



Defence Research and
Development Canada

Recherche et développement
pour la défense Canada



Decision rules for pictorial threat classification

David J. Bryant

Defence R&D Canada
Technical Report
DRDC Toronto TR 2009-126
July 2009

Canada

Decision rules for pictorial threat classification

David J. Bryant

Defence R&D Canada – Toronto

Technical Report

DRDC Toronto TR 2009-126

July 2009

Principal Author

Original signed by David J. Bryant, PhD

David J. Bryant, PhD

Approved by

Original signed by Justin Hollands, PhD

Justin Hollands, PhD

Head, Human Systems Integration Section

Approved for release by

Original signed by K. C. Wulterkens

K. C. Wulterkens

for Chair, Document Review and Library Committee

© Her Majesty the Queen in Right of Canada, as represented by the Minister of National Defence, 2009

© Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2009

Abstract

Two experiments examined the use of heuristic and analytic decision strategies in a simulated threat assessment task. Subjects learned to classify targets as friend or foe on the bases of uncertain cues (i.e., characteristics that were probabilistically associated with classification of a target as friend or foe). Subjects were then asked to classify targets that contrasted predictions of several decision rules, including a simple heuristic called Take-the-Best-for-Classification (TTB-C) that uses a single cue to classify targets and the Bayesian classification strategy that is based on formal statistic models. Results of Experiment 1 indicated that the mode of presentation (text versus picture) did not affect the tendency of subjects to use either decision strategy. Results of Experiment 2 indicated that exposure time of pictorial stimuli also did not affect the proportions of subjects employing TTB-C versus the Bayesian strategy. However, an unexpected but very large effect of the target set was observed in the second experiment. This effect may indicate that the interaction of the perceptual salience of cues with the diagnosticity of those cues is a predictor of strategy use. Future research will examine this possibility.

Résumé

On a examiné, dans le cadre de deux expériences, l'utilisation de stratégies de décision heuristique et analytique dans l'évaluation d'une menace simulée. Les sujets ont appris à classer les cibles comme étant amies ou ennemies sur la base de repères incertains (notamment, sur des caractéristiques probabilistes associées à la classification d'une cible amie ou ennemie). On a ensuite demandé aux sujets de classer les cibles qui mettaient en contraste les prédictions de plusieurs règles de décision, incluant la simple approche heuristique « ne garder que le meilleur en vue de la classification » (TTB-C), qui utilise un seul indice pour classer des cibles, et la stratégie de classification bayésienne qui repose sur des modèles de statistiques officielles. Les résultats de l'expérience numéro 1 indiquent que le mode de présentation (un texte par rapport à une image) n'a pas poussé les sujets à utiliser une stratégie de décision plutôt que l'autre. Les résultats de l'expérience numéro 2 indiquent que le temps d'exposition à une stimulation par l'image n'a pas eu d'incidence non plus sur les sujets qui ont utilisé l'approche TTB-C par rapport à la stratégie bayésienne. Toutefois, une conséquence inattendue, mais très importante de l'ensemble des cibles a été observée dans la deuxième expérience. Cette conséquence peut indiquer que l'interaction entre l'évidence perceptuelle des repères et la diagnosticité de ces repères est un prédicteur de la stratégie utilisée. D'autres recherches permettront d'examiner cette possibilité.

This page intentionally left blank.

Executive summary

Decision rules for pictorial threat classification

David J. Bryant; DRDC Toronto TR 2009-126; Defence R&D Canada – Toronto; July 2009.

Threat assessment is a basic military task in which one attempts to rapidly and accurately identify friendly, enemy and neutral forces. To better understand how humans make threat assessment judgments, we performed two experiments in which we manipulated the diagnosticity of available cues to contrast different theoretical models of human judgment.

Previous research [4] has examined four decision procedures that can serve as models of human judgment in a threat classification task. The Take-the-Best-for-Classification heuristic (TTB-C) is based on the premise that the single most valid cue can be used to make accurate threat classification judgments in a task environment in which that cue is highly predictive. Also considered were two variants of the additive rule, the Weighted Additive (WADD) and Unweighted Additive (ADD) rules. These calculate sums of cue values and select the alternative with the highest score. Finally, the Bayesian strategy makes probabilistic inference by applying Bayes' theorem to compute the conditional probabilities of friend and foe classifications for a target given the particular pattern of cue values present.

The aim of the current research was to examine factors that might predict when decision makers will employ heuristics versus a Bayesian decision strategy. The way information is presented in a decision making task often has a significant impact on the way people perform that task. For example, a pictorial format may facilitate use of compensatory additive or Bayesian decision procedures because the visual system has mechanisms to rapidly sum cues or compute probabilities. In addition, pictorial presentation may facilitate automatic (as opposed to deliberate) processing, suggesting that Bayesian computation is not necessarily cognitively demanding. In contrast, Glöckner and Betsch [41] have suggested that heuristics, such as TTB-C, are associated with deliberate processing and should be strongly affected by conditions that increase task demands or reduce available cognitive resources.

The purpose of Experiment 1 was to compare the performance of subjects who learned to classify contacts using textual versus pictorial cues. If, as Glöckner & Betsch [41] suggest, simultaneous availability of all cues is necessary to employ an automatic cue-integration procedure, subjects should be more likely to employ a Bayesian strategy when viewing pictures than when cues are provided textually. Text must be read sequentially, which would favour the use of a deliberate strategy, such as TTB-C. The aim of the Experiment 2 was to determine whether rapid presentation of pictorial stimuli affects the propensity of subjects to use a Bayesian decision strategy. If the Bayesian strategy depends on the recruitment of automatic perceptual mechanisms, it may be more evident in situations in which the deliberate use of an heuristic is made difficult.

Substantial proportions of subjects in Experiment 1 employed the TTB-C, ADD, and Bayesian rules in this experiment, and there was no significant difference between the text and pictorial conditions in the proportions of subjects employing the Bayesian and TTB-C rules. Results of

Experiment 2 indicated that exposure time to pictorial stimuli did not affect the proportions of subjects employing heuristic versus compensatory decision rules. Both heuristic and Bayesian rules are a viable strategy for both textual and pictorial stimuli. However, an unexpected but very large effect of the target set was observed. This suggests that the configuration of cues plays an important role in subjects' choice of a decision procedure. In one set, the uniform was the most predictive cue and it also seemed to be the most perceptually salient. In the second set, the helmet was most predictive but this cue was not considered to be as salient as the uniform. It may be that when a perceptually salient cue, or a cue with a pre-existing association to the classification task, is also the most diagnostic, subjects are able to quickly notice its relation to classification and use a simple rule such as TTB-C. In contrast, when a non-salient cue is most predictive, subjects do not have one cue that immediately stands out as a key predictor and so they tend to look at all cues, which suggests a compensatory and analytic decision rule.

Sommaire

Decision rules for pictorial threat classification

David J. Bryant; DRDC Toronto TR 2009-126; R & D pour la défense Canada – Toronto; Juillet 2009.

L'évaluation des menaces est une tâche militaire fondamentale où l'on tente de reconnaître, rapidement et avec précision, des forces amies, ennemies et neutres. Afin de mieux comprendre comment l'être humain s'y prend pour poser des jugements par suite de l'évaluation de menaces, nous avons mené deux expériences dans le cadre desquelles nous avons manipulé la diagnosticité des repères disponibles pour mettre en contraste différents modèles théoriques du jugement humain.

Des travaux de recherche antérieurs [4] ont permis d'examiner quatre processus décisionnels pouvant servir de modèles de jugement humain à l'intérieur d'une tâche de classification de menaces. L'approche heuristique consistant à « ne garder que le meilleur en vue de la classification » (TTB-C) repose sur l'hypothèse qu'un seul repère, hautement valide, peut servir à poser des jugements précis en matière de classification des menaces, dans un environnement opérationnel où ce repère est hautement prédictif. Deux versions de règle cumulative ont également été prises en considération : la règle pondérée cumulative (WADD) et la règle non pondérée cumulative (ADD). Ces règles calculent la somme de la valeur des repères et sélectionnent le mécanisme ayant le pointage le plus élevé. Enfin, la stratégie bayésienne réalise une inférence probabiliste en appliquant le théorème de Baye pour calculer les probabilités conditionnelles des classifications amies et ennemies d'une cible, compte tenu de la configuration particulière de la valeur des repères.

La présente recherche visait à examiner les facteurs pouvant prédire quand des décideurs allaient utiliser l'approche heuristique par rapport à une stratégie de décision bayésienne. La façon dont l'information est présentée dans une tâche de prise de décision a souvent des incidences importantes sur la façon d'exécuter cette tâche. À titre d'exemple, une présentation par images peut faciliter l'utilisation de procédures cumulatives compensatoires ou de décision bayésienne parce que le système visuel possède des mécanismes qui permettent d'additionner rapidement des repères ou de calculer des probabilités. De plus, la présentation par images peut faciliter le traitement automatique (par opposition à délibéré), suggérant que le calcul bayésien n'est pas nécessairement exigeant au plan cognitif. À l'opposé, Glöckner et Betsch [41] ont suggéré que les heuristiques, comme la règle TTB-C, sont associées à un traitement délibéré et qu'elles peuvent être fortement touchées par des conditions qui augmentent les exigences liées aux tâches ou réduisent les ressources cognitives disponibles.

L'expérience numéro 1 visait à comparer le rendement des sujets qui avaient appris à classifier les contacts à l'aide de repères textuels par rapport à des repères pictographiques. Si, comme le suggèrent Glöckner & Betsch [41], la disponibilité simultanée de tous les repères est nécessaire pour utiliser une procédure automatique d'intégration des repères, les sujets auront davantage tendance à utiliser une stratégie bayésienne lorsqu'ils voient des images que lorsque les repères apparaissent sous forme de textes. Un texte doit être lu par séquence, ce qui favorise l'utilisation d'une stratégie délibérée, comme la règle TTB-C. L'expérience numéro 2 visait à déterminer si la

présentation rapide d'une stimulation par l'image touche la propension des sujets à utiliser une stratégie de décision bayésienne. Si la stratégie bayésienne repose sur le recrutement de mécanismes perceptuels automatiques, elle peut s'avérer plus profitable dans des situations où l'utilisation délibérée d'une approche heuristique est difficile.

Un nombre important de sujets de l'expérience numéro 1 ont utilisé les règles TTB-C, ADD et bayésienne. Il n'y a pas eu de différence marquée entre les conditions textuelles et pictographiques dans le nombre de sujets qui ont utilisé les règles bayésienne et TTB-C. Les résultats de l'expérience numéro 2 indiquent que le temps d'exposition à une stimulation par l'image n'a pas eu d'incidence sur les sujets qui ont utilisé l'approche heuristique par rapport aux règles de décision compensatoire. Tant la règle heuristique que bayésienne constitue une stratégie viable pour la stimulation textuelle et pictographique. Toutefois, une conséquence inattendue, mais très importante de l'ensemble des cibles a été observée. Celle-ci suggère que la configuration des repères joue un rôle important dans le choix des sujets d'une procédure décisionnelle. Dans un ensemble, l'uniforme était le repère le plus prédictif et semblait également être le plus évident au plan perceptuel. Dans le deuxième ensemble, le casque était le repère le plus prédictif, mais pas aussi évident que l'uniforme. Il peut arriver que, lorsqu'un repère évident au plan conceptuel ou un repère ayant un lien préexistant avec une tâche de classification est également celui qui présente le diagnostic le plus probant, des sujets soient capables de remarquer rapidement sa relation avec la classification et utiliser une simple règle comme la règle TTB-C. À l'opposé, lorsqu'un repère non évident s'avère hautement prédictif, les sujets n'arrivent pas à voir un repère se démarquer immédiatement comme prédicteur clé. Ils ont alors tendance à regarder tous les repères, ce qui suggère une règle de décision analytique et compensatoire.

Table of contents

Abstract	i
Résumé	i
Executive summary	iii
Sommaire	v
Table of contents	vii
List of figures	ix
List of tables	x
Acknowledgements	xi
Introduction	1
Background.....	1
Analytic versus heuristic decision processes	1
Analytic Processes.....	2
Heuristic Processes.....	2
Decision Procedures for Threat Assessment.....	4
The Take-the-Best-for-Classification (TTB-C) Heuristic	4
Compensatory Rules.....	6
Bayesian Procedures.....	7
Results of Bryant (2007) [4].....	9
Research Question.....	10
Representational Format	10
Automatic versus Deliberate Classification.....	11
Overview of the Experiments.....	14
Experiment 1	15
Method.....	15
Subjects	15
Materials	15
Design	16
Procedure.....	18
Results.....	19
Training Session	19
Classification Strategy.....	20
Response Time	22
Discussion.....	22
Experiment 2	24
Method.....	24
Subjects	24

Materials	25
Design	26
Procedure	28
Results.....	28
Training Session	28
Classification Strategy	30
Response Time	30
Discussion.....	31
Conclusion.....	33
References	34
Annex A ..	39
Annex B: The Maximum Likelihood Method (MLM) for Bayesian Classification.....	45
List of symbols/abbreviations/acronyms/initialisms	49
Distribution list.....	2

List of figures

Figure 1: The Take-the-Best-for-Classification (TTB-C) Heuristic.....	5
Figure 2: Weighted Version of the Additive Rule for Classification	6
Figure 3: A Bayesian Network representation of the task used by Bryant [4]	8
Figure 4: Examples of pictorial stimuli	16
Figure 5: Classification Accuracy by Block in the Training Sessio.....	20
Figure 6: Examples of stimuli	25
Figure 7: Classification Accuracy by Block in the Training Session.....	30
Figure A1: The Take-the-Best-for-Classification Strategy	41
Figure A2: Bayesian Network for Friend-Foe Classification.....	42
Figure A3: The Weighted and Unweighted Additive Rules for Classificatio	44

List of tables

Table 1: Relative Frequencies of Cue Values for Friend and Foe Contacts.....	17
Table 2: Predicted Responses to Contact Types by Hypothesized Strategies	18
Table 3 Predicted Accuracy Levels by Contact Type	19
Table 4: Number of Subjects Classified as Using Hypothesized Decision Strategies	21
Table 5: Use of Decision Strategies in First and Second Conditions	22
Table 6: Relative Frequencies of Cue Values for Friend and Foe Contacts.....	26
Table 7: Predicted Responses to Contact Types by Hypothesized Strategies	27
Table 8: Predicted Accuracy Levels by Contact Type	28
Table 9: Number of Subjects Classified as Using Hypothesized Decision Strategies	31
Table A1: Relative Frequencies of Cue Values for Friend and Foe Contacts.....	39
Table A2: Predicted Responses to Contact Types by Hypothesized Strategies	40
Table A3: Summation of Cue Evidence for Cue Pattern 1,2,1,1 by the WADD and ADD Strategies	43
Table B1: Relative Frequencies of Cue Values for Friend and Foe Contacts	45
Table B2: Examples of Different Types of Items Used to Assess Decision Strategies.....	46

Acknowledgements

The experiments reported here were completed with the invaluable assistance of Elaine Maceda and Sabrina Kanani.

This page intentionally left blank.

1 Introduction

1.1 Background

Threat assessment is a basic military task in which one attempts to rapidly and accurately identify friendly, enemy and neutral forces. This task is performed in a wide range of operational environments, where the forces being assessed can be vehicles and platforms, such as aircraft, or individual persons, such as uniformed soldiers, guerrilla fighters, and civilians. In most operational settings, numerous data sources are available and can be considered to identify targets correctly. These data can include outputs of sensors (e.g., radar, electronic emission detection systems, Interrogate-Friend-or-Foe or IFF interrogation systems) or environmental data perceived directly by individuals.

Formally, threat assessment can be viewed as categorization based on uncertain cues. An operator or soldier must categorize one or more entities in the environment as friend, foe, or neutral using available perceptual cues and sensor information. This can be a difficult task when no cue provides certain classification. Uncertainty results from flawed human perception and ambiguous sensor data, as well as inherent uncertainty as to which characteristics are important, or diagnostic, to the identity of targets [1] [2]. This is especially true in asymmetric environments in which the enemy uses diverse equipment and attempts to blend into civilian populations, as well as in coalition operations in which allies may use different, unfamiliar equipment. Operating afield in unfamiliar nations can leave soldiers with limited knowledge of the kinds of information needed to distinguish neutral from potentially hostile factions.

To understand how humans make threat assessment judgments, we performed two experiments examining the cognitive processes people employ to classify targets as friend or foe. In these experiments, we manipulated the diagnosticity of available cues in order to contrast different classes of theoretical models of human judgment. Following on previous research [3] [4], the experiments assessed the analytic and heuristic processes employed by participants to categorize entities on the basis of probabilistic cues.

1.2 Analytic versus heuristic decision processes

A major question in decision research concerns the value of heuristics in relation to analytic procedures [5]. Whereas an heuristic is a simple decision procedure that offers the potential to quickly and easily resolve a specific decision, an analytic procedure promises an optimal solution at the cost of extensive computation and time. Both approaches have received empirical support and both can be successful approaches to decision making, and this precludes a simple conclusion that one is an inherently better approach than the other. Nevertheless, when considering a specific decision making task, one approach may be better when assessed in terms of not just accuracy but also in terms of the economical use of cognitive resources.

1.2.1 Analytic processes

The analytic approach is based on the premise that human decision making can be modeled in terms of formal processes predicted by *normative theories* of probability and logic. These theories have been popular in the development of support for military decision makers, perhaps because there is often an assumption that good decision making follows a rational approach in which decisions are based on expected outcomes and there is an attempt to select a course of action that will yield the optimal outcome.

Normative theories explain human judgment in terms of explicitly computable processes to take in information, code it symbolically, manipulate these symbolic representations, and generate some output. Analytic decision procedures based on these theories require some kind of formal comparison among decision alternatives using deliberate, procedural rules that quantify those alternatives. Numerous specific procedures for comparing alternatives are known, most of which can be computationally modeled. Many, for example, are based on Bayesian statistics and evaluate options in terms of base rates for different hypotheses and probabilities of the accuracy of different observations [6]. Other analytic strategies include subjective expected utility analysis, single feature difference, and elimination by aspects (see [7]).

To make a judgment, an analytic procedure generally specifies a number of dimensions along which to compare alternatives. Typically, these computations are based on compensatory algorithms in which all dimensions are weighted [8]. A popular general form is the linear compensatory model, which involves the computation of an overall score for each decision alternative based on the sum of relevant dimension values for each alternative, weighted by each dimension's importance [9]. Because the score of each alternative is based on all known dimensions, effects of large and small dimension values can compensate one another in determining the overall desirability of the alternative [10].

Analytic decision procedures are popular because they are designed to yield the optimal choice. The downside of such procedures, however, is that they must identify and compare all potential decision alternatives along all relevant dimensions. This means that analytic decision making entails extensive computations, even for fairly simple problems [11]. A comprehensive search for data to allow all comparisons is generally extremely time consuming, if not impossible, in real-world problems, given limitations of human knowledge and cognitive capacity. Moreover, it may not be possible to construct a complete representation of the problem space, including goals, the potential outcome values, and the probabilities of certain actions producing certain outcomes [12].

1.2.2 Heuristic processes

Although analytic models judge what is rational in terms of formal rules of logic and probability, the inherent limitations of human cognition have forced researchers to consider approaches that can be termed "bounded rationality" [13][14]. These approaches are based on the recognition that decision making mechanisms must work within the limits of time, knowledge, and computational power imposed by the situation and the decision maker him/herself [12][15][16]. Although there is a wide range of specific procedures that fall under the general definition of bounded rationality, the use of heuristics is assumed to be key to

implementing successful but cognitively plausible decision models. Heuristics are informal, intuitive strategies that specify simple steps, which are often based on pre-existing knowledge of probabilistic data, and are designed to work under a few general assumptions [17] [18].

The notion of *ecological* rationality leads to a focus in modeling decision making on the interaction of decision making mechanisms with environmental consequences. In line with this focus, Gigerenzer and the Adaptive Behavior and Cognition (ABC) Group [19] have developed the fast-and-frugal heuristic approach which seeks to develop models of cognition that are simultaneously plausible on psychological and ecological grounds, as well as being computationally specific (see [20]). Thus, fast-and-frugal heuristics are computationally simple procedures for making judgments with limited information that have been shown to be accurate and efficient solutions to certain judgment tasks [19][21]. The Take the Best (TTB) heuristic, for example, performs two-alternative choice tasks by determining the single cue dimension that both discriminates options and is believed to have the highest validity based on previous experience or learning (i.e., the cue is believed to offer the greatest conditional probability of indicating the correct choice given the cue's presence) [14][22].¹ In simulation studies with a variety of data sets drawn from psychology, economics, and other fields, TTB performs a choice task as accurately as more computationally intensive linear regression models [14][22]. In addition to achieving comparable accuracy, the TTB consistently exhibits a clear advantage over linear procedures in terms of frugality, consulting, on average, fewer cues and performing fewer computations than linear procedures.

A variety of studies have shown fast-and-frugal heuristics to be accurate and efficient solutions to certain judgment tasks. In simulation studies, TTB performs a choice task, such as choosing the larger of two cities, roughly as accurately as more computationally intensive linear regression models [14], a finding that has been replicated with 19 other data sets drawn from psychology, economics, and other fields and involving comparison along a variety of dimensions [22]. In addition to achieving comparable accuracy, TTB consistently exhibits a clear advantage over linear procedures in terms of frugality, consulting, on average, fewer cues and performing fewer computations than linear procedures. A non-compensatory heuristic such as TTB generally performs well when the task environment is itself structured such that the validity of cues falls off dramatically in a non-compensatory fashion [23].

Fast-and-frugal heuristics such as TTB can also provide plausible models of human decision making in tasks in which subjects are required to use probabilistically predictive cues to select an alternative (e.g., [24][25][26][27]). In these studies, only a subset (albeit a majority in some cases) of subjects can be classified as using TTB. Even under favourable conditions, subjects have frequently been observed to deviate from the principles of fast-and-frugal heuristics. Newell and Shanks [27], for example, found that the order in which subjects search for cues, in particular, often deviates substantially from what would be expected if decision makers were employing a fast-and-frugal heuristic (see also [26]). Often, a significant proportion of subjects seem to use more complex, compensatory procedures in

¹ It is often assumed in studies of fast and frugal heuristics that a person's internal representation of cue validities accurately reflects the ordering of cues according to their predictiveness [19], but this is not necessarily the case. The assumption that subjects in the experiments reported here acquired accurate cue validity knowledge is bolstered by the finding of Bryant [3] that subjects were able to report the relative validity of cues for stimuli similar given extensive threat classification training.

these experiments. As will be discussed in more detail, the propensity of subjects to employ TTB is affected by a range of factors, such as costs imposed on obtaining cues. Thus, it remains an open question as to the extent to which fast-and-frugal heuristics represent a general framework in which to understand human judgment.

1.3 Decision procedures for threat assessment

Given concerns that threat assessment is vulnerable to problems of information overload and uncertainty [1], the fast-and-frugal heuristic approach provides a potentially useful framework in which to study time- and information-stressed decision making. Fast-and-frugal heuristics may be a natural means to manage a heavy information load, in appropriate task environments. This appears to be true specifically for threat assessment, where surveys of experienced operators have indicated that operators do not consider or weigh all available data equally and that they employ decision making procedures that differ from those previously assumed [28].

To explore the potential of fast-and-frugal heuristics to model human threat assessment, Bryant [4] developed a simulated air threat assessment task in which to compare predictions of different decision models. In three experiments, subjects learned to classify simulated aircraft using four probabilistic cues, then classified test sets designed to contrast predictions of several compensatory and non-compensatory heuristics. Various “contacts” (simulated aircraft) were presented on a simulated radar screen for subjects to classify as either friend or foe based on the values of four cues. Each cue value had a specific probability of being associated with friend and foe contacts, with these probabilities determining the cue’s validity in classifying contacts.

To apply the fast-and-frugal heuristic approach to threat assessment, Bryant [3] [4] devised decision procedures specifically for the threat classification task. The following sections describe these in detail.

1.3.1 The Take-the-Best-for-Classification (TTB-C) heuristic

A variant of TTB, called Take-the-Best-for-Classification heuristic (TTB-C), was devised to perform the threat classification task.² Illustrated in Figure 1, TTB-C is based on the premise that the single most valid cue can be used to make accurate threat classification judgments in a task environment in which that cue is highly predictive. Unlike TTB, which chooses between two objects along a single dimension, TTB-C places a single object into one of two categories. Thus, TTB-C is simpler in some respects than TTB but it takes from TTB the basic search concept of locating the single best cue to make its decision.

Given an as-yet-unclassified contact, TTB-C begins by searching for the single most valid cue to serve as the basis for classification. In the experiment described in this report, all cues

² TTB-C is also derivable from the Lexicographic heuristic for two-alternative choice, which specifies the order in which cues are inspected according to a particular system (e.g., alphabetic, numeric). TTB-C is a specific instance of a Lexicographic heuristic in which relative cue validity serves as the system of cue search (see p. 82 in [19]).

associated with contacts will be available, so the most valid cue should be inspected. When the most valid available cue is located, the heuristic assesses which threat class has the greater probability of being true given the value of that cue and makes that threat class the output of the heuristic. The heuristic will be applied here to an experimental task in which subjects make a simplified two-category choice (friend or foe) but the heuristic could apply to threat classification with a larger set of threat classes. With the contact classified, the heuristic terminates. Should no valid cue be found, the heuristic can only guess.

TTB-C, as illustrated here, assumes that there exist one or more cues that have some non-random association to the threat class of contacts and that all, or some subset, of these cues can be inspected by the decision maker. Moreover, the decision maker must have acquired, through experience or training, knowledge of the relative validities of these cues. These, of course, are not minor assumptions but there is sufficient evidence that people can learn cue validities, even if their learning is imperfect [29] [30].

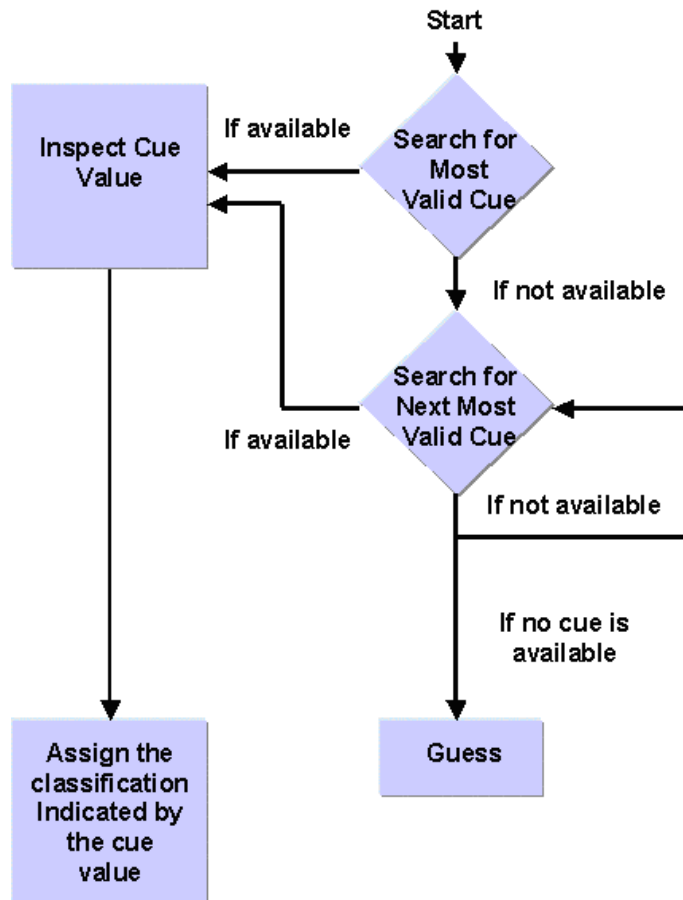


Figure 1: The Take-the-Best-for-Classification (TTB-C) Heuristic

1.3.2 Compensatory rules

Just as TTB-C is an adaptation of the TTB heuristic to the single-choice classification problem, other two-alternative choice decision strategies can be adapted. Among the decision strategies that have been examined are Franklin's Rule and Dawes' Rule. Franklin's rule is a procedure by which a decision maker calculates the sum of cue values weighted by the corresponding cue validities for each alternative and selects the alternative with the highest score [31]. Dawes' rule is similar and calculates the sum of unweighted cue values and selects the alternative with the highest score. Because both Franklin's and Dawes' Rules add bits of evidence for an alternative, they can be termed Additive Rules (for the sake of clarity, these rules will be referred to as the Weighted Additive and Unweighted Additive rules). Both employ all available cues although they do not compute probabilities to reach a decision.

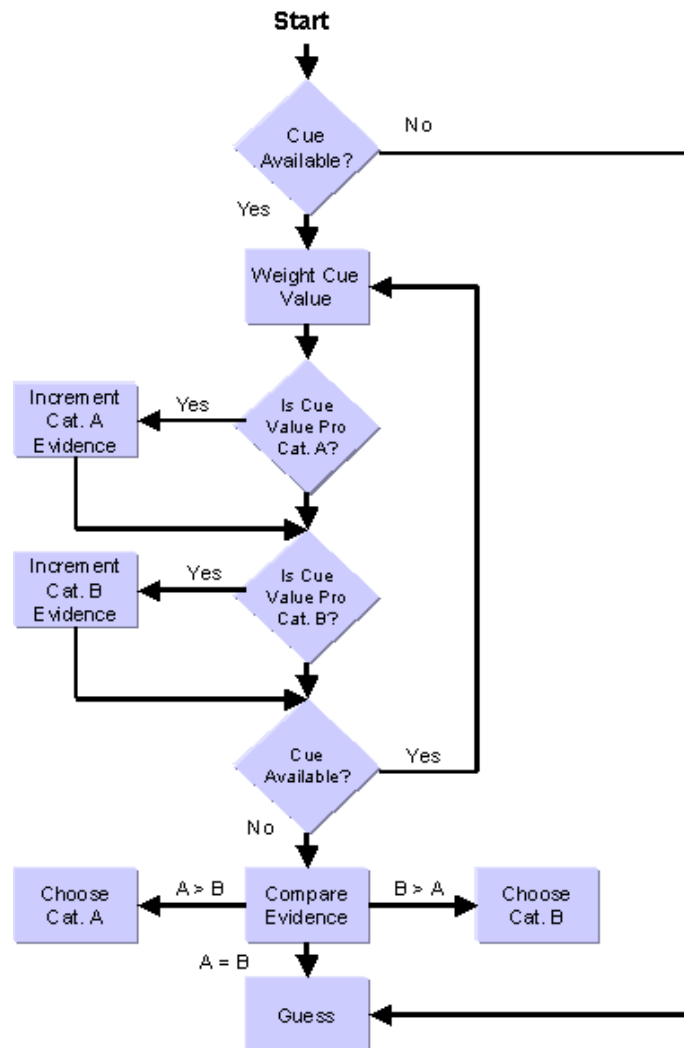


Figure 2: Weighted Version of the Additive Rule for Classification

Versions of the Additive Rules were formulated for the threat classification task. Unlike their progenitors, they do not compare cue values for two alternatives but rather examine each cue

value and assign evidence toward either friend or foe classification, depending on the associations of cue values to threat class. A running sum is maintained and, after all available cues have been inspected, used to place the contact in the friend or foe category. Figure 2 contains an illustration of Weighted Additive (WADD) Rule, which weights cues by their validity, adapted for threat classification. A classification version of the Unweighted Additive Rule (ADD) is performed just as illustrated in Figure 2 but without the weighting step following the selection of a cue. These rules use more information than TTB-C but are more generally useful because their accuracy is not limited to cases where a single cue is highly predictive.

1.3.3 Bayesian procedures

Another way to make a judgment on the basis of probabilistic cues is by means of a “naïve” Bayesian classifier. A naïve Bayes classifier is a system for making probabilistic inference by applying Bayes' theorem with a strong assumption of independence among cues (i.e., the probability with which one cue is associated to a classification outcome is completely unrelated to all other cues). That is, it assumes that the presence of any particular cue is unrelated to the presence of any other cue. In this procedure, a class of object is represented by a base rate (overall probability of an instance of that class occurring) and set of conditional probabilities that specify relationships of attributes to that class. Despite their simplifying assumption, naïve Bayes classifiers often work very well in complex real-world situations. Depending on the precise nature of the probability model, naïve Bayes classifiers can be trained very efficiently in a supervised learning setting. Thus, it is an appropriate model for learning the friend/foe classification used in Bryant [4].




Figure 3: A Bayesian Network representation of the task used by Bryant [4]

A Bayesian classifier can be instantiated by a number of different algorithms that calculate conditional probabilities. It can also be instantiated in a Bayesian network (or a belief network), which is a probabilistic graphical model that represents a set of variables and their probabilistic independencies. Thus, a Bayesian network can represent the probabilistic relationships between threat class (friend or foe) and predictive cues. Given a set of cues, the network can be used to compute the probabilities of the target being a friend or a foe.

A Bayesian network for the friend/foe task is shown in Figure 3. The top node represents the classification of a target as a friend (F). The case of a foe would be represented by the negation of friend (\bar{F}). Four nodes representing characteristics of the target, or cues (C_{1-4}), are connected to it according to their probabilistic association to the class of the target. Thus, each line linking a cue to the classification node is labeled by the conditional probability of the cue occurring given the classification of friend. Considering all cues as a set, the classification node represents the conditional probability of that class being true given the presence of the four linked cues. This is given by the formula (1):

$$P(F|C_1 \cap C_2 \cap C_3 \cap C_4) = \frac{P(C_1|F) \cdot P(C_2|F) \cdot P(C_3|F) \cdot P(C_4|F)}{P(C_1|F) \cdot P(C_2|F) \cdot P(C_3|F) \cdot P(C_4|F) + P(C_1|\bar{F}) \cdot P(C_2|\bar{F}) \cdot P(C_3|\bar{F}) \cdot P(C_4|\bar{F})} \quad (1)$$

Where:

$$P(C_j | \bar{F}) = 1 - P(C_j | F), j = 1 \text{ to } 4.$$

The Bayesian strategy was assumed to compute the conditional probabilities of friend and foe classifications given the particular pattern of cue values for a contact and select the alternative with the higher probability of being the correct classification. This is formally equivalent to the “profile memorization method,” which memorizes which option has the greater conditional probability of being correct for each cue configuration [19]. Martignon and Hoffrage [23] have described this method as the optimal Bayesian method for fitting known data.

1.3.4 Results of Bryant (2007) [4]

Bryant [4] examined the use of various decision procedures in three experiments, in which subjects learned to classify simulated aircraft using four probabilistic cues then classified test sets designed to contrast predictions of several compensatory and non-compensatory heuristics. Overall, results indicated that about half of the subjects who exhibited a classifiable, non-random strategy appeared to use a non-compensatory fast-and-frugal heuristic, but the other half used less frugal compensatory decision rules. Interestingly, the relative proportions of subjects exhibiting responses consistent with the fast-and-frugal heuristic versus other decision rules was largely unaffected by manipulations of time pressure and perceived cue uncertainty.

In a third experiment, Bryant [4] employed a severe time pressure manipulation that allowed subjects only four seconds in which to respond. This time limit made it extremely difficult to inspect all cues and was predicted to strongly favour the use of TTB-C because that decision rule only requires a single cue. The extreme time pressure condition was contrasted with a control condition that afforded sufficient time to examine all cues. Overall, a significant difference in the pattern of assigned decision strategies was observed between the high time pressure and control conditions. When time was severely limited, the majority of subjects employed either TTB-C or a guessing strategy, presumably because these strategies do not require time-consuming inspection of multiple cues. When time was relatively ample, subjects generally preferred strategies that made use of all available cues, either WADD or the Bayesian procedure. This finding indicated that a fairly extreme manipulation of time pressure is required to produce a marked shift in subjects’ preferences for strategy.

2 Research question

Research has determined that people may employ simple heuristics to perform cue-based tasks in some cases but employ a Bayesian strategy in other cases (e.g., [27]). Thus, a question arises as to what factors might predict when decision makers will employ heuristics versus a Bayesian strategy. Several factors have been found to affect whether subjects employ fast-and-frugal heuristics in experiments. The underlying information structure of the task, for example, is a key factor. A heuristic such as TTB is well-suited to a non-compensatory environment but less suited to a compensatory environment. That is, when there are a few highly predictive cues available, TTB can perform very well but when no such cues are available TTB is a poor strategy. Bröder [24] found that subjects were more likely to employ TTB when the task involved cues that were designed to create a non-compensatory environment with one highly weighted cue, rather than when cues created a compensatory environment.

Another factor influencing the use of heuristics is the cost, either in terms of monetary resources or cognitive effort, involved in searching for and obtaining cue information. Bröder [24] found 40% of subjects used TTB when the cost of cue information was relatively low and 60% when the cost was relatively high. Newell and Shanks [27] have also replicated the finding that a greater proportion of subjects employ TTB when cue costs are relatively higher.

Although some factors have been found to affect the likelihood of adopting a fast-and-frugal heuristic, other factors have been unexpectedly ineffective in changing subjects' decision strategies. Manipulations of time pressure have failed to substantially increase use of fast-and-frugal heuristics in all but the most extreme cases [4][24]. Similarly, manipulation of the experimental instructions [4], degree of knowledge of underlying cue validities [27], and number of cues [26] have failed to affect subjects' selection of decision strategy. Even simplifying the task environment and emphasizing the non-compensatory structure of a task does not dramatically increase the proportion of subjects using TTB [26].

2.1 Representational format

The way information is presented in a decision making task often has a significant impact on the way people perform that task (e.g., [32]). This is true for tasks involving reasoning about probabilities. Gigerenzer and Hoffrage [33], for example, found that people were better able to understand and use Bayesian conditional probabilities when problems were phrased in terms of absolute frequencies as opposed to abstract probabilities. It is likely that representational format affects the type of decision process one uses. Certainly, changes in format can affect the ease with which cues can be searched, which in turn affects the usefulness of compensatory procedures [34].

Bröder and Schiffer [25] have suggested that heuristics are sequential and thus best suited to situations in which information is received or processed in a sequential fashion – for example, through text. When pictorial stimuli are available, however, data are available simultaneously and can be processed in parallel. According to Bröder and Schiffer [25], heuristics are not suited to pictorial stimuli and, instead, people can more easily employ feature matching or

other perceptual, parallel process. To examine this possibility, Bröder and Schiffer [25] compared peoples' use of different strategies in a four-cue judgment task similar to that of Bryant [4] in which materials were presented either textually or in pictures. They observed greater use of a fast-and-frugal heuristic for textual items and greater use of a compensatory weighted additive rule for the pictorial items. This result supports the notion that people can use more complex, compensatory procedures when the presentation format allows parallel processing of data. It must be noted, however, that Bröder and Schiffer [25] also reported results in which no differences in strategy use were observed between textual and pictorial stimuli. They hypothesize that pictorial stimuli must comprise holistic images, not just separate pictures of features, to allow compensatory processing.

Pictorial format may facilitate use of compensatory or Bayesian decision procedures because the visual system has mechanisms to rapidly compute probabilities. Such mechanisms allow the visual system to infer what is in the environment with the imperfect sensory data it receives. Although the human visual system is a complex pattern recognition system, the images received by the eye are generally ambiguous because they are two-dimensional projections of a three-dimensional world. Thus, similar objects can create very different images, whereas very different objects can give rise to similar images [35]. Despite this ambiguity, however, the visual system does an excellent job of segregating and identifying discrete objects in the environment. Thus, Bayesian Perception Theory provides a model of visual perception based on a statistical analysis of sensory data [36].

The basic premise of Bayesian Perception Theory is that the visual system interprets sensory data to infer the most probable state of the world [35]. Essentially, the visual system computes the conditional probability of an hypothesized state of the world being true given the observed visual image on the basis of the priori probability of that state of the world (base rate), the posterior probability of the pattern of sensory data being encountered, and the probability of the sensory pattern being produced by the hypothesized state [37]. In this theory, visual perception depends on the integration of sensory cues according to a formal inference algorithm that adheres to Bayes' Theorem of conditional probabilities [38].

Substantial evidence supports the idea of Bayesian perception theory. Peoples' perceptual judgments have been found to be close to predictions of an "ideal" observer who infers objects based on Bayesian theory [35]. This approach, for example, has been successful in modeling depth perception, boundary detection, and grouping of elements belonging to the same surface.

Bayesian perception theory explains how people perceive objects in the environment, not how they identify or classify them [39]. Nevertheless, the integration of cues in classification judgment is analogous to the integration of sensory cues in visual perception. It is possible that people could recruit the Bayesian integration function used in perception to integrate probabilistic identity cues for classification.

2.2 Automatic versus deliberate classification

One might ask how it would be possible to recruit a Bayesian visual system to perform a classification task. After all, we have previously assumed that Bayesian computation is

cognitively demanding, which forms the rationale for predicting the use of heuristics when conditions limit time or otherwise restrict a person's ability to employ a demanding strategy. The answer lies in the distinction between deliberate and automatic processes.

Automaticity is often defined by three main criteria: insensitivity to intentional control, insensitivity to cognitive capacity limits, and absence of awareness [40]. With respect to this definition, an automatic process is one that a person performs without conscious control or awareness and which does not compete for limited cognitive resources. In contrast, a deliberate process is one in which processing is under conscious control, that the person has significant insight into (i.e., the person can describe how the process works), and is generally effortful and limited [41].

Ashby and colleagues [42] [43] [44] have argued that people possess two distinct systems for categorization. Specifically, they propose that people have access to both an explicit system that is deliberate and suited to learning rule-based class distinctions and an implicit system that is automatic and suited to learning how to integrate probabilistic cues to form categories. Evidence for this deliberate-automatic distinction comes from studies that have demonstrated that learning of rule-based classification schemes is strongly affected by cognitive load [45], task demands [46], and interfering tasks that compete for limited cognitive resources [42][47], whereas learning classification schemes based on integration of cues is largely unaffected by these factors.

Glöckner and Betsch [41] have suggested that heuristics, such as TTB-C, are associated with deliberate processing. As such, use of heuristics is strongly affected by conditions that increase task demands or reduce available cognitive resources. They argue, however, that people can access automatic processes to make classification judgments. In three experiments, Glöckner and Betsch [41] found that subjects could employ the WADD strategy to perform a cue-based decision task as long as information search was not restricted. Moreover, subjects' response times were very fast, suggesting that performance was based on automatic processing.

In bringing together the two independent lines of research on Bayesian perception and dual-systems of classification, it may be possible to better understand why some people employ heuristic decision rules in the threat assessment task whereas others employ a Bayesian procedure. This is consistent with Gigerenzer's concept of the "Adaptive Toolbox" in which people can choose an appropriate heuristic to suit a given problem [19]. According to the dual-system view, heuristics would be associated with deliberate rule-based processing. That is, to employ a heuristic such as TTB-C, one must consciously select a particular cue as the most valid and classify targets according to a simple rule. In contrast, the Bayesian procedure would be associated with automatic cue-integration. Subjects would not have to deliberately attempt to compute conditional probabilities but, rather, rely on automatic processes to integrate all available cues according to an algorithm that is consistent with Bayes' Theorem. Thus, different subjects could reveal different strategies depending on which classification system served as the basis for performance.

The threat assessment task used by Bryant [4] is technically an information integration task because optimal performance depends on integration of all four cues. Thus, it might be expected that subjects would employ an automatic classification system, yielding judgments

consistent with a Bayesian decision rule. Given the specific cue validities associated with cues, however, the maximal level of performance achievable with a heuristic such as TTB-C was almost the same as that of the Bayesian strategy. Thus, the task could be treated as a rule-based classification task with little discernable loss in accuracy. This allows subjects the option of approaching the task as either a rule-based or information-integration decision.

3 Overview of the experiments

The purpose of the experiments reported herein was to test the hypothesis that pictorial stimuli are more likely to elicit a Bayesian decision strategy from subjects than textual stimuli, which were primarily used in previous experiments [3][4]. This hypothesis is premised on the assumption that pictorial stimuli allow a simultaneous as opposed to a sequential scan of predictive cues. Bayesian Perception Theory suggests that cues can be automatically processed to compute the conditional probabilities of an item being a friend or foe, which can serve as the basis of classification judgments.

In the experiments reported herein, participants learned to classify targets as friend or foe and then were tested on the task. In the first experiment, contacts were pictures of stylized aircraft, whereas in the second experiment they were more realistic illustrations of dismounted infantrymen. In both cases, the test sets of contacts were designed to contrast predictions of several heuristics, including the TTB-C and Additive Rules developed specifically for the threat classification task, with a Bayesian strategy based on computation of the conditional probabilities of friend or foe classification given the particular pattern of cues.

The purpose of Experiment 1 was to compare the performance of subjects who learn to classify contacts using textual cues (i.e., cues indicated by verbal labels) versus those who learn to classify contacts using pictorial cues (i.e., cues indicated by diagrams that visually depict the cues). If, as Glöckner & Betsch [41] suggest, simultaneous availability of all cues is necessary to employ an automatic cue-integration procedure, then subjects should be more likely to employ a Bayesian strategy when viewing pictures than when cues are provided textually. Text must be read sequentially, which would favour the use of a deliberate strategy, such as TTB-C.

The aim of the Experiment 2 was to determine whether rapid presentation of pictorial stimuli affects the propensity of subjects to use a Bayesian decision strategy. If the Bayesian strategy depends on the recruitment of automatic perceptual mechanisms, it may be more evident in situations in which the deliberate use of an heuristic is made difficult. This prediction runs counter to previous hypotheses that use of a fast-and-frugal heuristic is encouraged by time constraints. The results of Glöckner and Betsch [41], however, suggest that this is the case only when information search is limited in such a way as to force a sequential strategy. Thus, Bryant [4] found that an extreme time limit resulted in greater use of TTB-C when subjects had to deliberately inspect each cue individually. In Experiment 2, pictorial stimuli will be presented very rapidly so that subjects will not be able to respond while viewing stimuli. Instead, after a stimulus has disappeared, they will have to rely on their immediate memory of the item to make a decision. This should encourage subjects to process stimuli perceptually, perhaps favouring an automatic Bayesian decision procedure.

Both experiments employ the method used by Bryant [4] to infer what decision procedure each subject employs. Because different classification procedures predicted different responses to the test items, it was possible to create a test set of items that elicited different predicted responses from the decision procedures under consideration. By assessing which procedure produced predictions most closely matching a subject's actual responses, it is possible to infer that the subject used a particular decision procedure.

4 Experiment 1

Subjects' responses to test items will be analysed by comparing their judgments to predictions of the competing decision procedures (heuristic and Bayesian) to identify what strategy each subject employs, and assess the consistency with which that strategy is employed. Contrasting results of the textual and pictorial cue conditions will allow us to determine whether subjects in these conditions exhibit different patterns of strategy use. In particular, we will learn whether subjects who make judgments with pictorial cues are more likely to use a Bayesian strategy, which would indicate a greater degree of information combination than heuristic use.

4.1 Method

4.1.1 Subjects

Subjects were 48 male and female volunteers who were employees of Defence Research and Development Canada - Toronto (DRDC Toronto), students conducting research at DRDC Toronto, or individuals recruited from local universities. All subjects were aged 18 and older, had normal or corrected-to-normal vision, and were unfamiliar with the specific hypotheses and stimulus configurations of the experiments. All received stress pay remuneration for participating.

This study, approved by the DRDC Human Research Ethics Committee (HREC), was conducted in conformity with the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans.

4.1.2 Materials

The experiment was conducted with Pentium Personal Computers (PCs) using the E-Prime experiment authoring software. The software presented instructions and stimuli, collected subject responses, and recorded data.

In the experiment, subjects learned to classify potentially hostile aircraft (contacts) as friend or foe. Each contact's identity was determined by the combination of four characteristics (cues) – the shape of the aircraft's nose (cone or round), the shape of the wing (swept or delta), the type of tail (flexed or raised), and the shape of the cockpit (oval or extended). Each contact's cue values will be generated according to a probability matrix (i.e., each cue value will have a specified probability of being associated with each class of contact, hostile and non-hostile). The contacts were presented in either pictorial or text format. The pictorial contacts, as illustrated in Figure 4, were simple line drawings. The text contacts consisted of a table that indicated the values of the four cues for that aircraft.




Figure 4: Examples of pictorial stimuli

Subjects performed two conditions, one involving textual presentation of stimuli and the other pictorial presentation. Consequently, two sets of 300 contacts (150 friend and 150 foe) were created for the training sessions and two sets of 100 contacts (50 friend and 50 foe) were created for the test sessions. The sets were counter-balanced across presentation conditions.

4.1.3 Design

Three variables were manipulated in this experiment. The first, varied within subjects, was the Cue Validity of each cue used to describe contacts in the training stimuli sets. To vary Cue Validity, each possible value of a cue (values 1 and 2) was probabilistically associated to friend and foe classifications such that each cue differed in diagnosticity. Thus, for one cue each possible value was paired with the friend or foe classification 90% of the time, for another cue 80% of the time and so on. Table 1 indicates the proportions of friend and foe contacts possessing each cue value for the four cues in the two training sets. A contact was created by, first, designating it a friend or foe, then assigning values to each of the four cues according to the probabilities in Table 1. For example, in Set 1 a friend would be assigned a value for cockpit (cue 1) of either “bubble” (10% chance) or “extended” (90% chance), a value for nose (cue 2) of either “cone” (40% chance) or “round” (60% chance) and so on. A foe in Set 1 would be assigned values for the same four cues but the probabilities for each value were reversed. Contacts in Set 2 were created in the same manner but probability values associated with each cue were different (see Table 1).

The second variable manipulated was the Contact Type in the test stimuli. Each test set was made up of patterns that offered contrasting predictions of the three contending classification strategies discussed previously; namely TTB-C, the Bayesian strategy, WADD and ADD. Eight cue patterns were identified for which at least one strategy offered a differing response than predicted by the other strategies. From these contacts, we created six Contact Types (A, B, C, D, E, and F) that distinguished the predicted accuracy of the possible strategies. The different item types are indicated in Table 2 with the predicted response of each decision strategy. Note that each cue pattern listed in Table 2 falls into a different Contact Type

depending on whether that pattern is associated with a friend or foe. Type A and B items elicit opposing predictions from TTB-C and the Bayesian strategy. Where TTB-C would predict that these patterns indicate a friend, the Bayesian strategy would predict they indicate a foe, and vice versa. Type C and D patterns elicit the same predictions from TTB-C and the Bayesian strategy but force both the Weighted and Unweighted Additive rules (WADD and ADD) to guess because equal numbers of cues suggest friend and foe classifications. Types E and F contacts distinguish the Weighted and Unweighted forms of the Additive strategies. Computation of target classifications by the TTB-C, Bayesian, WADD, and ADD strategies is illustrated in Annex A.

Table 1: Relative Frequencies of Cue Values for Friend and Foe Contacts

	SET 1							
	Cue 1 (Cockpit)		Cue 2 (Nose)		Cue 3 (Wing)		Cue 4 (Tail)	
	Value 1 (extended)	Value 2 (Bubble)	Value 1 (Round)	Value 2 (Cone)	Value 1 (Swept)	Value 2 (Delta)	Value 1 (Flexed)	Value 2 (Raised)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%
	SET 2							
	Cue 1 (Nose)		Cue 2 (Tail)		Cue 3 (Cockpit)		Cue 4 (Wing)	
	Value 1 (Round)	Value 2 (Cone)	Value 1 (Flexed)	Value 2 (Raised)	Value 1 (extended)	Value 2 (Bubble)	Value 1 (Swept)	Value 2 (Delta)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%

In the test set, each of the critical patterns was paired an equal number of times with friend and foe contacts. We predicted the levels of accuracy predicted by the hypothesized decision procedures for each Contact Type, shown in Table 3.

The third variable, varied within subjects, was the Presentation Format in which contact characteristics were presented to subjects in the training and test phases. In the Text cue condition, cues were presented textually; that is, subjects saw a chart listing the particular characteristics for each selected contact. In the Pictorial cue condition, cues for each contact were provided in pictorial form; that is, subjects saw a drawing of the contact that shows its particular characteristics.

Because subjects performed the experimental task in both a Text and Pictorial condition, two contacts sets were created using the same cue labels but with different underlying cue validity

structures. The two sets were fully counterbalanced with the two presentation conditions and the order in which they were performed.

Table 2: Predicted Responses to Contact Types by Hypothesized Strategies

Cue Pattern	Predicted Response of Strategy				Contact Types	
	TTB-C	Bayesian	WADD	ADD	Foe	Friend
1,2,1,1	Friend	Foe	Foe	Foe	B	A
2,1,2,2	Foe	Friend	Friend	Friend	A	B
1,1,1,1	Friend	Friend	Guess	Guess	D	C
2,2,2,2	Foe	Foe	Guess	Guess	C	D
1,1,2,2	Friend	Friend	Friend	Guess	F	E
1,2,1,2	Friend	Friend	Friend	Guess	F	E
2,1,2,1	Foe	Foe	Foe	Guess	E	F
2,2,1,1	Foe	Foe	Foe	Guess	E	F

Note: Cue pattern indicates the value (as 1 or 2) for each cue in order of cues listed in Table 1

4.1.4 Procedure

The experiment was divided into two sessions for the Pictorial and Text conditions, each with a training and test phase. In the training phase, subjects viewed 300 contacts, of which 150 were friends and 150 foes. Given the structure of cue information, some patterns were more likely to occur than others through a random generation of contacts and the training set contained a number of each pattern proportional to its expected frequency. The contacts were presented sequentially and a contact did not appear until the subject had made a response to the previous contact. For each contact, the subject made a classification judgment, indicating that the contact was either hostile or not hostile by pressing a labeled key on the computer keyboard. No other options were presented and subjects had to make a decision for each stimulus. After making his/her response, the subject was shown a message indicating whether the response was correct or not (accuracy feedback). Subjects received no initial information concerning the predictiveness of cues and all learning was accomplished through trial-and-error.

Table 3 Predicted Accuracy Levels by Contact Type

Heuristic	Contact Type*					
	A	B	C	D	E	F
TTB-C	0%	100%	0%	100%	0%	100%
Bayesian	100%	0%	0%	100%	0%	100%
WADD	100%	0%	Guess	Guess	0%	100%
ADD	100%	0%	Guess	Guess	Guess	Guess

Following the training phase, subjects were allowed a short break and then performed the test phase. The test phase followed the same procedure as the training phase with a few important differences. First, subjects received no feedback on the accuracy of their judgments. Second, subjects were given a 16 second time limit in which to make their judgment. If the subject did not respond within that time, the contact disappeared and a null response was recorded. Third, subjects were presented with only 100 contacts (10 each of type A, B, C, D, E, and F, and 40 randomly selected from all other patterns). Subjects always received test contacts in the same format (Pictorial or Text) as they had received during the training session.

4.2 Results

4.2.1 Training session

The contacts presented during the training session were divided into six blocks of 50 contacts each, based on the order of presentation (i.e., the first 50 contacts, the next 50, etc.). Accuracy scores (the percentage of contacts correctly classified as friend or foe) were calculated for each block for each subject to create mean accuracy scores, which are shown broken down by Presentation Format in Figure 5.




Figure 5: Classification Accuracy by Block in the Training Session

A two-way, within-subjects Analysis of Variance (ANOVA) revealed a significant effect of Training Block [$F(5,235) = 31.36$, $MSe = 0.006$, $p < .01$] but no significant main effect of Presentation Format [$F(1,47) < .001$, $MSe = 0.03$, $n.s.$], indicating that it was neither more nor less difficult to learn cue associations for text than pictorial contacts. There was also no significant interaction effect between the two factors [$F(5,235) = 0.65$, $MSe = 0.006$, $n.s.$].

Overall, subjects learned to classify contacts to a high degree of accuracy, which indicates learning of individual cue validities. Subjects' mean performance in the final block, however, was significantly less than the optimal levels predicted by TTB-C (90.0%), Bayesian (90.2%), and WADD (90.2%) ADD (84.8%) strategies. A t -test indicated that the combined Block 6 performance (.80) was significantly lower than the lowest predicted optimal (84.8%) [$t(48)=3.65$, $p < .01$].

4.2.2 Classification strategy

The test set was made up entirely of patterns that offered contrasting predictions of the contending classification strategies (see Table 2). Type A and B items, for example, elicit opposing predictions from TTB-C and the Bayesian strategy. Type C and D patterns elicit the same predictions from TTB-C and the Bayesian strategy but force the WADD and ADD rules to guess because equal numbers of cues, with equal combined weights, suggest friend and foe classifications.

In previous research [4], strategy use was assessed for each subject by comparing each subject's proportion of correct responses for each pattern type (A, B, C, D, E, and F) to the

predicted accuracy levels indicated by the various strategies. Across the six item types, it was possible in most cases to identify the strategy with which the subject's pattern of accuracies was most consistent. For this experiment, we adopted a method developed by Bröder and Schiffer [25][48] that performs essentially the same procedure but adds a more sophisticated method of evaluating the precise likelihood that a decision strategy produced a subject's pattern of responses.

Bröder and Schiffer's [25][48] Maximum Likelihood Method (MLM) chooses the best fitting model from TTB-C, Bayesian, WADD, ADD, and guessing based on the likelihood that a subject's responses were generated by each strategy. The MLM (also called Bayesian Classification Method) is explained in more detail in Annex B, which indicates the equations by which MLM computes, for each strategy, the conditional probability of the subject's responses being produced by that strategy. The power of MLM to discriminate between strategies depends on the numbers of discriminating items that do not yield guessing responses from one or more decision strategy. Thus, the test items were designed to contain those cue patterns that elicited contrasting predictions from the decision strategies.

Table 4: Number of Subjects Classified as Using Hypothesized Decision Strategies

Presentation Format	Decision Heuristic				
	TTB-C	Bayesian	ADD	WADD	Unclassifiable
<i>By Presentation Format</i>					
Text	17	15	13	2	1
Picture	23	12	9	2	2
<i>By Target Set</i>					
Set 1	22	11	11	3	1
Set 2	18	16	11	1	2

N = 48 for each presentation format & target set

Table 4 presents the number of subjects classified as using a given decision strategy. As can be seen, the proportions of subjects using each of the decision strategies were very similar in the Text and Picture conditions and a Pearson Chi-Square test revealed no significant difference between the two conditions [$\chi^2 = 2.29$, $df = 5$, n.s.]. In contrast to findings of Bryant [4], a large number of subjects were classified as using a Bayesian strategy in both presentation conditions. Bryant [4] had observed very few instances of Bayesian strategy use. More in line with those previous findings, large numbers of subjects employed TTB-C and the Unweighted Additive Rules. Only a few were classified as using the Weighted Additive and no subject was classified as using a guessing strategy.

The two target sets comprised different associations of cues with friend/foe classification. That is, although the same four cues were used in both sets, these cues were associated with different classifications and/or had different cue validities in the two sets. To assess the effect of the cue configuration of a target set, we collapsed strategy use across presentation format and separated it according to the target set. The numbers of subjects classified as using a

given strategy for each target set are shown in Table 4. The target set did not have a significant effect on subjects' choice of decision strategy [$\chi^2 = 2.66$, $df = 5$, *n.s.*].

Table 5: Use of Decision Strategies in First and Second Conditions

Condition Sequence	Decision Heuristic				
	TTB-C	Bayesian	ADD	WADD	Unclassifiable
First	20	12	12	3	1
Second	20	15	10	1	2

N = 48 for each presentation format

We looked at the consistency with which subjects used a particular strategy across the two Presentation Formats. Table 5 shows the number of subjects classified as using each of the decision strategies in the first and second sessions of the experiment (i.e. collapsed across Presentation Format). A total of 19 of the 48 subjects (40%) employed the same strategy in the Text and Picture conditions, meaning that most subjects (60%) did not employ the same strategy in the Text and Picture conditions. To assess whether the selection of a strategy in the second condition was affected by the strategy employed in the first, we analysed the probability of choosing the TTB-C strategy (coded as 1) versus a compensatory strategy (coded as 2) through a logistic regression.³ The strategy used in the first condition and the presentation format in the second condition were considered as predictors.⁴ Neither the strategy used in the first condition [$t(42) = 1.08$, *n.s.*] nor the presentation format [$t(42) = 1.70$, *n.s.*] emerged as significant predictors of the strategy used in the second condition.

4.2.3 Response time

Response times were measured from the time at which the test item appeared on the screen to the time at which the subject pressed either the friend or foe key on the computer keyboard. Although no predictions concerning response times were drawn from the decision strategies under consideration, mean response times were computed for subjects. Generally, subjects took a fair amount of time, on the order of 6 to 8 seconds, to indicate their decisions.

Because reaction times were not distributed normally, they were transformed by their natural log prior to analysis. Although a two-factor within-subjects ANOVA revealed no significant effect of Presentation Format [$F(1,47) = 0.58$, $MSe = 0.87$, *n.s.*], there were significant effects of Item Type [$F(5,235) = 4.78$, $MSe = 0.036$, $p < .01$] and the interaction of factors [$F(5,235) = 5.92$, $MSe = 0.029$, $p < .01$]. These effects seem to reflect the fact that mean response times were somewhat larger for item types C and D in the pictorial but not text presentation condition. Indeed, mean response times were approximately equal across all

³ Data of three subjects were excluded because they produced unclassifiable results in at least one session.

⁴ The presentation formats in the first and second sessions were linked and so only one need be coded in the regression analysis.

item types for text presentation. Although differences in mean response times across Item Type have been observed in previous experiments, there has been no consistency in which items elicit faster responses.

4.3 Discussion

Substantial proportions of subjects employed the TTB-C, ADD, and Bayesian rules in this experiment. The use of the Bayesian rule by so many subjects in the text condition was surprising given that only a few subjects were classified as Bayesian users in Bryant [4]. This finding may reflect the impact of differences in the experimental methodology employed in the current experiments or differences in the subject samples. This result suggests that use of the Bayesian rule is a viable strategy for both textual and pictorial stimuli. It had been predicted that pictorial presentation would favour the use of the Bayesian rule relative to textual presentation but this was not the case. There was no significant difference between the two conditions in the proportions of subjects employing the Bayesian and TTB-C rules. Simply presenting targets as pictures is not sufficient to make the Bayesian rule more appealing to subjects relative to textual presentation.

Pictorial presentation of contacts was intended to allow subjects to perceive all cues simultaneously and thus be able to use perceptual mechanisms to integrate them. Because the pictures remained on the screen until he or she responded, the subject could sequentially scan the picture to identify cues. To control for this possible search strategy, the next experiment contained a condition in which pictorial stimuli were presented for only a short duration then disappear before subjects can respond. This manipulation is intended to force subjects to rely on immediate memory for the items, which may favour an automatic Bayesian decision process if it is based on perceptual processing.

5 Experiment 2

In previous research [4], subsets of subjects were observed to employ both the analytic Bayesian procedure as well as heuristic procedures of varying degrees of complexity (i.e., Additive Rules, TTB-C). At a group level, however, the only factor to be observed to affect the relative proportions of subjects adopting a given decision procedure was time pressure at the time of test. Even then, a fairly extreme manipulation of time pressure was required to shift a larger proportion of subjects to employ the fast-and-frugal TTB-C heuristic as opposed to compensatory procedures. Thus, it was perhaps not so surprising that the manipulation of presentation format in the previous experiment did not yield any differences in the proportions of subjects using the various decision procedures.

The aim of Experiment 2 was to determine whether rapid presentation of pictorial stimuli affects the propensity of subjects to use a Bayesian decision strategy. It was expected that very rapid presentation will require subjects to perceptually encode pictorial stimuli and then analyse their perceptual memory of it. This manipulation was intended to promote the use of an automatic Bayesian decision procedure. Pictures, unlike texts, do not force subjects to engage in sequential search of cues. Thus, restricting viewing time was not expected to lead to more subjects adopting TTB-C.

In this experiment, the friend-foe classification task was reframed in terms of Combat Identification (CID) by dismounted infantry soldiers. CID is the capability to rapidly and accurately identify friendly, enemy and neutral forces, manage and control the battlespace, optimally employ weapons and forces, and minimize the risk of fratricide (the inappropriate engagement and potential wounding or killing of a friendly soldier or unit). Thus, CID is very similar to the air threat assessment task used in Experiment 1 and a formally equivalent task was developed for the second experiment.

5.1 Method

5.1.1 Subjects

Subjects were 48 male and female volunteers who were employees of DRDC Toronto, students conducting research at DRDC Toronto, or individuals recruited from local universities. All subjects were aged 18 and older, had normal or corrected-to-normal vision, and were unfamiliar with the specific hypotheses and stimulus configurations of the experiments. All received stress pay remuneration for participating.

This study, approved by the DRDC HREC, was conducted in conformity with the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans.

5.1.2 Materials

The experiment was conducted with Pentium PCs using the E-Prime experiment authoring software. The software presented instructions and stimuli, collected subject responses, and recorded data.

In the experiment, subjects learned to classify potentially hostile soldiers (contacts) as friend or foe. The contacts were presented in pictorial format, as illustrated in Figure 6.



Figure 6: Examples of stimuli

Each contact's identity was determined by the combination of four characteristics (cues) – the pattern and colour of the contact's uniform (Canadian Distinctive Pattern (CADPAT) or olive green), the presence of a face covering (black mask or no mask), the type of rifle held (C7 or AK-47), and the colour of the helmet (CADPAT or dark green). Each contact's cue values were generated according to a probability matrix (i.e., each cue value will have a specified probability of being associated with each class of contact, hostile and non-hostile). Unlike the completely fictional stimuli used in Experiment 1, some characteristics of the soldier stimuli (CADPAT uniform, C7 and AK-47 rifles) were likely known to at least some subjects. These subjects would have had pre-existing associations of these characteristics to friend and foe classifications (i.e., CADPAT with friend, AK-47 with foe) that could complicate learning. Previous experiments [3] [4], however, have shown that 300 training trials leads to very high classification performance levels and subjects should be able to acquire accurate representations of underlying cue validities despite any pre-existing associations. Nevertheless, this issue remains pertinent to future research.

Subjects performed two conditions, a long and a short presentation time. Consequently, two sets of 300 contacts (150 friend and 150 foe) were created for the training sessions and two sets of 100 contacts (50 friend and 50 foe) were created for the test sessions. The sets were counter-balanced across presentation conditions.

5.1.3 Design

Three variables were manipulated in this experiment. The first, varied within subjects, was the Cue Validity of each cue used to describe contacts in the training stimuli sets. To vary Cue Validity, each possible value of a cue (values 1 and 2) was probabilistically associated to friend and foe classifications such that each cue differed in diagnosticity. Thus, for one cue each possible value was paired with the friend or foe classification 90% of the time, for another cue 80% of the time and so on. Table 6 indicates the proportions of friend and foe contacts possessing each cue value for the four cues in the two training sets. A contact was created by, first, designating it a friend or foe, then assigning values to each of the four cues according to the probabilities in Table 6. For example, in Set 1 a friend would be assigned a value for uniform (cue 1) of either “CADPAT” (90% chance) or “Olive” (10% chance), a value for helmet (cue 2) of either “Canadian” (60% chance) or “Dark Green” (40% chance) and so on. A foe in Set 1 would be assigned values for the same four cues but the probabilities for each value were reversed. Contacts in Set 2 were created in the same manner but probability values associated with each cue were different (see Table 6).

Table 6: Relative Frequencies of Cue Values for Friend and Foe Contacts

	SET 1							
	Cue 1 (Uniform)		Cue 2 (Helmet)		Cue 3 (Rifle)		Cue 4 (Face Cover)	
	Value 1 (CADPAT)	Value 2 (Olive)	Value 1 (Can.)	Value 2 (Dark gr.)	Value 1 (C7)	Value 2 (AK-47)	Value 1 (None)	Value 2 (Covered)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%
	SET 2							
	Cue 1 (Helmet)		Cue 2 (Uniform)		Cue 3 (Face Cover)		Cue 4 (Rifle)	
	Value 1 (Can.)	Value 1 (Can.)	Value 1 (CADPAT)	Value 2 (Olive)	Value 1 (None)	Value 2 (Covered)	Value 1 (C7)	Value 2 (AK-47)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%

The second variable manipulated was Contact Type (for test stimuli). Each test set was made up of patterns that offered contrasting predictions for the three contending classification strategies discussed previously; namely TTB-C, the Bayesian strategy, and the Weighted and Unweighted Additive Rules. Eight cue patterns were identified for which at least one strategy offered a different response than the others. From these contacts, we created six Contact Types (A, B, C, D, E, and F) that distinguished the predicted accuracy of the possible

strategies. The different item types are indicated in Table 7 with the predicted response of each decision strategy. Note that each cue pattern listed in Table 7 falls into a different Contact Type depending on whether that pattern is associated to a friend or foe. Type A and B items elicit opposing predictions from TTB-C and the Bayesian strategy. Where TTB-C would predict that these patterns indicate a friend, the Bayesian strategy would predict they indicate a foe, and vice versa. Type C and D patterns elicit the same predictions from TTB-C and the Bayesian strategy but force the Unweighted Additive Rule to guess because equal numbers of cues suggest friend and foe classifications. Types E and F contacts distinguish the Weighted and Unweighted forms of the Additive strategies.

In the test set, each of the critical patterns was paired an equal number of times with friend and foe contacts. We predicted the levels of accuracy predicted by the hypothesized decision procedures for each Contact Type, shown in Table 8.

Table 7: Predicted Responses to Contact Types by Hypothesized Strategies

Cue Pattern	Predicted Response of Strategy				Contact Types	
	TTB-C	Bayesian	WADD	ADD	Foe	Friend
1,2,1,1	Foe	Friend	Friend	Friend	B	A
2,1,2,2	Friend	Foe	Foe	Foe	A	B
1,1,1,1	Foe	Foe	Guess	Guess	D	C
2,2,2,2	Friend	Friend	Guess	Guess	C	D
1,2,2,1 & 1,2,1,2	Friend	Friend	Friend	Guess	F	E
2,1,2,1 & 2,1,1,2	Foe	Foe	Foe	Guess	E	F

Note: Cue pattern indicates the value (as 1 or 2) for each cue in order of cues listed in Table 6

The third variable, varied within subjects, was the Exposure Time during which the picture of a contact was visible on the computer screen during the test phase. In the Control condition, contacts were presented for up to 16 seconds, during which subjects could indicate their response. The Control condition replicates the procedure used in previous studies [4] and provided ample time to subjects to inspect a contact. In the Brief Exposure condition, contacts were presented for 500 msec, followed by a visual mask for 500 msec. Subjects could enter their response only after the contact and mask had been presented.

Table 8: Predicted Accuracy Levels by Contact Type

Heuristic	Contact Type					
	A	B	C	D	E	F
TTB-C	100%	0%	100%	0%	100%	0%
Bayesian	0%	100%	100%	0%	100%	0%
WADD	0%	100%	Guess*	Guess*	100%	0%
ADD	0%	100%	Guess*	Guess*	Guess*	Guess*

* It is assumed a guessing strategy would yield 50% accuracy

5.1.4 Procedure

The experiment was divided into two sessions for the Control and Brief Exposure conditions, each with a training and test phase. In the training phase of both conditions, subjects viewed 300 contacts, of which 150 were friends and 150 foes. Given the structure of cue information, some patterns were more likely to occur than others through a random generation of contacts and the training set contained a number of each pattern proportional to its expected frequency. The contacts were presented sequentially and a contact did not appear until the subject had made a response to the previous contact. For each contact, the subject made a classification judgment, indicating that the contact was either hostile or not hostile by pressing a labeled key on the computer keyboard. No other option was presented and subjects had to make a decision for each stimulus. After his/her response, the subject was given accuracy feedback on their classification judgment. Subjects received no initial information concerning the predictiveness of cues and all learning was accomplished through trial-and-error.

Following the training phase, subjects were allowed a short break and then performed the test phase. The test phase followed the same procedure as the training phase with a few important differences. First, subjects received no feedback on the accuracy of their judgments. Second, subjects were presented with only 100 contacts (10 each of type A, B, C, D, E, and F, and 40 randomly selected from all other patterns). Finally, in the Brief Exposure condition, subjects saw a contact for 500 msec, followed by a coloured visual mask (random dot pattern), then made their judgment. In this and the Control condition, subjects were required to respond to all contacts. If, however, a subject had not responded within 16 seconds, a null response was recorded and the next contact presented.

5.2 Results

5.2.1 Training session

The contacts presented during the training session were divided into six blocks of 50 contacts each, based on the order of presentation (i.e., the first 50 contacts, the next 50, etc.).

Accuracy scores (the percentage of contacts correctly classified as friend or foe) were calculated for each block for each subject to create mean accuracy scores, which are shown broken down by Exposure Time Condition in Figure 7. Overall, subjects' mean accuracy in the final block was somewhat lower than that seen in previous experiments [4] and significantly less than the optimal levels of performance predicted by any of the decision models under consideration.

A within-subject ANOVA revealed a significant effect of Training Block [$F(5,230) = 12.61$, $MSe = .008$, $p < .01$] but no significant main effect of Presentation Time [$F(1,46) = 0.006$, $MSe = .06$, *n.s.*], which is expected because the training sessions were exactly the same in both cases. There was likewise no significant interaction effect between the two factors [$F(5,230) = 0.24$, $MSe < .001$, *n.s.*]. The training set used was examined as a categorical factor to determine whether one set might be easier to learn to classify, and, although the main effect of training set was not statistically significant [$F(1,46) = 0.96$, $MSe = .09$, *n.s.*], this factor did interact with Training Block [$F(5,230) = 2.27$, $MSe < .008$, $p < .05$]. The interaction reflects the finding that, when learning Set 1, subjects exhibited somewhat higher levels of accuracy in the early blocks (1-4) than that seen when learning Set 2. Subject's accuracy scores for the two sets were essentially equal on the final two blocks. This suggests that training Set 1 may have been easier to learn initially but that any advantage disappeared with extended practice.

A second ANOVA was performed on subjects' mean response times to contacts across blocks. This analysis revealed a significant effect of Training Block [$F(5,230) = 12.61$, $MSe = .008$, $p < .01$] as subjects tended to respond faster over the course of the training session. No other factor or interaction significantly affected response times in the training session.




Figure 7: Classification Accuracy by Block in the Training Session

As in the first experiment, subjects' mean performance in the final block, however, was significantly less than the optimal levels predicted by TTBC (90.0%), Bayesian (90.5%), WADD (89.0%), and ADD (85.0%) strategies. A *t*-test indicated that the combined Block 6 performance (.75) was significantly lower than the lowest predicted optimal (.85) [$t(48)=6.00, p<.01$].

5.2.2 Classification strategy

As in the first experiment, the test set was made entirely of patterns that offered contrasting predictions of the contending classification strategies (see Table 7). The MLM was again used to classify the decision strategy employed by each subject based on his or her sequence of responses to test items.

Table 9 presents the number of subjects classified as using a given decision strategy. As can be seen, the proportions of subjects using each of the decision strategies were very similar in the control and Brief Exposure conditions and a Pearson Chi-Square test revealed no significant difference between the two conditions [$\chi^2 = 2.91, df = 5, n.s.$]. In contrast to findings of Bryant [4], a large number of subjects were classified as using a Bayesian strategy in both conditions. More in line with those previous findings, large numbers of subjects employed TTBC and the ADD rule. Only a few were classified as using WADD and no subject was classified as using a guessing strategy.

Table 9: Number of Subjects Classified as Using Hypothesized Decision Strategies

Presentation Time	Decision Heuristic				
	TTB-C	Bayesian	ADD	WADD	Unclassifiable
<i>By Presentation Format</i>					
Control	19	16	6	0	7
Brief Exposure	21	18	5	1	3
<i>By Target Set</i>					
Set 1	35	4	4	1	4
Set 2	5	30	7	0	6

N = 48 for each presentation format & target set

The two target sets comprised different associations of cues with friend/foe classification. That is, although the same four cues were used in both sets, these cues were associated with different classifications and/or had different cue validities in the two sets. To assess the effect of the cue configuration of a target set, we collapsed strategy use across Exposure Time and separated it according to the target set. The numbers of subjects classified as using a given strategy for each target set are shown in Table 9. As can be seen, there is a striking difference in the numbers of subjects classified as using TTB-C versus the Bayesian strategy in Set 1 and Set 2 [$\chi^2 = 44.60$, $df = 5$, $p < .01$].

5.2.3 Response time

Response times were measured from the time at which the test item appeared on the screen to the time at which the subject pressed either the friend or foe key on the computer keyboard. Although no predictions concerning response times were drawn from the decision strategies under consideration, mean response times were computed for subjects. Generally, subjects required little time to make their responses, on the order of 1.2 to 1.4 seconds in the control condition and 650 to 800 msec in the Brief Exposure condition, to indicate their decisions.

Because reaction times were not distributed normally, they were transformed by their natural log prior to analysis. A mixed-design ANOVA revealed a significant effect of Exposure Time [$F(1,46) = 41.35$, $MSe = 1.11$, $p < .01$] and subjects were significantly faster to respond in the Brief Exposure than control condition. Subjects were not forced to respond faster in the Brief Exposure condition as they were allowed 16 seconds, just as in the control condition, after the test item had been briefly presented. The ANOVA also revealed a significant effect of Item Type [$F(5,230) = 5.49$, $MSe = 0.05$, $p < .01$], which reflected the fact that mean response times were somewhat faster for item types C and D in both the Brief Exposure and control conditions. Note that this result stands in contrast to the result observed in Experiment 1 in which items C and D tended to elicit *slower* responses. Variations in response time likely reflect random variation rather than a systematic effect. No other factors or interactions significantly affected response times.

5.3 Discussion

Rapid presentation of targets did not increase the proportion of subjects employing the Bayesian decision rule. Thus, it appears that static pictorial presentation itself does not promote the use of that or any other decision procedure, although this does not rule out the possibility that dynamic visual stimuli may elicit a preferred decision strategy. The absence of an effect of exposure time means that people can use both analytic and simple heuristic decision procedures equally as well when item appears very briefly as when they have essentially unlimited viewing time. Only extreme time pressure on cue search, such as employed in Bryant [4], seems to affect the relative use of a simple heuristic.

An unexpected but very large effect of the target set was observed in this experiment. This suggests that the configuration of cues plays an important role in subjects' choice of a decision procedure. In Set 1, the uniform was the most predictive cue and it also seemed to be the most salient. Some subjects spontaneously offered that they looked at this first and later informal surveys of researchers at DRDC Toronto also suggested that the uniform is seen as visually salient and an assumed a priori predictor of friend/foe. In Set 2, the helmet was most predictive and this cue was not considered to be as salient as the uniform. It may be that when a salient cue, or a cue with a pre-existing association to the classification task, is the most predictive, subjects are able to quickly notice its relation to classification and use TTB-C. In other words the salient, high validity cue suggests TTB-C as a decision rule. In contrast, when a non-salient cue is most predictive, subjects do not have one cue that immediately stands out as a key predictor and so they look at all cues to identify targets. This suggests a compensatory and analytic decision rule, either because subjects explicitly weigh all cues or because they acquire richer instances of contacts in memory, which supports a recognition-based decision rule that conforms to Bayesian predictions.

This explanation is consistent with some previous research on category learning. Various models of categorization (e.g., [49] [50]) have proposed that cue validity is the prime determinant of a cue's use in category learning. Nevertheless, some of these models admit that a cue's salience, either perceptual or conceptual (i.e., meaningfully/causally related to category membership), can play a role in the processes of learning categories and retrieving category information from memory [51]. Thus, the classic view of category learning is that the learner seeks cues that are most diagnostic and incorporates these in a representation of the category but may also select cues to category membership based on non-diagnostic characteristics (see Martin & Caramazza [52] for a review). Indeed, it has been found that completely irrelevant (i.e., non-diagnostic) cues can affect aspects of category performance such as judgments of category membership, ratings of category typicality, and the degree to which a person makes use of relevant cues [53] [54].

Martin and Caramazza [52] suggested in 1980 that people begin category learning by identifying features or cues that can be used to distinguish category members and non-members. Moreover, they argued that this process comprises a sequential search of cues, which a single cue at a time being evaluated for its usefulness. In their view, perceptually salient cues are most likely to be noticed early in the learning process and subsequently incorporated in the representation of the category. Thus, salient and predictive cues are more likely to be incorporated than predictive but non-salient cues, and even non-predictive cues may become associated with the category because they stood out early in the learning process.

In a series of experiments on perceptual category learning, Martin and Caramazza observed that the perceptual salience of cues did affect subjects' learning of categories even when salience was not aligned with cue validity. In particular, the salience of cues affected the order in which those cues were sampled.

The salience of a cue could also be determined by non-perceptual characteristics, such as pre-existing conceptual relationships between cues and specific categories [52]. In addition to being, perhaps, the most perceptually salient cue, the uniform may also have had a meaningful relationship to classification of a target as friend or foe. Subjects likely found it easier to learn and accept that the CADPAT uniform reliably signalled that a stimulus was a friend. The issue of conceptual salience will be explored in future research.

From this, it appears that the interaction of cue validity and cue salience should determine classification. When first encountering objects to be classified, such as the stimuli used in the experiments reported here, subjects may focus on one feature at a time to examine and evaluate as a category predictor. Subjects would be biased toward selecting salient cues first. In light of the target set effect seen in Experiment 2, it seems that if the first cue selected is highly predictive, most subjects recognize an heuristic such as TTB-C as a reasonable decision strategy. In contrast, when the first cue is not highly predictive, and the subject continues evaluating subsequent cues, he/she acquires knowledge of all cue validities and employs a Bayesian decision strategy.

6 Conclusion

The experiments reported here provide evidence that people can employ both a fast-and-frugal heuristic approach and a more complex Bayesian strategy to classify targets as friend or foe. Both provided accurate solutions to the experimental task, so it was hypothesized that the nature of the representational format might govern the decision procedure employed by subjects. Specifically, a Bayesian strategy was considered easier to employ when targets were depicted pictorially because all cues would be available for simultaneous processing, perhaps recruiting automatic perceptual process [41]. Textual presentation, in contrast, was expected to require sequential processing and be more consistent with the deliberate employment of a fast-and-frugal heuristic. The absence of any difference in strategy use between pictorial and textual presentation conditions in Experiment 1 is inconsistent with this hypothesis.

In an effort to replicate the results of Experiment 1, a speeded presentation technique was employed in Experiment 2 to make use of a deliberate strategy more difficult [41]. Contacts in the test phase were presented for only 500 msec each to dramatically limit subjects' time to inspect cues and thus make use of a deliberate heuristic difficult. It was hypothesized that this manipulation would favour the use of an automatic Bayesian strategy if it did, in fact, rely on the recruitment of automatic perceptual mechanisms. Contrary to this hypothesis, however, rapid presentation of test items did not affect subjects' propensity to employ either the Bayesian or TTB-C strategies.

Although the original hypothesis that subjects would be more likely to employ a Bayesian strategy for pictorial stimuli was not supported by the results, it was found that the target sets in Experiment 2 yielded a very large difference in decision behaviour. This finding raises the possibility that the salience of visual cues plays a role in determining decision strategy selection by subjects. Based on the relative salience of cues in the test sets, it may be that highly salient cues that are also highly predictive tend to lead subjects to employ an heuristic strategy such as TTB-C. In contrast, when the most predictive cues are not salient, subjects may be more likely to employ a Bayesian strategy. Subsequent research will examine the role of cue salience in threat assessment judgments.

References

- [1] Liebhaber, M. J., & Smith, C. A. P. (1999). *Tactical decision-making under stress: Baseline results*. Cited in Liebhaber, Kobus, & Feher (2002).
- [2] Cruz, E. E. (1996). *Fratricide in air-land operations*. Unpublished Master's thesis. Naval Postgraduate School, Monterey, CA.
- [3] Bryant, D. J. (2003). *Cue validity learning in threat classification judgments*. DRDC Toronto TR 2003-041, Defence R&D Canada – Toronto, Toronto, Canada.
- [4] Bryant, D. J. (2007). Classifying simulated air threats with Fast and Frugal Heuristics. *Journal of Behavioral Decision Making*, 20, 37-64.
- [5] Bryant, D. J., Webb, R. D. G., & McCann, C. (2003). Synthesizing two approaches to decision making in command and control. *Canadian Military Journal*, Spring, 29-34.
- [6] Klein, G. (1992). *Decision making in complex military environments*. Klein Associates, Fairborn, OH: Final Contract Summary Report, prepared for Naval Command, Control and Ocean Surveillance Center, San Diego, CA.
- [7] Yates, J. F. (1990). *Judgment and decision making*. Englewood Cliffs, NJ: Prentice Hall.
- [8] Hunt, E. (1994). Problem solving. In R. J. Sternberg (Ed.), *Thinking and problem solving* (pp. 215-232). San Diego, CA: Academic Press.
- [9] Hogarth, R. M. (1980). *Judgment and choice: The psychology of decision*. New York, NY: John Wiley.
- [10] Ben-Zur, H. (1998). Dimensions and patterns in decision-making models and the controlled/automatic distinction in human information processing. *European Journal of Cognitive Psychology*, 10, 171-189.
- [11] Fishburn, P. C. (1988). Normative theories of decision making under risk and under uncertainty. In Bell, D. E., Raiffa, H., & Tversky, A. (Eds.), *Decision making: Descriptive, normative, and prescriptive interactions* (pp. 78-98). New York, NY: Cambridge University Press.
- [12] Todd, P. M., & Gigerenzer, G. (2000). Precise of *Simple heuristics that make us smart*. *Behavioral and Brain Sciences*, 23, 727-780.
- [13] Gigerenzer, G. (1993). The bounded rationality of probabilistic mental models. In K. I. Manktelow & D. E. Over (Eds.), *Rationality: Psychological and philosophical perspectives* (pp. 284-313). New York: Routledge.
- [14] Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650-669.

- [15] Tietz, R. (1992). Semi-normative theories based on bounded rationality. *Journal of Economic Psychology*, 13, 297-314.
- [16] Todd, P. M. (2001). Fast and frugal heuristics for environmentally bound minds. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 51-70). Cambridge, MA: The MIT Press.
- [17] Sternberg, R. J. (1996). *Cognitive Psychology*. Fort Worth, TX: Harcourt Brace.
- [18] Colton, R. A., & Gange, R. H. (1993). *Including the human element in design of command and control decision support systems - The KOALAS concept*, Naval Research Lab., Washington, DC.
- [19] Gigerenzer, G., Todd, P. M., & The ABC Research Group (1999). *Simple heuristics that make us smart*. New York, NY: Oxford University Press.
- [20] Slegers, D. W., Brake, G. L., & Doherty, M. E. (2000). Probabilistic mental models with continuous predictors. *Organizational Behavior and Human Decision Processes*, 81, 98-114.
- [21] Gigerenzer, G., & Selten, R. (2001). Rethinking rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 1-12). Cambridge, MA: MIT Press.
- [22] Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In Gigerenzer, G., Todd, P. M., & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97-118). New York, NY: Oxford University Press.
- [23] Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 119-140). Oxford, UK: Oxford University Press.
- [24] Bröder, A. (2000). Assessing the empirical validity of the "take the best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1332-1346.
- [25] Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, 132, 277-293.
- [26] Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone "takes-the-best". *Organizational Behavior and Human Decision Processes*, 91, 82-96.
- [27] Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 53-65.

- [28] Liebhaber, M. J., Kobus, D. A., & Feher, B. A. (2002). *Studies of U.S. Navy air defense threat assessment: Cues, information order, and impact of conflicting data*. Technical Report 1888. SPAWAR Systems Center, San Diego, CA.
- [29] Brehmer, B. (1973). Single-cue probability learning as a function of the sign and magnitude of the correlation between cue and criterion. *Organizational Behavior and Human Decision Processes*, 9, 377-395.
- [30] Himmelfarb, S. (1970). Effects of cue validity differences in weighting information. *Journal of Mathematical Psychology*, 7, 531-539.
- [31] Reiskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics, and how can we tell? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 141-167). Oxford, UK: Oxford University Press.
- [32] Simon, H. A., & Hayes, J. R. (1976). The understanding process: Problem isomorphs. *Cognitive Psychology*, 8, 165-190.
- [33] Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102, 684-704.
- [34] Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of "success". *Journal of Behavioral Decision Making*, 17, 117-137.
- [35] Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, 13, 1-9.
- [36] Ernst, M. O. (2006). A Bayesian view on multimodal cue integration. In Knoblich, G., Thornton, I. M., Grosjean, M., & Shiffrar, M. (Eds.), *Human Body Perception From The Inside Out* (pp. 105-131). New York, NY: Oxford University Press.
- [37] Knill, D. C. (2007). Bayesian models of sensory cue integration. In K. Doya, S. Ishii, A. Pouget & R. P. N. Rao (Eds.), *Bayesian brain: Probabilistic approaches to neural coding* (pp. 189-206). Cambridge, MA: MIT Press.
- [38] Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences*, 10, 301-308.
- [39] Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271-304.
- [40] Hasher, L., & Zacks, R. T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, 108, 356-388.
- [41] Glöckner, A., & Betsch, T. (2008). Multiple-reason decision making based on automatic processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1055-1075.

- [42] Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, 8, 168-176.
- [43] Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition*, 30, 666-677.
- [44] Maddox, W. T., Ashby, F. G., Ing, A. D., & Pickering, A. D. (2004). Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition*, 32, 582-591.
- [45] Ashby, F. G., Ell, S. W., & Waldron, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition*, 31, 1114-1125.
- [46] Maddox, W. T., Lauritzen, J. S., & Ing, A. D. (2007). Cognitive complexity effects in perceptual classification are dissociable. *Memory & Cognition*, 35, 885-894.
- [47] Zeithamova, D., & Maddox, W. T. (2006). Dual-task interference in perceptual category learning. *Memory & Cognition*, 34, 387-398.
- [48] Bröder, A., & Schiffer, S. (2003). Bayesian strategy assessment in multi-attribute decision making. *Journal of Behavioral Decision Making*, 16, 193-213.
- [49] Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22-44.
- [50] Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53-79.
- [51] Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.
- [52] Martin, R. C., & Caramazza, A. (1980). Classification in well-defined and ill-defined categories: Evidence for common processing strategies. *Journal of Experimental Psychology: General*, 109, 320-353.
- [53] Edgell, S. E., Castellan Jr., N. J., Roe, R. M., Barnes, J. M., Ng, P. C., Bright, R. D., & Ford, L. A. (1996). Irrelevant information in probabilistic categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1463-1481.
- [54] Rein, J. R., Love, B. C., & Markman, A. B. (2007). *Feature relations and feature salience in natural categories*. In Proceedings of the 29th Meeting of the Cognitive Science Society (pp. 593-598), Nashville, TN.
- [55] Hu, X., & Batchelder, W. H. (1994). The statistical analysis of general processing tree models with the EM algorithm. *Psychometrika*, 59, 21-48.

Annex A

This Annex presents sample calculations of decision strategy predictions to illustrate how each decision strategy made judgments and how the six Contact Types (A, B, C, D, E, F) used in the test sets elicited different predictions from the various strategies.

Probabilistic structure of the training and test contacts

Contacts in Experiment 1 were created according to the probability matrix reproduced in Table A1, which indicates the proportions of friend and foe contacts possessing each cue value (1 or 2) for the four cues in the two training sets.

Table A1: Relative Frequencies of Cue Values for Friend and Foe Contacts

	SET 1							
	Cue 1 (Cockpit)		Cue 2 (Nose)		Cue 3 (Wing)		Cue 4 (Tail)	
	Value 1 (extended)	Value 2 (Bubble)	Value 1 (Round)	Value 2 (Cone)	Value 1 (Swept)	Value 2 (Delta)	Value 1 (Flexed)	Value 2 (Raised)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%
	SET 2							
	Cue 1 (Nose)		Cue 2 (Tail)		Cue 3 (Cockpit)		Cue 4 (Wing)	
	Value 1 (Round)	Value 2 (Cone)	Value 1 (Flexed)	Value 2 (Raised)	Value 1 (extended)	Value 2 (Bubble)	Value 1 (Swept)	Value 2 (Delta)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%

Each test set was made up of patterns that offered contrasting predictions of the three contending classification strategies discussed previously; namely TTB-C, the Bayesian strategy, WADD, and ADD. Eight cue patterns were identified for which at least one strategy offered a differing response than predicted by the other strategies. From these contacts, we created six Contact Types (A, B, C, D, E, and F) that distinguished the predicted accuracy of the possible strategies. The different item types are indicated in Table A2 with the predicted response of each decision strategy. Note that each cue pattern listed in Table 2 falls into a different Contact Type depending on whether that pattern is associated to a friend or foe. Type A and B items elicit opposing predictions from TTB-C and the Bayesian strategy.

Where TTB-C would predict that these patterns indicate a friend, the Bayesian strategy would predict they indicate a foe, and vice versa. Type C and D patterns elicit the same predictions from TTB-C and the Bayesian strategy but force the Unweighted Additive Rule to guess because equal numbers of cues suggest friend and foe classifications. Types E and F contacts distinguish the Weighted and Unweighted forms of the Additive strategies.

Table A2: Predicted Responses to Contact Types by Hypothesized Strategies

Cue Pattern	Predicted Response of Strategy				Contact Types	
	TTB-C	Bayesian	WADD	ADD	Foe	Friend
1,2,1,1	Friend	Foe	Foe	Foe	B	A
2,1,2,2	Foe	Friend	Friend	Friend	A	B
1,1,1,1	Friend	Friend	Guess	Guess	D	C
2,2,2,2	Foe	Foe	Guess	Guess	C	D
1,1,2,2	Friend	Friend	Friend	Guess	F	E
1,2,1,2	Friend	Friend	Friend	Guess	F	E
2,1,2,1	Foe	Foe	Foe	Guess	E	F
2,2,1,1	Foe	Foe	Foe	Guess	E	F

Note: Cue pattern indicates the value (as 1 or 2) for each cue in order of cues listed in Table 1

Example classification computations

This section illustrates how TTB-C, the Bayesian strategy, WADD and ADD compute classifications for a given target cue pattern. For this purpose, we refer to a Set 1 contact (see Table A1) composed of the following characteristics:

- Cue 1 (Cockpit) = Value 1 (Extended)
- Cue 2 (Nose) = Value 2 (Cone)
- Cue 3 (Wing) = Value 1 (Swept)
- Cue 4 (Tail) = Value 1 (Flexed)

This contact is represented as 1,2,1,1 in Table A2 and is a Type A item.

TTB-C

Figure A1 below illustrates TTB-C in the form of a flowchart. Given that all cues are available and that a cue value of 1 for Cue 1 is most strongly associated with a classification of Friend (see Table A1), we can compute the predicted response of TTB-C to cue pattern 1,2,2,2 in the following manner:

- Step 1: Search for most valid cue = Cockpit
- Step 2: Inspect cue value = Extended (1)
- Step 3: Assign classification (based on cue value's association to Friend or Foe) = Friend

Thus, TTB-C predicts that the target is a Friend.

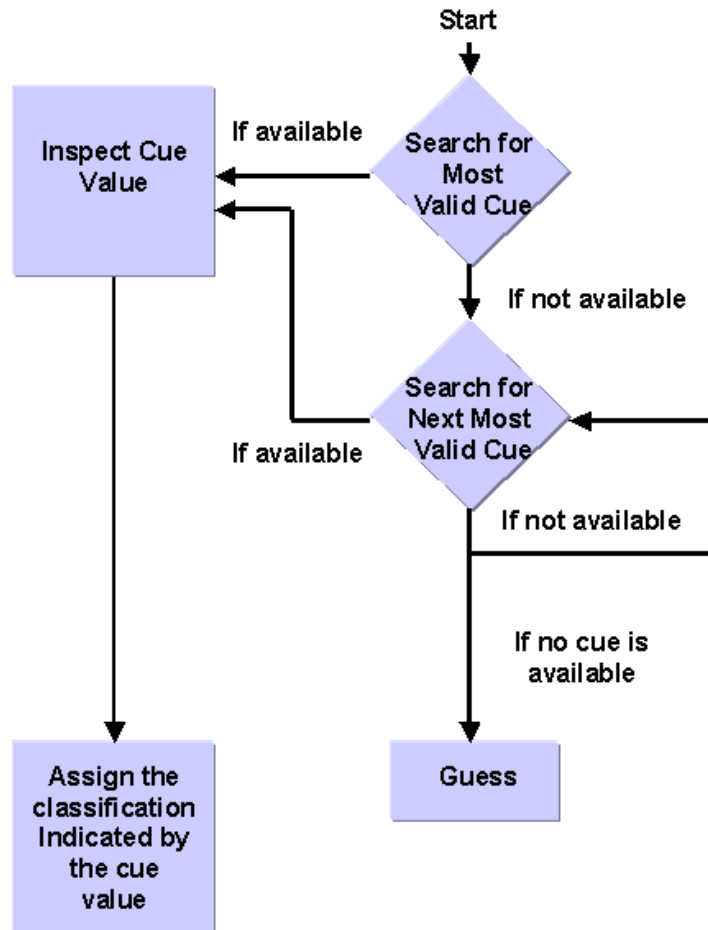


Figure A1: The Take-the-Best-for-Classification Strategy

Bayesian network

A Bayesian network for the friend/foe task is shown in Figure 3. The top node represents the classification of a target as a friend (F). The case of a foe would be represented by the negation of friend (\bar{F}). Four nodes representing characteristics of the target, or cues (C_{1-4}), are connected to it according to their probabilistic association to the class of the target. Thus, each line linking a cue to the classification node is labeled by the conditional probability of the cue occurring given the classification of friend. Considering all cues as a set, the

classification node represents the conditional probability of that class being true given the presence of the four linked cues. This is given by the formula (A1):

$$P(F|C_1 \cap C_2 \cap C_3 \cap C_4) = \frac{P(C_1|F) \cdot P(C_2|F) \cdot P(C_3|F) \cdot P(C_4|F)}{[P(C_1|F) \cdot P(C_2|F) \cdot P(C_3|F) \cdot P(C_4|F) + P(C_1|\bar{F}) \cdot P(C_2|\bar{F}) \cdot P(C_3|\bar{F}) \cdot P(C_4|\bar{F})]}$$

(A1)

Where:

$$P(C_j | \bar{F}) = 1 - P(C_j | F), j = 1 \text{ to } 4.$$

The Bayesian strategy was assumed to compute the conditional probabilities of friend and foe classifications given the particular pattern of cue values for a contact and select the alternative with the higher probability of being the correct classification.




Figure A2: Bayesian Network for Friend-Foe Classification

Based on Equation A1, we can compute the conditional probabilities of the target being a Friend or Foe given that the cue pattern is 1,2,2,2 in the following manner:

Given:

$P(C_1 = 1 F) = .90$	$P(C_1 = 1 \bar{F}) = .10$
$P(C_2 = 2 F) = .40$	$P(C_2 = 2 \bar{F}) = .60$
$P(C_3 = 2 F) = .30$	$P(C_3 = 2 \bar{F}) = .70$
$P(C_4 = 2 F) = .20$	$P(C_4 = 2 \bar{F}) = .80$

$$P(F|1,2,1,1) = \frac{P(C_1=1|F) \cdot P(C_2=2|F) \cdot P(C_3=1|F) \cdot P(C_4=1|F)}{[P(C_1=1|F) \cdot P(C_2=2|F) \cdot P(C_3=1|F) \cdot P(C_4=1|F) + P(C_1=1|\bar{F}) \cdot P(C_2=2|\bar{F}) \cdot P(C_3=1|\bar{F}) \cdot P(C_4=1|\bar{F})]}$$

$$P(F|1,2,1,1) = \frac{.90 \cdot .40 \cdot .30 \cdot .20}{[.90 \cdot .40 \cdot .30 \cdot .20] + [.10 \cdot .60 \cdot .70 \cdot .80]} = .39$$

and

$$P(\bar{F}|1,2,1,1) = \frac{P(C_1=1|\bar{F}) \cdot P(C_2=2|\bar{F}) \cdot P(C_3=1|\bar{F}) \cdot P(C_4=1|\bar{F})}{[P(C_1=1|F) \cdot P(C_2=2|F) \cdot P(C_3=1|F) \cdot P(C_4=1|F) + P(C_1=1|\bar{F}) \cdot P(C_2=2|\bar{F}) \cdot P(C_3=1|\bar{F}) \cdot P(C_4=1|\bar{F})]}$$

$$P(\bar{F}|1,2,1,1) = \frac{.10 \cdot .60 \cdot .70 \cdot .80}{[.90 \cdot .40 \cdot .30 \cdot .20] + [.10 \cdot .60 \cdot .70 \cdot .80]} = .61$$

Thus,

$$P(\bar{F}|1,2,1,1) > P(F|1,2,1,1)$$

Therefore, the Bayesian network predicts that the target is a Foe.

ADD and WADD

The Additive Rules examine each cue value and assign evidence toward either friend or foe classification, depending on the associations of cue values to threat class. A running sum is maintained and, after all available cues have been inspected, used to place the contact in the friend or foe category. Figure A3 contains an illustration of the Weighted Additive Rule (WADD), which weights cues by their validity, adapted for threat classification. A classification version of the Unweighted Additive Rule (ADD) is performed just as illustrated in Figure A3 but without the weighting step following the selection of a cue.

Table A3 presents the computation of evidence for a classification of friend or foe for the cue pattern 1,2,1,1. Both WADD and ADD sum greater evidence for a classification of foe based on the associations, weighted and unweighted, of cue values to the respective classes.

<i>Table A3: Summation of Cue Evidence for Cue Pattern 1,2,1,1 by the WADD and ADD Strategies</i>							
				Evidence Summation			
				WADD		ADD	
Cue	Cue Value	Cue Value Association	Weight	Friend	Foe	Friend	Foe
1	1	Friend	.9	.9	0	1	0
2	2	Foe	.6	0	.6	0	1
3	1	Foe	.7	0	.7	0	1
4	1	Foe	.8	0	.8	0	1
			Total	.9	2.1	1	3
			Classification	Foe		Foe	

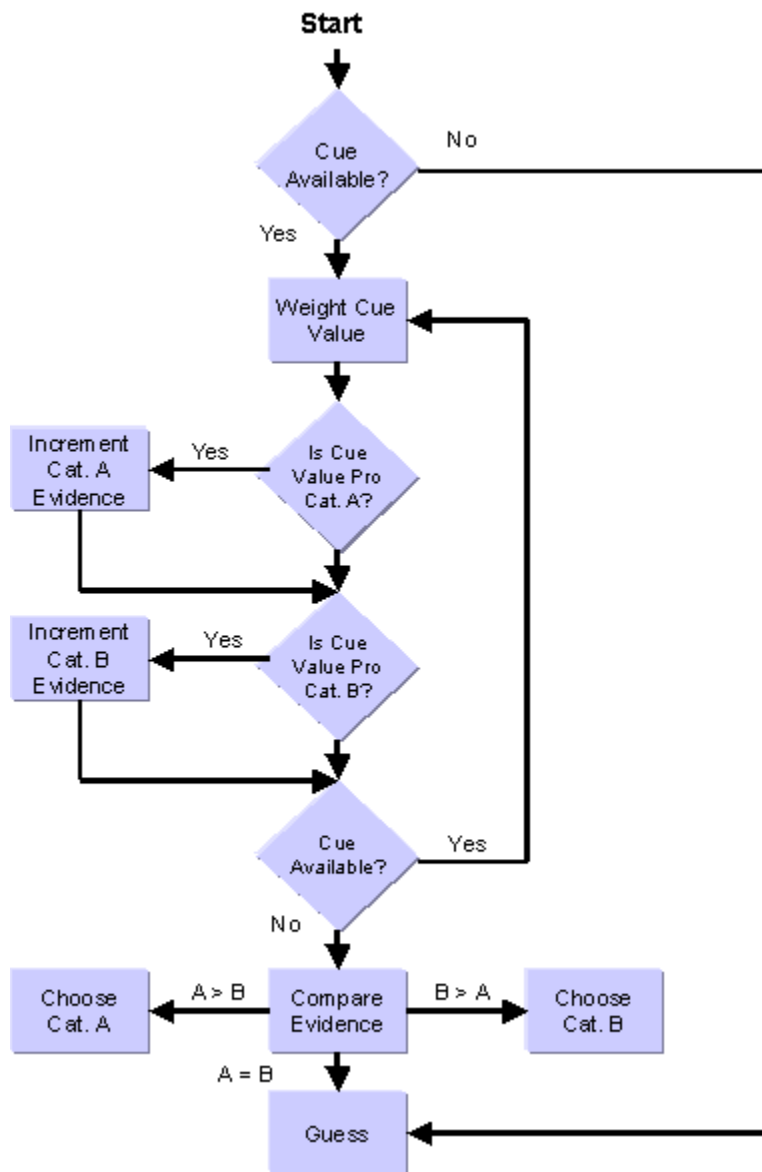


Figure A3: The Weighted and Unweighted Additive Rules for Classification

Annex B: The Maximum Likelihood Method (MLM) for Bayesian classification

The Maximum Likelihood Method (MLM), sometimes also termed the Bayesian Classification Method (BCM) method, was developed by Bröder and Schiffer [25] [48] as a means to assess which of a potential set of decision strategies was most likely employed by a subject in a multiple cue-based decision task. In short, the method they developed determines the conditional probability that a subject's sequence of responses for multiple decisions would occur given the use of a specified decision strategy. By determining this probability for a set of strategies, one is able to identify the most likely strategy to have produced the subject's observed responses.

The MLM is applied to a decision task in which items belonged to one of two possible types (friend or foe) and were described by values along four binary cues. The validity of each cue as a predictor of item type is a variable. For the purpose of illustrating the MLM, consider the cues presented in Table B1 along with the validity of each as a predictor of friend or foe.

Table B1: Relative Frequencies of Cue Values for Friend and Foe Contacts

	SET 1							
	Cue 1 (Cockpit)		Cue 2 (Nose)		Cue 3 (Wing)		Cue 4 (Tail)	
	Value 3 (extended)	Value 1 (Bubble)	Value 3 (Round)	Value 1 (Cone)	Value 3 (Swept)	Value 1 (Delta)	Value 3 (Flexed)	Value 1 (Raised)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%
	SET 2							
	Cue 1 (Nose)		Cue 2 (Tail)		Cue 3 (Cockpit)		Cue 4 (Wing)	
	Value 3 (Round)	Value 1 (Cone)	Value 3 (Flexed)	Value 1 (Raised)	Value 3 (extended)	Value 1 (Bubble)	Value 3 (Swept)	Value 1 (Delta)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%

Given four binary cues, there are 16 possible cue configurations that can be associated with friend or foe. All decision strategies under consideration make the same predictions for some of these configurations. That is, there are cue configurations for which each strategy will predict friend or foe. There are, however, a subset of items that elicit different predictions

from at least two strategies. Bröder and Schiffer's MLM makes use of these items. In particular, there are three critical item types j ($j = 1, 2, \text{ or } 3$) for the purpose of assessing decision strategy, which are listed in Table B2. Type 1 items elicit a prediction from the TTB-C strategy that is different from all others (types 1a and 1b simply reflect different cue configurations in which TTB-C makes a prediction opposite from the other strategies). Type 2 items elicit the same predictions from TTB-C and a Bayesian strategy but cannot be solved by either the Weighted or Unweighted Additive Rules, which can only guess. Type 3 items elicit the same predictions from TTB-C, the Bayesian strategy, and the Weighted Additive Rule but elicit guessing from the Unweighted Additive Rule.

<i>Table B2: Examples of Different Types of Items Used to Assess Decision Strategies</i>						
	Item Type					
	1		2		3	
Attribute/ Strategy	1a	1b	2a	2b	3a	3b
Cockpit	Extended	Bubble	Extended	Bubble	Extended	Bubble
Nose	Cone	Round	Round	Cone	Cone	Round
Wing	Swept	Delta	Swept	Delta	Swept	Delta
Tail	Flexed	Raised	Flexed	Raised	Raised	Flexed
Prediction of Strategies						
TTB-C	Friend	Foe	Friend	Foe	Friend	Foe
Bayesian	Foe	Friend	Friend	Foe	Friend	Foe
Unweighted Additive	Foe	Friend	Guess	Guess	Guess	Guess
Weighted Prose	Foe	Friend	Guess	Guess	Friend	Foe

The method determines the decision procedure that has the greatest likelihood of producing the data based on the predictions of the candidate set of procedures under consideration. In the experiments described in this report, those procedures are TTB-C, Bayesian, WADD, ADD, and Guessing. The method makes the assumption that subjects generate responses to test items according to one of these procedures is. It also assumes that subjects have a certain probability, ϵ , of making an error and generating a response not predicted by the procedure being used.

The likelihood of a subject's observed data vector (i.e. sequence of responses to test items) is calculated by the following formula (B1):

$$L(n_{jk}, n_{jk}^{\varepsilon} | k, \varepsilon_k, n_j) = \prod_{j=1}^3 \binom{n_j}{n_{jk}} \times (1 - \varepsilon_k)^{n_{jk}} \times \varepsilon_k^{(n_j - n_{jk})} \quad (\text{B1})$$

Where,

n_j = number of items of each type j presented in an experiment,

n_{jk} = number of choices in item type j that were predicted by strategy k ,

n_{jk}^{ε} = number of choices in item type j not predicted by strategy k , such that $n_{jk}^{\varepsilon} + n_{jk} = 1$,

ε_k = error probability of choosing the option not conforming to strategy k .

Thus, equation A1 gives the likelihood $L(n_{jk}, n_{jk}^{\varepsilon} | k, \varepsilon_k, n_j)$ that the observed data vector \mathbf{n} is equal to $(n_{jk}, n_{jk}^{\varepsilon})$, given strategy k , and unknown error probability ε_k . The unknown error term can be estimated by fitting the corresponding joint multinomial model to the frequency data [55], or by applying the formula given in Equation B2:

$$\hat{\varepsilon} = \frac{\sum_{j=1}^3 n_{jk}^{\varepsilon}}{\sum_{j=1}^3 n_j} \quad (\text{B2})$$

The following adjustments are applied to these formulae. When $k = \text{WADD}$, the index j in Equation B2 only runs from 1 to 2 because a person using WADD must guess for item type 2 and $\varepsilon_k = 0.5$ in Equation A1 for the case $k = \text{WADD}$ and $j = 2$. When $k = \text{ADD}$, the index j in Equation A2 only runs from 1 to 1 because a person using ADD must guess for item types 2 and 3 and $\varepsilon_k = 0.5$ in Equation A1 for the cases $k = \text{ADD}$ and $j = 2$ and $k = \text{ADD}$ and $j = 3$. For $k = \text{Guessing}$, no parameter estimation is necessary and all error probabilities are set to 0.5 in Equation A1.

To classify a subject's decision strategy, a likelihood ratio (Equation B1) is computed for every strategy and the vector classified as being produced by the particular strategy if the likelihood ration in favour of this strategy is larger than 1. Otherwise, the vector remains unclassified.

The power of MLM to discriminate between strategies depends on the numbers of discriminating items that do not yield guessing responses from one or more decision strategy. Thus, discriminating the WADD and ADD strategies may be more difficult than

discriminating the TTB-C and Bayesian strategies. Because the Guessing model has no free parameter, the other models will fit better than the random model in almost all cases.

List of symbols/abbreviations/acronyms/initialisms

ADD	Unweighted Additive Rule
ANOVA	Analysis of Variance
CADPAT	Canadian Distinctive Pattern
CID	Combat Identification
DRDC	Defence Research & Development Canada
HREC	Human Research Ethics Committee
IFF	Interrogate-Friend-Foe
MLM	Maximum Likelihood Method
MSe	Mean Square Error
PC	Personal Computer
TTB	Take-the-Best
TTB-C	Take-the-Best-for-Classification
WADD	Weighted Additive Rule

Distribution list

Document No.: DRDC Toronto TR 2009-126

LIST PART 1: Internal Distribution by Centre

2

2 TOTAL LIST PART 1

LIST PART 2: External Distribution by DRDKIM

1 Library and Archives Canada

DJCP

1 LCol Rick Johnston (Richard.Johnston@forces.gc.ca)

1 Cdr Robert Edwards (Robert.Edwards2@forces.gc.ca)

DRDC Valcatier

1 Dr. Eloi Bossé (Eloi.Bosse@drdc-rddc.gc.ca)

DST C4ISR 4

1 Dr. Benny Wong (Benny.Wong@forces.gc.ca)

CTC

1 Col Steve Bowes (Steve.Bowes@forces.gc.ca)

1 Maj Michael Standish (Michael.Standish@forces.gc.ca)

DGLS/COS Land Ops

1 Shawn Hoag (Shawn.Hoag@forces.gc.ca)

SJS

1 LCol Alan Smith (Alan.Smith@forces.gc.ca)

1 Maj John Whalen (John.Whalen2@forces.gc.ca)

ADM(S&T)

1 LCol Peter Scott (Peter.Scott@forces.gc.ca)

DSTA

1 Luc Lefebvre (Luc.Lefebvre@forces.gc.ca)

DRDC Atlantic

1 Dr. Mark Gammon (Mark.Gammon@forces.gc.ca)

63 TOTAL LIST PART 2

15 TOTAL COPIES REQUIRED

UNCLASSIFIED

DOCUMENT CONTROL DATA <small>(Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)</small>		
1. ORIGINATOR (The name and address of the organization preparing the document, Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's document, or tasking agency, are entered in section 8.) Publishing: DRDC Toronto Performing: DRDC 1133 Sheppard Ave. W., Toronto, ON, M3M Toronto 3B9 Monitoring: Contracting:		2. SECURITY CLASSIFICATION <small>(Overall security classification of the document including special warning terms if applicable.)</small> UNCLASSIFIED
3. TITLE (The complete document title as indicated on the title page. Its classification is indicated by the appropriate abbreviation (S, C, R, or U) in parenthesis at the end of the title) Decision Rules for Pictorial Threat Classification (U) (U)		
4. AUTHORS (First name, middle initial and last name. If military, show rank, e.g. Maj. John E. Doe.) David J. Bryant		
5. DATE OF PUBLICATION <small>(Month and year of publication of document.)</small> July 2009	6a NO. OF PAGES <small>(Total containing information, including Annexes, Appendices, etc.)</small> 64	6b. NO. OF REFS <small>(Total cited in document.)</small> 55
7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of document, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.) Technical Report		
8. SPONSORING ACTIVITY (The names of the department project office or laboratory sponsoring the research and development – include address.) Sponsoring: Tasking:		
9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant under which the document was written. Please specify whether project or grant.) 15au	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)	
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document) DRDC Toronto TR 2009-126	10b. OTHER DOCUMENT NO(s). (Any other numbers under which may be assigned this document either by the originator or by the sponsor.)	
11. DOCUMENT AVAILABILITY (Any limitations on the dissemination of the document, other than those imposed by security classification.) Unlimited distribution		
12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, when further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.)) Unlimited announcement		

UNCLASSIFIED

UNCLASSIFIED

DOCUMENT CONTROL DATA

(Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)

13. **ABSTRACT** (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

(U) Two experiments examined the use of heuristic and analytic decision strategies in a simulated threat assessment task. Subjects learned to classify targets as friend or foe on the bases of uncertain cues (i.e. characteristics that were probabilistically associated with classification of a target as friend or foe). Subjects were then asked to classify targets that contrasted predictions of several decision rules, including a simple heuristic called Take-the-Best-for-Classification (TTB-C) that uses a single cue to classify targets and the Bayesian classification strategy that is based on formal statistic models. Results of Experiment 1 indicated that the mode of presentation (Text versus picture) did not affect the tendency of subjects to use either decision strategy. Results of Experiment 2 indicated that exposure time of pictorial stimuli also did not affect the proportions of subjects employing TTB-C versus the Bayesian strategy. However, an unexpected but very large effect of the target set was observed in the second experiment. This effect may indicate that the interaction of the perceptual salience of cues with the diagnosticity of those cues is a predictor of strategy use. Future research will examine this possibility.

(U) On a examiné, dans le cadre de deux expériences, l'utilisation de stratégies de décision heuristique et analytique dans l'évaluation d'une menace simulée. Les sujets ont appris à classer les cibles comme étant amies ou ennemies sur la base de repères incertains (notamment, sur des caractéristiques probabilistes associées à la classification d'une cible amie ou ennemie). On a ensuite demandé aux sujets de classer les cibles qui mettaient en contraste les prédictions de plusieurs règles de décision, incluant la simple approche heuristique « ne garder que le meilleur en vue de la classification » (TTB-C), qui utilise un seul indice pour classer des cibles, et la stratégie de classification bayésienne qui repose sur des modèles de statistiques officielles. Les résultats de l'expérience numéro 1 indiquent que le mode de présentation (un texte par rapport à une image) n'a pas poussé les sujets à utiliser une stratégie de décision plutôt que l'autre. Les résultats de l'expérience numéro 2 indiquent que le temps d'exposition à une stimulation par l'image n'a pas eu d'incidence non plus sur les sujets qui ont utilisé l'approche TTB C par rapport à la stratégie bayésienne. Toutefois, une conséquence inattendue, mais très importante de l'ensemble des cibles a été observée dans la deuxième expérience. Cette conséquence peut indiquer que l'interaction entre l'évidence perceptuelle des repères et la diagnosticité de ces repères est un prédicteur de la stratégie utilisée. D'autres recherches permettront d'examiner cette possibilité.

14. **KEYWORDS, DESCRIPTORS or IDENTIFIERS** (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

(U) decision making; classification; combat identification

UNCLASSIFIED

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale



www.drdc-rddc.gc.ca

