number of red aircraft killed and the number of surviving blue aircraft can also be used as an operational objective in the commander's intent. According to Mansikka et al. (2021), the number of missiles required for a kill is another typical air combat performance indicator. The efficient use of resources, including missiles, is especially desirable for defensive operations, where a defender typically possesses inferior material resources compared to an attacker (Holmes, 1995).

Based on these air combat performance indicators, let us define three attributes for measuring the value of a COA: the probability of killing red aircraft (PK), the probability of survival for blue aircraft (PS), and the efficiency of air-to-air missile usage of blue flights (EM). Each COA can be ranked with respect to these three attributes. It is obviously impossible to accurately calculate the actual probability of a complex event, such as the survival of the blue aircraft for a given COA. However, this is not required for establishing a relative ranking of the COAs with respect to each attribute. These attribute-specific rankings are determined with a rule-based approach. Three rules are defined for each attribute, as shown in Table 2. The preference order is established by first dividing the COAs into more and less preferred groups according to the first rule, after which the second and third rule are used to generate smaller subgroups. The attributes and the associated rules are in no way absolute but they are sufficient in that they can be used to compare the COAs.

Table 2: Attribute rules. Lead and trail refer to leading and trailing blue flights.

Attribute	Rule 1	Rule 2	Rule 3
Kill probability (PK)	Shoot often	Prefer $\text{Fox}_{\text{close}}$	Lead shoots from a closer range than trail.
Survival probability (PS)	Prefer Fox_{long}	Shoot often	Trail shoots from a closer range than lead.
Missile efficiency (EM)	Shoot seldom	Prefer $\text{Fox}_{\text{close}}$	Lead shoots from a closer range than trail.

3.2 Additive Value Function

In this thesis, the attributes are assumed to be mutually preferentially independent (Keeney et al., 1993). Mutual preferential independence implies that the preference over the levels of an attribute does not depend on the levels of the other attributes. Thus, an additive value function can be used to evaluate the overall value of a COA. The additive value function is a weighted sum of the scores of a COA with respect to each attribute. Multiplicative and other, more complex, non-additive forms of multi-attribute value functions (see, e.g., Dyer and Sarin, 1979) are not considered in this thesis.

The set of decision alternatives is now denoted by $X = \{x^1, x^2, ..., x^{2401}\}$, where x^j is the j-th COA. Assuming that there are n attributes instead of the exemplary three, the overall value of the COA $x \in X$ is

$$v(x) = \sum_{i=1}^{n} w_i v_i(x_i)\,, \tag{1}$$

where x_i is the measurement level of COA x with respect to attribute i, v_i is the single attribute value function of attribute i, and w_i is the attribute weight of attribute i. The attribute weights are normalized, such that $w_i \in [0, 1]\,\forall\, i \in \{1, 2, ..., n\}$ and

$\sum_{i=1}^{n} w_i = 1$. The measurement level x_i is a direct measure of the COA with respect to the attribute i, e.g., the average number of missiles required for a kill in the case of the efficiency attribute EM. Defining the measurement levels can be difficult for some attributes, e.g, PK and PS, because there are no directly measurably quantities related to them. This problem is avoided with the rule-based approach discussed in Section 3.1.

The single attribute value function v_i in Equation (1) converts the measurement level x_i into a normalized score $v_i(x_i) \in [0, 1]$ with respect to attribute i. Finding an explicit form for each single attribute value function is not an easy task, as it usually requires input from the DM or from another expert in the field. Salo and Hämäläinen (2001) refer to this task as score elicitation and present examples of ordinal and cardinal methods for accomplishing it. Cardinal methods of score elicitation include directly rating the alternatives on a $[0, 100]$ scale and eliciting ratio comparisons about value differences. A simple, ordinal method is to rank the alternatives with respect to each attribute and linearly score them by giving the least preferred alternative a zero and the most preferred alternative a one. Numerous other score elicitation methods can be found in the literature (see, e.g, Keeney et al., 1993; Salo and Hämäläinen, 2010).

3.3 Incomplete Preference Information

Attribute weights describe the relative importance of attributes. The elicitation of exact point estimates for the attribute weights is challenging in the selection of COAs due to the imprecision of the commander's intent. Complete preference information is successfully used in many applications and numerous techniques have been developed for the elicitation of exact point weights (Weber, 1987). However, the modeling of incomplete preference information allows us to select COAs even if we cannot obtain exact weights. Incomplete preference information on the weights can be given in many forms, ranging from ordinal preferences, e.g., $w_1 \geq w_2 \geq w_3$ or 'attribute 1 is more important than attribute 2 which is more important than attribute 3', to interval-valued ratio statements, e.g., $3 \leq \frac{w_1}{w_2} \leq 5$ or 'attribute 1 is three to fives times more important than attribute 2'. These types of preference statements are discussed in detail by Salo and Hämäläinen (1992, 1995, 2010).

Incomplete preference statements impose linear constraints on the set of all possible attribute weights. The set of feasible weights consists of the combinations of attribute weights that satisfy those linear constraints (Salo and Hämäläinen, 1992). For the sake of example, let us inspect the set of feasible weights when working with the three COA attributes defined previously. The weights of attributes PK, PS and EM are denoted by w_{PK}, w_{PS}, and w_{EM}, respectively. Before we gain any information on the commander's intent, the set of feasible weights, denoted by S, is equal to the unconstrained set of weights $W = \{(w_{PK}, w_{PS}, w_{EM}) \,|\, w_{PK} + w_{PS} + w_{EM} = 1, w_i \geq 0, i = PK, PS, EM\}$. By contrast, complete information on the attribute weights, e.g., $(w_{PK}, w_{PS}, w_{EM}) = (0.4, 0.3, 0.3)$, reduces the set of feasible weights to a singleton, i.e., $S = \{(0.4, 0.3, 0.3)\} \subset W$.

Having discussed the two edge cases, i.e., no information and complete information, let us now consider an example with two interval judgments about the importance of the attributes. We have received the following verbal information from the commander: 'Kill probability is one to three times as important as survival probability and one to two times as important as missile efficiency'. This preference statement is formulated mathematically as $1 \leq \frac{w_{PK}}{w_{PS}} \leq 3$ and $1 \leq \frac{w_{PK}}{w_{EM}} \leq 2$. Alternatively, interval judgments of this type can be written in the form $I_{i,j} = [l_{ij}, u_{ij}]$, where l_{ij} is the lower bound and u_{ij} is the upper bound for the ratio of the i-th attribute weight and the j-th attribute weight (Salo and Hämäläinen, 1992). Using this notation, the preference information is compactly written as $I_{PK,PS} = [1, 3]$ and $I_{PK,EM} = [1, 2]$. Figure 2 visualizes the linear constraints imposed by the aforementioned interval judgments and the resulting set of feasible weights S in a ternary plot.

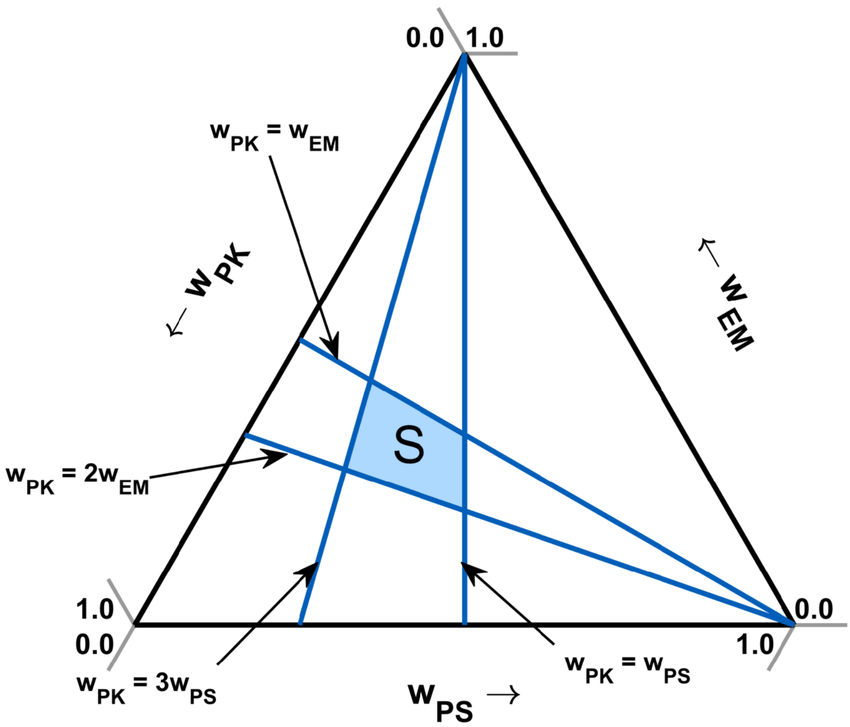

Figure 2: Ternary plot of the linear constraints imposed by the interval judgments $I_{PK,PS} = [1, 3]$ and $I_{PK,EM} = [1, 2]$. The resulting feasible region S is highlighted.

Interval judgments allow the commander to express his preferences as approximate ratio statements instead of exact estimates (Salo and Hämäläinen, 1995). This method falls into the category of preference programming methods. Preference programming can be characterized as the development of interactive decision support methods that accommodate incomplete preference information (Arbel, 1989). Salo and Hämäläinen (2010) present a variety of preference programming methods that are viable options for weight assessment in the context of this thesis. However, it is not necessary to implement support for all of these weight assessment methods directly into the MADA model developed in this thesis. We can accomplish a sufficient level of adaptivity in weight assessment by supporting two types of input: linear constraints and attribute-specific interval restrictions, because most of the methods showcased by Salo and Hämäläinen (2010) output one of the two.

3.4 Comparing Overall Value Intervals

The comparison of COAs with incomplete information differs essentially from the case of complete information. If the commander has exact estimates for the attribute weights, the comparison of decision alternatives is straightforward because an exact overall value for each COA is determined by the additive value function (1). By contrast, the use of incomplete preference information leads to each COA having a range of possible overall values on the set of feasible weights, making the comparison more difficult. Weber (1987) presents two concepts of dominance for comparing decision alternatives under incomplete preference information, both of which are based on representing the overall values as intervals.

The overall value interval of a COA is denoted by $V(x) = [\underline{v}(x), \bar{v}(x)]$, where $\underline{v}(x)$ is the lower bound and $\bar{v}(x)$ is the upper bound for the possible overall values of the COA. The attribute weights are allowed to vary within the constraints obtained from preference information (Salo and Hämäläinen, 1992). After the calculation of overall value intervals with the additive value function, the COAs can be compared based on pairwise dominance, the less restrictive dominance criterion introduced by Weber (1987). The COA x^j dominates the COA x^k in S, which is formally written as $x^j \succ_S x^k$, if and only if

$$\begin{cases} V(x^j) \geq V(x^k), \text{ for all } w \in S, \\ V(x^j) > V(x^k), \text{ for some } w \in S. \end{cases} \tag{2}$$

In other words, the conditions for pairwise dominance (2) hold when the overall value of x^j is greater than or equal to the overall value of x^k for all attribute weights in the feasible set and strictly greater for some. Pairwise dominance thus implies that the dominating COA is never worse than and sometimes better than the dominated COA in light of the commander's preferences.

The COA x is non-dominated if and only if there are no other COAs that dominate it in the set of feasible weights S, i.e.,

$$x \in X_{ND} \iff \nexists\, x' \in X : x' \succ_S x\,, \tag{3}$$

where $X_{ND} \subset X$ denotes the set of non-dominated decision alternatives. Therefore, each non-dominated COA can be considered effective because there are no strictly better alternatives available in view of the commander's preferences. Furthermore, Equation (3) implies that all dominated COAs should be left out of consideration because for each dominated COA $x \in X \setminus X_{ND}$ there exists at least one strictly better alternative. The concept of pairwise dominance thus allows us to exclude a number of ineffective, dominated COAs from the analysis and determine the set of preferred, non-dominated COAs.

The pairwise dominance relation of two COAs is determined by solving a set of linear optimization problems (see, e.g., Salo and Hämäläinen, 1992). The conditions for pairwise dominance (2) can be rearranged in the following manner:

$$V(x^j) \geq V(x^k), \text{ for all } w \in S \iff \min_{w \in S} [V(x^j) - V(x^k)] \geq 0 \tag{4}$$

$$V(x^j) > V(x^k), \text{ for some } w \in S \iff \max_{w \in S} [V(x^j) - V(x^k)] > 0\,. \tag{5}$$

The dominance relations of two COAs can thus be established by examining their minimum and maximum value differences. COA x^j dominates COA x^k if their minimum value difference (4) is non-negative and their maximum value difference (5) is positive. Conversely, if the minimum is negative and the maximum is non-positive then x^j is dominated by x^k. Lastly, if neither of these statements is true then neither COA dominates the other. This leads us to two linear optimization problems with the same constraints, but different objectives:

$$\min_{w \in S} \; V(x^j) - V(x^k) \tag{6}$$

$$\max_{w \in S} \; V(x^j) - V(x^k) \tag{7}$$

$$\text{s.t.} \quad Aw \leq b, \tag{8}$$

$$\sum_{i=1}^{n} w_i = 1, \tag{9}$$

$$w_i \geq 0. \tag{10}$$

The objective functions (6) and (7) stem directly from the rearranged conditions for pairwise dominance, (4) and (5), respectively. The first linear constraint (8) contains the m-by-n matrix A, the n-dimensional attribute weight vector w and the m-dimensional vector b, where m is the number of constraint rows imposed by the preference information and n is the number of attributes. The type of the preference statements dictates the definition of matrix A and vector b. Continuing with our example of interval judgements $I_{PK,PS} = [1, 3]$ and $I_{PK,EM} = [1, 2]$, the associated linear constraints would be given as

$$A = \begin{bmatrix} -1 & 1 & 0 \\ 1 & -3 & 0 \\ -1 & 0 & 1 \\ 1 & 0 & -2 \end{bmatrix} \text{ and } b = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

The elements of b are non-zero in the case of strict rankings, e.g., $w_{PK} - w_{PS} \geq 0.2$ (Salo and Hämäläinen, 2010). The remaining constraints (9) and (10) ensure that the attribute weights remain within the bounds of their definition, i.e., $w_i \in [0, 1]$.

Establishing every pairwise dominance relation for 2401 COAs by brute force would require us to solve $\binom{2401}{2} \cdot 2 = 5\,762\,400$ linear optimization problems. Fortunately, the number of the problems can be reduced with a few simple steps. First, the precise structure of the complete dominance relation graph (see, e.g., Weber, 1987), i.e., which COA dominates which, is not of any interest to the commander nor is it necessary for completing the decision process. It is sufficient to know only whether a COA is dominated or non-dominated. Therefore, if an alternative is established to be dominated, we can immediately exclude it from all further computations. Second, we can define and utilize an additional form of dominance, which is introduced by Weber (1987) as absolute dominance. According to this dominance criterion, COA x^j dominates COA x^k absolutely if the lower bound of the overall value of x^j is greater than the upper bound of the overall value of x^k, i.e., $\underline{v}(x^j) > \bar{v}(x^k)$. Unlike pairwise dominance, the absolute dominance of two COAs can be visually deduced from the overall value intervals. To showcase this fact, the overall value intervals of five arbitrarily defined COAs are plotted in Figure 3. The minimum possible value of x^5 is clearly greater than the maximum possible value of x^3, i.e., $\underline{v}(x^5) > \bar{v}(x^3)$. Thus, x^5 dominates x^3 absolutely. The fact that x_2 dominates x_1, x_3 and x_4 can be confirmed only by solving the associated linear optimization problems.

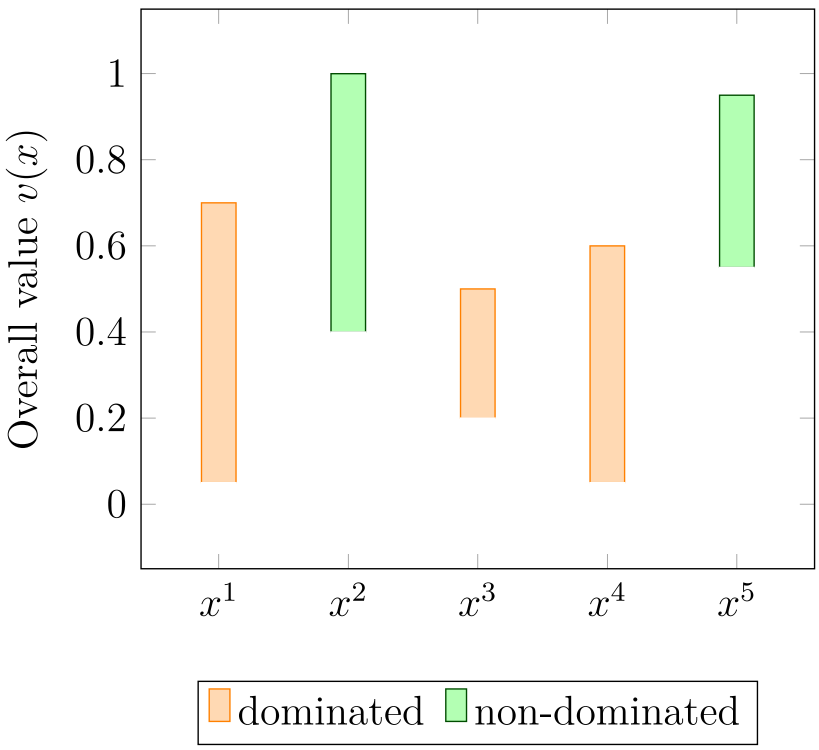

Figure 3: Overall value intervals of five COAs x^j, $j \in \{1, 2, 3, 4, 5\}$.

Absolute dominance is more restrictive and easier to identify in comparison to pairwise dominance. Moreover, the set of absolutely dominated COAs is a subset of the set of dominated COAs (Weber, 1987). Therefore, it is sensible to determine the absolute dominance structure of COAs before computing each pairwise dominance relation. Thus, the pairwise dominance relation of two COAs needs to be computed only when their overall value intervals overlap (Salo and Hämäläinen, 1995), e.g, in the case of COAs x^1 and x^2 in Figure 3. This allows us to further reduce the amount of pairwise dominance computations, which in turn increases the potential of the model for real-world use because the number of COAs can be large, as discussed in Section 2.

If the preference information acquired from the commander results in too many non-dominated COAs, we can elicit additional preference statements which will impose stricter linear constraints on the set of feasible weights. Assuming that the new preference statements adhere to certain rules of consistency (see, e.g., Salo and Hämäläinen, 2010), all existing dominance relations are preserved and, usually, new ones are established. Therefore, the set of non-dominated COAs X_{ND} stays the same or becomes smaller when exposed to additional preference statements.

When we have exhausted the available preference information, the set of non-dominated alternatives can not be reduced further in any objective manner. Decision rules are commonly used to derive a decision recommendation when there are several non-dominated alternatives left. Salo and Hämäläinen (2010) present a variety of decision rules, ranging from maximization of expected value to weight centralization. The following four principles are examples of decision rules that are based on overall value intervals of COAs (Salo and Hämäläinen, 2001):

1. *Maximax rule*: Choose the COA $x \in X_{ND}$ with the greatest upper bound of the overall value interval, i.e., $\bar{v}(x) \geq \bar{v}(x'), \forall x' \in X_{ND}$.
2. *Maximin rule*: Choose the COA $x \in X_{ND}$ with the greatest lower bound of the overall value interval, i.e., $\underline{v}(x) \geq \underline{v}(x'), \forall x' \in X_{ND}$.
3. *Central values*: Choose the COA $x \in X_{ND}$ with the greatest mid-point of the overall value interval, i.e., $[\underline{v}(x) + \bar{v}(x)] \geq [\underline{v}(x') + \bar{v}(x')], \forall x' \in X_{ND}$.
4. *Minimax regret*: Choose the COA $x \in X_{ND}$ with the smallest maximum regret, which indicates the largest difference between $v(x)$ and the value of other COAs, i.e., $\max_{x^* \in X_{ND}}[\bar{v}(x^*) - \underline{v}(x)] \leq \max_{x^* \in X_{ND}}[\bar{v}(x^*) - \underline{v}(x')], \forall x' \in X_{ND}$.

Out of the above four, Salo and Hämäläinen (2001) recommend the use of central values and minimax regret. However, these decisions rules should never be followed without discretion. The use of decision rules to compare alternatives with overall value intervals is somewhat counter-intuitive because they reduce the amount of information in the model. Furthermore, no conclusive statements can be made on the suitability of decision rules in specific contexts (Salo and Hämäläinen, 2010). Therefore, a sound approach would be to use multiple decision rules in parallel to establish an order of preference for the non-dominated COAs. The final decision as to which COA will be used in practice should always be made by the AOC, which can utilize additional information, e.g., geographical restrictions and positions of blue and red forces. Figure 4 summarizes the decision making process in its entirety.

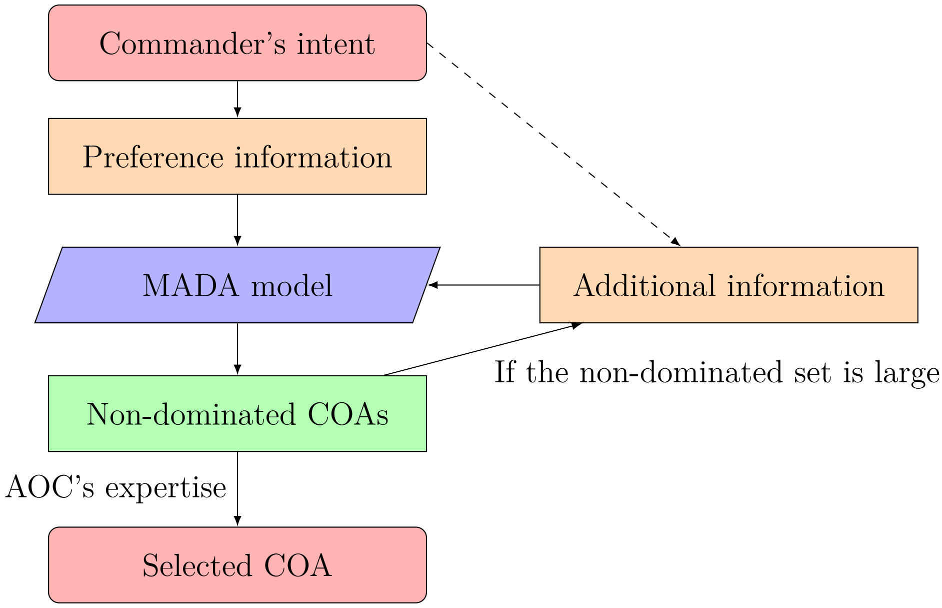

Figure 4: Flowchart of the decision making process.

3.5 Core Index

Even if the set of non-dominated COAs is large, it might be possible to comment on the effectiveness of individual TTPs. We accomplish this by applying robust portfolio modeling methodology to the problem of selecting an effective COA. Robust portfolio modeling is the application of preference programming methods to portfolio selection problems (Liesiö et al., 2007). In portfolio selection, the decision maker aims to choose an effective combination of several project proposals. This problem description is not dissimilar to that of COA selection. Indeed, we can consider COAs to be portfolios consisting of four projects, which are TTPs in this case. Liesiö et al. (2007) introduce an index that can be used to convey the preferability of an individual project. This index is called the core index and it describes how often a project appears in the non-dominated portfolios.

Let $T = \{L, M, C, LM, LC, MC, LMC\}$ be the set of available TTPs in air combat. The letters L, M and C refer to launch ranges Fox_{long}, Fox_{med} and $\text{Fox}_{\text{close}}$, respectively. The core index of TTP $t \in T$ with regard to the set of feasible weights S, denoted by $\text{CI}(t, S)$, is defined as

$$\text{CI}(t, S) = \frac{|\{x \in X_{ND}(S) \,|\, t \in x\}|}{|X_{ND}(S)|},$$

where $X_{ND}(S)$ is the set of non-dominated COAs with regard to S and $|\{\cdot\}|$ denotes the number of COAs in the respective set (Liesiö et al., 2007). Here, $t \in x$ indicates that the TTP t appears at least once in the COA x. If a TTP is included in all non-dominated COAs, its core index is 1 and it is called a core TTP. Conversely, if a TTP does not appear in any of the non-dominated COAs, its core index is 0 and it is called an exterior TTP. TTPs that do not fit into either of these categories are called borderline TTPs. The core index of a borderline TTP is strictly greater than zero but less than one (Liesiö et al., 2007).

The core or exterior status of a TTP remains unchanged when additional preference information is elicited (Liesiö et al., 2007). By definition, a non-dominated COA must include every core TTP and can not include any exterior TTP. Moreover, if we obtain complete preference information, the COA with the greatest overall value also includes all previously identified core TTPs. Therefore, the MADA model can provide decision support even if the number of non-dominated COAs is large by recommending TTPs with a high core index.

3.6 Implementation

The MADA model developed in this thesis is implemented in MATLAB R2020a. Microsoft Excel spreadsheets are used as input and output files for convenience. The input spreadsheet defines how the COAs are labeled, what attributes are used to compare them, and how the COAs are ranked with respect to each attribute. All of the experiments in Section 4 use the same labeling of COAs, i.e., the COA $x^j \in X = \{x^1, x^2, ..., x^{2401}\}$ always represents the same combinations of TTPs of four blue flights. The attributes and the associated rules remain as described in Section 3.1. The rules are used to established a ranking of COAs with respect to each of the three attributes. These rankings are converted into scores $v_i(x_i) \in [0, 1]$ in MATLAB. Different forms of the single attribute value function v_i are examined in the experiments.

The output spreadsheet contains overall value intervals which are computed with the additive value function (1) according to the given preference information. The preference information can be modified in MATLAB to observe its effect on the overall value intervals and dominance relations of COAs. As discussed in Section 3.4, the pairwise dominance relations are established by solving linear optimization problems, specifically with the *linprog* function from MATLAB's Optimization Toolbox. The MATLAB program formats all relevant data into the output spreadsheet, including a list of non-dominated COAs, which can then be sorted according to the decision rules introduced in Section 3.4, if necessary.

4 Experiments

In this section, the MADA model is tested with varying forms of incomplete preference information. The effect of the preference information on the number of non-dominated COAs is of particular interest. First, we examine a scenario where the incomplete preference information is given as an ordinal ranking of the relative importance of attributes. The shape of the single attribute value function is also varied. In the second experiment, we test the model with attribute weight interval restrictions that converge towards complete preference information, i.e., exact point weights. The exact point weights, towards which the weight intervals converge, is varied and the results are compared. A linear single attribute value function is used invariably throughout the second experiment.

4.1 Experiment I: Preference Order and Attribute Scoring

4.1.1 Effect on the number of non-dominated COAs

The first scenario to be examined involves incomplete preference information in the form of ordinal statements, specifically weak rankings (see, e.g., Salo and Hämäläinen, 2010). Weak rankings of attribute weights are of the form $w_i \geq w_j$. Here, two weak rankings are used to establish a preference order of attributes, e.g., $w_{PK} \geq w_{PS} \geq w_{EM}$. With three attributes, there exist a total of six unique preference orders. The overall value intervals and dominance relations of COAs are computed for the six preference orders. The sets of feasible weights that correspond to these preference orders are visualized in a ternary plot in Figure 5.

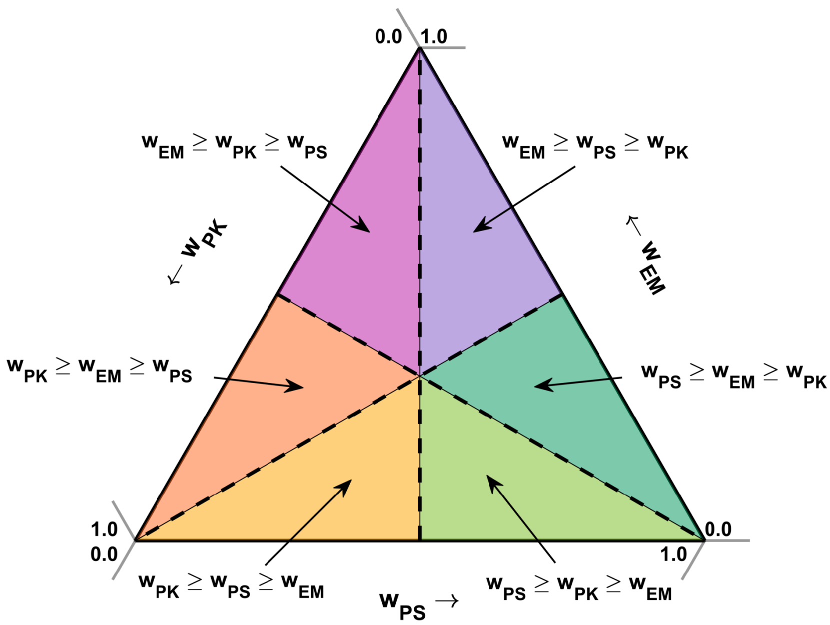

Figure 5: Ternary plot of the feasible regions defined by ordinal preference information.

In addition to the varying preference statements, we study the effect of the shape of the single attribute value function v_i on the number of non-dominated COAs. In our experiments, the measurement levels x_i exist only in the form of ordinal rankings. The simplest shape of the single attribute value function in this case is linear, where the COA on rank r receives a score $v_i(x_i)$ of $1 - \frac{r-1}{m-1}$ with $m = 2401$ being the number of COAs. Thus, the least preferred (rank 2401) COA with respect to attribute i receives a score of zero and the most preferred (rank 1) COA a one. In addition to linear scoring, we test a categorical single attribute value function. In categorical scoring, the COAs are divided into c categories based on their ranking. The categories are then scored linearly, i.e., all COAs in category j receive a score of $1 - \frac{j-1}{c-1}$, where category 1 is the most preferred and category c is the least preferred. Thus, the difference in score between two neighbouring categories is $\frac{1}{c-1}$. Figure 6 visualizes three distinct shapes of the single attribute value function, corresponding to linear and categorical scoring.

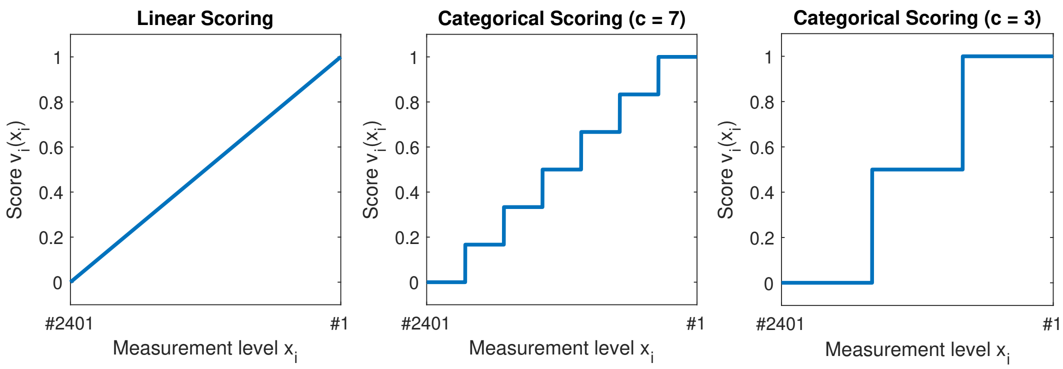

Figure 6: Three shapes of the single attribute value function. c denotes the number of categories in categorical scoring. The measurement levels x_i are ordinal rankings, where #1 denotes the best and #2401 the worst COA with respect to attribute i.

Table 3 presents the results of varying the preference order of attributes and the shape of the single attribute value function. Computing the overall value intervals and dominance relations of the COAs for a given preference order takes around 15 seconds on average with MATLAB R2020a running on a desktop computer with an Intel Core i5-9600K processor. The first observation to be made is that employing categorical scoring results in a constant number of non-dominated COAs. By contrast, when linear scoring is used the number of non-dominated COAs appears to fluctuate around the arithmetic mean of approximately 12.7. The preference orders $w_{PS} \geq w_{PK} \geq w_{EM}$ and $w_{EM} \geq w_{PS} \geq w_{PK}$ yield the most and least non-dominated COAs, respectively. There are no clear deviations in the results. However, there is a clear change in the behaviour of the MADA model when the shape of the single attribute value function is altered from linear to categorical. Unsurprisingly, the number of categories c seems to correlate with the number of non-dominated COAs. Because the number of non-dominated COAs is relatively low in this experiment, it is not necessary to analyse the core indices of individual TTPs.

Table 3: Number of non-dominated COAs with ordinal preference information. The weights of attributes probability of kill, probability of survival and efficiency of missiles are denoted by w_{PK}, w_{PS} and w_{EM}, respectively.

Preference information	Linear scoring	Categorical scoring (c=7)	Categorical scoring (c=3)
$w_{PK} \geq w_{PS} \geq w_{EM}$	11	3	59
$w_{PK} \geq w_{EM} \geq w_{PS}$	12	3	59
$w_{PS} \geq w_{PK} \geq w_{EM}$	18	3	59
$w_{PS} \geq w_{EM} \geq w_{PK}$	13	3	59
$w_{EM} \geq w_{PK} \geq w_{PS}$	15	3	59
$w_{EM} \geq w_{PS} \geq w_{PK}$	7	3	59

To gain an understanding of why categorical scoring leads to fixed numbers of non-dominated COAs, we study the three non-dominated COAs provided by scoring based on a division into seven categories. These three COAs are labeled as x^{558}, x^{939} and x^{957}. Inspecting the overall value intervals of x^{558}, x^{939} and x^{957} clarifies the situation: All three COAs have an overall value interval of $[1, 1]$. Their overall value is the maximum, i.e., 1, for any set of attribute weights. In this thesis, COAs that obtain a score of 1 with respect to each attribute are referred to as globally optimal COAs. Thus, the COAs x^{558}, x^{939} and x^{957} are globally optimal when categorical scoring with $c = 7$ is used. The preference information has no effect on the number of non-dominated COAs in this case, because the globally optimal COAs dominate all other alternatives.

Globally optimal COAs result naturally from the use of ordinal measurement levels x_i combined with the categorical single attribute value functions. For example, let us consider a division of COAs in to three categories with respect to each attribute. If the COA x^* is ranked in the top third with respect to every attribute, the score $v_i(x_i^*)$ is equal to one for every i. Thus, the lower and upper bounds of the overall value interval of x^* are also equal to one, i.e., $\underline{v}(x^*) = \bar{v}(x^*) = 1$. Let us now continue the previously discussed case of $c = 7$ and inspect the globally optimal COAs x^{558}, x^{939} and x^{957}. The TTPs constituting these three COAs are listed in Table 4 alongside the measurement levels of the COAs. The flights of a COA are labeled positionally with numbers 1, 2, 3 and 4, where 1 refers to the leading flight and 4 to the last flight. The measurement levels x_i are attribute-specific rankings, where #1 denotes the best and #2401 the worst COA with respect to attribute i. With 2401 COAs, the top seventh consist of the $2401/7 = 343$ best COAs.

The globally optimal COAs x^{558}, x^{939} and x^{957} are ranked in the top seventh with respect to every attribute. These three COAs balance the probability of killing the target and the probability of surviving while maintaining a good missile efficiency, at least from the perspective of the MADA model developed in this thesis. We

Table 4: Globally optimal COAs when categorical scoring with $c = 7$ is used. The TTPs are labeled with combinations of the letters L, M and C, which denote the launch ranges Fox_{long}, Fox_{med} and $\text{Fox}_{\text{close}}$, respectively. The measurement levels of COAs are ordinal rankings, where #1 denotes the best and #2401 the worst COA.

COA	Constituting TTPs				Measurement levels		
	1	2	3	4	PK	PS	EM
x^{558}	LM	LC	M	LC	#180	#316	#340
x^{939}	M	MC	LM	L	#321	#241	#128
x^{957}	M	MC	LMC	LC	#237	#145	#110

will not attempt to evaluate these COAs from a practical point of view, as that would require considerable expertise in the field of air combat. However, we can still try to recognize patterns in the TTPs of the non-dominated COAs. Indeed, in all three COAs, the TTPs of the first and third flights include Fox_{med}, the TTP of the second flight includes $\text{Fox}_{\text{close}}$ and the TTP of the fourth flight includes Fox_{long}. No conclusive statements can be made based on this observation, but it is interesting nonetheless.

4.1.2 Alternative scoring functions

In order to further understand the effects of different shapes of the single attribute value function, we inspect the set of non-dominated COAs with no preference information on the relative importance of attributes. More specifically, the behaviour of the number of non-dominated COAs with categorical scoring is studied as the number of categories c is increased. Two examples have already been presented, namely $c = 3$ and $c = 7$. The number of non-dominated COAs is plotted as a function of c in Figure 7. For comparative purposes, the number of non-dominated COAs is 52 if linear scoring is used and no preference information is given.

At first, the number of non-dominated COAs decreases rapidly as the number of categories is increased. At $c = 2$ the number of non-dominated COAs is 260. It reaches a minimum of 1 at $c = 8, 9, 10$ and then begins to increase. At the minimum point, the single non-dominated COA is the previously inspected x^{957}. The increase of non-dominated COAs is not monotonic because we are dealing with the dominance relations of a discrete set of decision alternatives. The number of non-dominated COAs is reduced temporarily when an increase in the number of categories leads to a decrease in the scores of some of the previously non-dominated COAs. We can expect the number of non-dominated COAs to eventually reach 52 because the categorical single attribute value function approaches the linear function as $c \to \infty$.

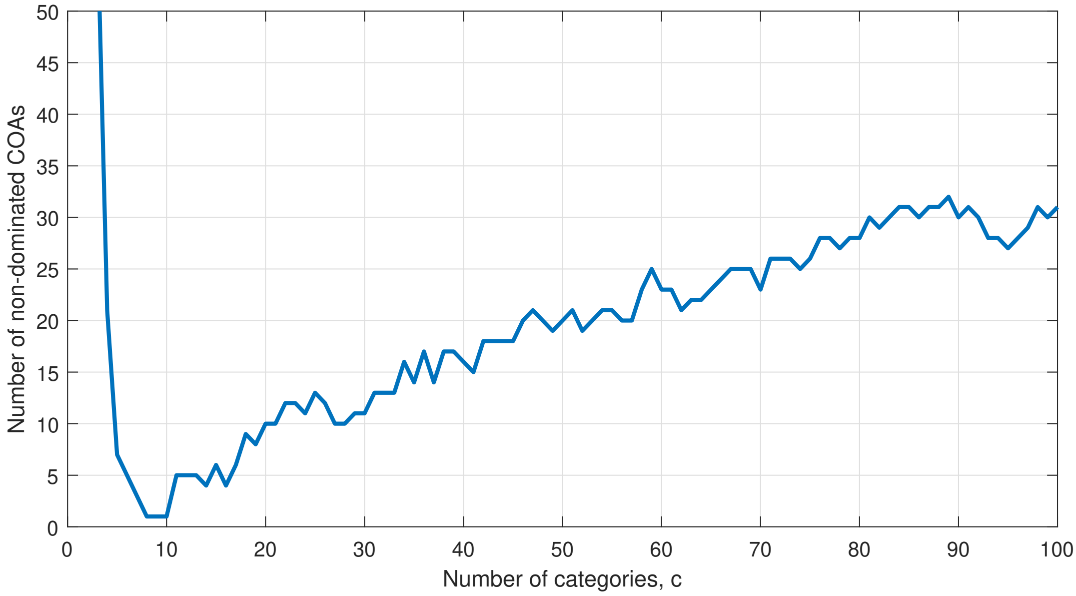

Figure 7: Number of non-dominated COAs as a function of the number of categories, when categorical scoring is used and no preference information is given on the relative importance of the attributes PK, PS and EM.

4.2 Experiment II: Convergence of Weight Intervals

4.2.1 Effect on the number of non-dominated COAs

The second experiment studies the behaviour of the set of non-dominated COAs as the set of feasible weights is restricted progressively. More specifically, the preference information is given in terms of attribute weight interval restrictions that convergence towards exact point weights. The scores of COAs are determined by the linear single attribute value function, as it does not produce any globally optimal COAs. Thus, the incomplete preference information actually has an effect on the non-dominated COAs. The experiment is performed with four different points of convergence: $(w_{PK}, w_{PS}, w_{EM}) = (1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. The first three points correspond to a convergence towards the attributes PK, PS and EM, respectively. The last point corresponds to a convergence towards the center of the weight region. In addition to examining the number of non-dominated COAs, we study the core indices of individual TTPs, defined in Section 3.5.

The linear optimization problem (6) for determining pairwise dominance relations contains the lower bound $w_i \geq 0$ for the weights of all attributes i. The convergent attribute weight intervals are implemented by defining an increasing lower bound $w_i \geq l$, $l \in [0, 1]$ for the appropriate weight. For example, when we perform the experiment with the point of convergence $(w_{PK}, w_{PS}, w_{EM}) = (1, 0, 0)$, the appropriate constraint is $w_{PK} \geq l$. The lower bound l is increased with a step size of $\epsilon = 0.005$ on every iteration. When the point of convergence is $(w_{PK}, w_{PS}, w_{EM}) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, the lower bound is applied to all three attribute weights. In this case, the step size must be reduced to one third of the original ϵ because it is added to three lower bounds instead of just one. Figure 8 provides visual examples of the convergent weight interval experiment. In the examples, the lower bounds are $w_{PK} \geq l = 0.3$ for convergence towards PK and $w_i \geq l/3 = 0.1\,, i = PK, PS, EM$ for convergence towards the center of the weight region. The area of an equilateral triangle is proportional to the square of the length of its sides. Thus, the area of the feasible region S is $(1 - 0.3)^2 = 49\%$ of the entire weight region in both examples.

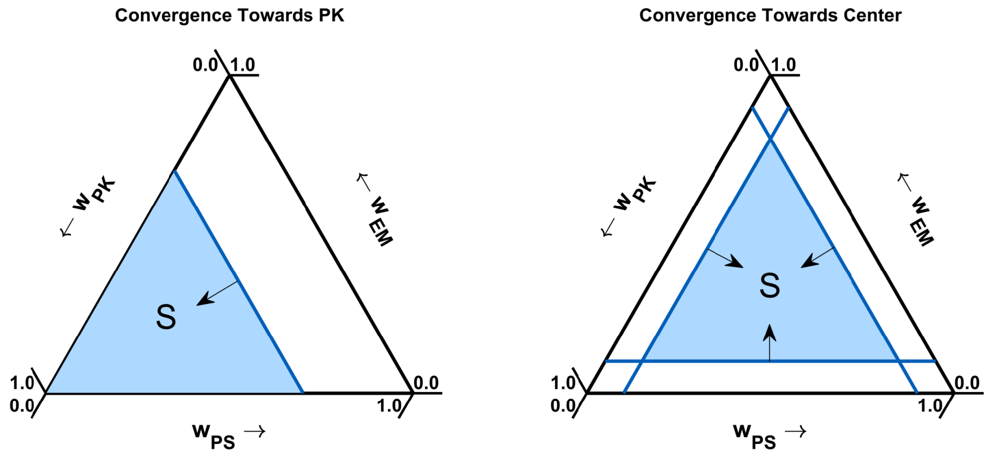

Figure 8: Examples of the second experiment visualized in ternary weight plots. The restrictions are $w_{PK} \geq 0.3$ for the left plot and $w_i \geq 0.1\,, i = PK, PS, EM$ for the right plot. The arrows indicate the convergence direction of the restrictions.

The behaviour of the number of non-dominated COAs is studied with respect to the area of the feasible region. The area is given is terms of percentages of the unconstrained weight region and 0% is defined to indicate the final singleton weight point. Figure 9 presents the the progression of the number of non-dominated COAs for the alternative points of convergence mentioned previously. In this experiment, the computation times of determining the non-dominated COAs ranged from seconds to just over a minute, depending on the size of the feasible region.

The number of non-dominated COAs behaves similarly in the cases of convergence towards the three corners of the weight region, i.e., $(w_{PK}, w_{PS}, w_{EM}) = (1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$. However, when the point of convergence is set to be in the center of the weight region, the number of non-dominated COAs decreases more rapidly. When the area is decreased to 40% of the entire weight region, there are only three non-dominated COAs left. By contrast, the number of non-dominated COAs at this point is 19, 23 and 17 for convergence towards PK, PS and EM, respectively. The faster decrease for convergence towards the center implies that there are only a few COAs that balance the three attributes exceptionally well. The single non-dominated COA for the exact weights $(w_{PK}, w_{PS}, w_{EM}) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is x^{957}, which was identified in the first experiment as globally optimal in the case of categorical scoring with $c = 7$.

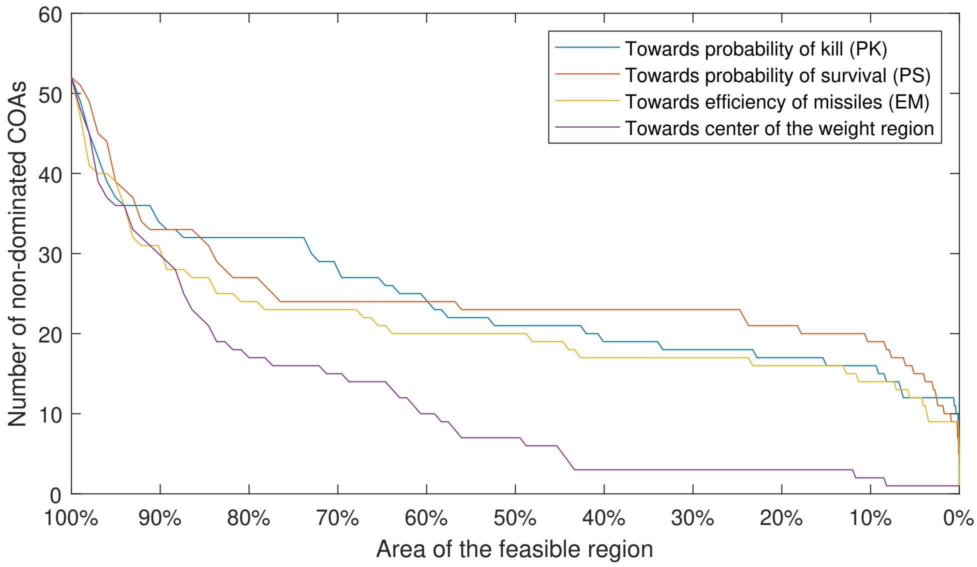

Figure 9: Number of non-dominated COAs as a function of the area of the feasible region with different points of convergence for the weight intervals.

4.2.2 Effect on the core indices of individual TTPs

The study of core indices is motivated by the possibility of identifying core TTPs that could be recommended even if the number of non-dominated COAs is large. As a reminder, the core index of a TTP $t \in T = \{L, M, C, LM, LC, MC, LMC\}$ indicates the fraction of non-dominated COAs that include the TTP t. Figure 10 presents the values of core indices before any preference information is given. The core indices of all TTPs lie within the interval $[0.4, 0.6]$, which indicates that the non-dominated COAs are balanced relatively well.

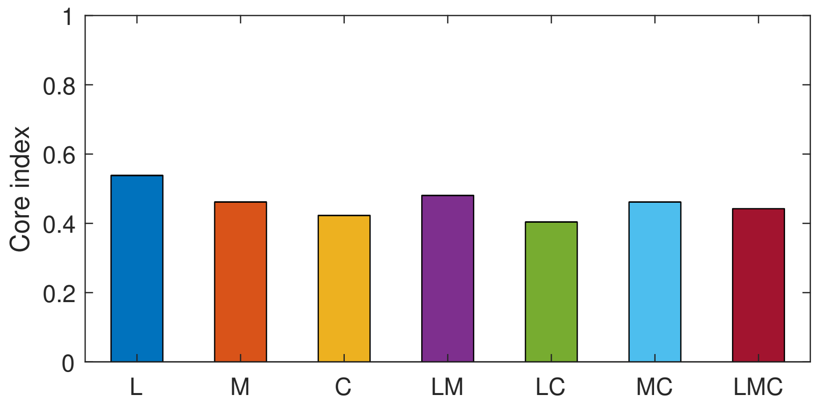

Figure 10: Core indices of TTPs when no preference information is given (area of the feasible region at 100%) and linear scoring is used. The letters L, M and C denote the launch ranges Fox_{long}, Fox_{med} and $\text{Fox}_{\text{close}}$, respectively.

The progression of core indices is visualized in Figure 11. Changing the point of convergence leads to noticeable differences in the core indices of TTPs. When the weight intervals converge towards PK, the use of the launch range $\text{Fox}_{\text{close}}$ is frequent in the non-dominated COAs. TTPs C and MC have the highest core indices at almost every point in the graph. The TTP L, which utilizes the launch range Fox_{long} and the corresponding geometry alternative R_{max}, has the lowest core index. Weighting the probability of kill more heavily than the other attributes should promote the use of TTPs that allow multiple missile launches and minimize the distance to the target. These findings are thus sensible from a practical point of view.

Convergence towards PS produces sensible core indices as well. When PS is weighted more heavily, the core index of TTP L is significantly higher than the rest. The second highest core index belongs to LMC, the TTP that utilizes all three launch ranges. Contrary to the results for convergence towards PK, the core index of C is decreased as the feasible region moves towards PS. This is expected because a flight employing R_{min} geometry compromises the ability to kinematically defeat incoming missiles and the TTP C only uses the closest launch range. Thus, using TTP C leads to a decreased probability of survival for the blue flight. The results for convergence towards EM are more difficult to interpret. In this case, LM has the highest core index until the area is only 11%, at which point the core indices of M and MC surpass it. We will not attempt to form a speculative explanation for these findings.

Convergence towards the center of the weight region results in core indices that clearly deviate from the rest. At the previously mentioned point of 40% area, TTPs L and LM are identified as external. Simultaneously, LMC becomes a core TTP. However, these results follow directly from the rapid decrease in the number of non-dominated COAs. The core index plots remain constant starting from the point of 40% area, where the number of non-dominated COAs is only three, until the number of non-dominated COAs decreases again around 10% area. Moreover, because the number of non-dominated COAs is low, core indices do not provide as much additional value to the decision maker. However, we can conclude that the non-dominated COAs consist of TTPs that are sensible with regard to the given preference information.

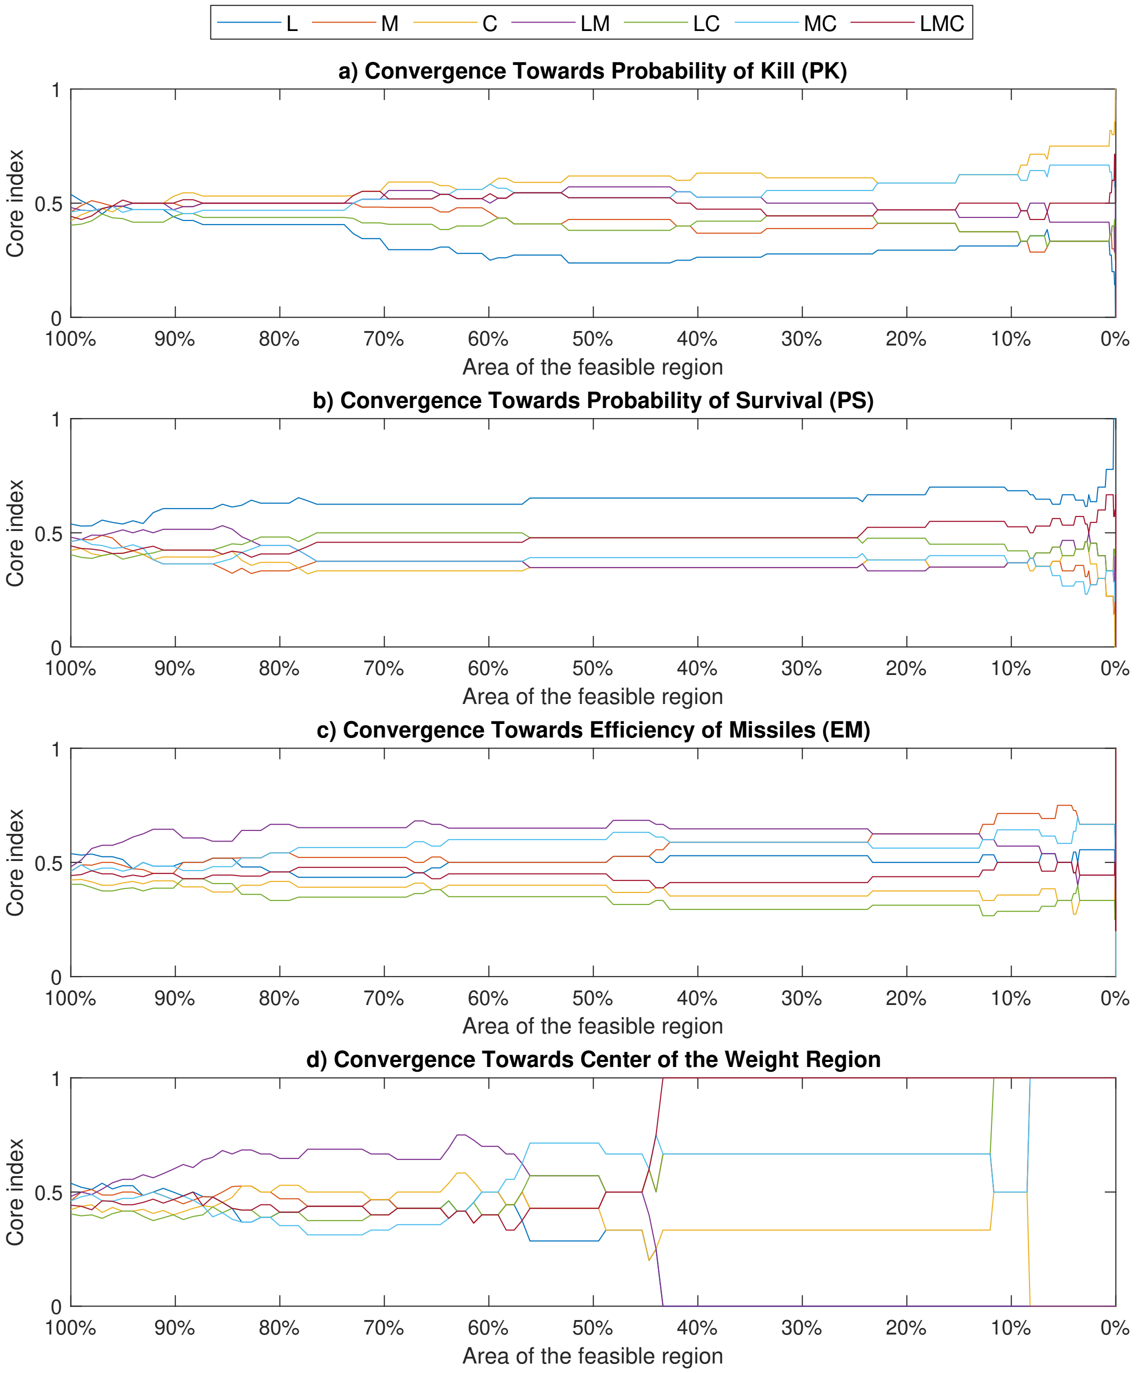

Figure 11: The core indices of TTPs as functions of the area of the feasible region. The results are presented for convergence towards a) probability of kill, b) probability of survival, c) efficiency of missiles and d) the center of the weight region. The TTPs are labeled with combinations of the letters L, M, and C, which denote the launch ranges $\mathrm{Fox}_{\mathrm{long}}$, $\mathrm{Fox}_{\mathrm{med}}$ and $\mathrm{Fox}_{\mathrm{close}}$, respectively.

5 Conclusion

In this thesis, we develop a multi-attribute decision analysis model for the selection of air combat tactics. After defining the vital concepts of air combat and discussing the current practices in COA selection, the thesis describes the MADA model and how it can be used to improve the process of selecting a COA. The COAs are evaluated and ranked with respect to three attributes: probability of kill, probability of survival and efficiency of missiles. An additive value function is used to determine overall value intervals for the COAs based on incomplete preference information elicited from the commander's intent. The overall value intervals are compared using the concept of pairwise dominance. We are thus able to identify non-dominated COAs, which are considered effective with respect to the given preference information.

We perform two extensive experiments to study the effect of varying preference information on the set of non-dominated COAs identified by the model. The results provided by the experiments lead us to the three main conclusions. First, incomplete preference information, e.g., weak ranking of attributes, suffices to reduce the number of COAs in consideration to feasible levels. Second, the non-dominated COAs identified by the model consist of TTPs that are generally sensible with respect to the given preference information. Finally, the shape of the single attribute value function has a significant impact on the set of non-dominated COAs.

The MADA model is capable of identifying a small set of effective COAs from thousands of alternatives. The model also provides robustness throughout the chain of command as it does not rely on verbal interpretations that are prone to unintentional modifications. Nevertheless, in its current state, the model has plenty of room for improvement. First, the model is based on a simplified version of air combat tactics. The simplified TTPs have only two constituting factors, namely geometry and launch range alternatives. Second, many situational factors are left unconsidered. For example, the surface-to-air capability of the opponent can impose geographical restrictions for the blue flights. Lastly, close attention must be paid to selecting single attribute value functions in order to obtain valid recommendations from the model.

The MADA model developed in this thesis fulfills its intended purpose, but it also opens up interesting avenues for future research. The problems associated with the selection of air combat tactics call for collaboration with air combat experts. Feedback from air combat experts is especially beneficial for fine-tuning the single attribute value functions. Simulations of air combat could be used to rule out COAs that do not adhere to the geographical restrictions of the situation. Regardless of the opportunities for further development, we believe that the model can already be considered a successful application of multi-attribute decision analysis and incomplete preference information.

References

AJP-3.3. Allied Joint Doctrine for Air and Space Operations. *Allied Joint Publication-3.3*, NATO Standardization Office, 2016.

A. Arbel. Approximate articulation of preference and priority derivation. *European Journal of Operational Research*, 43(3):317–326, 1989.

CNATRA P-825. All Weather Intercept (AWI), Flight Training Instruction, Advanced NFO T-45C/VMTS. *CNATRA P-825 (Rev. 02-17) PAT*, United States Navy, Chief of Naval Air Training, 2017.

J. S. Dyer and R. K. Sarin. Measurable multiattribute value functions. *Operations Research*, 27(4):810–822, 1979.

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

J. M. Holmes. *The counterair companion: short guide to air superiority for joint force commanders.* Air University Press, Maxwell Air Force Base, Alabama, 1995.

JP 1-02. Department of Defense Dictionary of Military and Associated Terms. *Joint Publication 1-02*, United States Armed Forces, Joint Chiefs of Staff, 2016.

JP 3-30. Joint Air Operations. *Joint Publication 3-30*, United States Armed Forces, Joint Chiefs of Staff, 2019.

R. L. Keeney, H. Raiffa, and R. F. Meyer. *Decisions with multiple objectives: preferences and value trade-offs.* Cambridge University Press, New York, 1993.

J. Kokkala, K. Berg, K. Virtanen, and J. Poropudas. Rationalizable strategies in games with incomplete preferences. *Theory and Decision*, 86(2):185–204, 2019.

J. Liesiö, P. Mild, and A. Salo. Preference programming for robust portfolio modeling and project selection. *European Journal of Operational Research*, 181(3):1488–1505, 2007.

J. Liesiö, P. Mild, and A. Salo. Robust portfolio modeling with incomplete cost information and project interdependencies. *European Journal of Operational Research*, 190(3):679–695, 2008.

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

V. Mattila and K. Virtanen. Ranking and selection for multiple performance measures using incomplete preference information. *European Journal of Operational Research*, 242(2):568–579, 2015.

J. Mustajoki, R. P. Hämäläinen, and A. Salo. Decision support by interval SMART/SWING—incorporating imprecision in the SMART and SWING methods. *Decision Sciences*, 36(2):317–339, 2005.

A. Salo and R. P. Hämäläinen. Preference assessment by imprecise ratio statements. *Operations Research*, 40(6):1053–1061, 1992.

A. Salo and R. P. Hämäläinen. Preference programming through approximate ratio comparisons. *European Journal of Operational Research*, 82(3):458–475, 1995.

A. Salo and R. P. Hämäläinen. Preference ratios in multiattribute evaluation (PRIME)-elicitation and decision procedures under incomplete information. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 31 (6):533–545, 2001.

A. Salo and R. P. Hämäläinen. Preference programming–multicriteria weighting models under incomplete information. In *Handbook of Multicriteria Analysis*, pages 167–187. Springer, 2010.

A. Salo and A. Punkka. Ranking intervals and dominance relations for ratio-based efficiency analysis. *Management Science*, 57(1):200–214, 2011.

T. Tervonen, J. Liesiö, and A. Salo. Modeling project preferences in multiattribute portfolio decision analysis. *European Journal of Operational Research*, 263(1): 225–239, 2017.

M. Weber. Decision making with incomplete information. *European Journal of Operational Research*, 28(1):44–57, 1987.