



This electronic thesis or dissertation has been downloaded from the University of Bristol Research Portal, <http://research-information.bristol.ac.uk>

Author:

Kowalski, Pawel B

Title:

Multi-source heterogeneous intelligence fusion

General rights

Access to the thesis is subject to the Creative Commons Attribution - NonCommercial-No Derivatives 4.0 International Public License. A copy of this may be found at <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>. This license sets out your rights and the restrictions that apply to your access to the thesis so it is important you read this before proceeding.

Take down policy

Some pages of this thesis may have been removed for copyright restrictions prior to having it been deposited on the University of Bristol Research Portal. However, if you have discovered material within the thesis that you consider to be unlawful e.g. breaches of copyright (either yours or that of a third party) or any other law, including but not limited to those relating to patent, trademark, confidentiality, data protection, obscenity, defamation, libel, then please contact collections-metadata@bristol.ac.uk and include the following information in your message:

- Your contact details
- Bibliographic details for the item, including a URL
- An outline nature of the complaint

Your claim will be investigated and, where appropriate, the item in question will be removed from public view as soon as possible.

Multi-Source Heterogeneous Intelligence Fusion

By

PAWEŁ BARTŁOMIEJ KOWALSKI



Department of Engineering Mathematics
UNIVERSITY OF BRISTOL

A dissertation submitted to the University of Bristol in accordance with the requirements of the degree of DOCTOR OF PHILOSOPHY in the Faculty of Engineering.

NOVEMBER 2021

Word count: fifty-two thousand one hundred and thirty seven

ABSTRACT

The ever-increasing amount of data means that today's decision-makers and analysts are often faced with an overwhelming amount of information, arriving in a variety of formats from multiple sources. This work addresses three interrelated challenges within information fusion - context exploitation, explainability and situation awareness and demonstrates their application to maritime situational awareness.

In uncertain reasoning, scenario context can be defined as information relevant to the reasoning task at hand but not directly involved with it. This context can be exploited in the construction of the model. The benefit of exploiting context in the construction of the model is two-fold. On the one hand, its explicit representation can provide better insight into the problem. On the other hand, it may improve the expressiveness of the model and result in an inference that better represents available knowledge. Context exploitation can also improve the interpretability of the model and can be used to generate better explanations.

A reasoning process built upon evidential networks is transparent by design. Still, it can be not easy to follow due to many variables and the complex information combination process. As such, tools for explanation generation and tracking the origin of beliefs, uncertainty and conflict in evidential reasoning systems are proposed. Finally, we must consider the problem of fusion on a higher level in the Joint Directors of Laboratories (JDL) framework, where uncertain relations between multiple entities need to be tracked. An extension of the conceptual graphs framework to allow uncertainty modelling is proposed, along with appropriate information fusion strategies. These three notions are put together, and their relationship is demonstrated in two maritime domain scenarios: situation awareness of multiple vessels involved in illegal fishing and underwater infrastructure threat assessment focused on the impact of partially reliable sources.

DEDICATION AND ACKNOWLEDGEMENTS

I would like to thank my supervisor, Trevor Martin, for his support, expertise and guidance. This work was funded and supervised by Thales Group, and particular thanks go to my industrial supervisor David Harvey for valuable suggestions, as well as to Claire Laudy from Thales for discussions and comments. Furthermore I would like to acknowledge the support from NATO Centre for Maritime Research and Experimentation, where I performed significant portion of the work on this thesis. In particular I would like to thank Anne-Laure Joussetme, who supervised my work during my time at CMRE and whose guidance and advice was imperative for the completion of this thesis. I would also like to thank other staff at CMRE, especially other members of the DKOE project, including Nadia Ben-Abdallah and Maximilian Zocholl the extensive discussions with whom helped shape this thesis.

I cannot possibly list all the people who supported me in some way or another during research and writing of this thesis. Above all, I must thank my parents and my grandmother for their unwavering belief. They were there both when things were easy and when they were difficult. Knowing that there always has been someone to fall back on may have been the difference between success and failure. I wish to specifically thank Ula, Kasper and Michał for their support, but also all my friends in the UK, in Poland and in Italy who helped me at one time or the other throughout these last several years.

This work was supported by Engineering and Physical Sciences Research Council [Grant number: EP/I028153/1] and Thales Group. A portion of this work was supported by NATO CMRE as part of the DKOE project

AUTHOR'S DECLARATION

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED: DATE:

TABLE OF CONTENTS

	Page
List of Tables	xi
List of Figures	xiii
1 Introduction	1
1.1 Main notions and their relations	1
1.2 Thesis structure	3
1.3 Key contributions	5
1.4 List of research outputs	5
1.4.1 Poster presentations	5
1.4.2 First-author conference publications	6
1.4.3 Journal publications	6
2 Background	7
2.1 The basic principles and generalised models of information fusion	7
2.2 Uncertainty modelling formalisms	9
2.2.1 Set based and interval based approaches	9
2.2.2 Probability theory	10
2.2.3 Possibility and fuzzy sets	12
2.2.4 Dempster-Shafer belief functions	14
2.2.5 Interpretations of belief functions	21
2.3 Information quality and types of uncertainty	21
2.3.1 Source quality and information quality	21
2.3.2 Types of uncertainty and modelling of uncertainty	22
2.4 Graphical models	24
2.4.1 Conceptual graphs	24
2.4.2 Bayesian networks	25
2.4.3 Markov networks	26
2.4.4 Markov Logic networks	26
2.4.5 Valuation-based systems and evidential networks	27
2.5 Context awareness	29
2.5.1 Context in artificial intelligence	29

2.6	Summary	31
3	Reasoning with partially reliable sources and context	33
3.1	Context in information fusion	33
3.1.1	Categories of context	34
3.2	Models for partially reliable sources	34
3.2.1	Contextual belief correction	39
3.3	Reasoning with context	40
3.3.1	The contextual reasoning model	41
3.3.2	Illustration: Vessel behaviour analysis example	45
3.4	Properties of contextual reasoning model	47
3.4.1	Context-of only - contextual belief correction	47
3.4.2	Context-for problem	47
3.4.3	Context-for source	48
3.4.4	Combined models	51
3.4.5	Interpretability	51
3.4.6	Further examples of source behaviour and contextual reasoning models	54
3.5	Dependence relationships between sources	57
3.5.1	Consonant sources	57
3.5.2	Non-consonant reports	60
3.5.3	Idempotent combination of evidence using VBS	60
3.5.4	Possibly independent sources and context-aware independence	66
3.6	Summary	68
4	Provenance of information, conflict and uncertainty	71
4.1	Relation to contextual reasoning, explainability and interpretability	71
4.2	Theoretical foundations of provenance	72
4.2.1	Provenance of information	72
4.2.2	Mathematical formalisation of information fusion and decision making	72
4.2.3	Generalised measures for the provenance of decision, belief and information	74
4.2.4	Provenance of conflict	75
4.2.5	Origin of uncertainty	76
4.3	Origin of information and decision in belief functions framework	76
4.3.1	Simple contribution measures and the problem with preference	77
4.3.2	Generation of explanations in the theory of evidence	78
4.3.3	Distance-based contribution measures	79
4.3.4	Generation of BBAs	88
4.3.5	Randomly generated seed BBAs	93
4.4	Information fusion as cooperative game theory - Shapley value	96
4.5	Contribution measures - the preference property	98
4.5.1	Dempster's rule of combination and maximum of belief decision making	99

4.5.2	Dempster’s rule of combination and other decision-making approaches . . .	105
4.5.3	Other combination rules	106
4.6	Contribution to conflict	106
4.6.1	Conflict contribution measures	108
4.7	Contribution to uncertainty	110
4.7.1	Measures of second-order uncertainty	111
4.7.2	Example: combination methods facilitating nonspecificity	112
4.8	Discussion and summary	113
5	Situation awareness using conceptual graphs	115
5.1	Modelling of situations	115
5.1.1	JDL perspective	116
5.1.2	Situation awareness	116
5.1.3	Situations and context	116
5.2	Conceptual graphs, information graphs and ontologies	117
5.2.1	Uncertainty management	118
5.2.2	Motivation	118
5.3	Uncertainty management for partial conflict resolution	118
5.3.1	Set-based approach	119
5.3.2	Uncertain information fusion	120
5.3.3	Multi-source aggregation scenario using uncertain information fusion	121
5.4	Partial reliability ratings	123
5.5	Comparison of the two models	125
5.6	Modal uncertainty	127
5.7	Fusion strategies	128
5.8	Reasoning and association	130
5.9	Relation to belief function framework	130
5.9.1	Uncertain graphs as belief functions	131
5.9.2	Representation as evidential networks	132
5.9.3	Existential valuations	136
5.10	Markov logic networks	136
5.11	Discussion and challenges	137
6	Maritime Situation Awareness scenarios	139
6.1	Scenario 1: multi-entity illegal fishing	139
6.1.1	Scenario overview	139
6.1.2	Rules and sources of information	140
6.1.3	Multi-source reasoning implementation and results	142
6.1.4	Explaining the results	149
6.2	Scenario 2: Infrastructure protection and threat assessment	153
6.2.1	Threat assessment multi-source evidential network	153

6.2.2	Inference results	157
6.2.3	Explaining the impact of source reports	159
6.2.4	Explaining the source model	162
6.3	Summary	166
7	Discussion	169
7.1	Thesis summary and next steps	169
7.2	Further research and relation to other areas within Artificial Intelligence	172
	Bibliography	175

LIST OF TABLES

TABLE	Page	
3.1	Examples of the six categories of contextual information in MSA domain, adapted from [53] and [54]	35
3.2	The six reasoning models	51
3.3	Masses m_1 through to m_5 , unnormalized and normalized	53
3.4	Discord and nonspecificity for normalized and unnormalized masses m_1 , m_2 and m_6	53
3.5	Multi-valued mapping Γ describing the interpretation of each value of the report R given the hypothesis H	55
3.6	Combination of the two dependent sources using a variety of combination rules: conjunctive rule of combination (CRC), disjunctive (DRC), Dubois-Prade conjunctive and disjunctive rules for dependent sources and cautious rule of combination; numerical results provided for $\lambda = 0.8$ and $\mu = 0.7$	59
3.7	Combination of the two dependent sources using a variety of combination rules: numerical results provided for $\lambda = 0.8$ and $\alpha = 0.25$	60
3.8	Comparison of XQR model-based combination for different values of D and S (assuming $s_1 = s_2$) with combination using Dempster's, cautious and disjunctive rules	67
4.1	Mass assignments and initial reliability of sources S_1 and S_2	87
4.2	Highly focused seed basic belief assignment used	91
4.3	Coefficients of determination R^2 (focused BBA, divergent mode)	91
4.4	Coefficients of determination R^2 (focused BBA, perturbation mode)	92
4.5	Basic probability assignment used as the seed BBA	93
4.6	Coefficients of correlation r and determination R^2	93
4.7	Mean μ and standard deviation σ of r for the Monte Carlo simulation with randomly generated BBAs	94
4.8	Proportion of cases where the preference property is not satisfied for the different contribution measures using Dempster's rule of combination and maximum of belief decision	100
4.9	Proportion of cases where the preference property is not satisfied for the different contribution measures and various decision making methods using Dempster's rule of combination	106

4.10	Proportion of cases where the preference property is not satisfied for the different contribution measures and various decision making methods using Dempster's, Dubois-Prade or PCR rules of combination	107
4.11	Conflict contribution measures for masses m_1 to m_4	110
4.12	Specificity and nonspecificity contribution measures for masses m_1 to m_4	112
5.1	Semantic evidence combination results	124
5.2	Semantic evidence combination results including necessity ratings	125
5.3	Probability distribution on outcomes of the maritime surveillance scenario using sources 1 - 3 and the original model from [107]	126
5.4	Probability distribution on outcomes of the maritime surveillance scenario using the set-based uncertainty model for partial conflict resolution	126
5.5	Nonspecificity and average consistency values for the family of candidate fusion outcome graphs	129
5.6	Semantic evidence combination results represented as a belief function	132
5.7	Evidential network joint distribution	134
5.8	Belief distribution with one of the masses having a degree of necessity	135
5.9	Belief distribution with one of the masses having a degree of necessity with fusion performed using conceptual graphs	135
5.10	Belief distribution with one of the masses having a degree of necessity with fusion performed using Dubois-Prade conjunctive-disjunctive rule	136
5.11	Weighted rules for the Markov Logic Network	137
5.12	Markov Logic Network inference results	137
6.1	Inference on Illegal fishing \times Fishing \times Illegal Activity for <i>ARAMIS</i>	143
6.2	Inference on joint belief distribution of Illegal activity (<i>ARAMIS</i>) \times Illegal activity (<i>CHEKHOV</i>)	146
6.3	Marginals obtained from the inferred joint belief distribution on $IA(C) \times RV(A, C) \times IF(A) \times F(A) \times IA(A)$	147
6.4	Inference of behaviour of vessels B and D. There is no evidence supporting any implication of illegal activity for either vessel	147
6.5	Variables of the threat assessment evidential network	154
6.6	Model of solution to threat assessment expressed in natural language rules	155
6.7	Logical relationships in the threat assessment valuation network (not including source models), with parametrised probability values α_i . Values in brackets are the ones used in the current implementation	156
6.8	List of sources	157

LIST OF FIGURES

FIGURE	Page
1.1 The relationships between the three main focus areas of this thesis	2
2.1 Marginalization and vacuous extension (from [19])	16
2.2 Three basic types of uncertainty - typology from [36]	23
2.3 Uncertainty typology (or dictionary) adapted from [38]	23
2.4 An example of a conceptual graph describing the destination of fishing vessel Marcel .	25
2.5 Structure of a Bayesian network for identification of critical situations between two ships	25
3.1 Reporting process	35
3.2 The basic XQR model. The source report is driven by the variable of interest and the source quality. The source quality itself may be affected by the problem variable (akin to contextual belief correction)	38
3.3 An evidential network, showing an implementation of the XQR model	38
3.4 Extension of Steinberg's model of exogenous and endogenous variables, accounting for variables outside the context of the problem [52]	40
3.5 System-level view of the reasoning system with a set of inputs and outputs. Information flow determines the difference between context variable and a variable of interest . . .	42
3.6 A variable of interest X and an associated report R provided by a source which behaviour is determined by Q . The context variable C may affect either or both Q and X	43
3.7 An evidential network, showing an implementation of the $XQRC$ model	43
3.8 Ternary plot showing the relationship between different contextual models for $\lambda = 0.2$, $\mu = 0.6$, $\gamma_1 = 0.5$ and $\gamma_2 = 0.3$	52
3.9 Valuation network for inference on the report provided by the aerial UAV for vessel detection and classification model	55
3.10 Two reports R_1 and R_2 such that the report R_2 is conditionally independent on X given R_1	58
3.11 A valuation network implementing idempotent combination of reports R_1 and R_2 . . .	61
4.1 Sensitivity of belief and plausibility to discounting of a mass, illustrating projection of Imp_i onto $\widehat{Bel}_i = \widehat{Pl}_i$	80
4.2 Illustration of the triangle formed by the three masses	81
4.3 Contribution of m_1 to belief in ω_1 using $N = 3$ and $b = 0.4$	84

4.4	Contribution of m_1 to belief in ω_1 using $N = 5$ and $b = 0.4$	85
4.5	Contribution of m_1 to belief in ω_1 using $N = 3$ and $b = 0.7$	86
4.6	Contribution of m_1 to belief in ω_1 using $N = 5$ and $b = 0.7$	87
4.7	Contribution of m_1 to belief in ω_1 for different values of b using $r_1 = r_2 = 1$ and $N = 3$.	88
4.8	Contribution of m_1 to belief in ω_1 for different values of b using $r_1 = r_2 = 1$ and $N = 5$.	89
4.9	Sensitivity analysis of the final fusion outcome with respect to the reliability of information sources	90
4.10	Change in quality of decision when a single body of evidence is removed under Dempster-Shafer rule of combination combination, using highly focused seed BBA, divergent mode	92
4.11	Change in quality of decision when a single body of evidence is removed under Dempster-Shafer rule of combination, using highly focused seed BBA, high perturbation	93
4.12	Change in quality of decision when a single body of evidence is removed under Dempster-Shafer rule of combination combination, using highly focused seed BBA, low perturbation	94
4.13	Monte Carlo histogram using $\hat{C}(m_1, m A)$ (top) and benchmark (bottom)	95
4.14	Valuation network	97
4.15	Bivariate hexagonal binning showing density of bba pairs satisfying and not satisfying the preference property for <i>Bel</i> as a function of the distance and the conflict between them, using maximum of belief decision making and Dempster's combination	101
4.16	Bivariate hexagonal binning showing density of bba pairs satisfying and not satisfying the preference property for distance-based measure as a function of the distance and the conflict between them, using maximum of belief decision making and Dempster's combination	102
4.17	Set of scatter plots for decision making using maximum of belief, combination using Dempster's rule and the belief measure being used as the contribution measure	103
4.18	Set of scatter plots for decision making using maximum of belief, combination using Dempster's rule and the belief measure being used as the contribution measure	104
4.19	Histogram of bba pairs satisfying and not satisfying the preference property for <i>Bel</i> , where $0.2 < d_J(m_1, m_2) < 0.45$ and $0.1 < m_J(\emptyset) < 0.4$ for different values of $S_{A \subseteq B}$	105
4.20	Histogram of bba pairs satisfying and not satisfying the preference property for different values of $S_{A \subseteq B}$	105
5.1	Excerpt from a maritime concept vocabulary	119
5.2	Conceptual graph representation without uncertainty modelling	119
5.3	Conceptual graph representation with uncertainty modelling	120
5.4	Conceptual graph representation with uncertainty modelling	120
5.5	Conceptual graph representation with uncertainty modelling	120
5.6	Excerpt from a maritime concept vocabulary	121
5.7	Basic knowledge model of the topic of interest	122
5.8	Graph G_1 generated from the first source	122
5.9	Graph G_2 generated from the second source	122

5.10	Graph $G_1 \otimes G_2$ generated using disjunctive fusion of both concepts in G_1 and G_2	122
5.11	Graph G_{123} , the result of combination of the three sources $G_1 \otimes G_2 \otimes G_3$ using an appropriate fusion strategy	123
5.12	Graph G_4 with $Nec(G_4) = 0.3$ and $Pos(G_4) = 0.8$	124
5.13	Graphs resulting from combination of G_2 and G_4 assuming that G_4 is necessary (5.13a) and that G_4 is possible (5.13b)	125
5.14	Simplified model for the maritime surveillance scenario from [107], the geographical location has been removed	125
5.15	Evidential network for possible implementation of trivial type hierarchy	132
5.16	Valuation network showing the concepts and relations in the model for Marcel analysis example	133
6.1	A single vessel conceptual graph shown for <i>ARAMIS</i>	141
6.2	Valuation network for assessment of illegal activity and illegal fishing	142
6.3	A single vessel conceptual graph shown for <i>ARAMIS</i> obtained through inference on AIS data with $P = 0.44442$	144
6.4	A single vessel conceptual graph shown for <i>BOHUN</i> obtained through inference on AIS data with $P = 0.9$	144
6.5	A single vessel conceptual graph shown for <i>CHEKHOV</i>	144
6.6	Excerpt from the greater conceptual graph, showing the different possible situations regarding rendez-vous of vessel <i>CHEKHOV</i> . The case when neither of the two sources is true is not shown here, but it corresponds to 6.5 only ($P = 0.0625$)	145
6.7	Valuation network for assessment of illegal activity for a cargo vessel	145
6.8	The full conceptual graph showing the most likely world including all the four vessels .	148
6.9	Sensitivity spaces for inference on illicit activity of <i>CHEKHOV</i>	151
6.10	Sensitivity spaces for inference on illicit activity of <i>CHEKHOV</i> using Dempster's rule of combination exclusively	152
6.11	Shapley values for contribution of the sources to the solution from Figure 6.8	152
6.12	Simplified, high-level relationships between inference targets (green) and directly relevant or auxiliary information (cyan).	153
6.13	The evidential network implementing the multi-source threat assessment reasoning .	156
6.14	Amount of evidence <i>supporting</i> (belief, lower bound), amount of evidence <i>consistent with</i> (plausibility, upper bound) and amount of evidence <i>against</i> THREAT (top graph) and its three components (bottom graph), for each of the four vessels	158
6.15	Sensitivity spaces for the impact of the individual reports on inference on THREAT . .	160
6.16	Sensitivity spaces for the impact of INTENT on inference on TREAT	161
6.17	Sensitivity spaces for the impact of the individual reports on INTENT	161
6.18	Source behaviour as variation of α and β , with directions corresponding to maximum positive changes of belief, plausibility, specificity and consonance	163
6.19	Sensitivity spaces for belief and plausibility with two-dimensional quality valuation and different directions of source uncertainty shift for source 2	164

6.20 Sensitivity spaces for inference on Threat where the parameters describing source behaviour are varied	165
6.21 The comparison of sensitivity of the inference result to the report with (blue) and without (red) the source model included	166

INTRODUCTION

The ever-increasing amount of data means that today's decision-makers and analysts are often faced with an overwhelming amount of information, arriving in a variety of formats from multiple sources. This thesis begun with a very open-ended research question—*what are the ways in which the sense-making in a general heterogeneous multi-intelligence scenario can be improved?* Is very clear that no single answer exists, but rather there are a variety of different aspects of the problem each of which can be investigated by itself. As such, this thesis addresses three major challenges which arise in the field of information fusion. The first is concerned with selection of information to use. This includes modelling the relationships between various variables of interest and understanding the performance of the different sources which provide information. Some information may exist which is relevant to a problem, but does not address it directly—such information is considered the *context* of the problem, and its exploitation is one of the main areas of focus of this thesis. Another challenge is understanding the fusion process. If its outcome is some decision it is important that the sources which contribute to it the most can be identified. This notion of identifying the origin of information is called *provenance* and itself, as well as more general generation of explanations, or *explainability*, is another major aspect of this thesis. Finally we must consider that the fusion problem often goes beyond identifying the true value of some variable, but often requires uncovering uncertain relationships between multiple entities. This problem of *situation assessment* is the third of the main notions discussed here. These three areas are presented in context of *maritime situational awareness* and discussed in light of relevant examples.

1.1 Main notions and their relations

As discussed above, this thesis addresses three key challenges associated with heterogeneous information fusion, but all of them are addressed through the lens of handling uncertain information. There is little doubt that handling information that is uncertain or imperfect to different

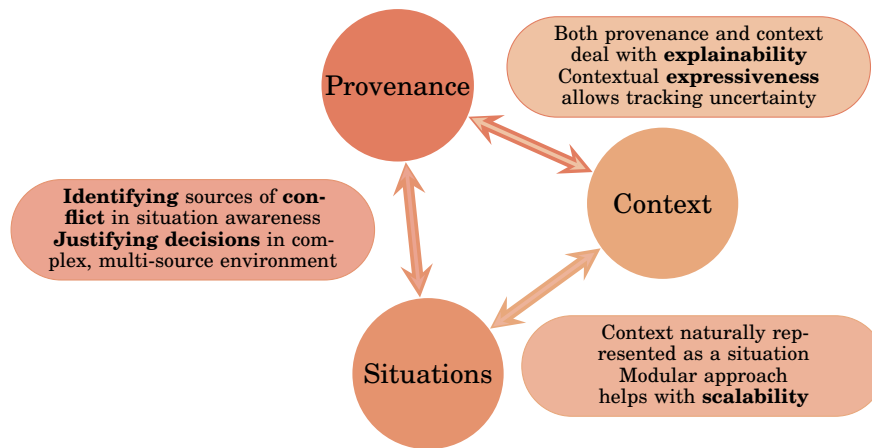


Figure 1.1: The relationships between the three main focus areas of this thesis

degrees is one of the most important, if not the most important, challenge within the field of information fusion. Questions of contextual reasoning, explainability, and situation awareness need to be addressed, taking into account the key issue of uncertainty. As such multiple methods of uncertainty representation are discussed at length. However, throughout this thesis, the bulk of work is focused on the belief function theory and usage of evidential networks for reasoning. Alternative approaches to uncertainty modelling, in particular probability and possibility theories, are employed at times, but the main focus remains with Dempster-Shafer theory of evidence and its extensions. The literature review and background discussion on uncertain reasoning are the focus of Chapter 2.

Figure 1.1 shows the relationships between the three main areas of interest introduced in the previous section. First, let us discuss these concepts in a little more detail. The first, context, as discussed earlier, is concerned with information which is relevant, but not directly involved in the fusion process. A typical example of a context variable in many maritime fusion problems is weather. While it may not be a part of the problem itself, it is likely to affect both the behaviour of the agents as well as the performance of the available sensors. This extended modelling of performance makes it possible to increase the expressiveness of the information source model, which thus can make it possible to track the origin of uncertainty better.

The provenance—and more generally the provision of explanations—can be considered one of the main challenges when it comes to deployment of automated decision support systems, in particular in safety-critical situations. In order for users to be able to employ such systems for decision-making, they need to trust it. This trustworthiness, among others, can be facilitated by explainability and the main method of provision of explainability in information fusion systems is by provenance—by showing which sources of information contribute to what extent to the particular decision as well as by showing how credible they are through their contributions to conflict. Context feeds into explainability directly by explicit modelling of additional variables, making it possible to assess the impact of contextual variables and source reports independently. However, additionally, as mentioned above, it can improve the tracking of uncertainty.

Finally we go back to the third challenge, the modelling of multi-entity situations. This is

particularly relevant to making sense of semantic information, where many entities or agents can be described with a variety of relations between them. In maritime situational awareness this can be illustrated with multiple vessels collaborating, possibly performing rendez-vous or being connected through some other property (the owner or the country). Alternatively this can also relate to a single entity with a variety of properties which need to be investigated in more detail. This example of multiple vessels which share an owner is an illustration of a particular context relevant in this scenario—the owner of the vessels may not be part of the fusion problem (e.g. the vessels' intentions), but is relevant to it and the resulting relationship between the vessels can be used to make inference about their intentions. Furthermore, as the complexity of the scenario increases it becomes even more important to be able to produce explanations and track the impact of individual information sources.

1.2 Thesis structure

In an uncertain reasoning scenario, context can be defined as a set of information relevant to the reasoning task at hand but not directly involved with it. This context can be exploited in the construction of the model. The benefit of this is two-fold: on the one hand, its explicit representation can provide better insight into the problem, and on the other hand, it may improve the expressiveness of the model and result in an inference which better represents available knowledge. Context exploitation can also improve the interpretability of the model and can be used to generate better explanations. The work on context stems from the existing models for partially reliable sources within evidence theory. Existing implementations of contextual belief correction are consistent with the notion of *context-of* some specific problem (the set of all information relevant to a particular problem), the approach used in this PhD work builds on the notion of *context-for*, additional information used to either improve the characterisation of sources or constrain the problem domain. A detailed discussion of the notions of *context-of* and *context-for* takes place in Chapter 3. While the notion of context has already been addressed to some extent in the existing literature on information fusion, it has not been to date formalised in the light of belief functions theory. The aforementioned contextual belief correction does not interpret context the way most other literature does. A more in-depth analysis of context, especially when discussing its impact on source modelling, forms a natural link with the problem of non-independent sources. These topics are discussed at length in Chapter 3 of this thesis. Most of the work from Chapter 3 was carried out at the Centre for Maritime Research and Experimentation as part of the Data Knowledge and Operation Effectiveness project under the supervision of Anne-Laure Jousset.

The second major challenge identified in information fusion is that of provenance or explainability. Chapter 4 discusses this problem. There are three main properties of the fusion result which may be tracked through the process. These include the information or decision itself, the conflict as well as uncertainty. Identifying the relative contribution of a source to belief in a hypothesis has been previously done using sensitivity spaces. In this thesis, some alternative approaches are proposed, using geometric interpretations of belief functions or interpreting the fusion process as a cooperative game and exploiting the Shapley value. This is similar to recent

developments in explainability methods in machine and deep learning, where the Shapley value is used as an explainability tool. Some properties such a contribution measure should satisfy are proposed, and the degree to which different metrics comply with the properties of the ideal contribution measure is used as a benchmark.

Furthermore, identifying the origin of and handling conflict is a major issue in information fusion, with several authors discussing methods of conflict management as well as computing the relative contribution to the conflict. Finally, in some information fusion systems, it can be useful to identify the relative impact of individual sources on the overall uncertainty

This notion is extended to encompass different source models, such as truthfulness and relevance as well as more complex models of source behaviour, by considering the different dimensions of source reliability to be a scalar field. This is one of the major links between the notions discussed in Chapter 3 and in Chapter 4. Furthermore, the additional expressiveness provided by contextual reasoning is itself a method of obtaining a greater degree of explainability or interpretability.

Chapter 5 addresses the notion of situation awareness and dealing with the fusion of information pertaining to multiple objects with some relationships between them. A degree of situation awareness is obtained by extending the notion of *conceptual graphs*, a graphical representation of logical relationships, to uncertain conceptual graphs, such that each concept node may have an uncertain value. A fusion strategy allowing for two different levels of reliability of a source providing the conceptual graph is developed—with the upper level, corresponding to plausibility or possibility, implying that some of the information provided is true, and the lower corresponding to the situation where all information provided by the source has to be true. A method of representing conceptual graphs as valuation networks has been proposed, allowing uncertain reasoning in these situations. Furthermore, links between Markov Logic Networks, valuation networks and uncertain conceptual graphs have been briefly investigated. The modular approach exploited throughout this thesis through the use of valuation networks and utilisation of context for determination of the scope of reasoning supports the scalability, an important aspect of multi-source, multi-target uncertain fusion.

Finally, two scenarios within the maritime situation awareness domain have been proposed to demonstrate the applicability of the methods developed. The first one deals is a situation awareness scenario with multiple sources reporting on multiple targets. Contextual information is used to model the sources, and different provenance techniques are used to support decision-making. The second scenario pertains to the assessment of a threat a vessel may pose to underwater infrastructure. More emphasis is placed on the explanation of the impact of individual sources and the assessment of the impact of source imperfections. These scenarios are presented in Chapter 6. The second scenario from Chapter 6 was developed at the Centre for Maritime Research and Experimentation as part of the Data Knowledge and Operation Effectiveness project with Anne-Laure Jousset and Maximilian Zocholl.

1.3 Key contributions

The key contributions of this thesis are as follows: in Chapter 3, a formal extension of the behaviour-based source correction model is proposed to include context. This is framed in relation to existing models of contextual reasoning and the notions of *context-for* and *context-of*. This makes it possible to also frame existing work on contextual belief correction in relation to the notions of *context-of* source and *context-of problem*. As such, this work creates a bridge between two different areas of research in information fusion whilst also standing in its own right as a method for contextual reasoning, increasing the expressiveness of the model while allowing better explainability. Furthermore, some methods of dealing with sources with unknown dependency relations are proposed and discussed.

In Chapter 4, a new method for assessing the contribution of a single body of evidence to the belief in a proposition is proposed through the exploitation of the geometric interpretation of belief functions and the geometric distance. Several properties a contribution measure should satisfy are proposed, which apply to contributions to information, conflict and uncertainty. One of these is not easily satisfiable and is instead used as a benchmark to compare this new contribution measure to other possible measures. The properties of this distance-based contribution measure are investigated, and it is shown to have relatively good performance regardless of the decision-making method used and belief combination method. This is unlike other measures, which tend to perform significantly better in some configurations but significantly worse in others.

Chapter 5 extends the notion of conceptual graphs to allow for embedding uncertainty in conceptual graphs. A fusion strategy allowing for two different levels of reliability of a source providing the conceptual graph is developed, and the notion of modal uncertainty embedded in graphs is briefly addressed. Links between the uncertain conceptual graph formalisms and Markov Logic Networks and valuation networks are established.

In Chapter 6, the explainability methods from Chapter 4 are framed in relation to source models discussed in Chapter 3. Interpretation of the different dimensions of source reliability to be a scalar field makes it possible to construct sensitivity spaces that can be used for the determination of the impact of source uncertainty on the reasoning.

1.4 List of research outputs

1.4.1 Poster presentations

1. BELIEF 2018 Compiègne - Provenance Across Belief Combinations with VAST 2018 use case, P.Kowalski, T.Martin
2. LANL Data and Information Fusion Conference 2019, Maritime Situation Awareness Workshop 2019, BFAS Summer School 2019 - Arms smuggling investigation: context-aware uncertain information fusion with evidential networks, P. Kowalski, A.-L. Jousset

1.4.2 First-author conference publications

1. 2017 13th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD) - Assessment of the contribution of a source of information to the final belief in a proposition, P.Kowalski, T.Martin
2. 2018 21st International Conference on Information Fusion (FUSION) - Provenance Across Evidence Combination in Theory of Belief Functions, P.Kowalski, T.Martin
3. 2019 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA) - Embedding Uncertainty in Conceptual Graphs for Semantic Information Fusion, P.Kowalski, T.Martin
4. 2020 IEEE 23rd International Conference on Information Fusion (FUSION) - Explainability in threat assessment with evidential networks and sensitivity spaces, P. Kowalski, M. Zocholl, A.-L. Josselme
5. 2021 IEEE 24th International Conference on Information Fusion (FUSION) - Investigating suspicious vessel behaviour in light of context, P. Kowalski, A.-L. Josselme
6. 2021 IEEE 24th International Conference on Information Fusion (FUSION) - Towards measuring information value in a multi-intelligence context, A.-L. Josselme, T.Wickramaratne, P. Kowalski

1.4.3 Journal publications

1. Information Fusion - Explaining the impact of source behaviour in evidential reasoning, P. Kowalski, M. Zocholl, A.-L. Josselme [accepted, special issue on XAI]
2. International Journal of Approximate Reasoning - Contextual Reasoning about Partially Reliable Sources with Evidential Networks, P. Kowalski, A.-L. Josselme [expected submission mid-2021]

BACKGROUND

“Probability is the most important concept in modern science, especially as nobody has the slightest notion what it means ”

— Bertrand Russel (*probably*)

The problem of information fusion is the assessment of the true state of the world given a set of sources providing information. The information in question may be uncertain, vague or imperfect in another way. The sources may present this information in different formats, using a variety of uncertainty representation frameworks. Finally, there may be multiple ways in which these information pieces may be combined to construct the image of the real world. In this chapter, some of the tools used within the field of information fusion are introduced. The first section is concerned with the concept of information fusion itself, its basic principles and generalised models. In the subsequent section, the concept of information quality is discussed with a focus on different types of imperfections and uncertainties. Furthermore, the uncertainty formalisms and models are addressed and discussed in detail before discussing the available tools for knowledge representation which may, to a different extent, implement the uncertainty models discussed. Finally, we address some of the more advanced concepts within information fusion: graphical models, ontology-based fusion and Markov logic networks.

2.1 The basic principles and generalised models of information fusion

Information fusion is a type of aggregation process, the purpose of which is to extract truthful knowledge from a number of incomplete or otherwise uncertain information sources. The following section discusses the key properties and principles of information fusion, as described by Dubois

et al. [1]. These properties are *unanimity*, *information monotonicity*, *consistency enforcement*, *optimism*, *fairness*, *insensitivity to vacuous information*, *commutativity* and *minimal commitment*. Note that as these are general properties, they may not necessarily be applicable to specific settings, particularly when considering imprecise and uncertain fusion.

The basic postulate of *unanimity* implies that the information sources unanimously agree on has to be assumed to be true - thus, the weakest form implies that all the information agreed to be *possible* remains possible, and all the information assumed to be *impossible* by all the sources remains impossible post-fusion (this is also referred to as possibility and impossibility preservation respectively). *Information monotonicity* implies that for two (or more) non contradicting sets of information sources, the fusion of a more informative set should result in more informative results. A consistent piece of information is a piece of information such that it does not contain a contradiction. Similarly information items are mutually consistent if either their supports overlap (weak mutual consistency) or their cores do (strong mutual consistency). *Consistency enforcement* means that fusion of individually consistent inputs should result in a consistent output (or at least, the result of their combination is not empty), *optimism* implies that (in the absence of information to the contrary) it should be assumed that all sources are equally reliable in agreement with their mutual consistency. *Fairness* implies that each input should propagate to the fusion output, *insensitivity to the vacuous information* means that vacuous information should not contribute to the final result and *minimal commitment* implies that the final result should contain as little information as possible whilst complying with the remaining properties.

Finally, when discussing the principles of fusion, it is important to differentiate between information fusion and similar yet different techniques. Primarily the difference between preference aggregation and information fusion is that the latter describes the state of the world based on a set of imperfect inputs, whereas the former describes the *ideal* world in accordance with sources representing individuals or criteria. Consider a simple example of weather. For the sake of information fusion, if two sources report the temperature to be *hot* and *cold* respectively, the purpose of fusion is to resolve this inconsistency and determine the real temperature. If two individuals are trying to decide on a holiday destination, a *mild* climate may be an acceptable trade-off between hot and cold - however, such a middle-ground approach is very unlikely to yield the appropriate answer in an information fusion scenario.

From a more formal viewpoint, Dubois et al. [1] proposed a general setting for information fusion. A generalised representation of information items is proposed, such that T_i represents an *information item*, which may be a subset of possible worlds, a plausibility relation, a fuzzy set, a mass assignment and similar. A *vacuous item* is an information item expressing total ignorance.

Let $\mathcal{T} = \{T_1 \dots T_n\}$ be the set of possible information items of a certain format. Again, according to [1] a n-ary *fusion operation* is a mapping $f^n : \mathcal{T}^n \rightarrow \mathcal{T}$. When it is not ambiguous we replace f^n with f (see¹).

First of all we consider a set of worlds, or possibilities Ω , one of which is true. Each information item T can be characterized as having the following features: its *support*, $S(T) \subseteq \Omega$, its core

¹Note the change from Dubois' article where the arity of the fusion operator is denoted by subscript

$C(T) \subseteq \Omega$ and an induced *plausibility ordering*. The notion of plausibility ordering corresponds to that of potential surprise[2] - the more plausible a state of affairs is, the less surprising its presence is. Alternatively a state of affairs is plausible if there is little indication that it would not be true. The *support* $S(T)$ corresponds to the set of all worlds, which are not impossible, whereas the core $C(T)$ corresponds to the set of all fully plausible worlds. This definition can be generalised to different frameworks for representing information items, such as sets, plausibility relations, Dempster-Shafer mass assignments and others, as is shown throughout this chapter.

2.2 Uncertainty modelling formalisms

The question of how uncertainty should be represented has been a subject of a long-time debate. Nowadays, there is little question that the theory of probability formulated by Rev. Thomas Bayes is the one most commonly used and well-understood throughout the world; taught in primary schools and adapted to be used in cutting-edge Artificial Intelligence solutions. At a deeper inspection, however, it is clear that the theory of probability is not the be-all and end-all solution to uncertainty quantification (even though some Bayesians would certainly disagree)[3]. For one, the commitment to numbers means that any two events must be comparable; on another, there is no way of actually quantifying ignorance or second-order uncertainty. Furthermore, when combining probability with logic, further issues arise, such as Lewis triviality (discussed later in this section). In this section, we explore the different approaches to expressing - and quantifying - uncertainty, including but not limiting ourselves to the probability theory and discuss how they can be used in an information fusion framework.

2.2.1 Set based and interval based approaches

Most representations of uncertainty begin with the concept of *possible worlds*, sometimes called *elementary outcomes*, a *frame of discernment* or in probability theory the *sample space*. Generally, these concern all the possible outcomes of the event under consideration. In one of the most common known random events, tossing of a six-sided die, we consider six possible worlds, each realisation corresponding to the side on which the die may land. We will sometimes refer to the set of likely or known worlds as *events* or *propositions*, and to the individual worlds as *hypotheses*.

Thus we may have the most coarse method of representing uncertainty. Consider some (large) set of possible worlds W and an agent expressing her knowledge on the true state of the world. The size of the set of the worlds $W' \subseteq W$ that the agent considers possible can be considered a qualitative measure of her uncertainty. Even though there is no method of quantifying the relative likelihood of each possible world, we still are able to model the agent's uncertainty to some degree. For example, consider the problem of identifying a seaborne vessel. One possible set of worlds may be $W = \{\text{military}, \text{commercial}, \text{leisure}\}$. It is clear that the granularity of this set of worlds is arbitrary. Clearly, we could be performing this identification on a lower level - the type of military vessel, the activity of the commercial vessel or even the Maritime Mobile Service Identity (MMSI) of the vessel itself.

An agent may describe her knowledge as "The vessel is not a military vessel", defining the possible worlds as $W' = \{commercial, leisure\}$. The agent considers *possible* every U such that $U \cap W' \neq \emptyset$ and *knows* every U such that $W' \subseteq U$. To illustrate if we consider a refinement of the frame of possible worlds such that $commercial = \{fishing, cargo\}$, the agent considers $U = fishing$ to be *possible*. On other hand if we consider a coarsened set of worlds $W = \{military, non - military\}$ the agent *knows* $U = non - military$

This concept of set-based uncertainty can be easily extended to the case where the set is non-discrete. Consider the case where the variable of interest is the vessel's speed. Consider here that the possible worlds lie within the range $W = [0, 60]$ knots. An agent describes his assessment of the vessel speed as being between 7 and 9 knots - $W' = (7, 9)$. As such, the agent considers the speed of 8 knots possible and knows that the speed is within $U = [6, 10]$.

2.2.2 Probability theory

Bayesian probability is undoubtedly the best-known approach to uncertainty representation and for drawing inferences in the presence of uncertainty. Continuing with the possible worlds semantics, we can consider probability to be a measure describing the likelihood that a world w out of the set of possible worlds $W = \{w_1, w_2, \dots, w_n\}$ is the real world. The agent assigns a number to each of the possible worlds; the greater the number, the greater the likelihood. In some cases, not all the possible worlds may be assigned a probability value; it is, however, assumed that the subset of W to which a probability is assigned satisfies closure properties.

Definition 2.1. An algebra over W is a set \mathcal{F} of subsets of W that contains W and is closed under finite union and complementation (if $U \in \mathcal{F}$ and $V \in \mathcal{F}$ then so are $U \cup V \in \mathcal{F}$ and $\bar{U} \in \mathcal{F}$)

This algebra \mathcal{F} over W is the domain of the probability measure. Conventionally, its range is the interval $[0, 1]$. Hence the following definition:

Definition 2.2. A *probability space* is a tuple (W, \mathcal{F}, P) , where \mathcal{F} is defined as above and P is a *probability measure* on W , a function $P : \mathcal{F} \rightarrow [0, 1]$. P also satisfies the following axioms:

P1. $P(W) = 1$

P2. $P(U \cup V) = P(U) + P(V)$ if U and V are disjoint elements of \mathcal{F}

A crucial theorem in probabilistic modelling is the well-known Bayes' theorem, describing the probability of event based on a set of conditions related to the said event. It is usually stated as follows:

$$(2.1) \quad P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

where $P(A|B)$ denotes the conditional probability of A given B .

An important aspect of Bayesian probability is that it does not allow for explicit modelling of ignorance. This can often be a disadvantage in information fusion, where the fact that there is information missing may in itself be information. Typically this lack of ignorance is modelled using

the principle of indifference, where if no knowledge on probabilities of some events is provided, they are assumed to be equally likely.

2.2.2.1 Justifying probability

This chapter has been opened by an alleged quote from Bernard Russell, claiming that "nobody has the slightest notion what it (probability) means". If we choose to quantify belief using probability, it is important to understand the meaning behind these numbers or, at the very least, justify their use. The classical approach to applying probability would be to reduce a situation to a number of elementary outcomes and assume each of them to be equally likely (principle of indifference). For instance, if it is known that given some geo-temporal conditions, 30% of vessels in an area are fishing ships, it can be safely assumed that if a vessel is observed on radar, in the absence of any other information, the likelihood of it being a fishing vessel is $P(\text{fishing}) = 0.3$, and by complementing $P(\text{not fishing}) = 0.7$ (as this particular vessel is equally likely to be any "elementary ship", and 30% of these are fishing vessels). The problem with this approach is that it is often not possible to find the set of "equally likely elementary outcomes", in particular when we do not have full knowledge about the situation. Some alternatives are relative frequencies or subjective assessments. The issue with the former is that it is often hard or impossible to extend it to all probabilistic situations, in particular complex, uncommon events. This is the attempt of constructing an objective assessment of probability. The subjective alternative assumes that there is no such thing as the objective notion, and it is simply assigned by an individual. Some attempts to argue that such selection should satisfy the axioms $P1$ and $P2$ have been made, most famous in terms of betting behaviour. Although the discussion of the philosophical argumentation for subjective probability may be out of the scope of this thesis, it is discussed again later in this chapter when discussing probability transformations of belief functions.

2.2.2.2 Bayesian fusion

In Bayesian statistics, probability is interpreted as a degree of belief, and hence all quantities involved in the fusion task are interpreted as being subjectively random. This allows for symmetric treatment of all the information involved. In this sense, any known value provided by an information source is a consequence of the quantity of interest adopting some "true" but unknown value. Hence the uncertainties are modelled probabilistically using the *degree of belief* interpretation. The transformation of sources of information to achieve a homogeneous mathematical representation can be performed using a variety of techniques, such as the application of the maximum entropy principle [4]. The fusion itself is performed through an application of the Bayes theorem (2.1).

As discussed previously, a downside of Bayesian methods is a requirement for a prior probability distribution that may be unknown. If this is not the case, a variety of vague (uninformative) prior distributions has been proposed in literature [5]. If the information delivered by the fusion sources dominates the fusion task, this can be relatively simple, as the posterior is relatively stable against

different choices of the prior. However, if this is not the case, sensitivity analysis typically needs to be performed using a variety of priors.

Consider the following simple example where a sensor provides an information on an object, classifying it either as type A or type B , so that the quality of interest $z = \{A, B\}$. N sensors provide this information, each with a reliability rating r , which is used to model uncertainty - $r \times 100\%$ of the time the sensor will rely correct information, and the remainder it will provide no useful information. Thus the information sources can be modelled as pairs $d_i = \{t_i, r_i\}$, where t_i is the reported object type. Thus using these definitions and the rules of probability the following holds:

$$P(d_i|z) = P(t_i, r_i|z) = P(t_i|r_i, z)P(r_i|z)$$

And following from the above definition of reliability:

$$P(t_i|r_i, z) = r_i \cdot \mathbb{1}_{z=t_i} + (1 - r_i) \frac{\mathbb{1}_{z=t_i} + \mathbb{1}_{z \neq t_i}}{2}$$

Bayes theorem can be used to derive:

$$P(z|d_1 \dots d_N) = \frac{P(d_1 \dots d_N|z)P(z)}{P(d_1 \dots d_N)} \propto P(d_1 \dots d_N|z) = \prod_{(i=1)}^N P(t_i|r_i, z) = \prod_{(i=1)}^N \left(\frac{1+r_i}{2} \mathbb{1}_{z=t_i} + \frac{1-r_i}{2} \mathbb{1}_{z \neq t_i} \right)$$

thus producing an expression for probability of the variable of interest given a set of information sources $d_i = \{t_i, r_i\}$ [6].

Until recently, a major weakness of the Bayesian statistics was that the calculation of posterior distribution was not computationally viable; as the number of variables increases, the calculation can become computationally infeasible. This has changed with the development of the Monte Carlo Markov Chain (MCMC) theory - a Markov chain which converges to a unique stationary distribution, which also is the posterior distribution of interest [7] [8].

2.2.3 Possibility and fuzzy sets

Possibility measures are another method of assigning numbers to possible worlds. First introduced by Lofti Zadeh and further developed by Didier Dubois and Henri Prade, it is both an extension of fuzzy sets theory and fuzzy logic as well as an alternative to probability theory for uncertainty modelling. In Zadeh's view it was meant to provide graded semantics for natural language statements, whereas Duibois and Prade developed it as a complete uncertainty model. A possibility measure Pos assigns a value in $[0, 1]$ to each subset of W , satisfying the following axioms:

Pos1. $Pos(\emptyset) = 0$

Pos2. $Pos(W) = 1$

Pos3. $Pos(U \cup V) = \max(Pos(U), Pos(V))$ if U and V are disjoint

The critical difference between the possibility and probability axioms is that the union of two disjoint sets $U \cup V$ yields the maximum in possibility theory and the sum of appropriate measures

in the probability theory. Furthermore, it can be shown that *Pos3* holds when the sets are not disjoint, whereas this is not true for the corresponding probability axiom.

Furthermore in possibility theory there exists a dual called necessity measure *Nec*. It is defined as follows

$$(2.2) \quad Nec(U) = 1 - Pos(\bar{U})$$

Fuzzy logic is a concept strongly related to possibility theory. It is a form of multi-valued logic, where the truth values may take the form of any number between 0 and 1. This is in contrast to Boolean logic, which enforces "crisp" values - only values of 0 or 1 are allowed.

The purpose of such fuzzy logic can be illustrated with the *sorites paradox*[9] (paradox of the heap). Consider a heap of sand from which individual grains are being removed. Whilst it would seem obvious that removing a relatively small number of grains does not turn the heap into a non-heap, it is also clear that at some unspecified point, the remaining grains do not constitute a real heap. In fuzzy logic, the membership function can be specified by essentially providing a lower bound on "definitely a heap"(1) and an upper bound on "not a heap at all"(0). If the number of grains is between these two bounds, the truth value is also larger than 0 but less than 1.

Fuzzy data fusion can be used in situations where vague or partial data is fuzzified using a gradual membership function and then standard fuzzy logic fusion rules can be used. Simple fuzzy fusion rules are either conjunctive rules (typically used when fusing data provided by equally reliable and homogeneous sources) such as intersection or product of fuzzy sets or disjunctive rules such as addition or intersection of fuzzy sets (used if at least one of the sources is unreliable, or the data is conflicting) [10]. The following equations represent fuzzy set intersection, product (conjunctive), sum and union (disjunctive):

$$(2.3) \quad \mu_1^\cap(x) = \min[\mu_{F_1}(x), \mu_{F_2}(x)] \quad \forall x \in X$$

$$(2.4) \quad \mu_2^\cap(x) = \mu_{F_1}(x) \cdot \mu_{F_2}(x) \quad \forall x \in X$$

$$(2.5) \quad \mu_1^\cup(x) = \max[\mu_{F_1}(x), \mu_{F_2}(x)] \quad \forall x \in X$$

$$(2.6) \quad \mu_2^\cup(x) = \mu_{F_1}(x) + \mu_{F_2}(x) - \mu_{F_1}(x) \cdot \mu_{F_2}(x) \quad \forall x \in X$$

To illustrate this, consider, again, a medical diagnostics scenario with a patient calling a medical helpline for advice. For the sake of simplicity, assume that the symptoms of some pathology (e.g. risk of a stroke) are severe headaches and high systolic blood pressure. Both these terms are fuzzy terms, and the former cannot be reliably measured in a crisp manner, whereas the latter is likely to vary between individual patients. If a patient reports "mild to severe" headaches and "higher than usual" blood pressure, the former could be interpreted as being a "severe headache" to a degree of 0.7, and the latter as "high blood pressure" to a degree of 0.6. Using a conjunctive fusion rule from equation 2.3, this could put the patient at risk of a stroke of 0.6. Based on this, the consulting doctor could decide that the patient is not in need of emergency medical attention but should see a GP as soon as possible.

Some adaptive fuzzy fusion rules have been developed as a compromise between strictly conjunctive and disjunctive fusion approaches. An example of such adaptive fusion rule has been

proposed by D. Dubois and H. Prade for fusing data from two sources with a varying degree of conflict:

$$(2.7) \quad \mu_{Adaptive}(x) = \max \left[\frac{\mu_i^\cap(x)}{h(\mu_{F_1}(x), \mu_{F_2}(x))}, \min [1 - h[\mu_{F_1}(x), \mu_{F_2}(x)], \mu_i^\cup(x)] \right] \forall x \in X$$

Where $h(\mu_{F_1}(x), \mu_{F_2}(x))$ measures the degree of conflict between the membership functions and $\mu_i^\cup(x)$ and $\mu_i^\cap(x)$ are the desired fuzzy fusion rules[11].

Fuzzy fusion is well suited to modelling the membership of a target in a vaguely defined class, as opposed to modelling uncertain membership in a well-defined class. However, similarly to Bayesian probability, it does require prior knowledge of membership functions for different fuzzy sets. [10].

The fusion of possibility distribution is based on the set-theoretic or logical view of information. The selection of fusion methods is dependent on how agreeable data sources are and what is known about their reliability; the basic conjunctive and disjunctive fusion rules discussed in the context of fuzzy fusion are appropriate only in a very limited number of scenarios[12]. Generally, if the purpose of fusion is to visualise the proportion of sources that treat some situation as possible, arithmetic mean or another type of mean of possible values may be most appropriate [12]. Possibility fusion is argued to be most appropriate in environments where no statistical data is available and in the fusion of heterogeneous data sources [11]. Although it is less popular than evidence theoretic or probabilistic methods, and hence less research on the topic exists, it has been shown to be capable of achieving comparable results [13]. Some areas where possibilistic data fusion approaches were proven to be successful include autonomous robot localisation in partially known environments[14][15] and risk assessment [16]. However, it is also argued that this method of fusion is not appropriate for merging conflicting evidence, especially when one piece comes from the general knowledge base and another is a piece of evidence directly conflicting the prior data [12].

2.2.4 Dempster-Shafer belief functions

The Dempster-Shafer theory[17][18] can be considered both as a generalisation of probability (by assigning the probability mass to sets of outcomes rather than to individual outcomes) and of possibility theory (through the induction of upper and lower probabilities).

Mathematically the theory of evidence is formalised as follows. Consider a set of possible, exclusive and exhaustive worlds, named *frame of discernment*, $\Omega = \{\omega_1, \omega_2, \dots, \omega_n\}$. The belief mass distribution (or Basic Probability Assignment – BPA, or mass function) is defined as

$$(2.8) \quad m : 2^\Omega \rightarrow [0, 1]$$

$$(2.9) \quad \sum_{A \subseteq \Omega} m(A) = 1$$

The mass associated with the empty set, $m(\emptyset)$, can represent hypotheses outside the frame of discernment, relaxing the exhaustivity hypothesis. Under the *closed world assumption*, we have $m(\emptyset) = 0$ while the *open world assumption* allows $m(\emptyset) \neq 0$. Some special cases of belief functions include the *simple support function*, the *vacuous belief function* and the *dogmatic belief function*.

Definition 2.3 (Simple support belief function). A simple support belief function is a belief function such that there exists only one set with non-zero mass which is not the universal set. As such a simple support mass function for $X \in \Omega$ with weight P is:

$$(2.10) \quad m(A) = \begin{cases} 1 - P & \text{if } A = X \\ P & \text{if } A = \Omega \\ 0 & \text{otherwise} \end{cases}$$

The notation X^P is sometimes used to denote such a belief function. Note that the *weight* refers to the mass on the universal set rather than the focus.

Definition 2.4 (Vacuous belief function). A vacuous belief function is a belief function with mass of 1 on the universal set, i.e. $m(\Omega) = 1$

Definition 2.5 (Dogmatic belief function). A belief function is *dogmatic* if $m(\Omega) = 0$

Finally, we may consider the probability distribution a special case of a belief function where all the mass is assigned to singletons.

Hence the *plausibility* and *belief* functions can be deduced, which may be interpreted as upper and lower probabilities, or, consistently with possibility theory as corresponding to the possibility and necessity measures, respectively

$$(2.11) \quad Bel(A) = \sum_{B:B \subseteq A} m(B)$$

$$(2.12) \quad Pl(A) = \sum_{B:B \cap A \neq \emptyset} m(B)$$

Belief functions on joint frames Just like in the case of probability theory, belief functions can be expressed on joint frames. Consider two frames of discernment $\Omega_X = \{x_1, x_2, \dots, x_n\}$ and $\Omega_Y = \{y_1, y_2, \dots, y_n\}$. We may consider the joint frame of discernment Ω_{XY} to be the Cartesian product of the two frames: $\Omega_X \times \Omega_Y = \{x_1 \times y_1, x_1 \times y_2, \dots, x_n \times y_n\}$

A joint belief mass for two variables X and Y is defined over $\Omega_{XY} = \Omega_X \times \Omega_Y$ such that $m_{XY} : 2^{\Omega_{XY}} \rightarrow [0, 1]$. This mass describes the joint partial knowledge over variables X and Y , representing the relationship between them.

The operations of *vacuous extension* and *marginalization* represent coarsening and focusing of knowledge respectively (see Figure 2.1). The vacuous extension of the mass m_X defined on the frame Ω_X to the joint frame Ω_{XY} , with $A \in \Omega_{XY}$ and $B \in \Omega_X$ is defined as follows:

$$(2.13) \quad m^{\Omega_X \uparrow \Omega_{XY}}(A) = \begin{cases} m^{\Omega_X}(B) & \text{if } A = B \times \Omega_Y \\ 0 & \text{otherwise} \end{cases}$$

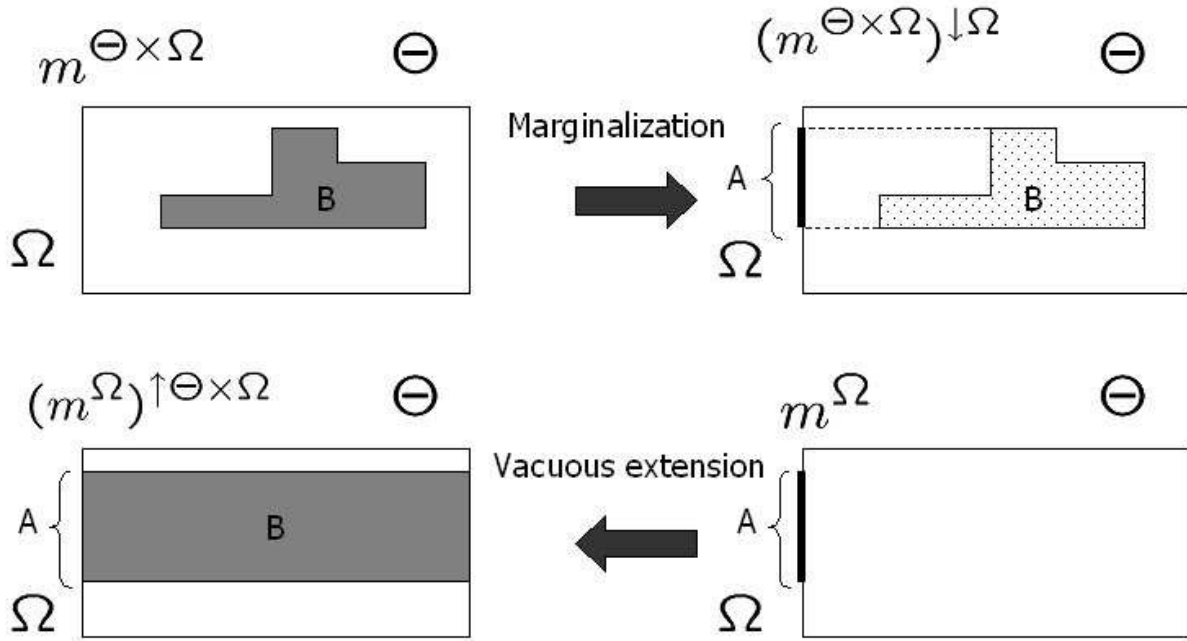


Figure 2.1: Marginalization and vacuous extension (from [19])

Marginalisation in the belief function framework is defined as follows, with $A \in \Omega_X$ and $B \in \Omega_{XY}$:

$$(2.14) \quad m^{\Omega_{XY} \downarrow \Omega_X}(A) = \sum_{B \subseteq \Omega_{XY}, B \downarrow \Omega_X = A} m_{XY}(B)$$

where $B \downarrow \Omega_X$ is the set projection of $B \subseteq \Omega_{XY}$ over Ω_X .

A property of the joint belief functions which is of particular interest is that they be used to encode uncertain logical relations through the truth set. Consider some logical formula ϕ , which holds with probability p . It can be encoded by the simple support belief function

$$(2.15) \quad m^\phi(A) = \begin{cases} p & \text{if } A = \{B : \phi(B)\} \\ 1-p & \text{if } A = \Omega_X \times \Omega_Y \\ 0 & \text{otherwise} \end{cases}$$

where $\{B : \phi(B)\}$ is the set of all $B \in \Omega_X \times \Omega_Y$ satisfying the formula ϕ .

For example, given frames $\Omega_X = \{x_1, x_2\}$ and $\Omega_Y = \{y_1, y_2\}$, a material implication $x_1 \rightarrow y_2$ can be written as $\neg(x_1 \wedge \neg y_2)$. Thus the focal set supported by the mass function describing this relation is $\{x_1 \times y_2, x_2 \times y_1, x_2 \times y_2\}$.

Discounting The problem of making inference on true state of the world given some source is discussed in more detail in the next chapter. However an operation which is often used in belief function theory is that of discounting. This process, sometimes referred to as Shafer's discounting redistributes a proportion of the mass assigned to the non-universal set to the universal set. It can be considered the most basic implementation of the reliability (or relevance) of the source. The

mass m is discounted by the factor α as follows

$$(2.16) \quad m^\alpha(A) = \begin{cases} \alpha m(A) & \text{if } A \neq \Omega \\ \alpha m(A) + (1 - \alpha) & \text{if } A = \Omega_X \end{cases}$$

Combination of evidence and normalisation The conjunctive rule of combination (CRC) combining different sources of information pertaining to the same frame of discernment was proposed in [20]. Using notation consistent with above, the joint mass is obtained as follows

$$(2.17) \quad m_J^C = (m_1 \odot m_2)(A) = \sum_{A=B \cap C} m_1(B)m_2(C)$$

The mass attributed to the empty set can be considered the conflict between the two masses:

$$(2.18) \quad m_J^C(\emptyset) = \sum_{\emptyset=B \cap C} m_1(B)m_2(C)$$

As such mass function can be normalized by forcing $m(\emptyset) = 0$ if the frame of discernment is considered to be exhaustive

$$(2.19) \quad m^\wedge(A) = \begin{cases} \frac{1}{1-m(\emptyset)}m(A) & \text{if } A \neq \emptyset, \\ 0 & \text{if } A = \emptyset \end{cases}$$

The normalized CRC is called Dempster's rule of combination and is denoted by $m_J^{DS} = m_1 \oplus m_2$.

Some other rules of combination exist². Probably one of the most widely known, albeit rarely used, is the disjunctive combination rule (DRC), proposed by Smets [21]. As per its name, the joint mass function is obtained as follows:

$$(2.20) \quad m_J^D = (m_1 \oslash m_2)(A) = \sum_{A=B \cup C} m_1(B)m_2(C).$$

One of the main criticisms of Dempster's rule is that it provides highly counterintuitive results when the sources of information are conflicting. Zadeh's example of disagreeing experts is probably the most well-known instance of such criticism. Essentially, when two sources of information are primarily conflicting, but there exists a tiny intersection, Dempster's rule will assign the entire mass to this intersection through normalisation. This problem is avoided when using the unnormalised CRC but at the cost of relaxing the closed-world assumption.

Yager's rule is another solution to this problem whereby the conflicting mass is assigned to the universal set (ignorance) rather than redistributed across all the focal sets [22]. The difference between using the CRC and Yager's rule can be attributed to the presumed cause of conflict between the two sources. If we use CRC, we assume that the two sources are reliable, and the disagreement comes from modelling issues, such as an incorrect frame of discernment. In this case, assigning the mass to the empty set corresponds to the belief that the true state of the world lies outside of the frame of discernment considered. In the case of Yager's rule, conflict resolution

²To the extent that it has become a running joke in the belief functions community that every aspiring belief functions researcher must propose their own rule for combining two belief functions

is obtained by assuming that the disagreement stems from the unreliability of the sources, and as such, the conflict is treated as vacuous information.

Whereas Dempster's rule requires the sources to be reliable, truthful and independent, un-normalised CRC allows the sources to be conflicting and disjunctive rule can be applied when it is known that some of the sources are not reliable or truthful (or rather at least one of the sources is reliable and truthful). Dubois and Prade[23] propose a generalised rule to the case when the sources are not independent. This variation of conjunctive rule (and, by extension, other set-operation based rules) can be used if the relationship between the two dependent mass functions is known. This means, the two mass functions to be combined are treated as if they were defined on different frames of discernment Ω and Ω' and the joint mass $m^{\Omega \times \Omega'}$ defining the dependency relation between the two masses m_1 and m_2 is known such that $m_1 = m^{\downarrow \Omega}$ and $m_2 = m^{\downarrow \Omega'}$

From this we obtain the Dubois-Prade conjunctive and disjunctive rules as follows:

$$(2.21) \quad m_J^{DPC}(A) = \sum_{A=B \cap C} m(B \times C)$$

and

$$(2.22) \quad m_J^{DPD}(A) = \sum_{A=B \cup C} m(B \times C)$$

with $B \in \Omega$ and $C \in \Omega'$.

Finally, another approach to handling non-independent sources is the cautious combination rule proposed by Denoeux [24]. This rule is different to Dubois and Prade's in that it does not require the knowledge of the dependency relation between the two masses. The idempotency of this rule is achieved by drawing from the canonical decomposition of a belief function, a process by which some basic belief assignments can be decomposed into a set of simple support functions and the Least Commitment Principle (LCP). It can be interpreted as using the two mass functions to generate a set of constraints and subsequently selecting the least informative mass.

The canonical decomposition of a non-dogmatic belief function is performed with the use of the *commonality* function q :

$$(2.23) \quad q(A) = \sum_{B \supseteq A} m(B)$$

which can be used to compute the weight of each *generalized simple support mass function*, with $w \in [0, \infty)$. This implies that the mass on the unique non-universal focal set may be negative, in which case it is no longer a real mass function, but can be interpreted as absorption of beliefs [25]:

$$(2.24) \quad w(A) = \prod_{B \supseteq A} q(B)^{(-1)^{|B|-|A|+1}}$$

and the cautious rule of combination defines the joint BBA as one with the following weight function

$$(2.25) \quad w_{1 \otimes 2} = \min(w_1(A), w_2(A)) \quad \forall A \in \Omega$$

and thus the mass can be obtained

$$(2.26) \quad m_1 \circledast m_2 = \bigoplus_{A \in \Omega} A^{w_1 \circledast 2}$$

2.2.4.1 Dissimilarity measures in the theory of evidence

A dissimilarity measure between BBAs can represent the degree of (dis)similarity between two different bodies of evidence. [26] provides a comprehensive review of various dissimilarity measures. Some, such as Jousselme's distance, are defined based on the geometrical interpretation of a BBA [27], whilst others use an alternative approach.

Geometrical interpretation of BBAs A body of evidence can be seen as a discrete random variable with its values being the members of 2^Θ and with the probability distribution m . While this interpretation is useful in many situations, allowing the usage of the random sets framework in the context of the theory of evidence [28], a given BBA \mathcal{B}, m has fixed values $m(A) \forall A \in \mathcal{B}, \mathcal{B} \subseteq 2^\Theta$. Hence it is possible to neglect the random aspect of the body of evidence \mathcal{B}, m and use the following geometric interpretation.

A vector space \mathcal{E}_{2^Θ} is generated by all the subsets of Θ . Therefore it is possible to define a BBA as a special case of a vector of \mathcal{E}_{2^Θ} , \vec{m} with coordinates $m(A_i)$ such that $\sum_{i=1}^{2^N} m(A_i) = 1; m(A_i) \geq 0$. Note that this definition does not impose the closed-world assumption.

Jousselme's distance Jaccard index, or Jaccard similarity coefficient is a statistic used for comparing the similarity (or diversity) of sample sets.

$$(2.27) \quad S(A, B) = \frac{|A \cap B|}{|A \cup B|}$$

$S(A, B) = 0$ implies that the sets have nothing in common; $S(A, B) = 1$ means they are identical. The Jousselme distance is an L_2 Euclidean distance with the Jaccard's matrix used as the weighting matrix. It is defined on \mathcal{E} as follows

$$(2.28) \quad d_J(m_1, m_2) = \sqrt{\frac{1}{2}(\vec{m}_1 - \vec{m}_2)^T \mathbf{Jac}(\vec{m}_1 - \vec{m}_2)}$$

where \mathbf{Jac} is the metric (weighting) matrix, a $2^N \times 2^N$ matrix with its elements $\mathbf{Jac}(A, B) = \frac{|A \cap B|}{|A \cup B|}$ [29], proven to be a strict distance matrix in [30].

Other geometric measures of dissimilarity exist and they are discussed at length in [26]

2.2.4.2 Decision making under belief functions

The general decision making problem is making a decision δ based on $m(\cdot)$. In its most classical variant, we make an assumption that $\delta \in \Theta$, i.e. the decision problem consists of choosing a specific element of the frame of discernment $\hat{\theta}$, which represents the set of solutions to the problem at hand. Note that in a more realistic scenario, the decision making problem would be based on the set of actions that would not necessarily be related by a one-to-one mapping to the frame of discernment. Many decision criteria have been proposed in the literature with varying levels of complexity and adoption. In this section, some of the most commonly used ones are presented.

1. Maximum of belief (pessimistic scheme) This decision making procedure is based off the so-called pessimistic scheme, where the element of the frame of discernment with the largest belief value is chosen

$$(2.29) \quad \hat{\theta} = \underset{i}{\operatorname{argmax}} Bel(\theta_i)$$

2. Maximum of plausibility (optimistic scheme) If a more optimistic - less cautions - approach is to be adopted the element of th frame of discernment with the maximum of plausibility is chosen (i.e. with the least amount of evidence directly against it)

$$(2.30) \quad \hat{\theta} = \underset{i}{\operatorname{argmax}} Pl(\theta_i)$$

3. Maximum of probability As the two previous approaches can be seen to be either overly cautious or optimistic a more balanced attitude is usually taken. In order to do this the basic belief assignment is transformed into a probability measure $P(\cdot)$ compatible with the belief interval $[Bel(\theta_i), Pl(\theta_i)]$ using one of many probabilistic transformations, such as the commonly used pignistic probability transformation[31]

$$(2.31) \quad P_{Bet}(\theta_i) = \sum_{\theta \in 2^\Theta} m(\theta) \times \frac{|\theta \cap \theta_i|}{|\theta_i|}$$

thus making it possible to choose $\hat{\theta}$

$$(2.32) \quad \hat{\theta} = \underset{i}{\operatorname{argmax}} P(\theta_i)$$

Other popular transformations of BBA to probability mass include taking the mass of singletons directly and normalising (thus ignoring partial ignorance) and plausibility-based methods [32].

Another theoretical approach, which is referenced later in this paper is to make decisions by using a strict distance metric $d(\cdot, \cdot)$ between two BBAs. The methods of implementation of these distances are discussed in the next section.

$$(2.33) \quad \hat{\theta} = \underset{\theta}{\operatorname{argmin}} d(m, m_\theta)$$

where m is the basic belief assignment based on which the decision is made, and m_θ is the BBA focused on θ i.e. $m(\theta) = 1$.

In their work on decision-making using distances, Dezert and Han [33] have proposed a quality indicator $q(\hat{\theta})$ defined as follows:

$$(2.34) \quad q(\hat{\theta}) = 1 - \frac{d(m, m_{\hat{\theta}})}{\sum_{\theta \in 2^\Theta \setminus \theta} d(m, m_\theta)}$$

which evaluates how good the decision is compared to all other possible decisions - in the case of classical decision problem to all other singletons in the frame of discernment. It can be seen that this indicator will reach its maximum when the BBA $m(\cdot)$ is fully focused on the decision and the higher the value of $q(\hat{\theta})$ the more confidence can we have in the decision $\hat{\theta}$.

2.2.5 Interpretations of belief functions

The original two interpretations of belief functions are that of upper and lower limits on probabilities (Dempster's interpretation) which was then rephrased as the sum of evidence by Shafer. However, there are several interpretations of belief functions that are distinct from Dempster-Shafer's Theory. Let us recall that the original Dempster's theory considered belief functions to be lower bounds of traditional probability distributions[20]. Shafer abandoned this view leaning towards the interpretation of belief functions as subjective opinion, with its meaning consistent with notions of "evidence", "doubt", or "support". This was taken further by Philippe Smets, who has proposed a different approach to the combination of belief functions in response to Lotfi Zadeh's criticism.

Transferable Belief Model The cornerstone of Smets' Transferable Belief Model[31] (TBM) is the division of reasoning levels into the *credal* level, dealing with beliefs and implemented using belief functions and the *pignistic* level, on which the decisions are made. The transformation between the two is performed with the pignistic probability transformation (Equation 3). The key difference in the treatment of belief functions, however, is the relaxation of the exhaustivity principle. The conjunctive combination rule without normalisation is used, and the probability mass on the empty set is allowed, corresponding to an unexpected outcome - an event outside the known frame of discernment. This interpretation of mass on the empty set became popular even among authors working outside TBM and not necessarily considering the transfer of beliefs between the credal and pignistic levels. This approach is used throughout this thesis unless stated otherwise.

2.3 Information quality and types of uncertainty

All information to be dealt with originates from some type of sources, be they *hard*, such as sensors or *soft* such as human reports. All kinds of sources tend to have some degree of information quality and imperfections associated with them. The individual sources may be more or less reliable, and the information they provide may or may not be credible. Furthermore, other dimensions of information quality exist, such as ambiguity, incompleteness, imprecision and others.

It is common to classify uncertainty as either epistemic or aleatoric, which roughly correspond to knowledge-related and randomness-related, respectively. The former can be reduced, either through the inclusion of additional or better information sources, change in the fusion mechanism or similar means, whereas the latter cannot. In the framework, we use the latter broadly corresponds to imprecision and the former to ignorance, as defined later in this section.

2.3.1 Source quality and information quality

It is self-explanatory that the concepts of information quality and source quality are very closely related as a source is mostly characterised by the quality of information it generates. Conversely, information of unknown quality can be characterised according to its originating source. Regard-

less, the two terms are not strictly equivalent. STANAG 2511[34] [35] distinguishes between the two, and each piece of information is jointly and independently characterised according to its source quality and information quality. The former, *source reliability*, is determined by the history of the past information provided by the source and the *credibility* of the information itself is determined by its relationship with other available pieces of information on the same topic and their respective qualities. As such, the reliability of the source is designated by a letter from A to F, where values A to E denote decreasing confidence in the source, and F denotes ignorance regarding its quality. Similarly, the quality of information is assessed on a numeric scale of 1 to 6, where values 1 to 5 denote decreasing credibility, and 6 implies the lack of ability to judge the truth.

2.3.2 Types of uncertainty and modelling of uncertainty

A large number of typologies exist for the classification of types of information quality dimensions in general or uncertainty specifically. Here we recall two such attempts at classification, one from Generalised Information Theory [36] [37] and a dictionary of different terms relating to data imperfection produced by Smets [38]. Note that the latter is consistent with vocabulary in [39], which follows from [40]. It is important to bear in mind that even considering these typologies, the exact meanings and definitions of some terms may vary. For instance, it is clear that in Smets typology in Figure 2.3, the meaning of "uncertainty" is somewhat different to its interpretation in Figure 2.2. Furthermore, it may seem at first glance that in Figure 2.3 terms such as uncertainty and imprecision are mutually exclusive; it is important to note that this is not the case. According to Smets, the distinction between imprecision, inconsistency and uncertainty is related to the model of possible worlds. An imprecise statement is satisfied by several worlds, and an inconsistent one is satisfied by none. An uncertain statement imposes ordering on the several worlds which satisfy it; as such uncertain and imprecise statements are closely related. These terms tend to become easier to distinguish when considered under specific uncertainty representation formalism - although different formalisms cover only some subsets of the variety of information quality measures discussed.

2.3.2.1 Crisp imprecision or nonspecificity

Crisp imprecision can be considered the first level of uncertainty expression and describes the (in)ability of the source to distinguish between its elements and providing a set of values it believes to be true. It is modelled by the source, providing a set of hypotheses within which the true hypothesis lies or, for numeric information, an interval containing the true value. For example, a classifier may be unable to distinguish between a tugboat and a fishing vessel. Similarly, all sensors produce imprecise readings according to their specifications. In crisp sets non-specificity can be quantified using Hartley function[41] and by its generalizations in different uncertainty modelling formalisms[42].

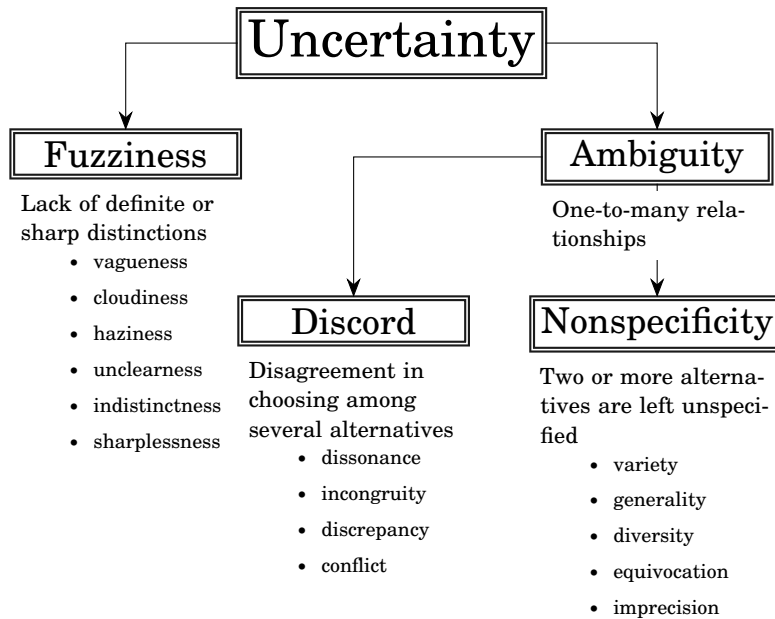


Figure 2.2: Three basic types of uncertainty - typology from [36]

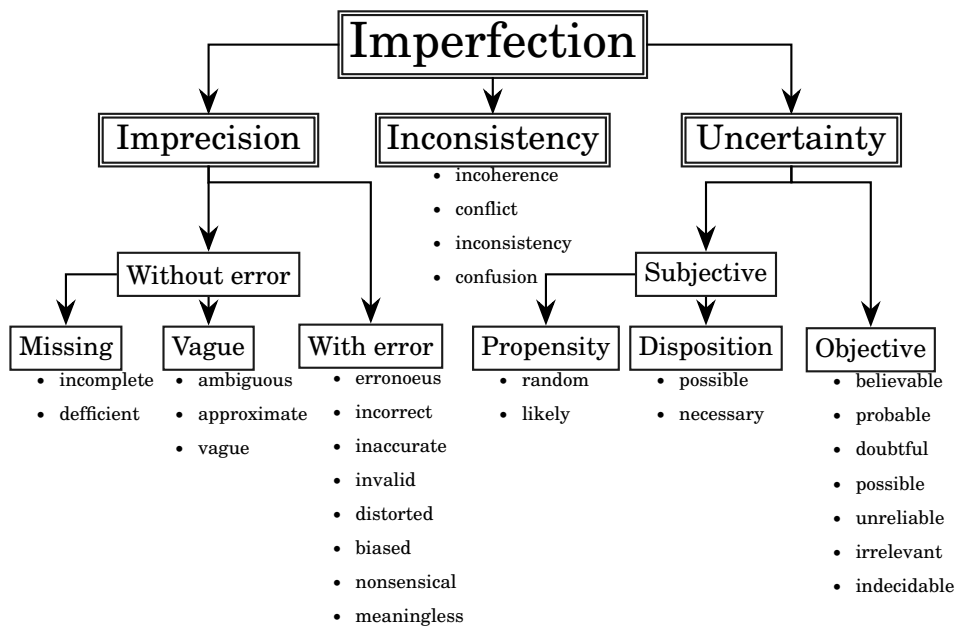


Figure 2.3: Uncertainty typology (or dictionary) adapted from [38]

2.3.2.2 Fuzziness

Fuzziness is a type of imprecision related to but distinct from nonspecificity. This is typically, albeit not exclusively, associated with human sources. For example, a human informant may have spotted a "large ship" rather than providing numeric values on its gross tonnage or length. This semantic vagueness means that the statement "ship X is large" usually is not either true or false but rather has some degree of truth assigned to it. This can be modelled using formalisms such as fuzzy set theory and possibility theory [43] [44]. Note that the measures of nonspecificity and

fuzziness are distinct and independent from one another. Furthermore, the two measures have a different impact on the overall information content. Whereas it is self-explanatory that any reduction in nonspecificity implies an increase in information content, the converse is not always true.

2.3.2.3 Discord

We consider discord to be the second level of uncertainty measures, modelled with monotone capacities such as probabilities, possibilities, belief measures and similar. Monotone measures are defined as functions g such that for any $A, B \in \Omega$ such that $A \subseteq B$ then $g(B) \geq g(A)$; they are often used to model some confidence degree in an event being true or false. Furthermore, the non-additive measures such as belief or possibility induce second-order uncertainty, defining an interval within which the true confidence degree lies. An objective probability distribution could be induced from a vessel density map, whereas a subjective one could be produced by a human agent expressing their beliefs.

The uncertainty content associated with such a measure can be quantified using appropriate measures such as Shannon entropy in the case of probability distributions or its appropriate generalisations in the theories of evidence and possibility. In some formalisms, such as the theory of evidence, several alternative measures may exist [45].

2.4 Graphical models

2.4.1 Conceptual graphs

Conceptual graphs are a framework that essentially allows for the representation of first-order predicate logic in graphical form [46]. This formalism is a model which encompasses a basic ontology and graph structure. There exist multiple extensions and re-imaginings of this model, allowing for definitions of rules and constraints; equalities between concepts and others.

Concepts are used to represent different entities to be modelled as part of a knowledge base. Each concept is labelled using two components: the conceptual type and a marker. The type defines the generalised entity, whereas the marker is used to identify a specific object in the world. The relation nodes indicate the relations established between different entities in the represented situation.

The vocabulary (or type hierarchy) is a simple ontology used to define the conceptual and relation types used to label concept and relation nodes, respectively. This can be formalised as $\mathcal{V} = (T_C, T_R, \text{markers})$, where T_C denotes the partially ordered set of concept types, T_R the partially ordered set of relation types and *markers* the set of individual markers (names) used to label the concept nodes. Furthermore $*$ denotes the *generic marker*. The three sets are assumed to be pairwise disjoint.

Therefore a basic conceptual graph G can be defined by a 4-tuple $G = (C_G, R_G, E_G, I_G)$, where C_G, R_G, E_G is a finite undirected and bipartite multigraph with two types of nodes - concept and relation denoted by C_G and R_G respectively, and E_G being the set of edges.

The final element of the 4-tuple. l_G is the naming function of the nodes and edges such that each concept node c is labelled with a pair $l_G(c) = t(c) \in T_C, m(c) \in \text{markers} \cup \{*\}$ and each relation node r is labelled by $l_G(r) \in T_R$.



Figure 2.4: An example of a conceptual graph describing the destination of fishing vessel Marcel

Conceptual graphs are discussed in greater detail in Chapter 4. Figure 2.4 shows an example of a conceptual graph, describing a relation of a *FishingVessel* (concept type) named *Marcel* (marker), which has a *hasDestination* relation with another concept node of concept type *Destination* with a generic marker.

This section discusses various approaches for graphical representation of relationships between uncertain variables and appropriate reasoning

2.4.2 Bayesian networks

Bayesian networks are tools for modelling with uncertain beliefs based on the probability theory described above, and in particular, they are useful for representing problems involving many related hypotheses. Such a network consists of two components: a directed acyclic graph (DAG) and a set of conditional probabilities - these can be thought of as the qualitative and quantitative components of a Bayesian network, respectively. The structure of a Bayesian network corresponds to the conditional independence between the propositional variables represented by the nodes[47]. An example Bayesian network from [48] is shown in Figure 2.5.




Figure 2.5: Structure of a Bayesian network for identification of critical situations between two ships

The above example illustrates a potential Bayesian network identifying some of the factors, which may determine whether a vessel poses a threat to another one. The conditional independence relations are such that the likelihood of warning is conditionally independent on the distance between the vessels and the collision likelihood, given that the degree of threat is known.

Some weaknesses of Bayesian networks are due to the principles of Bayesian probability modelling. Specifically, all the hypotheses and relationships are fixed in advance, and it is only the evidence that varies across problems. As multi-source information problems may involve uncertain numbers of entities, they may produce exponential sets of association hypotheses[49].

2.4.3 Markov networks

Markov networks, or Markov random fields, are undirected graphical models. It is described by a graph, where the nodes represent random variables. It can be considered similar to a Bayesian network in its representation of dependencies between variables, the main difference being that Markov networks are undirected and may be cyclic. As such, they can be applied to a wider range of problems where no natural directionality is associated with the dependencies, and some dependencies can be expressed by MRFs, which cannot be expressed by Bayesian Nets (and vice versa). A major downside is that the exact computation of probabilities for MRF's is NP-hard in the general case, and approximation techniques need to be used.

Whereas in Bayesian networks, conditional probability distributions are associated with each node, probability in Markov networks is computed using potentials over cliques. A potential is a function mapping every combination of variables in a clique to a positive real number. A probability of a world is then proportional to the product of all the potentials corresponding to the values taken by variables in that world.

Unlike in the directed Bayesian case, we are not concerned with how one variable is generated from the other, but rather the level of coupling between the dependent variables in the graph. As such, rather than needing to define how some variable A is constructed from B, we define the strength of interactions of these variables. More formally:

Definition 2.6. A Markov Random Field(MRF) is a probability distribution p over a set of variables $X = \{x_1, \dots, x_n\}$ with corresponding frames of discernment $\{\Omega_1, \dots, \Omega_n\}$ defined by an undirected graph G , the nodes of which correspond to the variables in X .

$$(2.35) \quad p(x_1, \dots, x_n) = \frac{1}{Z} \prod_{c \in C} \phi_c(x_c)$$

where C is the set of cliques and ϕ is a non-negative function over the variables in the clique (such that x_c is defined over $\prod_c \Omega_c$)

Furthermore the partition function Z is used for normalization:

$$(2.36) \quad Z = \sum_{x_1, \dots, x_n} \prod_{c \in C} \phi_c(x_c)$$

2.4.4 Markov Logic networks

A Markov logic network is a tool for bridging probabilistic representation and first-order logical reasoning. A relatively recent development proposed by Richardson and Domingos [50] uses a weighted first-order logic knowledge base as a template for the construction of a Markov Random Field. First Order Logic (FOL) is a collection of formal logic systems which unlike propositional logics allows for quantifiers and variables. As such it is possible to create statements (or rules) which generalise over many possible entities.

The reasoning behind weighting formulas in a knowledge base is that while many formulas are typically true in the real world, they may not always be true and, as such, are going to only capture a fraction of relevant knowledge or need to be excessively specific by including all the

possible exceptions. A classical example of such a first-order logic rule is "birds fly" or, in FOL $\forall x, bird(x) \implies flies(x)$. This rule can have an infinite number of *groundings*, expressions where variables are replaced with constants. In MLN we consider a finite set of constants, or values that x can take. Consider that we live in a world where only birds are seagulls, sparrows and penguins. In this case the earlier formula has 3 groundings and it actually holds true in two of them, as we know that $flies(sparrow)$, $flies(seagull)$ and $\neg flies(penguin)$.

Formally we have:

Definition 2.7. A Markov logic network L is a set of pairs (F_i, w_i) , where F_i is a formula in first-order logic and w_i is a real number. Together with a finite set of constants $C = \{c_1 \dots c_n\}$ it defines a Markov network $M_{L,C}$ as follows:

- $M_{L,C}$ contains one binary node for each possible grounding of each predicate appearing in L . The value of that node is 1 if the ground atom is true and 0 otherwise.
- $M_{L,C}$ contains one feature for each possible grounding of each formula F_i in L . The value of the feature is 1 if the formula is true and 0 otherwise and the weight of the feature is w_i

The probability of a world specified by the ground Markov network $M_{L,C}$ is given by:

$$(2.37) \quad p(x_1 \dots x_n) = \frac{1}{Z} \exp\left(\sum_i w_i n_i(x)\right)$$

which clearly follows from the exponential form of Equation 2.35.

A set of assumptions is in place to ensure that the set of possible worlds for (L, C) is finite and that $M_{L,C}$ represents a unique, well-defined probability distribution over these worlds. These assumptions are that of unique names (different constants refer to different objects), domain closure (only objects in the domain are representable using constant and function symbols in (L, C)) and known functions (for every function in L its value applied to any possible arguments is known and is in C). In particular, this last assumption makes it possible to replace functions by their values when grounding formulas.

2.4.5 Valuation-based systems and evidential networks

An evidential network is a graphical model representing knowledge about a set of uncertain variables modelled using belief functions [51]. Nodes in the graph represent uncertain variables, while the valuations represent uncertain relationships between subsets of variables or uncertain evidence about a single variable. It is the realisation of a Valuation-Based System (VBS) [51] (or valuation network), consistent with the evidence theory, where valuations are encoded with mass functions.

A valuation algebra is a tuple $(\mathcal{X}, \mathcal{V}, \odot, \downarrow)$ where \mathcal{X} is a set of uncertain variables of interest, \mathcal{V} is a set of valuations, \odot is an aggregation operation between valuations and \downarrow is a focusing operation.

Let X and Y be two uncertain variables defined over the frames Ω_X and Ω_Y respectively. A joint mass function for X and Y is defined over $\Omega_{XY} = \Omega_X \times \Omega_Y$ such that $m_{XY} : 2^{\Omega_{XY}} \rightarrow [0, 1]$. This mass function describes the joint partial knowledge over variables X and Y , representing the relationship between them.

The conjunctive combination operation is implemented using a modified version of Equation (2.17). If m_X and m_Y are the two joint mass functions for variables X and Y respectively, with $A \in \Omega_{XY}$, $B \in \Omega_X$ and $C \in \Omega_Y$, then:

$$(2.38) \quad (m_X \odot m_Y)(A) = \sum_{B \cap C = A} m_X^{\Omega_X \uparrow \Omega_{XY}}(B) m_Y^{\Omega_Y \uparrow \Omega_{XY}}(C)$$

where $m_X^{\Omega_X \uparrow \Omega_{XY}}$ denotes the *vacuous extension*, of the mass function m_X defined over the frame Ω_X to the joint frame Ω_{XY} , with $A \in \Omega_{XY}$ and $B \in \Omega_X$:

$$(2.39) \quad m^{\Omega_X \uparrow \Omega_{XY}}(A) = \begin{cases} m^{\Omega_X}(B) & \text{if } A = B \times \Omega_Y \\ 0 & \text{otherwise} \end{cases}$$

Marginalisation corresponds to focusing of knowledge and is defined as follows, with $A \in \Omega_X$ and $B \in \Omega_{XY}$:

$$(2.40) \quad m^{\Omega_{XY} \downarrow \Omega_X}(A) = \sum_{B \subseteq \Omega_{XY}, B \uparrow \Omega_X = A} m_{XY}(B)$$

where $B \uparrow \Omega_X$ is the set projection of $B \subseteq \Omega_{XY}$ over Ω_X .

The fusion problem is solved by extending all mass functions to the joint space, and marginalising over the variable of interest. Let denote by $Y \in \mathcal{X}$ the variable of interest, by Ω the joint frame of all variables, and let m be a valuation of \mathcal{V} , then the result of the fusion problem is:

$$(2.41) \quad m^{\Omega_Y} = \left(\odot_{m \in \mathcal{V}} m \uparrow \Omega \right) \downarrow \Omega_Y$$

In practice, axioms for local computation implemented by a VBS avoid ever computing the joint and rather use belief propagation schemes within binary joint trees [51].

Joint belief functions can be used to encode uncertain logical relations through the truth set. Consider some logical formula ϕ which holds with probability p . It can be encoded by the simple support belief function:

$$(2.42) \quad m(A) = \begin{cases} p & \text{if } A = \{B : \phi(B)\} \\ 1 - p & \text{if } A = \Omega_X \times \Omega_Y \\ 0 & \text{otherwise} \end{cases}$$

where $\{B : \phi(B)\}$ is the set of all $B \in \Omega_X \times \Omega_Y$ satisfying the formula ϕ . In this paper we use the following propositional logic symbols: \leftrightarrow denotes equivalence, \rightarrow denotes material implication, \wedge and \vee denote logical conjunction and disjunction respectively and \neg denotes negation. For example, given frames $\Omega_X = \{x_1, x_2\}$ and $\Omega_Y = \{y_1, y_2\}$, a material implication $x_1 \rightarrow y_2$ can be written as $\neg(x_1 \wedge \neg y_2)$. Thus the focal set supported by the mass function describing this relation is $\{x_1 \times y_2, x_2 \times y_1, x_2 \times y_2\}$.

2.5 Context awareness

The notion of context itself has been defined by various authors in different ways. Generally speaking, from an information fusion perspective, context is information that does not directly characterize the focal element of the problem but rather its surrounding or situation. The approach to the context in this thesis is most similar to that from [52], where context variables are exogenous to the problem, but the endogenous variables may be dependent on them. Some possible context types may include user context (requirements), physical context (time of the day, weather, infrastructure), device context (available information sources and their quality) and more[53][54].

Several works on aligning uncertain reasoning and information fusion with contextual information exist. In the probabilistic framework, the Markov Logic Networks (MLN) can be used to incorporate context [55]. Furthermore, in probabilistic reasoning, there exist the concepts of contextual independence and hence contextual belief networks [56]. In evidence theory, we have the work of Mercier et al. on contextual discounting and, more generally, on contextual information correction [19]. Gundersen discussed the relation between context and situation awareness, distinguishing between agent and situation context; however, with little interest in handling uncertainty and conflict [57]. All in all, the purpose of the context within information fusion is largely adaptation and explainability, as discussed in a recent survey by Snidaro et al. [58].

In all of these approaches, the exact meaning of context varies and is not necessarily consistent with its usual definition. In the existing work using MLN context is treated just like any other piece of evidence, and relatively little emphasis is placed on what makes context contextual. The contextual belief networks share the shortcomings of other Bayesian approaches, such as dealing with incomplete information, modelling ignorance and logical reasoning. Contextual discounting in belief functions framework does not deal with what we would consider the context in an information fusion sense; its context is actually the true state of the world. In this chapter, a new model for contextual reasoning with partially reliable sources is proposed, placing significant emphasis on understanding what context is and how it can be exploited within the belief function framework.

2.5.1 Context in artificial intelligence

The notion of context has largely entered the computer science mainstream as part of the concept of pervasive or ubiquitous computing [59], as part of which the paradigm of concept-aware systems was developed. The underlying concept is that pervasive computing provides services to the user regardless of the environment (as opposed to the 20th-century vision of localized computing), whereas context-awareness ensures that the right service is provided to the user at the appropriate time and place, adapting its behaviour without explicit user intervention [60]. The exact meaning of context, however, does not have a single universally agreed-upon definition. Regardless, the key ideas expressed in the variety of definitions converge. In general, the notion of context can be summarized as information that is relevant to the problem at hand but does not explicitly define it. This could be a political situation in an arms trade investigation, legal regulations when tracking an illegal fishing ship or meteorological data for assessment of radar data quality.

Some most widely accepted definitions consistent with these generalised notions have been proposed by Dey [61]: "Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves." and Brézilion[62] "Context is what contains a problem solving without intervening in it explicitly." Winograd[63] comments on Dey's definition and criticizes its broadness in reference to *any* information. His definition "The context is a set of information. This set is structured, it is shared, it evolves and serves the interpretation" restricts it by enforcing structure and relevance to the application. McCarthy [64] proposed logical representation of context as formal objects. Without going into the detail of this representation, the basic relation asserts a proposition in some context; however, the assertion itself is always realized in a context. From this, some interesting consequences follow. First of all, context is always relative to another context (contextual information cannot exist without some context of its own). From this, it follows that contexts have infinite dimension and hence cannot be described completely. Finally, with several contexts exist in a discussion, there is a common context above them [62]

Another approach takes into account two alternative paradigms of context derived from the very definition of that word in the English language. *Context* can mean either "that which unites or binds together" or "that which surrounds and gives meaning to something else". Based on this Gong [65] proposes a distinction between *context-for* (CF) and *context-of* (CO). With reference to some situation X, *context-of* X may refer to all the things within the space of X, whereas *context-for* X describes all the external factors which may affect X. In maritime domain we may consider *context-of* illegal fishing - in such context, we are concerned with fishing quotas, vessels' fishing capabilities, marine protected areas (MPAs) and such. On the other hand *context-for* such a problem may include season or meteorological conditions. In [66] Steinberg discusses the difference between the two, placing more emphasis on the type of reasoning. In this case, CF starts with a specific problem (such as inferencing or decision) and attempts to discover additional information to resolve uncertainties. CO starts with the perceived situation to derive information about the situation. As such, the purpose of CO reasoning is to assess expectation; the objective of CF is to handle goal-driven reasoning.

Steinberg also focuses on context from a more ontological or semantic perspective [66][67]. Situations are defined in terms of entities, relationships and activities in a method very similar to ontologies or conceptual graphs. In this framework, "context is a situation that provides information that can be used either a) to condition expectations or b) to improve the understanding of a given inference or planning/control problem".

Pomerol and Brézilion[68] briefly discuss the relationships between knowledge and context, identifying *contextual knowledge* as a subset of all knowledge. Both general and contextual knowledge can be explicit or implicit and can be made explicit; the particular differences include that contextual knowledge is this part of knowledge that is useful for describing the nature state preceding decision and hence may have several realizations. Thus while knowledge is fixed, the procedural context may change.

2.6 Summary

The purpose of this chapter was to provide an extensive review of different concepts and tools associated with the fusion of uncertain information. We have first discussed the generalised framework for information fusion as proposed by Dubois et al. Subsequently, we have discussed the imperfections in data and the different frameworks for uncertainty representation. The Dempster-Shafer Theory, or the theory of belief functions and its variants, were discussed extensively. We have also briefly discussed how the various uncertain reasoning formalisms could be aligned with the theory of belief functions through the utilisation of the notion of random sets. Subsequently, we briefly addressed the problem of information representation and storage and discussed the main graphical reasoning models - Bayesian Nets, Markov Logic Networks and, in particular, Evidential Networks. Not all of the concepts discussed are used throughout the remainder of this thesis. Still, they are discussed here to position the work done throughout this thesis in relation to literature as well as to make it easier to see the possible avenues for future work. In particular, the notions of valuation networks and belief function theory are used extensively throughout the remainder of this thesis. However, we are also concerned with the generalised notions of uncertain information fusion, especially in Chapter 4. Markov logic networks, conceptual graphs, and possibility theory are again discussed in Chapter 5 for their use in situational awareness.

REASONING WITH PARTIALLY RELIABLE SOURCES AND CONTEXT

Disclaimer: The contents of this chapter have been largely produced in collaboration with NATO STO Centre for Maritime Research and Experimentation (CMRE), under the supervision of Anne-Laure Jousselme as part of the Data Knowledge and Operational Effectiveness project. A significant portion of this content has been published in the CMRE technical report "Multi-source contextual reasoning for vessel behavioural analysis"

Information fusion, as discussed in the previous chapter, is concerned with the combination of information from multiple sources to get a better estimate of some variable of interest. However, in most cases, the sources do not provide information on the variable of interest directly but rather describe some observation of it or some other variable that can be used for reasoning about the true state of the variable of interest. As such, models for reasoning with partially reliable sources or reasoning generally are an important part of the body of knowledge pertaining to uncertain information fusion. The importance of Contextual Information (CI) within the world of Information Fusion (IF) stems from the fact that IF systems are typically designed to work in "well-defined conditions", exploiting both observed data and a priori knowledge. In some cases, this world behaviour can be very complex and cannot be defined sufficiently without constraining the problem domain using some form of contextual information. In maritime situational awareness, some examples include assessing the quality of a radar reading within the context of weather conditions or identifying events that may be meaningless without context (such as illegal fishing within the geographical context of marine protected areas).

3.1 Context in information fusion

As in the case of context nomenclature in the field of information fusion remains somewhat ambiguous and inconsistent. As such we are not concerned with the distinction between *knowledge*,

data, sensor and *information* fusion and in this thesis exclusively use the latter term. It is widely agreed on that the Joint Directors of Laboratories (JDL) fusion model has gained most acceptance in the information fusion community - with several revisions and modifications proposed by different authors[69][70]. Snidaro et al. [71] provide a detailed review of context-based information fusion applications across the different levels of JDL framework. Some notable examples include: contextual sensor characterization (JDL Level 0) [72], contextual maritime data association and filtering (JDL Level 1) [73][74], maritime situational awareness (JDL Level 2) [75][76] as well as a multi-level (JDL Levels 0-3) context-based harbour surveillance application [77].

In many cases, contextual information in fusion can be interpreted as a constraint, consistent with one of the definitions provided earlier. These constraints can range from hard or physical to soft or procedural, with geographical context (maps or routes) being an example of the former and ship blacklisting an example of the latter [71]. However, there are also situations where the context may be interpreted as additional features, semantics or situation elements that do not directly reduce uncertainty in the situation space. A typical example would be the detection of anomalous situation, which changes the context and increases the number of possibilities. Another situation where context does not necessarily constrain the problem space is when the context is used to determine information quality which may, again, increase uncertainty in the problem space [53][78].

3.1.1 Categories of context

There clearly exists a vast diversity of categories and applications of context in the field of information fusion. The work of Razzaque [54] addresses the problem of classification of contextual information. Six conceptual contextual parameters are identified. The author justifies the need for this with several advantages, including better modelling of information quality, improvements in source selection, context refinement and context manipulation in order to make the transition from low-level to high-level processing and fusion more straightforward. We would add to this list also the explainability and interpretability of both the reasoning process itself and its outcomes. The proposed categories are *user context*, *physical context*, *network context*, *activity context*, *device context* and *service context*. Examples of each of these six categories at JDL levels 1 to 3 in the Maritime Situation Awareness domain are shown in Table 3.1

3.2 Models for partially reliable sources

The information which needs to be handled in many fusion problems is provided by a number of sources - ranging from most simple physical sensors to machine learning classifiers and unstructured human intelligence. A physical sensor may be simply unreliable or may have been tampered with; classifiers performance is determined by a confusion matrix, and information provided by human witnesses may be vague or even untruthful. These sources vary not only in accuracy and reliability but also potentially in truthfulness and, in some cases, may display more complex behaviours.

Table 3.1: Examples of the six categories of contextual information in MSA domain, adapted from [53] and [54]

Context	JDL Level 1	JDL Level 2	JDL Level 3
User	Analyst tracking targets	VTS operator	Intelligence officer
Physical	Coast lines, sea state	Anchorage ares, channels	Critical infrastructure
Network	Network of radars	Communications network	Intelligence analysts
Activity	Usual tracks or routes	Season, routes, patterns	Recent piracy incidents
Device	Radars, satellites	AIS coverage	Availability of sources
Service	Location, heading	Anomalies of interest	Standard identities

We need to consider the disconnect between the source report, the source knowledge and the true state of the world. This introduces two levels of noise or uncertainty associated with a source report - the difference between the true state of the world and the source perception of it and the potential disconnect between what the source knows and the contents of its report. In general, the former can be considered the source reliability and the latter the truthfulness (albeit this is a simplified view as there may be other types of uncertainty involved - such as a human source unable to express its knowledge due to semantics). This general view is shown in Figure 3.1 [79].




Figure 3.1: Reporting process

When sources provide unreliable information about the domain of interest, it is necessary to somehow incorporate this reliability in the uncertainty representation. In the theory of evidence, this is typically done using one of many discounting operations, most often the classical discounting proposed by Shafer. However, while this does allow for the inclusion of the reliability assessment in the uncertainty structure, it is treated as part of the source report rather than an external variable or metaknowledge.

Some authors have discussed modelling this more explicitly. One of the earlier works discussing partially reliable sources or witnesses in uncertainty modelling context is that of Bovens and Hartmann [80], where a witness, whose reliability REL is a random variable provides a report REP on some hypothesis HYP . This approach naturally translates into evidential reasoning - if the witness is unreliable, she or he provides no information whatsoever on the variable of interest. Hence Haenni and Hartmann [81] extended this model to postulate a more general approach based on Dempster-Shafer theory.

Consider the simple example of an Automatic Identification System (AIS) broadcast over a lossy channel. The AIS is a transponder on board of ships that automatically emits vessels' information but that can easily be switched off by vessels' captains. Useful to avoid collisions,

the AIS reveals vessels positions that vessels' masters may wish to keep hidden for some periods of time. If the signal is broadcasted it may or may not be received depending on the conditions, however it cannot be possibly received if it is never broadcasted. The variable *HYP* corresponds to the AIS broadcast whereas the variable *REP* corresponds to the reception.

Under Haenni and Hartmann model this source report is generated based on the true state of the hypothesis *HYP* and the source reliability *REL*. There exists some probabilistic relationship between *HYP* and *REP*, which is well-defined but can be different depending on the reliability of the source. Haenni and Hartmann list an extensive collection of possible source behaviours and a complete discussion is beyond the scope of this chapter. However the most complete model relevant to us in this example is the *Perfect / Indicator* model. If the source is reliable, it provides perfect information, otherwise it can generate false positives and false negatives, based on some random variables *P* and *Q* which determine the likelihood of false positive or false negative.

$$\Sigma_{(PD)} = \left\{ \begin{array}{l} Rel \implies (Hyp \iff Rep) \\ \neg Rel \implies (Hyp \implies (P \iff Rep)) \\ \neg Rel \implies (\neg Hyp \implies (Q \iff Rep)) \end{array} \right\}$$

For the AIS example we assume that it is impossible for a false positive to be provided (we are only concerned with missing vessels rather than spoofed vessels) and as such *Q* is always *False*, whereas *P* is some random variable with its distribution being part of the model.

A more general Belief Based Correction (BBC) model was proposed by Pichon et al. [82]. It is based on the idea that the source can be in one of the possible states $Q = \{q_1..q_n\}$, which defines the inference on *X*, the true state of the world given the source report *Y*. Formalized with multi-valued mapping

$$(3.1) \quad \Gamma_A : Q \rightarrow X \forall A \subseteq Y$$

This can be represented by the categorical mass function:

$$(3.2) \quad m_{\Gamma}^{Q \times X \times Y} \left[\bigcup_{q \in Q, A \subseteq Y} (\{q\} \times A \times \Gamma_A(q)) \right] = 1$$

which can be combined with the other pieces of evidence: a mass function m^Y on *Y* provided by the source and the mass function m^Q on *Q* describing the metaknowledge on the source. The mass $m_{\Gamma}^{Q \times X \times Y}$ has vacuous marginals on *Q*, *Y*. The marginal on *X* is vacuous if the entire frame of *X* is within the image of Γ .

For the AIS example we can consider two source behaviour models: one where the channel is lossless and one where it is lossy.

The variable of interest *X* denotes whether the signal is transmitted $\Omega_X = \{t, \neg t\}$, and the report variable *Y* denotes whether the signal is received or not $\Omega_Y = \{r, \neg r\}$. The channel can lose packets with some likelihood and hence has two behaviour models $\Omega_Q = \{l, \neg l\}$. Thus the source is a perfect source if the channel is lossless ($Q = \neg l$) or always-negative if it loses packets ($Q = l$). Therefore the multi-valued mapping Γ maps the report and source behaviour to a set of hypotheses:

$$\begin{aligned}
(3.3) \quad & \Gamma_r(l) = t \\
& \Gamma_r(\neg l) = t \\
& \Gamma_{\neg r}(l) = \{t, \neg t\} \\
& \Gamma_{\neg r}(\neg l) = \neg t
\end{aligned}$$

which can be represented as a categorical mass function m_Γ :

$$m_\Gamma \left(\left\{ \begin{array}{l} l \times r \times t, l \times \neg r \times \Omega_X \\ \neg l \times r \times t, \neg l \times \neg r \times \neg t \end{array} \right\} \right) = 1$$

Unlike Haenni and Hartmann's model, where the variables of interest are all binary, there is no constraint on the cardinality of X and Y ; furthermore, the masses X , Y and Q can be non-Bayesian as such BBC model makes it possible to define a wide variety of source behaviours as long as they can be captured in a mass function on the joint frame $Q \times X \times Y$.

Although the BBC model is proven to be equivalent to an evidential network implementation with three variables, the general underlying assumption is that the source behaviour is entirely determined by metaknowledge. Even though it may be uncertain, it is not considered to be an output of some reasoning process. As such, the accuracy of the source is determined only by the prior metaknowledge and possibly the true state of the world, in a manner similar to the contextual information correction concept introduced in [19].

Here we use the BBC model as the main reasoning tool; however, consider the "source behaviour" variable Q to describe not only metaknowledge but rather be affected by other elements of the model. We use X to denote the variable of interest and R to denote the source report, corresponding respectively to X and Y in BBC and to Hyp and Rep in Haenni and Hartmann's model. Throughout the rest of this thesis, this model is referred to as the XQR model. The XQR and the BBC model are mathematically equivalent for the purpose of describing source behaviours and making inference on some variable of interest. The difference between the two is semantic: the BBC model places emphasis on the relationship between the source report and the true state of the world in light of some metaknowledge, whereas the XQR model treats the information about the source and information provided by the source as two equally important sources of information. Furthermore, the former is not constrained to be static metaknowledge; it could be an output of a reasoning process or dynamically obtained otherwise. As such, while an evidential network is a possible representation of the BBC model, the XQR model is primarily considered to be a module forming part of a larger evidential network. This places more emphasis on reasoning about the source quality but also makes the model more suitable for an extension. BBC is defined by Γ , whereas XQR is defined by m . The causal relationships within this model are captured by the DAG in Figure 3.2.

This DAG captures the high-level causal relationships. The source report is considered to be affected both by the true state of the world and the source quality. The contextual belief correction of Mercier et al. behaviour is captured on this graph by the dashed link between X and Q . However,

this does not mean that the value of Q can be inferred from X . While this graph captures the causal relationships, it does not determine the reasoning implementation. The mass function obtained from the multi-valued mapping Γ_A is used to connect the three variables for reasoning purposes.




Figure 3.2: The basic XQR model. The source report is driven by the variable of interest and the source quality. The source quality itself may be affected by the problem variable (akin to contextual belief correction)

The XQR model is formally defined using a valuation network. The network consists of three variables X , Q and R defined on arbitrary frames Ω_X , Ω_Q and Ω_R . The variable X corresponds to the variable of interest, R corresponds to a report provided by some source and Q describes the state the source is in. The relationship between the three is captured by the valuation on the joint frame m $\Omega_X \times \Omega_Q \times \Omega_R$. This valuation is defined by the mass m_{XQR} such that it has vacuous marginals i.e. $m_{XQR}^{\downarrow X} = \Omega_X$, $m_{XQR}^{\downarrow Q} = \Omega_Q$ and $m_{XQR}^{\downarrow R} = \Omega_R$. The source report is described by a unary valuation on Ω_R and the available information about the source state can be described by a unary valuation on Ω_Q or obtained otherwise (e.g. as a result of some other reasoning process by adding a valuation on the joint frame of Ω_Q and another added variable). The resulting valuation network is shown in Figure 3.7.




Figure 3.3: An evidential network, showing an implementation of the XQR model

To continue with the AIS example, the variable of interest X denotes whether the signal is transmitted $\Omega_X = \{t, \neg t\}$, and the report variable R denotes whether the signal is received or not $\Omega_R = \{r, \neg r\}$. The channel can lose packets with some likelihood and hence has two behaviour models $\Omega_Q = \{l, \neg l\}$. Thus the source is a perfect source if the channel is lossless ($Q = \neg l$) or always-negative if it loses packets ($Q = l$). This source behaviour model can be expressed by a joint mass function m over the joint space $\Omega_X \times \Omega_R \times \Omega_Q$:

$$m \left(\left\{ \begin{array}{l} l \times r \times t, l \times \neg r \times \Omega_X \\ \neg l \times r \times t, \neg l \times \neg r \times \neg t \end{array} \right\} \right) = 1$$

which is the same mass function as in the BBC model.

3.2.1 Contextual belief correction

In the theory of belief functions the concept of context is typically associated with *contextual discounting* or more generally *contextual belief correction*, a method of belief function correction proposed by Mercier[19] and a generalization of the classical discounting process, as defined in Section 2.2.4.

Classically the purpose of discounting is to take into account the relative reliability of a source of evidence by transferring a portion of the belief mass provided by the source to the universal set representing ignorance, such that:

$$\begin{aligned} m^\alpha(A) &= \alpha m(A) \forall A \subset \Omega \\ m^\alpha(\Omega) &= m(\Omega) + (1 - \alpha)m(A). \end{aligned}$$

The contextual discounting process proposed in [19] reasonably assumes that the reliability of the source is not going to be the same for every element in the frame of discernment. For a trivial example, consider a classifier recognising a vessel type on frame $\Omega = \{\text{CARGO}, \text{TANKER}, \text{PASSENGER}\}$. This classifier may be far more accurate in identifying the passenger ship than one of the other two types. Basic extensions of this model include coarsening - where a particular set of outcomes may be misidentified - and reinforcement, where a source is assumed to be more accurate for a specific target.

In existing works, the notion of context used here is related to the true state of the world, meaning the true label of the object observed—a finite set of “contexts” partitions the set Ω , leading to coarsening operations. The reliability of the source is thus defined for each context $A \subseteq \Omega$, representing the degree to which the source may be reliable, knowing that the true state of the world lies in A . The model has been extended to other correction mechanisms (e.g., reinforcement) while an interpretation within the BBC model has been provided in [83].

The knowledge held by the agent is computed by:

$$(3.4) \quad m_{Ag}^\Omega \left[m_{Ag}^{\Omega_X}, m_{Ag}^{\Omega_Q} \right] = \left(\bigodot_l^L m_{Ag}^{\Omega_Q}[\theta_l] \uparrow^{\Omega_X \times \Omega_Q} \bigodot m_{Ag}^{\Omega_X}[\{Q\}] \uparrow^{\Omega_X \times \Omega_Q} \right) \uparrow^{\Omega_Q}$$

where $m_{Ag}^{\Omega_X}$ is the information provided by the source (which is conditional on the source reliability) and $m_{Ag}^{\Omega_Q}[\theta_L]$ is a set of L masses describing the agent’s information about the reliability of the source given a particular context θ_L . This equation corresponds to first deconditioning the agent’s knowledge of reliability conditioned on the context and then combining with the source report conditioned on the reliability.

In general, these methods are generalized by the BBC model and, by extension, by the *XQR* model.

As the “context” as referred to here is actually the true state of the variable of interest, it is questionable whether this model captures context as per definitions discussed earlier, as most definitions consider the context to be information that does not explicitly define the problem. However, it can be argued that the sensing process takes space in the context of the true state of

the world, which is consistent both with the approach of Mercier et al. [84] and with the distinction into *context-of* and *context-for*. Alternatively, we can consider that the true state of the world is the *context-for* estimation of the reliability of the source. With this in mind, this interpretation is consistent with McCarthy’s formalism $ist(c, p)$ - the proposition that the source is reliable (or unreliable) holds true with greater (or lower) likelihood given some context.

3.3 Reasoning with context

Figure 3.4 illustrates the partitioning of the problem space. First let us consider the set of all variables, which are somehow related to the fusion problem at hand, that we will partition to reflect the notions of *context-of* and *context-for* introduced in Section 2.5.1. These can be partitioned into the set of endogenous variables (or problem variables) $\mathcal{X} = \{X_1 \dots X_n\}$ and the set of exogenous variables. The former is then further partitioned into contextual variables $\mathcal{C} = \{C_1, \dots, C_m\}$ and excluded variables E_i that we decide not to consider in our model. The contextual variables \mathcal{C} are considered to provide *context-for* the problem described by \mathcal{X} .

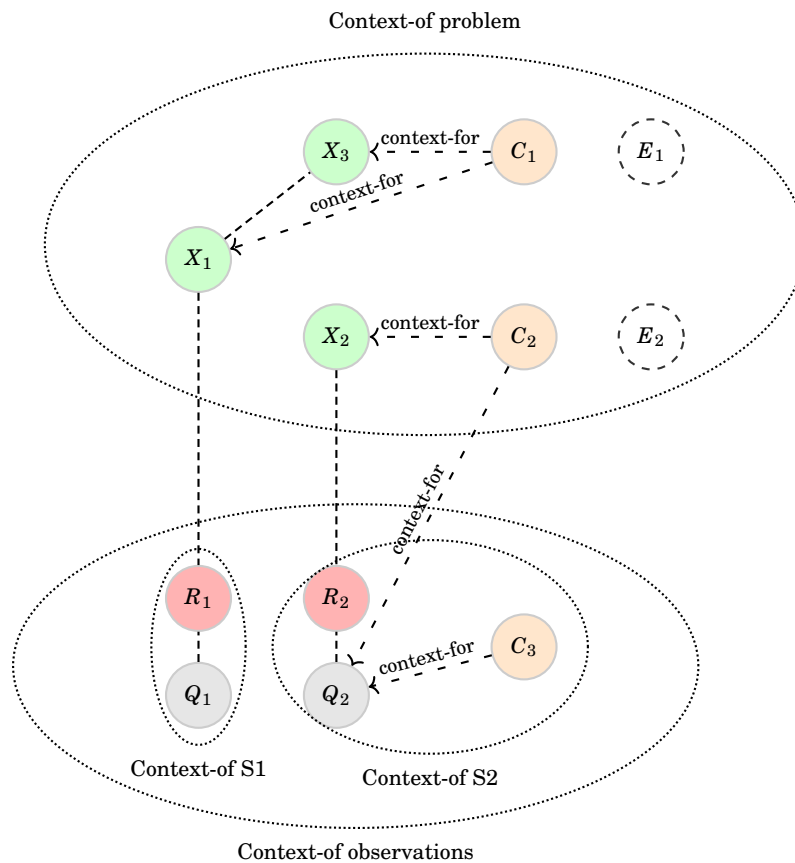


Figure 3.4: Extension of Steinberg’s model of exogenous and endogenous variables, accounting for variables outside the context of the problem [52]

This partitioning is consistent with the one proposed by Steinberg and Bowman [52]: what the authors call “universe of discourse” is called “*context-of* the problem” here. Rather than a vocabulary shift, introducing the notion of *context-of* enables considering and distinguishing

between more variables than the original model of [52]. Indeed, in a fusion problem involving multiple partially reliable sources, we are not only concerned with the variables related to the problem itself. As discussed in the previous section, we assume that the problem variables cannot be observed directly, but through some observational model as described earlier. Hence, according to the *BBC* and the *XQR* model a set of variables $\mathcal{R} = \{R_1, \dots, R_n\}$ describes the observations of variables \mathcal{X} provided by some partially reliable sources which behaviour is described by a set of variables $\mathcal{Q} = \{Q_1, \dots, Q_n\}$. We consider these variables \mathcal{R} and \mathcal{Q} to lie outside the *context-of* the problem but rather within the *context-of* the source. In a similar vein to *context-of* the problem, some endogenous variables exist within the *context-of* the sources which provide *context-for* the sources (such as the variable C_3 in Figure 3.4). The definition of two *contexts-of* (the problem and the sources) allows a formal connection between Steinberg and Bowman’s model and the *BBC* model.

Furthermore, we may also consider that some exogenous variables within the *context-of* the problem, may provide some *context-for* the source in addition to, or instead of, providing *context-for* the problem (such as the variable C_2 in Figure 3.4). As such, a contextual variable C is a variable such that it is exogenous to the problem (and as such would not appear in the *XQR* model from the previous section), but it is relevant or useful to the problem by providing some *context-for*.

From McCarthy’s notion of infinite contexts, we can always consider how each context considered exists within a greater context. In particular, while each source exists within its own context (*context-of* S_i), a greater *context-of* observations exists, which encapsulates *contexts-of* all sources. Furthermore this entire partitioning takes place within the *context-of* the model, which encapsulates all the variables and thus is the universe of discourse.

3.3.1 The contextual reasoning model

Here we consider the case where an exogenous variable may either constrain or provide information on the values of the problem variable or the case where it may affect the behaviour of the partially reliable source. The context is only of interest to the extent to which it may provide clues for understanding the entities of interest. On the modelling level, it may be unclear what the differences are in handling context and problem variables. We propose some assumptions regarding the relationship between the variables in the previously defined *XQR* model and a new contextual variable C . Although these assumptions will not be applicable for all scenarios, they provide a good starting point and can be relaxed as required.

Problem variables are observed indirectly ($X \rightarrow R$) Focus variables (and focus variables only) are observed by partially reliable sources, and they are only observed indirectly. As such, any estimate of a focus variable needs to be obtained through inference. A sound argument stands that the same should apply for context variables; however, this is constrained as it could make the model grow infinitely (as mentioned earlier, there is no context without context).

Source quality is not affected by the variable of interest $X \not\rightarrow Q$ The source quality (or, more specifically, the source behaviour selection) is not affected by the variable of interest. This

does not mean that the source is equally reliable for every possible state of the variable of interest, but rather that this relationship is represented by the source behaviours instead. For example, the strength of the signal does not affect channel interference - but the source can still be more accurate if the signal is strong.

Source quality can be affected by exogenous variables ($C \rightarrow Q$) The source quality selection can be affected by contextual variables - for example, high traffic density can cause channel interference for the signal discussed earlier. This is the major difference between this model and the other partially reliable source models discussed in the previous section, as the source behaviour may vary with context.

Context affects the variable of interest but not vice versa $C \rightarrow X, X, Q \not\rightarrow C$ By definition, context can affect the state of the variable of interest. It can be argued that there are many situations in which a context and variable of interest can be swapped. However, by definition of the variable types, context needs to be useful, and problem variable is the target of inference - doing this inference in reverse would make the context the problem variable. Consider the example from [67], where aircraft speed may be the context for its type, and its type may be the context for its speed. These are two different problems, and the selection of problem and context variables depends on operational needs. If both type and speed need to be estimated, then they are both problem variables.




Figure 3.5: System-level view of the reasoning system with a set of inputs and outputs. Information flow determines the difference between context variable and a variable of interest

Another way to view the differences between the context and focus variables in this model is to consider the flow of information. Treating the entire reasoning model as a system, we consider the metaknowledge and the relationships between variables as internal elements of the system. Then we can consider the external information on source reports, source quality and context as system inputs and the variables of interest as its outputs.

The consequence of this approach is that it constraints one of the key properties of this model: a variable of interest cannot be observed directly and can only be obtained as a result of inference (even if the inference is trivial, such as a perfect, reliable source). Furthermore, only variables of interest are considered outputs.

With this in mind, we can extend the XQR model discussed in the previous section with the fourth variable type C , which represents the contextual information, as per 3.6. In the belief functions framework, this is modelled as an evidential network, as shown in Figure 3.7.




Figure 3.6: A variable of interest X and an associated report R provided by a source which behaviour is determined by Q . The context variable C may affect either or both Q and X




Figure 3.7: An evidential network, showing an implementation of the XQRC model

It is important to note that there is not necessarily a strict distinction between *problem variable* and *context variable* as both represent the different aspects of the true state of the world, and the divide is mainly if not entirely determined by the problem statement. In this case, we distinguish between X and C through the assumption that the variable or variables X are the variables of interest and as such on themselves should directly affect the report (i.e. the report is conditionally independent on C provided X and Q). In other words, once a variable is reported on, it becomes a focus variable.

Recall from an earlier discussion on context that from an information fusion perspective, the definition of context implies that it must be useful. It can be shown that if the context is not relevant, then the XQRC model is equivalent to the XQR model and, by extension, to the BBC scheme.

On the following pages, for the sake of readability, $A^{\uparrow B}$ is a shorthand for $A^{\uparrow\Omega_B}$ and similarly $A^{\downarrow B}$ is a shorthand for $A^{\downarrow\Omega_B}$, where B is a variable, and Ω_B is its frame of discernment.

Proposition 3.1 (Relevance of context). *The XQRC model cannot be rewritten as XQR model unless $(m_{CX} \oplus m_C^{\downarrow X})^{\downarrow X} = \Omega_X$ and $(m_{CQ} \oplus m_C^{\downarrow Q})^{\downarrow Q} = \Omega_Q$ or the postulate that m_{XQR} has vacuous*

marginals is relaxed

Proof. Provided that m_{CQ} has vacuous marginal on C and $m_{XQR}^{\uparrow X \times Q} = \Omega_{X \times Q}$

$$(3.5) \quad (m_{CQ} \oplus m_C^{\uparrow Q})^{\downarrow C} = m^C$$

$$(3.6) \quad (m_{CX}^{\uparrow X} \oplus m_{CX}^{\uparrow Q})^{\downarrow X} = m_{CX}^{\uparrow X}$$

Note that $m_{CQ}^{\downarrow C} = \Omega_C$ and m_{XQR} has vacuous marginals on all singletons. Simplifying the XQRC model to an XQR model is tantamount to removing the variable C by combining all the valuations which contain it and aligned with the remaining variable set $\{X, Q, R\}$. This is then combined with the original valuation on $\{X, Q, R\}$, m_{XQR} . This results in the new mass \hat{m}_{XQR}

$$(3.7) \quad \hat{m}_{XQR} = ((m_{CX}^{\uparrow Q} \oplus m_C^{\uparrow X \times Q} \oplus m_{CQ}^{\uparrow X})^{\downarrow (X \times Q) \downarrow R}) \oplus m_{XQR}$$

such that:

$$(3.8) \quad \hat{m}_{XQR}^{\downarrow X} = \left[(m_{CX}^{\uparrow Q} \oplus m_C^{\uparrow X \times Q} \oplus m_{CQ}^{\uparrow X})^{\downarrow (X \times Q) \downarrow R} \oplus m_{XQR} \right]^{\downarrow X} = \\ = \left[(m_{CX} \oplus m_C^{\uparrow X})^{\downarrow X \downarrow R \times Q} \oplus m_{XQR} \right]^{\downarrow X} = (m_{CX} \oplus m_C^{\uparrow X})^{\downarrow X}$$

Hence $\hat{m}_{XQR}^{\downarrow X} = \Omega_X \iff (m_{CX} \oplus m_C^{\uparrow X})^{\downarrow X} = \Omega_X$, and the context is irrelevant to X . Same holds for $(m_{CQ} \oplus m_C^{\uparrow Q})^{\downarrow Q}$ and the vacuous marginal on Q . ■

The implication of this proposition that if the context is not relevant to the uncertain reasoning problem, the two models are equivalent. This is in line with some definitions of context discussed earlier, namely the property of context relevance.

Another important property of this model is that the context information is not included twice in the event that a single context affects the reasoning problem in more than one way. This is ensured by the valuation network model but can be explicitly proven.

Proposition 3.2 (Inclusion-exclusion (No double counting of context)). *If m_{CX} or m_{CQ} is vacuous the impact of context can be marginalized on Q or X respectively and included in \mathcal{M} . Let these be denoted by $\hat{m}_{CQ} = m_{CQ} \oplus m_C^{\uparrow Q}$ and $\hat{m}_{CX} = m_{CX} \oplus m_C^{\uparrow X}$. If the condition does **not** hold, \hat{m}_{CQ} and \hat{m}_{CX} cannot be both included in this set, unless m_C is vacuous.*

Proof. Let \mathcal{M} denote the set of all the valuations in the XQRC model, and $m_J = \bigoplus_{m \in \mathcal{M}} m$ denote the full joint mass. Let $\mathcal{M}_1 = \{m_{XQR}, m_Q, m_R\}$ be the set of valuations included in the XQR model. Consider also the sets of variables $\mathcal{V} = \{X, R, Q, C\}$ and $\mathcal{V}_1 = \{X, R, Q\}$. Assume m_C is not vacuous.

$$(3.9) \quad m_J = \bigoplus_{m \in \mathcal{M}} m^{\uparrow \mathcal{V}} = \bigoplus_{m \in \mathcal{M}_1} m^{\uparrow \mathcal{V}} \oplus m_{CX}^{\uparrow \mathcal{V}} \oplus m_{CQ}^{\uparrow \mathcal{V}} \oplus m_C^{\uparrow \mathcal{V}}$$

If m_{CQ} is vacuous this becomes:

$$(3.10) \quad m_J = \bigoplus_{m \in \mathcal{M}} m^{\uparrow \mathcal{V}} = \bigoplus_{m \in \mathcal{M}_1} m^{\uparrow \mathcal{V}} \oplus \hat{m}_{CX}^{\uparrow \mathcal{V}} \oplus m_{CQ}^{\uparrow \mathcal{V}} = \bigoplus_{m \in \mathcal{M}_1} m^{\uparrow \mathcal{V}} \oplus \hat{m}_{CX}^{\uparrow \mathcal{V}}$$

And the context variable can be removed

$$(3.11) \quad m_J^{\mathcal{V}_1} = \bigoplus_{m \in \mathcal{M}_1} m^{\uparrow \mathcal{V}_1} \oplus \hat{m}_{CX}^{\uparrow \mathcal{V}_1}$$

The same applies to \hat{m}_{CQ} if m_{CX} is vacuous. If neither m_{CX} nor m_{CQ} are vacuous then the context valuation m_C is included twice:

$$(3.12) \quad \bigoplus_{m \in \mathcal{M}_1} m^{\uparrow \mathcal{V}} \oplus \hat{m}_{CX}^{\uparrow \mathcal{V}} \oplus \hat{m}_{CQ}^{\uparrow \mathcal{V}} = \bigoplus_{m \in \mathcal{M}_1} m^{\uparrow \mathcal{V}} \oplus m_{CX}^{\uparrow \mathcal{V}} \oplus m_C^{\uparrow \mathcal{V}} \oplus m_{CQ}^{\uparrow \mathcal{V}} \oplus m_C^{\uparrow \mathcal{V}} \neq m_J$$

provided that m^C is not vacuous. ■

In other words we can consider the base model \mathcal{M}_1 and its extensions $\mathcal{M}_X = \mathcal{M}_1 \cup \hat{m}_{CX}$ and $\mathcal{M}_Q = \mathcal{M}_1 \cup \hat{m}_{CQ}$, where m_{CQ} and m_{CX} respectively are vacuous. Clearly the full model $\mathcal{M} \neq \mathcal{M}_Q \cup \mathcal{M}_X$ as in that case the impact of context would be included twice.

Relation to contextual belief correction Recall from the discussion of contextual belief correction that in CBC, the agent's contextual knowledge is provided as a set of BBA's conditioned on a different context. Both relationships between context C and the variable of interest X and quality Q can be interpreted as being obtained from a set of contextually conditional beliefs. In this case we can consider $m_{CX} = \bigodot_i^C m_X[c_i]^{\uparrow C \times X}$ and $m_{CQ} = \bigodot_i^C m_Q[c_i]^{\uparrow C \times Q}$. As such, this remains consistent with McCarthy context formalisation - a belief that holds true in some context c . However, the introduction of a new contextual variable makes it possible to consider that variable to be *context-for* in Gong's interpretation, as opposed to *context-of* provided by the CBC.

3.3.2 Illustration: Vessel behaviour analysis example

Recall the example of the AIS broadcast discussed first in Section 3.2.

The introduced contextual variable C represents weather and takes values in $\Omega_C = \{W_g, W_b\}$, corresponding to good and bad weather states, respectively. We consider two uncertain, semantic expressions describing the impact the contextual information on inference regarding the AIS transmission status:

1. Bad weather reduces the likelihood that the vessel master will disable the AIS transmission (for better prevention of collisions with other ships)
2. Bad weather reduces the likelihood that the broadcasted AIS signal will be received correctly.

3.3.2.1 Context as uncertain information source - context-for problem

The statement *Bad weather reduces the likelihood that the vessel master will disable the AIS transmission* (1) can reasonably be only interpreted as evidence or constraint on the true state of the world.

The logical relationship can be written as a material implication: $(C = W_b) \implies (X = t)$. Note, that this corresponds to a belief in $X = t$, conditional on $C = W_b$. This context will provide no

constraint on X in the event that the weather is good. Furthermore uncertainty must be included: the captain is *less likely* to disable the AIS; this relationship is by no means certain. As such the truth table associated with the above logical expression can be encoded in the focal set defined over $\Omega_C \times \Omega_X$:

$$F_1 = \left\{ \begin{array}{l} W_b \times t \\ W_g \times \Omega_X \end{array} \right\}.$$

The relationship between the context C and the variable of interest X can be then represented by the belief function m_c such that

$$\begin{aligned} m_{cX}(F_1) &= \gamma_1 \\ m_{cX}(\Omega_C \times \Omega_X) &= 1 - \gamma_1 \end{aligned}$$

for some value of $\gamma_1 \in (0,1)$, which represents the degree of belief in the hypothesis that this relationship holds.

Treating context as evidence is the most natural interpretation of such information. In general, it could be argued that any relevant context of a situation can be interpreted in this manner. For instance, we previously discussed how context might affect quality in this particular scenario (*Bad weather reduces the likelihood that the broadcasted AIS signal will be received correctly* (2)); we could interpret this instead as direct evidence on the AIS being broadcasted ($C = W_b$) \implies ($X = t$). Note that this is the same logical expression as in the case of the statement (1).

In many reasoning systems where the reporting models and quality of the source are not modelled explicitly, this could be the default method. The logic behind this approach is rooted in the probabilistic thinking model - a reduction in our belief in the hypothesis that the AIS is not being broadcasted can be treated as a belief in its converse. However, this approach fails to take advantage of the power of evidential reasoning, which allows for explicit modelling of ignorance.

3.3.2.2 Context and information quality - context-for source

The alternative way in which context can be utilised for inference in this model is by affecting the source quality. Once again we reinterpret the statement *Bad weather reduces the likelihood that the broadcasted AIS signal will be received correctly* (2) and this time represent it to support the belief in $Q = l$. This is represented by the logical expression ($C = W_b$) \implies ($Q = l$) and hence encoded in the mass function defined over $\Omega_C \times \Omega_Q$:

$$\begin{aligned} m_{cQ}(F_2) &= \gamma_2 \\ m_{cQ}(\Omega_C \times \Omega_Q) &= 1 - \gamma_2 \end{aligned}$$

where the focal set F_2 corresponds to the truth table of the above logical expression:

$$F_2 = \left\{ \begin{array}{l} W_b \times l \\ W_g \times \Omega_Q \end{array} \right\}.$$

The $XQRC$ model allows distinguishing between context-for source and for the problem. By that, it allows explicitly encoding the independent behaviours of the source (here the AIS sensor)

and the relevant agent with the situation of interest (here the ship captain). Both may impact the reception of the AIS signal, which in both cases is caused by the context (here the weather). The contextual variable here lies within the context of the problem, and as such, can act both as a context-for problem (affecting variables within the context of the problem) and context-for source. We could contrast this with a variable in context-of source acting as a context-for source, which is not related to the context of the missing vessel - such as electrical malfunctions near the AIS receiver, which may impact its quality.

3.4 Properties of contextual reasoning model

In this section, some topics associated with inference on contextual information are discussed. These, in particular, are concerned with modularity and scalability, expressiveness and interpretability or explainability.

To consider the added value of contextual reasoning, we will compare the original model where the context-for the problem is not considered whatsoever to implementations utilising context-for problem as well as context-for source. The rationale behind this is that utilisation of context-for problem only corresponds simply to the inclusion of additional sources - and while it may bring about some advantages of contextual reasoning, it will not fully exploit the added value.

3.4.1 Context-of only - contextual belief correction

In the most basic model, the *context-for* the problem, i.e. the weather variable is not considered at all. We still operate in the context-of the AIS channel behaviour, which can be represented either using Mercier et al. contextual belief correction or the XQR model as we have chosen to do.

Consider the AIS example with observations: $m^R(\neg r) = 1$ and $m^Q(l) = \lambda$, $m^Q(\neg l) = 1 - \lambda$, which means that the AIS signal was not received and AIS loss rate is λ .

Let m_1 be the joint $m_{XQR} \oplus m^Q \oplus m^R$

$$m_1(\{\neg l \times \neg r \times \neg t\}) = 1 - \lambda$$

$$m_1(\{l \times \neg r \times \Omega_X\}) = \lambda$$

This joint valuation covers the space of three variables. It can be marginalized on the variable of interest X to finish the inference process.

$$m_1^{\downarrow X}(\neg t) = 1 - \lambda$$

$$m_1^{\downarrow X}(\Omega_X) = \lambda$$

3.4.2 Context-for problem

Using context-for problem only allows better inference (compared to no contextual variables); and leverages the modularity and scalability advantages provided by valuation networks. This, to some extent, helps the interpretability of the model.

This corresponds to the inclusion of additional two mass functions, the joint m_{CX} representing the impact of weather as context-for the internal variable of interest (or alternatively, using semantics similar to these in CBC, the conditional knowledge about X given the context) and m^C . As these masses can be marginalized to provide information on X directly, this corresponds closely to the method that would be used to include this additional *context-for* variable in the absence of a specific contextual reasoning model.

$$\begin{aligned} m_{CX}(\{W_b \times t, W_g \times t, W_g \times \neg t\}) &= \gamma_1 \\ m_{CX}(\Omega_C \times \Omega_X) &= 1 - \gamma_1 \end{aligned}$$

Combined with some observation on context $m_C(W_b) = \mu$, $m_C(\Omega_C) = 1 - \mu$ and marginalized on X this corresponds to an additional piece of evidence on X .

$$\begin{aligned} (m_{CX} \oplus m^C)^{\downarrow X}(t) &= \gamma_1 \lambda \\ (m_{CX} \oplus m^C)^{\downarrow X}(\Omega_X) &= 1 - \gamma_1 \lambda \end{aligned}$$

Which can be combined with the previous result $m_2 = m_1 \oplus (m^C \oplus m_{CX})^{\downarrow Q}$

$$\begin{aligned} m_2^{\downarrow X}(\emptyset) &= \gamma_1 \mu (1 - \lambda) \\ m_2^{\downarrow X}(t) &= \gamma_1 \mu \lambda \\ m_2^{\downarrow X}(\neg t) &= (1 - \gamma_1 \mu)(1 - \lambda) \\ m_2^{\downarrow X}(\Omega_X) &= (1 - \gamma_1 \mu) \lambda \end{aligned}$$

The improvement in inference is straightforward, as it is possible to leverage a greater amount of information - in a manner similar to simply exploiting another source.

Consider that we can simply add these two masses as a module. Scalability is a major issue evidential reasoning due to the additional memory and computational costs associated with the growing frame of discernment; this issue is exacerbated in multi-dimensional belief functions due to the rate at which the sizes of frames of discernment can increase. Shenoy-Shafer's architecture of valuation networks [51] partially addresses this problem by exploiting binary join trees, which keep the instantaneous size of a single frame of discernment under control, as discussed briefly in Section 2.4.5. Instead of marginalizing the joint in the previous subsection, we could add the other contextual reasoning mass m_{CX} instead, or possibly another source report.

This makes it possible to obtain separate inferences with and without this additional context for the problem, which in turn makes it possible to accurately assess the impact of the context on the solution.

3.4.3 Context-for source

Inclusion of context-for sources in addition to the context-for problem allows a greater expressiveness (through diffidence). As per the previous case, this corresponds to the inclusion of the masses corresponding to the m_{CQ} , as well as the report on contextual variable m^C .

As per the previous case, we have mass function m_{CQ}

$$\begin{aligned} m_{CQ}(\{W_b \times l, W_g \times l, W_g \times \neg l\}) &= \gamma_2 \\ m_{CQ}(\Omega_C \times \Omega_Q) &= 1 - \gamma_2 \end{aligned}$$

This can be directly combined with m_1 such that $m_3 = m_1 \oplus m_{CQ}$. In this scenario, we cannot obtain an inference on X from context only.

$$\begin{aligned} m_{CQ} \oplus m_1(\{W_g \times \neg l \times \neg r \times \neg t\}) &= \gamma_2(1 - \lambda) \\ m_{CQ} \oplus m_1(\{l \times \neg r \times \Omega_X \times \Omega_C\}) &= \gamma_2\lambda \\ m_{CQ} \oplus m_1(\{\neg l \times \neg r \times \neg t \times \Omega_C\}) &= (1 - \gamma_2)(1 - \lambda) \\ m_{CQ} \oplus m_1(\{l \times \neg r \times \Omega_X \times \Omega_C\}) &= (1 - \gamma_2)\lambda \end{aligned}$$

Combined with the same observation on context as before, and marginalized on X

$$\begin{aligned} m_3^{\downarrow X}(\emptyset) &= \gamma_2\mu(1 - \lambda) \\ m_3^{\downarrow X}(\neg t) &= (1 - \gamma_2\mu)(1 - \lambda) \\ m_3^{\downarrow X}(\Omega_X) &= \lambda \end{aligned}$$

Expressiveness An interesting property of this approach to contextual information exploitation is the increased expressiveness. In particular consider the marginal on X : $m_3^{\downarrow X}$ and compare it to the same marginal prior to the introduction of context-for source: $m_1^{\downarrow X}$. The relation between the two has been shown through a sequence of conjunctive combination operations, however it can also be shown that there does not exist a *proper* mass m such that $m \oplus m_1^{\downarrow X} = m_3^{\downarrow X}$. Note that $m_1(\Omega_X) = m_3(\Omega_X)$

Proof. For any two mass functions m_1 and m_2

By the conjunctive rule of combination:

$$(m_1 \oplus m_2)(\Omega) = \sum_{\Omega=A \cap B} m_1(A)m_2(B) = m_1(\Omega)m_2(\Omega).$$

Therefore $(m_1 \oplus m_2)(\Omega) = m_1(\Omega) \iff m_2(\Omega) = 1$.

Also by CRC: If $m_2(\Omega) = 1$ and m_2 is a proper mass function then: $m_J = m_1$.

Therefore no proper mass function exists satisfying $m_1(\Omega) = (m_1 \oplus m_2)(\Omega)$ and $m_1 \neq (m_1 \oplus m_2)$

■

An interesting point to note is that the normalized $m_2^{\downarrow X}$ is a discounted $m_1^{\downarrow X}$. Thus as expected, this particular simple example of contextual reasoning implements the discounting operation (contextual discounting since it only affects the source report $R = \neg r$ and would have no effect on mass assigned to $R = r$).

Context-for source as diffidence Canonical decomposition and decomposition of belief functions [25] can be used to support interpretability. Consider the difference between m_1 and m_3 . We have shown that there does not exist a proper bba m such that $m_1 \oplus m = m_3$. We can find the *pseudo* bba m_{3-1} such that $m_{3-1} \oplus m_1 = m_3$ which is one way of evaluating the "impact" of valuation m_{CQ} . The commonality function is:

$$\begin{aligned} q_{2-1}(\Omega) &= 1 \\ q_{2-1}(\neg t) &= 1 - \gamma_2 \mu(1 - \lambda) \\ q_{2-1}(\emptyset) &= 1 \end{aligned}$$

from which we naturally obtain the pseudo-bba m_{2-1}

$$\begin{aligned} m_{2-1}(\Omega) &= 1 \\ m_{2-1}(\neg t) &= -\gamma_2 \mu(1 - \lambda) = -m_2(\emptyset) \\ m_{2-1}(\emptyset) &= \gamma_2 \mu(1 - \lambda) = m_2(\emptyset) \end{aligned}$$

The negative mass can be interpreted as a diffidence mass as referred to in [85] or belief absorption [25] and the mass on an empty set corresponds to inconsistency or conflict. For Dubois et al. [85] diffidence represents bias or doubt. In this case, we observe a mass function inducing diffidence, where it can correspond to negative evidence on source quality and as such as a form of doubt about the information this source provides. Overall the combination $m_1 \oplus m_{2-1}$ corresponds to transfer of mass from $m_1(\neg t)$ to $m_2(\emptyset)$.

An interesting remark is that the overall information available about the variable of interest is reduced, seemingly conflicting with the principle of information monotonicity in uncertain information fusion. To recall Dubois [1] states that given two sets of non-disagreeing (consistent) agents with providing less information than the other, fusing the less informative set should not provide a result more informative than fusing the other.

If we consider the two sets of agents, or in this model the set of valuations, to be the AIS report only $A_1 = \{m, m^R\}$ and the set of AIS report and contextual information $A_2 = \{m, m^R, m^C, m_{CQ}\}$ they seem non-conflicting. It is, however important to note that they are both non-informative in the absence of meta-information on the source quality. When the two sets are augmented with the quality valuation m^Q , which places a non-zero mass on $Q = l$, the set A_1 remains consistent, whereas set A_2 does not. Specifically $(\bigcirc A_2)(\emptyset) \neq 0$. As such, despite the set A_2 containing more information than A_1 (as $A_1 \subset A_2$), it is acceptable that its fusion result will be less informative without violating the information monotonicity rule. However, in an implementation without conflict propagation, i.e. normalizing after every step in the valuation network, this conflict will not be visible at the final node. Furthermore, despite inconsistencies within the set A_2 , the normalized result of inference on X will be n -consistent [86] for any n , meaning that inconsistencies in the knowledge base do not propagate into the output.

3.4.4 Combined models

Finally, we will consider models including both context-for source and context-for problem in different forms

Context-for both source and problem We may also consider the case where both context-for source and context-for problem are taken into account, $m_4 = m_3 \oplus (m^C \oplus m_{CX})^{\uparrow Q}$:

$$\begin{aligned} m_4^{\downarrow X}(\emptyset) &= (1 - \lambda)\mu[\gamma_2 + (1 - \gamma_2)\gamma_1] \\ m_4^{\downarrow X}(t) &= \gamma_1\mu\lambda \\ m_4^{\downarrow X}(\neg t) &= (1 - \lambda)[1 - \mu + \mu(1 - \gamma_1)(1 - \gamma_2)] \\ m_4^{\downarrow X}(\Omega_X) &= \lambda[1 - \mu + \mu(1 - \gamma_1)] \end{aligned}$$

Treating context-for source as context-for problem Furthermore, let us recall how in absence of the contextual model, we could consider all contextual rules as evidence on X or as context-for problem. As such, the statement *Bad weather reduces the likelihood that the broadcasted AIS signal will be received correctly* can be interpreted as $C \rightarrow X$.

Thus consider the mass function $m_D^{C \times X}$ capturing this knowledge:

$$\begin{aligned} m_D^{C \times X}(\{W_b \times t, W_g \times t, W_g \times \neg t\}) &= \gamma_2 \\ m_D^{C \times X}(\Omega_C \times \Omega_X) &= 1 - \gamma_2 \end{aligned}$$

Thus let $m_5 = m_1 \oplus m_D^{C \times X}$ and $m_6 = m_1 \oplus m_D^{C \times X} \oplus m_{CX}$, corresponding to the cases where the context-for source is interpreted as context-for problem.

With this, we have a total of six possible models, as displayed in Table 3.2.

Table 3.2: The six reasoning models

Mass	Description
m_1	No context-for
m_2	Context-for problem only
m_3	Context-for source context only
m_4	Both CF-source and CF-problem
m_5	CF-source interpreted as CF-problem
m_6	As m_5 , plus CF-problem

3.4.5 Interpretability

There are several ways in which the methods laid out in this chapter help interpretability and explainability of the results. One of the most obvious ones is the modularity and scalability discussed earlier. The modular nature of the model means that it is possible to compute and store an inference on the variable of interest at different stages of the computation and use that to track the changes in the marginal.

Consider $\lambda = 0.2$, $\mu = 0.6$, $\gamma_1 = 0.5$ and $\gamma_2 = 0.3$. The belief masses $m_1, \dots, m_6^{\downarrow X}$ computed with these parameters are displayed in Figure 3.8 and in Table 3.3.

For conciseness from this point onward, m_n refers to $m_n^{\downarrow X}$ unless specified otherwise. The relatively non-intensive computation of the intermediary masses allows straightforward analysis. It can be easily seen that the introduction of context-for source acts as a discounting operation ($m_1 \rightarrow m_3$), whereas using context to place constraints on the variable of interest is tantamount to the introduction of another information source ($m_1 \rightarrow m_2$). Combining both (m_4) increases the belief in $X = t$. Although it makes sense provided one understands the model (due to conflict reduction), it may seem counter-intuitive at first as a (contextual) discounting operation increases the degree of belief in some hypothesis.

The final mass m_5 represents treating the contextual information on quality as constraining the variable of interest instead. It is of interest as it represents a more classical approach to information aggregation, and as discussed earlier, this could be the default method in some systems not fully utilising the evidential framework ability to model ignorance. As such, the ignorance $m_5(\Omega)$ is the lowest.




Figure 3.8: Ternary plot showing the relationship between different contextual models for $\lambda = 0.2$, $\mu = 0.6$, $\gamma_1 = 0.5$ and $\gamma_2 = 0.3$

Table 3.3: Masses m_1 through to m_5 , unnormalized and normalized

	m_1	m_2	m_3	m_4	m_5	m_6
\emptyset	0	0.24	0	0.144	0	0.312
t	0	0.06	0.079	0	0	0.06
$\neg t$	0.8	0.56	0.74	0.656	0.77	0.488
Ω	0.2	0.14	0.18	0.2	0.23	0.14

Table 3.4: Discord and nonspecificity for normalized and unnormalized masses m_1, m_2 and m_6

	Unnormalized		Normalized	
	$NS(\cdot)$	$m^{(2)}(\emptyset)$	$NS(\cdot)$	$m^{(2)}(\emptyset)$
m_1	λ	0	λ	0
m_3	λ	C_\emptyset	$\frac{\lambda}{1-K}$	0
m_5	$\lambda(1-\gamma_2\mu)$	$C_\emptyset + 2\lambda\mu\gamma_2(1-\gamma)(1-\gamma_2\mu)$	$\frac{\lambda(1-\gamma_2\mu)}{1-K}$	$\frac{2\lambda\mu\gamma_2(1-\gamma)(1-\gamma_2\mu)}{(1-K)^2}$

Decomposition of belief functions We have discussed how the transfer of mass associated with the inclusion of external conflict can be interpreted as a diffidence function. However, we can continue with the decomposition-based approach for a better understanding of the results.

Similar analysis can be performed for the changes from m_1 to m_4 by computing m_{4-1} and then performing canonical decomposition to view the individual components. These can be shown to be m_{3-1} as computed earlier (the component induced by m_{CQ}) and a simple support function on t induced by m_{CX} .

Difference between context-for source and problem interpretation Consider the two interpretations m_{CQ} and $m_D^{C \times X}$ of the same "opinion" on the impact of context again. The former affects the quality and implements contextual reasoning as discussed throughout this chapter; the latter can be thought of as direct inference. Consider the difference between $m_3 = m_1 \oplus m_{CQ} \oplus m_C(W_b)$ and $m_5 = m_1 \oplus m_D^{C \times X} \oplus m_C(W_b)$, i.e. the interpretations of the same piece of contextual information as either context-for source or problem. We can use the uncertainty quantification methods from Section 2.3 to discuss the difference between the two approaches in more detail.

By inspection we can see that the only difference between m_5 and m_3 is some mass being transferred from the universal set Ω (m_3) to t (m_5). By definition this implies an increase in discord and reduction in ambiguity. Table 3.4 shows the nonspecificity ($NS(m) = \sum m(A) \log_2 |A|$, which becomes $NS(m) = m(\Omega)$ when $|\Omega| = 2$) and $m^{(2)}(\emptyset)$ values for normalized and unnormalized masses m_1, m_3 and m_5 . Note that in unnormalized case a significant proportion of auto-conflict is due to the original mass on the empty set. For clarity let $K = m(\emptyset) = (1-\lambda)\mu\gamma_2$ and $C_\emptyset = 2K - K^2$

There exist some interesting albeit expected relationships. Particularly it can be seen that between m_1 and normalized m_3 there is no increase in $m^{(2)}\emptyset$ which acts as a measure of internal conflict; in the unnormalized case this is equivalent to the explicit conflict $m(\emptyset)$ instead. However, an increase in ambiguity can be seen, as we know m_3 to be less informative than m_1 . On the contrary, in the case of m_5 , the ambiguity decreases, but at the cost of increasing the internal

conflict. In the unnormalized case, the measure is composed of both $m(\emptyset)$ and the additional internal conflict.

This brings us to another important result from the decision making perspective: since m_3 remains a simple support function on $\neg t$, the only non-universal focal set which can have a non-zero belief mass is $\neg t$ - assuming one of the usual decision making approaches (such as TBM maximisation of pignistic probability [31]), no vacuous decision and $\lambda, \gamma_2, \mu < 1$, m_3 will always yield the decision $X = \neg t$. In the case of m_5 however there exist values of λ, γ_2, μ which result in the decision $X = \neg t$.

3.4.6 Further examples of source behaviour and contextual reasoning models

In this section, additional source reasoning models exhibiting more complex behaviours are discussed. First of all we consider a report provided by an UAV which is generated using one of two sensors with different properties. The context affects the choice of the sensor, as well as the reasoning. In addition, we also discuss two human intelligence examples - a report obtained by questioning a witness and a report provided by a trained analyst.

3.4.6.1 Vessel detection and classification

The following example is very loosely based on the typical sensors available to an aerial UAV. Consider an aerial UAV used for marine surveillance, running a detection and classification task using two sensors (e.g. EO and IR). We assume that the classification task exploits data from one or the other but not both - for example, for some reason, the construction does not allow simultaneous usage of both sensors. The two modes of operation can be referred to as mode 1 and mode 2. On top of the different operational models, let there be a chance of a false negative in mode 1 and a chance of false positive in mode 2. Context variables such as time of day and weather can affect either or both mode selection and error rate.

Let the variable of interest X denote the vessel type and be defined on the frame $\Omega_X = \{c, f, o, n\}$ denoting *cargo*, *fishing*, *other* and *none*. Let the report be produced on the same frame so that $\Omega_R = \Omega_X$. The different modes of operation are described as follows: in mode 1, it is possible to distinguish all the different types with a chance of false negative, whereas in mode 2, it is possible to distinguish cargo from others; however, a chance of false positive exists. This model describes a typical situation where a high-quality sensor can be used in good conditions, but a less reliable sensor is available when the conditions are subpar. This implies total of 4 different behavioural models - h_{11} (mode 1, normal), h_{12} (mode 1, erratic), h_{21} (mode 2, normal) and h_{22} (mode 2, erratic). Therefore the quality-behaviour variable, Q takes its values in the frame of discernment $\Omega_Q = \{h_{11}, h_{12}, h_{21}, h_{22}\}$

Furthermore, it is reasonable that there should be a relation between C and X . Some possible relationships could be due to ships, in general, less likely to be in the surveyed location in bad weather, fishing ships less likely to be at sea at night. This makes this model consistent with the generalized model shown in Figure 3.6. The relationship between the variables Q , R and X is

determined by the conditional mass assignment m_S which follows from Equation 3.2 provided Γ from Table 3.5.

Table 3.5: Multi-valued mapping Γ describing the interpretation of each value of the report R given the hypothesis H

Γ		$Q \in \Omega_Q$			
		h_{11}	h_{12}	h_{21}	h_{22}
$R \in \Omega_R$	c	c	c	c	$\{c, n\}$
	f	f	f	$\{f, o\}$	$\{f, o, n\}$
	o	o	o	$\{f, o\}$	$\{f, o, n\}$
	n	n	$\{c, f, o, n\}$	n	n

The context variables C_1 and C_2 which refer to weather (C_1) and the time of day (C_2) are defined on the frames $\Omega_{C_1} = \{g, b\}$ and $\Omega_{C_2} = \{d, n\}$, denoting *good*, *bad*, *day* and *night* respectively.

As such, we denote the following relationships between the variables:

$$C_1 = d \rightarrow Q \in \{h_{11}, h_{12}\}$$

$$C_1 = n \rightarrow Q \in \{h_{21}, h_{22}\}$$

$$C_2 = g \rightarrow Q \in \{h_{11}, h_{21}\}$$

$$C_2 = b \rightarrow Q \in \{h_{12}, h_{22}\}$$

These can be encoded using valuations m_{C_1Q} and m_{C_2Q} . For simplicity sake, let us assume that only C_1 has an effect on X such that

$$C_1 = d \rightarrow X = \neg f$$

and this relationship is modelled by valuation m_{CX} . The model using these four valuations corresponds to the valuation network shown in Figure 3.9




Figure 3.9: Valuation network for inference on the report provided by the aerial UAV for vessel detection and classification model

3.4.6.2 Human intelligence examples

In general, contrary to physical sensors, human intelligence cannot be physically modelled. Some behaviour models could be generated by analysts for human intelligence reports, e.g. accuracy regarding the frame of discernment, likelihood of truth-telling etc. For instance, provided a witness report or other kind of human intelligence source, an analyst could elicit a set of potential behaviour models, potential probability distribution and/or set of variables affecting the selection of the model. However, when the source report accessible to us is already a "full analyst report", there is little processing that can be done.

In this case, it is assumed that the frame of discernment of the agent report is consistent with the frame of discernment of variable of interest. Otherwise (if the agent reports on some other variable), some relationship between the two must be derived lest the report is insignificant to the problem at hand. Although the 'behaviour model' of the analyst is unknown it seems reasonable to treat her as a rational agent reporting directly on a variable of interest. Therefore, the only natural way to model the quality of the agent in this situation is via Shafer's discounting model.

One key problem with an "analyst" type of source is that, whatever its report, it is a result of another unknown inference process. This means that there is going to be a risk of double counting as it cannot be easily verified whether the information used to produce the analyst report is or is not overlapping with the information available to the other VBS. Some discussion on this topic takes place in a later section. Here we discuss simple human intelligence examples, one including a witness report undergoing appropriate modelling and the other being a simple report provided by an analyst.

HUMINT example 1: Simple witness report Consider a witness reporting on a type of ship she or he has spotted. Let the witness report on the frame $\Omega_R = \{c, f, o\}$ and the frame of interest be $\Omega_X = \{c, f, o, n\}$, referring to *cargo*, *fishing*, *other* and *none*. An analyst may read the witness statement and come to the following conclusions:

1. It is possible that the witness is entirely correct and truthful $\Gamma_A = A$ for all $A \in \Omega_R \cap \Omega_X$
2. There is some chance that the witness is lying either about the type of the ship or about having seen the ship at all $\Gamma_A = \neg A$ (note that the set $\neg A$ includes *none* for every A in Ω_R)
3. There is a chance that the witness is truthful, but has misclassified a *tanker* (which belongs to the *other* class) as *cargo* $\Gamma_c = \{c, o\}, \Gamma_A = A$ if $A \neq c$
4. There is a chance that regardless of what the witness says, he or she does not actually describe the vessel of interest $\Gamma_A = \{\Omega_X\}$ for all $A \in \Omega_R$

And the analyst provides some belief distribution on $Q \in \{H_1, H_2, H_3, H_4\}$ corresponding to the four behavioural hypotheses above. Using these four behaviour classes, inference can take place as per the previous examples

HUMINT example 2: Analyst report This is an example of a more complex analyst report coming with justification. Consider the following report: *"The images taken by the UAV have been analysed, and a vessel has been certainly spotted in area XYZ. There is strong evidence that it is a cargo vessel (60%). We can verify that it certainly is not a fishing vessel."* The key difference between this and the previous one is that some information about the inference process is provided. As such it makes sense that this source could be readily combined with the witness statements discussed earlier, whereas double-counting could occur if it were to be combined with the UAV reports. However, if new evidence is obtained regarding the quality of the UAV report, it should also affect the quality of the analyst report in question (assuming that this evidence has not been available to the analyst).

3.5 Dependence relationships between sources

Another area of interest when discussing source behaviours is their potential inter-dependency. In this section, we discuss possible scenarios where data cross-contamination may occur, namely when the sources to be combined are not independent. This section is laid out as follows. First, we consider trivial examples with consonant belief assignments provided by the reports - i.e. one of the reports is a strictly less informative variant of the other. Different combination methods, taking into account the dependencies and otherwise, are compared. At the next stage, we consider a more complex case where the strict information content ordering no longer holds. We propose a method of idempotent combination of masses using an appropriately built evidential network exploiting the conjunctive rule of combination only. This method is easier to implement than the cautious rule, as the cautious rule cannot directly be used in a valuation network, but at the cost of being less than or equally informative. We extend this concept to the sources where partial dependency may or may not exist. Furthermore, we comment on the general fundamental property that the result of fusion should be more informative than either of the sources and discuss whether it should be applicable in all the cases, in particular when the sources are known to be dependent.

3.5.1 Consonant sources

Here the basic examples of interdependent reports are considered. To begin with, consider 2 sources providing reports R_1 and R_2 reporting variable of interest X . The most typical example of double counting of information occurs when the same piece of information is taken into account twice for inference resulting in bias and overly confident result. This can be generalized to the case where R_2 is not necessarily identical to R_1 , but it is conditionally independent from X given R_1 . This general model is shown in Figure 3.10. Note that this model includes a dashed and a dotted line. The dashed line represents the "true" model, i.e. the process by which this report is actually generated. In a fusion scenario, the knowledge of conditional independence of R_2 from X given R_1 means that for inference X is conditionally independent on R_2 given R_1 . If the purpose is to make inference on X , R_2 could be safely ignored. As such, we must assume that it is not known

which of the two reports is the original or alternatively which of the two is a better approximation of X . As such, for the purpose of inference on X , the model with the dotted line is considered.




Figure 3.10: Two reports R_1 and R_2 such that the report R_2 is conditionally independent on X given R_1

Let us illustrate this first with a simplified version of the AIS example used throughout this chapter. In this trivial model the frame of the variable of interest remains unchanged $\Omega_X = \{On, Off\}$ but the frame of discernment of the report is set to be the same as the frame of the variable of interest - both R_1 and R_2 are now defined on the frame Ω_X and $R_i = On \leftrightarrow X = On$, $R_i = Off \leftrightarrow X = Off$. This means the report can be treated as a perfect source with no frame alignment.

In the most simple example where the two sources are identical, we can consider R_2 to be inference on X using the same model relating R_1 and X . In other words, R_2 is not a real report, but rather a result of some reasoning process based on R_1 . If the two were to be combined, the result would be overly confident and clearly biased. As such, only one of the reports can be used for the combination using the conjunctive rule. The selection of the source plays no effect as the information they provide is identical.

This becomes more complicated when the general case is taken into account, and the information provided by R_2 is not necessarily identical to that provided by R_1 . This could be interpreted as R_2 being the inference on X using a different behaviour model. This is illustrated in the first part of this section using two simple examples. First, we consider one where the two reports are consonant and secondly one where they are not. Furthermore, it is shown that both of these examples can be modelled using the XQR source model

In the first example, consider the second report, R_2 , being a *more cautious* interpretation of the first report, e.g. due to Shafer discounting.

The following R_1 corresponds to $R = off$ in the original model from Section 3.3.2. To keep the results concise, we change the meaning of λ to refer to the probability of correct transmission rather than loss rate in this example.

$$m_{R_1} = \begin{cases} \lambda & R_1 = off \\ 1 - \lambda & R_1 = \Omega \end{cases}$$

Consider now the report R_2 which is the outcome of Shafer discounting of the initial report

using discounting rate μ . As such we obtain the following mass on R_2 .

$$m_{R_2} = \begin{cases} \mu\lambda & R_2 = \text{off} \\ 1 - \mu\lambda & R_2 = \Omega \end{cases}$$

Since R_2 is conditioned on R_1 consider the joint distribution induced by R_1 and R_2 on $\Omega \times \Omega$. Since R_2 is a less committed interpretation of R_1 it is only informative if R_1 is informative.

$$m_{R_1 \times R_2} = \begin{cases} \lambda(1 - \mu) & R_1 = \text{off} \times R_2 = \text{off} \\ \lambda\mu & R_1 = \text{off} \times R_2 = \Omega \\ 1 - \lambda & R_1 = \Omega \times R_2 = \Omega \end{cases}$$

At this point, let us discuss the different methods of combining the two non-distinct bodies of evidence. As we know and as is shown later in this section, Dempster's conjunctive rule of combination results in a mass function more committed than either of the two sources, a behaviour which is incorrect for non-distinct bodies of evidence. The disjunctive combination rule, on the contrary, produces an under-committed result. One straightforward solution is to simply choose one or the other. Furthermore, Dubois and Prade [23] proposed variations of conjunctive rule (and by extension, other set-operation based rules), which can be used if the relationship between the two dependent mass functions is known, i.e. if the joint mass function can be obtained from the marginals. Finally, we have the cautious rule (Equation 2.26), which can be applied even when the relationship between the sources is not known. The comparison of all these rules is shown in Table 3.6

Table 3.6: Combination of the two dependent sources using a variety of combination rules: conjunctive rule of combination (CRC), disjunctive (DRC), Dubois-Prade conjunctive and disjunctive rules for dependent sources and cautious rule of combination; numerical results provided for $\lambda = 0.8$ and $\mu = 0.7$

	m_{R_1}	m_{R_2}	CRC	DRC	DP-C	DP-D	Cautious
<i>off</i>	λ	$\lambda\mu$	$1 - [(1 - \lambda\mu)(1 - \lambda)]$	$\lambda^2\mu$	λ	$\lambda\mu$	λ
	0.8	0.56	0.912	0.448	0.8	0.56	0.8
Ω	$1 - \lambda$	$1 - \lambda\mu$	$(1 - \lambda\mu)(1 - \lambda)$	$1 - \lambda^2\mu$	$1 - \lambda$	$1 - \lambda\mu$	$1 - \lambda$
	0.2	0.44	0.088	0.552	0.2	0.44	0.2

It is clear that for this example, all the combination rules, bar the disjunctive rule of combination, can be generalized as a weighted conjunctive combination rule $m_1^\alpha \oplus m_2^\beta$ where the superscript corresponds to Shafer discounting rate. This, in turn, naturally corresponds to the source model discussed earlier (with one behaviour model (perfect source) on some report variable inducing R_1 and the other inducing R_2). This is useful as it makes it possible to perform a fusion of dependant sources without the implementation of alternative combination rules.

The reason why this is possible in this case is that the secondary source R_2 is entirely driven by R_1 and, as such, does not contain any information not already present in R_1 . The cautious rule of combination can be interpreted as an implementation of the inclusion-exclusion principle in

terms of confidence and diffidence components of belief functions: confidence components of the two belief functions are combined with the diffidence equal to their intersection, as such, in the case where R_1 contains R_2 , this is simply equivalent to the selection of one over the other.

3.5.2 Non-consonant reports

Unfortunately, and as expected, it can be shown that this trivial result does not necessarily extend to cases where the two sources provide unique information. The following example can be used to illustrate this. Consider the same source 1 inducing report R_1 as per the previous example. The second dependent source R_2 is obtained by moving a proportion of mass α from *off* to *on*. Hence the joint mass

$$m_{R_1 \times R_2} = \begin{cases} \lambda\alpha & R_1 = \textit{off} \times R_2 = \textit{on} \\ \lambda(1-\alpha) & R_1 = \textit{off} \times R_2 = \textit{off} \\ 1-\lambda & R_1 = \Omega \times R_2 = \Omega \end{cases}$$

Table 3.7: Combination of the two dependent sources using a variety of combination rules: numerical results provided for $\lambda = 0.8$ and $\alpha = 0.25$

	m_1	m_2	CRC	DRC	DP-C	DP-D	Cautious
\emptyset	0	0	$\alpha\lambda^2$	0	$\alpha\lambda$	0	$\frac{\lambda^2\alpha}{1-\lambda+\lambda\alpha}$
	0	0	0.16	0	0.2	0	0.4
<i>on</i>	0	$\alpha\lambda$	$\alpha\lambda(1-\lambda)$	0	0	0	$\frac{\lambda\alpha(1-\lambda)}{1-\lambda+\lambda\alpha}$
	0	0.2	0.04	0	0	0	0.1
<i>off</i>	λ	$(1-\alpha)\lambda$	$1-\dots$	$(1-\alpha)\lambda^2$	$(1-\alpha)\lambda$	$(1-\alpha)\lambda$	$\frac{\lambda(1-\lambda)}{1-\lambda+\lambda\alpha}$
	0.8	0.6	0.76	0.48	0.6	0.6	0.4
Ω_X	$1-\lambda$	$1-\lambda$	$(1-\lambda)^2$	$1-(1-\alpha)\lambda^2$	$1-\lambda$	$1-\lambda+\alpha\lambda$	$\frac{(1-\lambda)^2}{1-\lambda+\lambda\alpha}$
	0.2	0.2	0.04	0.52	0.2	0.4	0.1

We can see some significant differences between the different combination rules and furthermore see that it is not necessarily possible to replace the rules with a series of conjunctive combinations. This brings up the question of the optimal way to discount or otherwise model the reliability of the sources in such a way that they can be combined using conjunctive combination rules without overstating the evidence. By the optimum method, it is implied to retain the maximum amount of information from each source (as it is trivial to perform this kind of combination in a suboptimal way by simply ignoring all but one source).

For one the two reports R_1 and R_2 can be considered two *interpretations* of the same *source report*, consistent with the partially reliable source model. In order to be able to combine them when the exact relationship is unknown a different approach is required.

3.5.3 Idempotent combination of evidence using VBS

In this section it is shown that it is possible to combine two pieces of evidence in an idempotent manner using consecutive conjunctive combinations, which can be implemented in VBS. This




Figure 3.11: A valuation network implementing idempotent combination of reports R_1 and R_2

result is noted by Pichon et al. in [87], where it is stated that the weighted average of several sources is a special case of the behaviour-based correction model. We consider the two reports, R_1 and R_2 to be two *interpretations* of a single source report.

Consider the valuation network displayed in Figure 3.11. We have the two reports represented by variable nodes R_1 and R_2 and defined in valuations m_{R_1} and r_2 . X is, per usual, the variable of interest. The variable Q represents the "modes of operation" of the singular source. Here let it be defined $Q \in \{1, 2\}$, corresponding to selection of either of the two reports and driven by the valuation m_Q . Consider $X \in \Omega_X = \{X, \neg X\}$, $R_1 \in \Omega_1 = \{X_1, \neg X_1\}$ and $R_2 \in \Omega_2 = \{X_2, \neg X_2\}$. The masses m_1 and m_2 , defined on the frames of discernment $R_1 \times X \times Q$ and $R_2 \times X \times Q$ respectively represent the following relations:

$$\begin{aligned} R_1 = X_1 &\leftrightarrow X = X && \text{if } Q = 1 \\ R_2 = X_2 &\leftrightarrow X = X && \text{if } Q = 2 \end{aligned}$$

Now it can be shown that this combination (the inference on X) is idempotent if m_Q is dogmatic.

General solution Let us show the general solution of the most basic realization of such a network.

The mass functions m_1 and m_2 have their entire mass assigned to a single focal set F_1 and F_2 respectively such that

$$F_1 = \left\{ \begin{array}{l} 1 \times X_1 \times X \\ 1 \times \neg X_1 \times \neg X \\ 2 \times \Omega_1 \times \Omega_X \end{array} \right\} \quad F_2 = \left\{ \begin{array}{l} 2 \times X_2 \times X \\ 2 \times \neg X_2 \times \neg X \\ 1 \times \Omega_2 \times \Omega_X \end{array} \right\}$$

Note that for sake of saving space the notation used includes Cartesian products of the sets when appropriate, e.g. $1 \times \Omega_X = \{1 \times X, 1 \times \neg X\}$.

Let the masses m_{R_1} and m_{R_2} be defined as follows

$$m_{R_1}(\theta) = \begin{cases} r_1 & \text{if } \theta = X_1 \\ \bar{r}_1 & \text{if } \theta = \neg X_1 \\ r_1^\omega & \text{if } \theta = \{X_1, \neg X_1\} \end{cases} \quad m_{R_2}(\theta) = \begin{cases} r_2 & \text{if } \theta = X_2 \\ \bar{r}_2 & \text{if } \theta = \neg X_2 \\ r_2^\omega & \text{if } \theta = \{X_2, \neg X_2\} \end{cases}$$

And the mass on Q :

$$m_Q(\theta) = \begin{cases} s_1 & \text{if } \theta = 1 \\ s_2 & \text{if } \theta = 2 \\ s_\omega & \text{if } \theta = \{1, 2\} \end{cases}$$

Which after combination yields the following marginal on X

$$\begin{aligned} m(X) &= r_1 r_2 + s_1(r_1 \bar{r}_2 + r_1 r_2^\omega) + s_2(\bar{r}_1 r_2 + r_1^\omega r_2) \\ &= r_1 r_2 + s_1 r_1 (1 - r_2) + s_2 r_2 (1 - r_1) \\ &= s_1 r_1 + s_2 r_2 + s_\omega r_1 r_2 \\ (3.13) \quad m(\neg X) &= \bar{r}_1 \bar{r}_2 + s_1(\bar{r}_1 r_2 + \bar{r}_1 r_2^\omega) + s_2(r_1 \bar{r}_2 + \bar{r}_1 r_2^\omega) \\ &= s_1 \bar{r}_1 + s_2 \bar{r}_2 + s_\omega \bar{r}_1 \bar{r}_2 \\ m(\Omega_X) &= s_1 r_1^\omega + s_2 r_2^\omega + \\ &\quad + s_\omega(r_1^\omega r_2^\omega + r_1 \bar{r}_2 + r_1 r_2^\omega + \bar{r}_1 r_2 + \bar{r}_1 r_2^\omega + r_1^\omega r_2 + r_1^\omega \bar{r}_2) \end{aligned}$$

Proof of idempotence In this section, we show that this combination is idempotent for $n = 2$ provided that m_Q is dogmatic.

Proof. Consider the dogmatic distribution m_Q such that $s_2 = 1 - s_1$ and $s_\omega = 0$. In this case Equation system 3.13 simplifies to:

$$\begin{aligned} m(X) &= s_1 r_1 + s_2 r_2 + s_\omega r_1 r_2 \\ &= s_1 r_1 + s_2 r_2 \\ (3.14) \quad m(\neg X) &= s_1 \bar{r}_1 + s_2 \bar{r}_2 \\ m(\Omega_X) &= s_1 r_1^\omega + s_2 r_2^\omega \end{aligned}$$

Which is a weighted sum, which is idempotent by definition. ■

Reduced informativeness in non-dogmatic case In the event that the mass m_Q is not dogmatic, i.e. $s_\omega \neq 0$ the result from Equation 3.13 holds. If the masses are identical, the process is no longer idempotent, but some of the mass is moved to the universal set (or to supersets of larger cardinality in the general case discussed later).

Consider the equal report masses m_{R1} and m_{R2} such that $r_1 = r_2 = r$, $\bar{r}_1 = \bar{r}_2 = \bar{r}$ and $r_1^\omega = r_2^\omega = r^\omega$.

Equation 3.13 becomes:

$$\begin{aligned} m(X) &= s_1 r + s_2 r + s_\omega r^2 \\ m(\neg X) &= s_1 \bar{r} + s_2 \bar{r} + s_\omega \bar{r}^2 \\ m(\Omega_X) &= (s_1 + s_2) r^\omega + \\ &\quad + s_\omega (r^\omega + 2(r\bar{r} + r r^\omega + \bar{r} r^\omega)) \end{aligned}$$

Since $s_\omega \neq 0$, $s_1 + s_2 < 1$ and $m(X) < r$. Now consider the extreme scenario where $s_1 = s_2 = 0$ and $s_\omega = 1$. Now we have $m(X) = r^2$, $m(\neg X) = \bar{r}^2$ and $m(\Omega_X) = (r^\omega)^2 + 2(r\bar{r} + r r^\omega + \bar{r} r^\omega) = 1 - r^2 - \bar{r}^2$.

As such for two identical reports m_R we have:

$$\begin{aligned} r^2 &\leq m(X) \leq r \\ \bar{r}^2 &\leq m(\neg X) \leq \bar{r} \\ r^{\omega^2} + 2(r\bar{r} + rr^{\omega} + \bar{r}r^{\omega}) &\geq m(\Omega_X) \geq r^{\omega} \end{aligned}$$

Thus this combination method is idempotent if and only if either the mass m_Q or the mass r is dogmatic and simple (in the generalized case). Also, in the extreme case ($s_{\omega} = 1$), it degenerates to pointwise multiplication with the remaining mass being assigned to the empty set.

Generalization to a larger frame of discernment Consider $X \in \Omega_X = \{X_1, X_2, \dots, X_N\}$. Let \mathcal{F} be the set of all focal sets in 2^{Ω_X} . Two reports R_1 and R_2 are determined by valuations r and \hat{r} and a "selector" variable Q is defined on $\Omega_Q = \{1, 2\}$ such that 1 is associated with R_1 and vice versa. Q is determined by valuation m_Q such that $m_Q(1) = s_1$, $m_Q(2) = s_2$ and $m_Q(\Omega_Q) = s_{\omega}$ as per previous paragraph. It can be shown that for any focal set $F \in \mathcal{F}$

$$(3.15) \quad m(F) = s_1 r(F) + s_2 \hat{r}(F) + s_{\omega} \sum_{F=F' \cup F''} r(F') \hat{r}(F'')$$

The exact same proof can be used here to show that this combination is idempotent if m_Q is dogmatic.

Proof. Let $r(F) = \hat{r}(F) \forall F \in \mathcal{F}$. Furthermore let $m_Q(1) + m_Q(2) = 1$ and hence $s_{\omega} = 0$ From Equation 3.15:

$$\begin{aligned} m(F) &= s_1 r(F) + s_2 r(F) \\ &= (s_1 + s_2) r(F) \\ &= r(F) \end{aligned}$$

■

3.5.3.1 Discussion and interpretation

It has been mentioned several times throughout this section that the problem of possibly interdependent sources is a special case of the source model (see Section 3.5.1) where the two reports are interpreted as behaviours of the same source. To show that this is true, consider that the BBC model requires a multi-valued mapping Γ for every behaviour hypothesis in Ω_Q and every possible source report in Ω_R to the true state of the world in Ω_X .

Consider the original AIS example modeled using the BBC model as per Section 3.3.2. Let the report be $m_R(Rec) = 0.8$, $m_R(\neg Rec) = 0.2$ (which could be reasonable for data gathered over some

short period of time). Now assume that the two hypotheses we are concerned are H_2 and H_3 and the distribution is such that $m_Q(H_2) = m_Q(H_3) = 0.5$. The inference is simple but for reference:

$$M \oplus m_Q^{|R \times X} \left(\left\{ \begin{array}{l} rec \times on \\ \neg rec \times \Omega_X \end{array} \right\} \right) = 0.5 \quad M \oplus m_Q^{|R \times X} \left(\left\{ \begin{array}{l} rec \times \Omega_X \\ \neg rec \times off \end{array} \right\} \right) = 0.5$$

Where M is the mass function induced by the multi-valued mapping Γ .

From here it is straightforward to compute the marginal on x : $x(on) = 0.4$, $x(off) = 0.1$, $x(\Omega_X) = 0.5$. Now consider an alternative interpretation, where we have two independent experts who *have access* to the AIS receiver and make their inference based on that. They provide the following masses r and \hat{r} , which are the results of applying H_2 and H_3 to the mass r

$$r : \left\{ \begin{array}{l} on : 0.8 \\ \Omega_X : 0.2 \end{array} \right. \quad \hat{r} : \left\{ \begin{array}{l} off : 0.2 \\ \Omega_X : 0.8 \end{array} \right.$$

Applying Equation 3.15 with $s_1 = s_2 = 0.5$ results in $x(on) = 0.4$, $x(off) = 0.1$, $x(\Omega_X) = 0.5$ which is the same result as using the previous method, showing the two interpretations to be equivalent.

It naturally follows that being consistent with the source model, this approach could be applied to sources providing different behaviour models, as discussed through the earlier sections of this report. Furthermore, it stands to reason that this would be consistent with context-aware approaches.

In [82] it has been suggested that the BBC model and, by extension, the XQR model is a special case of Haenni's [88] remark that the disjunctive rule can be recovered by forming an evidential network with two binary frames representing the reliability of the two sources and a categorical mass function inducing the disjunctive combination. The approach to combination proposed here is, too, an extension of Haenni's framework where the mass function relating the reliabilities of the two sources is not necessarily categorical. As we have discussed before, if $s_\omega = 1$, the disjunctive rule is induced, and the network is simplified to the one which yields Haenni's result.

Comparison with the cautious rule Let us compare this approach to using the cautious rule of combination. The key advantage from the implementation point of view is that it can be used directly in the BFM, where combination rules other than the conjunction are not supported.

The two approaches should be compared from a more theoretical rather than entirely implementation-oriented perspective. The underlying principle of the cautious rule is that if the two sources are equally reliable, one source claims that the target is in the set A and the other that it is in the set B we should be certain that it indeed lies within the intersection $A \cap B$. The cautious rule selects the least informative element satisfying the above condition.

On the contrary, utilizing the metaknowledge model for combination can be interpreted as a probabilistic combination of the sources - source A takes precedence with some probability; otherwise, source B takes precedence. The vacuous mass on source selection corresponds to a disjunctive combination. Furthermore, some degree of total agreement (pointwise multiplication of the masses) is transferred to the combination outcome regardless of the relative weights assigned

to sources. As such, it could be interpreted that, in general, provided two sources claiming that the target is in sets A and B respectively, we know for certain that it is plausible that the target is in the disjunction $A \cup B$, with some probabilities assigned to sets A and B respectively. Assignment of parameters s_n and s_w determines the degree to which mass is transferred to the sets A , B and $A \cup B$.

In information fusion, in general, it is usually assumed that provided two sources, the combination of them should be not less informative than either one. This is often relaxed for conflicting sources, and rules such as the disjunctive rule of combination are used. However, it could be argued that this should not necessarily apply to sources that are not independent. In fact, if both sources are otherwise equal, they should have similar information content. It is interesting to consider that in such a scenario, lack of information could be considered information itself - if two sources provide consonant beliefs, both being equally reliable and trustworthy, it is not necessarily optimal to commit to the more informative one.

Consider an example of two experts providing information, one of them being more cautious than the other - either one of them is overly cautious or the other is overly bold - but it is unknown which one it is. This is essentially the first example in this section, with consonant belief assignments, which we will now use to compare the approach outlined here to the other combination rules. The results were previously shown in Table 3.6, where all the combination rules for dependent sources defaulted to either of the two masses. Applying Equation 3.15 to this problem with weights $s_1 = s_2 = 0.5$ yields $m_J(\text{off}) = 0.68$ and $m_J(\{\text{on}, \text{off}\}) = 0.32$, which is the only result where the information content of the joint is greater than the less-informative report but less than the more-informative report. As such, it could be argued that it is the only case where the extra (meta)-information provided by the second source is leveraged (which could be interpreted as "Source 1 is overconfident") without ignoring the first report entirely. Shifting the weights makes it possible to make this report more or less informative.

This method of combination is not more informative than cautious rule of combination for any set of inputs. In general $w_J(A) \in [w(A), \hat{w}(A)]$ for any A in m_J where m_J is the result of combination. In comparison $w_\wedge(A) = \min(w(A), \hat{w}(A))$ where $m_\wedge(A)$ is the resultant mass produced by the cautious combination rule. As such $w_J(A) \geq w_\wedge(A)$ for any A and hence $m_\wedge \sqsubseteq_w m_J$, i.e. m_J is less informative than m_\wedge under w -ordering.

It also makes sense that the information content of the joint mass could be improved when considering the scenario where the sources may or may not be dependent. The behaviour of the cautious rule is such that it allows for overlap of information between the sources, but it will extract information from both nevertheless, whereas the approach outlined here performs more of an either-or combination. Allowing a hypothesis that the beliefs are indeed independent would mean that the combination is no longer idempotent and as such, it would be down to appropriate metaknowledge selection to ensure that no double counting of evidence takes place. This is discussed in the next section.

3.5.4 Possibly independent sources and context-aware independence

Let us discuss a scenario where it is unknown whether two pieces of information are or are not independent. For instance, a human intelligence source may provide some information, and the reasoning model for it may not be available. As an example, consider an agent providing a report on the frame of discernment from the earlier vessel identification example $\Omega = \{c, f, o\}$ which is consistent with the frame of the variable of interest. Furthermore, let there be another report R providing information on the same frame of discernment. Now there exists a possibility that this information was already included in the analyst report (and may have been disregarded on purpose due to the unknown reasoning model - as such information content analysis of the report may not be sufficient) or that it was not available at the time the report was generated. As such, it stands to reason there could be two models - one with this piece of information was included (and the sources are dependent) and one without.

Another interesting concept is that of a metainformation source, namely a piece of information affecting the behaviour of the source, which often cannot be expressed in the same frame of discernment Ω . This section, and some other more complex relationships between information sources, and their relation to the XQR model, are discussed.

3.5.4.1 Possibly independent information source

First, consider the simple scenario where a report is produced by some source that may or may not be independent. Consider R on $\Omega_R = \Omega$.

Possibly independent information source - some information source is providing evidence that the actual value of X is not A . This information can be readily represented as a valuation on X (a separate report, R_2). However, a double-counting problem occurs as it is unknown whether this piece of information has already been included in the analyst report R_1 .

One apparent solution is that briefly described at the end of the previous section - using the generalized metaknowledge model, we can independently model the hypotheses that R_1 and R_2 are dependent and that they are independent. In the former case, we can split it further into two more hypotheses using selection variable to allow for idempotent combination. This can be formalized using two quality variables: *dependence* $D = \{Dep, \neg Dep\}$ and *selection* $S = \{1, 2\}$, the operation of latter was discussed in detail in the previous section. The Cartesian combination of the two defines the joint frame $D \times S = \{1 \times Dep, 1 \times \neg Dep, 2 \times Dep, 2 \times \neg Dep\}$. However, the selection variable is irrelevant when the reports are independent, and this can be mapped to the following XQR model hypotheses

$$\begin{aligned}1 \times Dep &\rightarrow H_1 \\2 \times Dep &\rightarrow H_2 \\ \Omega_S \times \neg Dep &\rightarrow H_3\end{aligned}$$

where H_1 and H_2 correspond to the appropriate hypotheses from the previous section and H_3 corresponds to the simple conjunctive combination of the two reports.

Consider we assume a 50% probability that the reports are independent $D(\neg Dep) = 0.5$ and use equal weights $S(1) = S(2) = 0.5$. As such we obtain the following distribution on the hypotheses: $Q(\{H_1, H_3\}) = 0.25$, $Q(\{H_2, H_3\}) = 0.25$, $Q(H_3) = 0.5$. Note that in this case there is no difference for the combination process between mass assigned to $Q(\{H_2, H_3\})$ and $Q(H_2)$ (this is analogous to normal discounting process where no information on reliability implies non-reliability). Let report R1 be determined by the mass $r(o) = 0.5$, $r(\{o, f\}) = 0.2$, $r(\{o, c\}) = 0.2$, $r(\Omega_X) = 0.1$ and the report R2 by mass $r(\neg c = \{o, f\}) = 0.8$, $r(\Omega_X) = 0.2$. The marginal on X using the aforementioned quality descriptor variables is then $m(o) = 0.455$, $m(\{o, f\}) = 0.39$, $m(\{o, c\}) = 0.07$, $m(\Omega_X) = 0.085$

Let us take a step back and once more compare these results to these obtained using other combination methods - Dempster's rule, the cautious rule and for the different valuations on D and S

Table 3.8: Comparison of XQR model-based combination for different values of D and S (assuming $s_1 = s_2$) with combination using Dempster's, cautious and disjunctive rules

	r_1	r_2	Dempster \oplus	Cautious \wedge	Generalized metaknowledge		
					H_S^D		
					$H_{0.5}^{0.6}$	$H_{0.5}^0$	$H_0^{0.5}$
o	0.5	0	0.66	0.58	0.455	0.25	0.33
$\{o, f\}$	0.2	0.8	0.28	0.24	0.39	0.5	0.42
$\{o, c\}$	0.2	0	0.04	0.12	0.07	0.1	0.02
Ω_X	0.1	0.2	0.02	0.06	0.085	0.15	0.23

Where H_S^D denotes the VBS used in this section, with $S(1) = S(2) = S$ and $D(\neg Dep) = D$. Note that in the case where $D = 0$ this corresponds to the idempotent case discussed previously, and when $S = 0$, this becomes the weighted conjunctive-disjunctive rule, i.e. mass assigned to every focal set is the weighted sum of the results of the combination using Dempster's and disjunctive rule respectively. Furthermore, unsurprisingly the result for $H_{0.5}^{0.5}$ is the weighted sum of the result for $H_{0.5}^0$ and Dempster's rule.

3.5.4.2 Metainformation and context-aware fusion

The source of information takes the following form: *evidence which typically points at c is actually pointing at o*. It could hypothetically be driven by some piece of information telling us that, for example, *tanker (other) types are likely to be misclassified as cargo*.

Similarly to the previous case, there could be two behaviour hypotheses corresponding to the cases where this extra information has already been included in the source and the case where it has not. However, since this source does not provide information on X directly, the combination is less straightforward. The key difference here is that the source in question is meaningless in the absence of another report - it does not directly provide information on the variable of interest and, as such, is not a *report* in a sense used throughout this work. It would best be described as a contextual variable in this case, as it places some constraints on the focus but does not intervene directly. This also means that usual combination rules - whether the one discussed in the previous section, the cautious rule or the conjunctive rule cannot be applied directly.

Let the source be interpreted in such a way that mass assigned to the focal sets c and $\{c, o\}$ should be redistributed to the hypothesis o ¹. This could be modelled as a mapping $\Lambda : 2^{\Omega_X} \rightarrow 2^{\Omega_X}$ such that $\Lambda(\theta) = o$ if $\theta = c$ or $\theta = \{c, o\}$ and $\Lambda(\theta) = \theta$ otherwise.

Using the same logic as before we can reason that there should be two modes of inference: one where this piece of contextual information has already been taken into account and one where it has not. Let us denote the former hypothesis with H_1 and the latter with H_2 . This corresponds to the *dependence* variable D from the previous case. The mapping Γ from report-behaviour space to focus space can be then derived:

$$\begin{array}{ll} \Gamma_{\theta}(H_1) = \theta & \forall \theta \in \Omega_X \\ \Gamma_{\theta}(H_2) = o & \theta = c \\ \Gamma_{\theta}(H_2) = \theta & \theta \neq c \end{array}$$

Consider the initial report R_1 as per previous example: $r(o) = 0.5$, $r(\{o, f\}) = 0.2$, $r(\{o, c\}) = 0.2$, $r(\Omega_X) = 0.1$ and the assumed degree of dependence $Q(H_1) = d(Dep) = 0.6$ $Q(H_2) = d(\neg Dep) = 0.4$. The dependence value is chosen to reflect the information content of the previous example, which included $d(Dep) = 0.5$ and $r_2(\Omega_X) = 0.8$ This yields the following resultant mass function $m(o) = 0.58$, $m(o, f) = 0.24$, $m(o, c) = 0.12$, $m(\Omega_X) = 0.06$ - although this is the same as the result of cautious combination in the previous example this is coincidental and this result does not hold as parameters are varied.

3.6 Summary

This chapter addressed the challenge of problem modelling, and more specifically, source modelling with valuation network through two lenses. On the one hand, we have considered the challenge of modelling partially reliable sources and the role that contextual information can play in the design of accurate models. It has been shown that it is possible to align the different existing definitions of context in light of *context-of* and *context-for*. Subsequently, it is possible to frame the problem of context in a way consistent with contextual correction of belief functions and behaviour-based correction model for partially reliable sources. The introduction of context as a variable in a variant of the behaviour-based correction model has the natural benefit of increased explainability associated with the transparency-by-design of valuation networks, but also it is shown to improve the expressiveness of the model. On another, we have considered the problem of non-independent sources. Some existing literature on the problem was discussed, in particular combination methods for belief functions that take into account interdependencies between the sources, such as the cautious rule and Dubois-Prade rule non-independent sources. This is compared to using Dempster's rule only by considering the degree of inter-dependency between two sources as part of the behaviour-based correction model. Some questions remain open-ended.

¹A more interesting variant would be the case where this has no effect on the mass assigned to the focal sets $\{c, f\}$ and $\{c, f, o\}$, but it does not seem possible to implement in VBS and would pose other significant problems

Particular areas of interest which could be investigated in more detail include a more formal investigation of non-independent sources. The notion of meta-independence has been proposed in [87], but was not discussed in much detail. A more in-depth investigation into source independence as well as contextually conditional source independence is one of the key areas of further work. The question of conflict between sources was briefly mentioned; however, especially in light of other notions discussed in this chapter, it deserves a more in-depth discussion, as a conflict between the sources is an indicator of either incorrectly modelled sources (where their reliability is overestimated) or greater issues such as an incorrect model of the problem (e.g. non-exhaustive frame of discernment as per Zadeh's example). As this chapter addresses the intricacies of source modelling and the modelling of relationships between them, it is an area that could be addressed further. The problem of interpretation of temporal reports, where a single source is sequentially sampled is relevant to the work in this chapter and has not been addressed at all.

PROVENANCE OF INFORMATION, CONFLICT AND UNCERTAINTY

Provenance refers to the origin of *something*. In the context of multi-source information fusion, this term may be employed to refer to tracking the source of information, conflict or uncertainty. The first concept discussed here is the one that most intuitively comes to mind when considering the issue of provenance in fusion, namely the origin of the information. Identifying the origin of and handling conflict is a major issue in information fusion, with several authors discussing methods of conflict management as well as computing the relative contribution to the conflict. Finally, in some information fusion systems, it can be useful to identify the relative impact of individual sources on the overall uncertainty.

This chapter is structured as follows. First, the notions of explainability and interpretability are recalled, and the relation between this and the previous chapter is established. Subsequently, the theoretical foundations of provenance are introduced, focusing on the provenance of information, establishing the expected behaviours of a generalized provenance measure. In the next section, some provenance of information measures in the belief function framework are proposed and compared. This is followed by a discussion of measures of the provenance of conflict and nonspecificity.

4.1 Relation to contextual reasoning, explainability and interpretability

As discussed in the previous section, two main reasons for context exploitation exist - adaptability and explainability. The explainability provided by context exploitation can be further augmented by provenance measures, i.e. quantifying the relative contribution to the belief (or decision) of individual sources, which can be either source reports or pieces of contextual information. It has also been shown that the context-aware partially reliable source model allows improved expressiveness through the inclusion of what essentially are negative belief masses. As with

this approach individual sources can increase overall uncertainty, in such an information fusion system, the provenance of uncertainty will actually be important.

4.2 Theoretical foundations of provenance

4.2.1 Provenance of information

Information fusion is a specific aggregation process that aims to extract truthful knowledge from incomplete, uncertain or conflicting sources. In general terms, we could consider the outcome of the information fusion process as a type of knowledge base with multiple sources contributing towards it.

Now, for a certain piece of information within that knowledge base, we may be interested in tracking the origin of that particular piece of information - for example, if we believe that a patient suffers from a specific disease, is that belief-driven by a doctor's opinion, patient's self-assessment or a blood test? As such, we may interpret the provenance of information as the process of tracking the origin of a specific piece of knowledge within our knowledge base.

From this, we can discuss two very closely related concepts - the provenance of belief and the provenance of decision. The difference between belief and information is very subtle, and for the purpose of this discussion, we will treat belief as a more refined form of information. For instance, the provenance of information could be identifying the source, which gives us the information on possible distributions of seats in parliaments following an election - a wide array of possible outcomes. On the other hand, the provenance of belief would involve identifying the sources driving the belief in one specific outcome, e.g. one specific party winning by a landslide. Do note that these distinctions are blurry at best, and they are largely affected by granularity and "level" of data or information. However, it is important to remember that what is considered high-level information under one implementation or use case may become low-level data in another case.

Finally, we can discuss the provenance of the decision. Whilst the concept may be more complicated, it is easier to draw a line between itself and the other two. Going back to the scenario with identifying the disease troubling a patient, consider that a decision must be made regarding the further course of treatment. Whilst this decision could be a simple function of our degree of belief in the possible diagnoses, it is likely there will be other considerations such as side effects or cost.

4.2.2 Mathematical formalisation of information fusion and decision making

In this section, some mathematical terms are defined in such a manner that they may be applied regardless of the uncertainty formalism chosen for information fusion. This can be complemented with the mathematical definitions from [1].

Consider a decision making problem where an *action* is taken under uncertainty and it has some *consequence* dependent on the *state of the world*. Let $\Omega = \{\omega_1 \dots \omega_n\}$ be the frame of discernment, i.e. the set of all possible worlds and $\mathcal{C} = \{c_1 \dots c_n\}$ be the set of possible consequences. Finally the set $\mathcal{A} = \{a_1 \dots a_n\}$ is the set of all possible actions. Thus $a : \Omega \rightarrow \mathcal{C}$ In addition let us

assume that there exists a utility function mapping the effect of each consequence to a real-valued variable $u : \mathcal{C} \rightarrow \mathbb{R}$

In [1] a general setting for representation of information items is proposed, such that T_i represents an *information item*, which may be a subset of possible worlds, a plausibility relation, a fuzzy set, a mass assignment and so forth. This framework was summarised in Chapter 2.

Let $\mathcal{T} = \{T_1 \dots T_n\}$ be the set of possible information items of a certain format. Again, according to [1] a *n-ary fusion operation* is a mapping $f^n : \mathcal{T}^n \rightarrow \mathcal{T}$. When it is not ambiguous we replace f^n by f (see¹).

We may also define a knowledge base $\mathcal{K} \subseteq \mathcal{T}$, a set that contains all the information items based on which a decision is made. Additionally to account for different fusion mechanisms we define a set $\mathcal{F} = \{f_1 \dots f_n\}$. There exist multiple decision making formalisms which belong to the set \mathcal{D} , $d \in \mathcal{D} : \mathcal{T} \rightarrow \mathcal{A}$

We may define a *combine-and-decide* operation g as a composition of a fusion operation f and a decision operation d such that $d \circ f : \mathcal{K}^n \rightarrow \mathcal{A}$ - a mapping from a subset of the knowledge base to an action.

Furthermore let us, as per [1] further the information item T as having the following features: its *support*, $S(T) \subseteq \Omega$, its *core* $C(T) \subseteq \Omega$ and an induced *plausibility ordering*. The *support* $S(T)$ corresponds to the set of all worlds which are not impossible, whereas the *core* $C(T)$ corresponds to the set of all worlds which are fully plausible. This definition can be generalised to different frameworks for representing information items, such as sets, plausibility relations, Dempster-Shafer mass assignments and others.

Furthermore, we can define two additional terms which allow us to compare the relative informativeness of information sources: *information ordering* and *imprecision index*

Information ordering is a partial preorder relation (reflexive and transitive): $T \sqsubseteq T'$ means that T provides at least as much information as T' - it is more precise or more specific i.e. if $T \sqsubseteq T'$ then $S(T) \subseteq S(T')$ and $C(T) \subseteq C(T')$. This coincides with set inclusion and fuzzy set inclusion, but its meaning is less obvious for belief functions and plausibility relations. Note that for this ordering to be meaningful, a degree of mutual consistency is required.

Imprecision index is a measure $II(T)$ of how much information is contained in an information item - depending on the uncertainty representation model used, it can be reduced to some index of non-specificity. From this, it is possible to define information monotonicity as a property of the fusion operator f

Definition 4.1. Let information monotonicity be a property of the fusion operator f , which can be enforced only when the sources of information do not contradict one another

It can be defined as follows If $\forall i T_i \sqsubseteq T'_i$ and $\exists j : T_j \sqsubset T'_j$ then $f(T_1, \dots, T_n) \sqsubset f(T'_1, \dots, T'_n)$ if $S(T_1) \cap S(T_2) \dots \cap S(T_n) = \emptyset$

¹Note the change from Dubois' article where the arity of the fusion operator is denoted by subscript

4.2.3 Generalised measures for the provenance of decision, belief and information

In this chapter contribution measures are proposed to quantify the provenance of decision, belief and information.

The first contribution measure, quantifying the provenance of decision, can be defined as a multivariate mapping from a subset of knowledge base, an element in the knowledge base and an action to a real-valued variable. The purpose of this contribution is to assess the degree to which a single element in the knowledge base contributes to the selection of a particular action when a subset of the knowledge base is combined and used to make a decision.

$$(4.1) \quad c : 2^{\mathcal{K}} \times \mathcal{K} \times \mathcal{A} \rightarrow \mathbb{R}$$

Furthermore, we may also consider a simplified case where the decision making is limited to identifying the most likely element in the frame of discernment. In this case, $d : T \rightarrow \Omega$ and the following measure quantifies the provenance of belief.

$$(4.2) \quad c : 2^{\mathcal{K}} \times \mathcal{K} \times \Omega \rightarrow \mathbb{R}$$

Finally in order to quantify the provenance of information $d : T \rightarrow 2^{\Omega}$:

$$(4.3) \quad c : 2^{\mathcal{K}} \times \mathcal{K} \times 2^{\Omega} \rightarrow \mathbb{R}$$

Now we can propose the properties that a provenance measure should satisfy:

Proposition 4.1 (Preference). *Given fusion operator f and decision method d , knowledge base \mathcal{K} , where $K^* \subseteq \mathcal{K}$, $k_1, k_2 \in K^*$ and possible distinct actions $a_1, a_2 \in \mathcal{A}$ if*

$$d \circ f(K^*) = a_1$$

and

$$d \circ f(K^*/k_1) = a_2$$

and

$$d \circ f(K^*/k_2) = a_1$$

then

$$c(K^*, k_1, a_1) \geq c(K^*, k_2, a_1)$$

This can be understood as follows: if removing a piece of information k_1 is sufficient to change the decision from a_1 to something else, but removing k_2 does not have this effect, the contribution of k_1 to a_1 must be greater than contribution of k_2 to a_1 .

Proposition 4.2 (Non-degeneracy). *Given knowledge base \mathcal{K} , where $K^* \subseteq \mathcal{K}$ and $k_1 = k_2$*

$$c(K^*, k_1, a) = c(K^*, k_2, a) \forall a \in \mathcal{A}, K^* \in \mathcal{K} : k_1, k_2 \in K^*$$

Proposition 4.3 (Inclusion). *Given fusion operator f which enforces information monotonicity and decision method d , knowledge base \mathcal{K} , where $K^* \subseteq \mathcal{K}$ with $k_1, k_2 \in K^*$ such that $k_1 \sqsubseteq k_2$ and the sources are not contradicting then*

$$c(K^*, k_1, a) \geq c(K^*, k_2, a)$$

Proposition 4.4 (Vacuous insensitivity). *Given knowledge base \mathcal{K} and a vacuous information item $k^V \in \mathcal{K}$ then:*

$$c(K^*, k^V ma) = 0 \forall a \in \mathcal{A}$$

4.2.4 Provenance of conflict

The notion of conflict is strongly related to that of (in)consistency in classical logic. In the belief function framework, conflict is defined as the inconsistency of the conjunctive combination of two pieces of information. In a logical framework, formulas are considered conflicting if inclusion in a knowledge base causes it to become inconsistent, i.e. there does not exist an interpretation satisfying both of these formulas.

More formally, if we consider information items T_i to be formulas in some propositional language, the knowledge base K is a finite set of logical formulas. A belief base K is inconsistent if there exists a formula α such that $K \vdash \alpha$ and $K \vdash \neg \alpha$. Some measures quantifying the degree of inconsistency are based on the concept of *minimal inconsistency sets* [89].

Definition 4.2 (Minimal inconsistent sets). Given an inconsistent belief base K , its minimal inconsistent subsets MI are defined by:

$$MI(K) = \{K' \subset K \mid K' \vdash \perp \text{ and } \forall K'' \subset K', K'' \not\vdash \top\}$$

i.e. they are inconsistent subsets of the knowledge base, such that removal of any single element makes it consistent. As such, we can state that any elements in K' are in conflict. Furthermore, we can consider measures quantifying inconsistency to be quantifying conflict.

As such, similarly to the previous section, *consistency measure* is a mapping from a subset of the knowledge base to a real valued variable:

$$(4.4) \quad \phi : 2^{\mathcal{K}} \rightarrow \mathbb{R}$$

. and conflict between elements or subsets of the knowledge base can be defined as:

$$(4.5) \quad \hat{\phi} : f(k_1, k_2) \rightarrow \mathbb{R}$$

. where k_1, k_2 are disjoint subsets of \mathcal{K} and f is some fusion operator. This measure has some desirable properties (as per [90]): 1. Bounded 2. Extreme measures 3. Symmetry ($\hat{\phi}$) 4. Imprecision monotonicity 5. Vacuous insensitivity (ignorance is bliss) 6. Insensitivity to refinement

and the provenance of conflict is a multivariate mapping from a subset of the knowledge base and an element in the knowledge base to a real-valued variable

$$(4.6) \quad \kappa : 2^{\mathcal{K}} \times \mathcal{K} \rightarrow \mathbb{R}$$

Again, the purpose of this contribution measure is to assess how much a particular element of the knowledge base contributes to the conflict when some subset of this knowledge base is combined. As per previous case consider the following properties:

Proposition 4.5 (Preference). *Given a knowledge base \mathcal{K} , where $K^* \subseteq \mathcal{K}$, $k_1, k_2 \in K^*$, if*

$$\phi(K^*/k_1) \geq \phi(K^*/k_2)$$

then

$$\kappa(K^*, k_1) \leq \kappa(K^*, k_2)$$

Proposition 4.6 (Non-degeneracy). *Given a knowledge base \mathcal{K} , where $K^* \subseteq \mathcal{K}$, $k_1, k_2 \in K^*$, if*

$$\phi(K^*/k_1) = \phi(K^*/k_2)$$

then

$$\kappa(K^*, k_1) = \kappa(K^*, k_2)$$

Proposition 4.7 (Inclusion). *Given fusion operator f which enforces information monotonicity, knowledge base \mathcal{K} , where $K^* \subseteq \mathcal{K}$ with $k_1, k_2 \in K^*$ such that $k_1 \sqsubseteq k_2$*

$$\kappa(K^*, k_1) \geq \kappa(K^*, k_2)$$

Proposition 4.8 (Vacuous insensitivity (ignorance is bliss)). *Given knowledge base \mathcal{K} and a vacuous information item $k^V \in \mathcal{K}$ then:*

$$\hat{\phi}(v, k^V/v) = 0$$

4.2.5 Origin of uncertainty

Finally, we consider tracking the origin of uncertainty in the information fusion system. The concept of uncertainty refers to several different notions, as outlined in Chapter 2 Origin of uncertainty is quite similar to the origin of the conflict, as it can be argued that conflict within the information is a form of uncertainty as discussed in more detail in Section 2.3.2. As such particular focus here is placed on the non-conflicting types of uncertainty, fuzziness (vagueness) and nonspecificity. In the majority of information fusion systems inclusion of additional sources reduces uncertainty. However, there are some exceptions to this rule. In Dempster's rule of combination, a highly conflicting source can increase the size of the universal set due to conflict redistribution. As discussed in the previous section, more complex relationships between variables, such as contextual reasoning, may also increase uncertainty. Information about the quality of the sources may reduce the degree of certainty associated with the overall inference.

4.3 Origin of information and decision in belief functions framework

As in the remainder of this work, the main focus throughout the rest of this chapter will be on implementations under the belief functions framework.

4.3.1 Simple contribution measures and the problem with preference

A natural method of assessing the contribution of the individual source to the final belief in a proposition is to use one of the common measures within the belief function framework, such as the belief, plausibility, commonality or similar as the measure of contribution. Although generally, they are a good indication of relative contributions of the belief sources, it can be shown that they do not always satisfy all the proposed properties, even when selected in conjunction with the appropriate decision method and combination rule. For example, consider the most natural candidate, the belief measure Bel used in an information fusion system exploiting Dempster's rule of combination and the maximum of belief rule for singleton selection. The Bel is considered the most natural candidate as the greater a belief in a hypothesis from a single source of information, the more it will contribute to the overall belief in that hypothesis following information fusion. In other words, if we consider several sources of information, the one with the greatest belief in the hypothesis can be expected to contribute the most to the overall belief in that hypothesis. Other possible candidates, such as plausibility and commonality measures by definition, do not satisfy the vacuous insensitivity property (Proposition 4.4).

It is straightforward to show that there exist cases where the preference property is not satisfied, as we prove below:

Proof. Consider two mass functions m_1 and m_2 defined on the frame of discernment $\Omega = \{\omega_1 \dots \omega_n\}$ within a knowledge base \mathcal{K} , such that m_1 is a simple support function with weight $1 - \alpha$ for some $\omega \in \Omega$ and m_2 is a simple support with weight $1 - \beta$ for a focal set F such that $\omega \subset F$. Let

$$\hat{m} = \bigoplus_{m \in \mathcal{K}, m \neq m_1, m_2} m \text{ and } m_J = \bigoplus_{m \in \mathcal{K}} m = \hat{m} \oplus m_1 \oplus m_2$$

By definition of belief measure

$$Bel_{m_1}(\omega) = \alpha$$

$$Bel_{m_2}(\omega) = 0$$

.

For Bel to be a contribution measure satisfying the preference property using the maximum of belief decision and Dempster's rule of combination, if:

$$\operatorname{argmax}_i Bel_{m_J}(\omega_i) = \omega$$

and

$$\operatorname{argmax}_i Bel_{\hat{m} \oplus m_2}(\omega_i) = \omega$$

then

$$\operatorname{argmax}_i Bel_{\hat{m} \oplus m_1}(\omega_i) = \omega$$

.

Now consider \hat{m} with mass distribution such that mass γ is placed on $\bar{\omega} \in \Omega$ such that $\bar{\omega} \neq \omega, \bar{\omega} \notin F$ and the remaining $1 - \gamma$ is distributed among focal sets $\mathcal{G} = \{G_1, G_2 \dots G_n\}$ such that $F \cap G = \omega, \forall G \in \mathcal{G}$.

By Dempster's rule of combination:

$$Bel_{\hat{m} \oplus m_2}(\omega) \propto \beta(1 - \gamma)$$

and

$$Bel_{\hat{m} \oplus m_2}(\bar{\omega}) \propto (1 - \beta)\gamma$$

Hence if $\beta > \gamma$ then $\arg \max_i Bel_{\hat{m} \oplus m_2}(\omega_i) = \omega$ and similarly if $\alpha < \gamma$ then $\arg \max_i Bel_{\hat{m} \oplus m_1}(\omega_i) = \bar{\omega}$.

Thus for $\alpha < \beta$ the preference property is not satisfied ■

In less formal terms, this can be associated with the property of Dempster's rule, as well as the conjunctive rule of combination by which the impact of a mass m_1 on the resulting belief in a hypothesis can be equally driven by the mass assigned to that hypothesis as well as some of its supersets (provided that they are not included in m_2). This implies that in general unary measures cannot fully capture this.

It is important to note that the preference property is difficult to satisfy for all possible cases, and it will be a recurring problem throughout this chapter. However it would not seem reasonable to dismiss it in general, but rather it can be useful to treat it as a sort of benchmark as it is done throughout this chapter. In general we can consider a situation with two partially conflicting sources m_1 and m_2 , where the hypotheses ω_1 and ω_2 are both strongly supported by m_1 , such that ω_1 is slightly more likely than ω_2 . However ω_1 is deemed not plausible by m_2 and ω_2 is weakly supported. By the definition of preference, m_2 should have a greater contribution. However, it is unlikely to be able to find a contribution measure that will satisfy all possible cases. Although Dempster's conflict redistribution is part of the problem, the above result holds even in the case of the unnormalised conjunctive combination rule.

4.3.2 Generation of explanations in the theory of evidence

There exists some, albeit limited, work on explainability in the context of belief functions. Strat and Lawrence introduced the concept of sensitivity spaces for the provision of explanation abilities in systems using belief functions [91] [92]. This work was later expanded on by Xu and Smets [93]. In this section, the existing literature is discussed, and some of the notions are extended in line with requirements for a contribution measure, as outlined earlier in this chapter.

4.3.2.1 Sensitivity spaces

The method proposed for assessing the impact of individual bodies of evidence on the final belief in some proposition involved utilization of first derivatives of belief and plausibility measures with respect to the discounting rate (reliability).

For a hypothesis A and BoE i with reliability r_i

$$(4.7) \quad \widehat{Bel}_i(A) = \left. \frac{\delta Bel(A)}{\delta r_i} \right|_{r_i=1}$$

and

$$(4.8) \quad \widehat{Pl}_i(A) = \left. \frac{\delta Pl(A)}{\delta r_i} \right|_{r_i=1}.$$

These two measures are used to generate sensitivity spaces, which explain the marginal impact of each BoE on the hypothesis of interest. By extension, they can also be used as contribution measures.

The notion of preference we have proposed is closely aligned with how these derivatives are used to assess the impact of a piece of evidence, as removing a single body of evidence is tantamount to setting its reliability $r_i = 0$. This was discussed to a greater extent by Xu and Smets [93], where they show that computation of this derivation for unnormalised belief functions is equivalent to simply removing the body of evidence of interest. While in the majority of situations, the preference property as defined earlier is satisfied when an appropriate measure is selected given the decision-making method (e.g. sensitivity of belief for a maximum of belief), there exist exceptions to that behaviour, in particular when one of the masses is strongly inconsistent.

Strat and Lowrance considered the overall impact of a body of evidence on the overall inference for a particular hypothesis to be the distance from the origin on the $\widehat{Bel} - \widehat{Pl}$ space, i.e.

$$(4.9) \quad Imp_i = \left| \left(\widehat{Pl}_i(A), \widehat{Bel}_i(A) \right) \right| = \sqrt{\widehat{Pl}_i(A)^2 + \widehat{Bel}_i(A)^2}$$

while this captures the overall impact quite well, it does not consider the direction of that impact and as such it is unsuitable to be used as a contribution measure. A natural extension to make it suitable for use as a contribution measure is to consider its scalar projection onto the direction corresponding to maximum contribution, $\widehat{Bel} = \widehat{Pl}$, as per Fig 4.1 Let \vec{d} denote the vector

$$(4.10) \quad d_i \vec{d}(A) = \begin{bmatrix} \widehat{Bel}_i(A) \\ \widehat{Pl}_i(A) \end{bmatrix}$$

and we can consider the scalar projection of \vec{d}_i onto $\widehat{Bel}_i = \widehat{Pl}_i$ to be a contribution measure. For convenience let \vec{p} denote the unit vector parallel to $\widehat{Bel}_i = \widehat{Pl}_i$. Thus we have the contribution measure:

$$(4.11) \quad d_i = \|(\vec{d}_i \cdot \vec{p})\vec{p}\|$$

4.3.3 Distance-based contribution measures

Consider a scenario where two bodies of evidence m_1 and m_2 are combined using one of the combination rules. The identification process using one of the decision rules identifies the hypothesis $A \subseteq \Omega$ as true or most likely to be true. We propose a measure quantifying the relative contribution of an individual BBA to the decision. One possible method of constructing a provenance measure is to exploit the geometric interpretation of belief functions. As such, Jousselme's distance is used to define such a measure. Other distance metrics could be used here instead, but this distance is




Figure 4.1: Sensitivity of belief and plausibility to discounting of a mass, illustrating projection of Imp_i onto $\widehat{Bel}_i = \widehat{Pl}_i$

not only a true distance measure but also, regardless of its criticisms, relatively well understood and straightforward to implement. This measure has been proposed in [94] and further discussed in [95].

Given a mass assignment $m(S)$ and a proposition A let $m_A(S)$ be a mass assignment function such that:

$$(4.12) \quad m_A(S) = \begin{cases} \frac{m(S)}{\sum_{S \cap A \neq \emptyset} m(S)}, & \text{if } S \cap A \neq \emptyset \\ 0, & \text{otherwise} \end{cases}$$

Using the definition of a plausibility measure this can also be written as

$$(4.13) \quad m_A(S) = \begin{cases} \frac{m(S)}{Pl(A)}, & \text{if } S \cap A \neq \emptyset \\ 0, & \text{otherwise} \end{cases}$$

This *plausibility conditioned mass assignment* can be understood as the distribution of uncertainty under the assumption that A is fully plausible.

If the mass assignment $m(S)$ is the outcome of a fusion process $m = m_1 \oplus m_2$, the distance $d_J(m_A, m_1)$ is a measure of similarity between the contributing mass function m_1 and the resultant mass function m conditioned under the assumption that A is plausible. It can be interpreted as a measure of the contribution of m_1 to the final belief that A is true.

Thus we propose a contribution metric

$$(4.14) \quad C(m_1, m|A) = 1 - d_J(m_A, m_1)$$




Figure 4.2: Illustration of the triangle formed by the three masses

In the case where we want to compare relative contribution of several masses it is also appropriate to normalize this measure

$$(4.15) \quad \hat{C}(m_1, m|A) = \frac{C(m_1, m|A)}{\sum_{m_i \in M \setminus m_1} C(m_i, m|A)}$$

This measure can be shown to satisfy non-degeneracy and inclusion. Due to the inherent properties of distance measures, it cannot satisfy the vacuous insensitivity property.

Non-degeneracy Follows from triangle inequality ■

For the discussion of inclusion we use a very strong form of information monotonicity in order to reduce the internal conflict within the belief functions. Namely:

Definition 4.3 (d-commitment). Mass m_1 is more d-committed than m_2 ($m_1 \square_d m_2$) if m_2 can be obtained from m_1 via Shafer's discounting

By definition, this is a stronger commitment than specialization and hence pl-commitment or q-commitment.

Satisfaction of this property can be easily proven for $|\Omega| = 2$ and experimentally verified for larger frames of discernment

Inclusion Consider masses m_1 and m_α such that m_α is obtained from m_1 through Shafer's discounting with discounting factor α . Drawing from the geometrical interpretation of belief functions, all the masses $m_1 \oplus m_\alpha$ lie on a straight line between m_1 and $m_1^2 = m_1 \oplus m_1$. Similarly all for all possible values of α , m_α lies on the straight line between m_1 and m_v , where $m_v(\Omega) = 1$. Let \vec{m}_1 , \vec{m}_v and \vec{m}_1^2 denote the position vectors for the corresponding masses.

As such we can consider a triangle $\Delta m_1 m_v m_1^2$. The distance from any point on $m_1 \vec{m}_v$ to any point on $m_1 \vec{m}_1^2$ is always greater than the distance from m_1 to any point on $m_1 \vec{m}_v$ if the angle $\omega = \angle(m_v, m_1 m_1^2) \in [\frac{\pi}{2}, \frac{3\pi}{2}]$. Furthermore let us denote unit vectors \vec{m}_d and \vec{m}_r , corresponding to the directions from m_1 to m_1^2 and to m_v respectively. See Figure 4.2 for illustration For $\omega \in [\frac{\pi}{2}, \frac{3\pi}{2}]$ we need:

$$(4.16) \quad \langle m_d, m_r \rangle \leq 0$$

. Since **Jac** is used as the weighting matrix, we have

$$(4.17) \quad \langle m_1, m_2 \rangle = \sum_{A, B \subseteq \Omega} m_1(A) m_2(B) \frac{|A \cap B|}{|A \cup B|} = \vec{m}_1^T \mathbf{Jac} \vec{m}_2$$

with $\mathbf{Jac}(A, B) = \frac{|A \cap B|}{|A \cup B|}$

Thus we have:

$$(4.18) \quad \begin{aligned} \langle m_d, m_r \rangle &= (\vec{m}_v - \vec{m}_1) \mathbf{Jac} (\vec{m}_1^2 - \vec{m}_1) = \\ &= \begin{bmatrix} -m_1(A) & -m_1(B) & m_1(A) + m_1(B) \end{bmatrix} \begin{bmatrix} 1 & 0 & \frac{1}{2} \\ 0 & 1 & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & 1 \end{bmatrix} \begin{bmatrix} m_1^2(A) - m_1(A) \\ m_1^2(B) - m_1(B) \\ m_1^2(\Omega) - 1 + m_1(A) + m_1(B) \end{bmatrix} \\ &= \frac{1}{2} (m_1(A) - m_1(B)) [m_1(A) - m_1^2(A) + m_1^2(B) - m_1(B)] \\ &+ \frac{1}{2} (m_1(A) + m_1(B)) [m_1^2(\Omega) - m_1(\Omega)] \end{aligned}$$

Thus for $\langle m_d, m_r \rangle \leq 0$ it is sufficient that:

$$(4.19) \quad \frac{1}{2} (m_1(A) - m_1(B)) [m_1(A) - m_1^2(A) + m_1^2(B) - m_1(B)] \leq 0$$

and

$$(4.20) \quad \frac{1}{2} (m_1(A) + m_1(B)) [m_1^2(\Omega) - m_1(\Omega)] \leq 0$$

From the properties of Dempster's rule of combination we know that $m_1^2(\Omega) \geq m_1(\Omega)$, which satisfies Inequality 4.20. Furthermore we know that if $m_1(A) \geq m_1(B)$ then $m_1^2(A) \geq m_1^2(B)$ and $m_1^2(A) - m_1(A) \geq m_1^2(B) - m_1(B)$. Thus if $m_1(A) \geq m_1(B)$ then $m_1(A) - m_1^2(A) + m_1^2(B) - m_1(B) \leq 0$ and if $m_1(A) \leq m_1(B)$ then $m_1(A) - m_1^2(A) + m_1^2(B) - m_1(B) \geq 0$ thus satisfying Inequality 4.19 and by extension satisfying Inequality 4.16.

This result can be directly applied to the case of mass conditioned on the plausibility of A provided that $m_1(A) \geq m_1(B)$ ■

Vacuous sensitivity In the belief functions framework, a vacuous element is a basic belief assignment such that all of the mass is assigned to the universal set $v : m(\Omega) = 1$

For the vacuous insensitivity property to hold,

$$C(v, m|A) = 1 - d_J(m_A, v) = 0 \quad \forall m, A$$

which implies:

$$d_J(m_A, v) = 1 \quad \forall m, A$$

which is not satisfiable for any m, A ■

4.3.3.1 Performance of distance-based contribution measures

In [94] and [95] the properties of this measure were investigated. Some of the results presented in these papers are recalled here.

In [94] the rationale of this proposed metric was evaluated by investigating its suitability for sensitivity analysis.

In order to verify the rationality of the proposed metric, we perform the following simulations

Monte Carlo simulation of fusion of biased and unbiased random BBAs We consider two sources of information providing information in the form of mass assignment functions m_1 and m_2 over the same frame of discernment Ω . Source 1 provides a mass function biased in favour of the proposition A , whereas Source 2 provides a randomly generated unbiased mass assignment function.

Algorithms for generation of mass assignment functions The mass functions are obtained by generating a vector \vec{x} of 2^N random numbers in the interval $[0,1]$. In order to generate the unbiased mass assignment function the values x_i are normalised so that $\sum_{i \in 2^N} x_i = 1$. The i th value of the normalised vector is then associated to the focal elements ω_i so that $m(\omega_i) = x_i$. In order to generate the biased mass assignment function the vector \vec{x} is normalised so that $\sum_{i \in 2^N} x_i = 1 - b$, where b is a bias coefficient. The elements of this vector are then associated to the focal elements ω_i such that $m(\omega_i) = x_i$ $i \neq 1$ and $m(A) = x_1 + b$.

Experimental setting Two sets of experiments are performed

- The degree of reliability r of S1 is varied and S2 is assumed to be fully reliable ($r = 1$). This is performed for discrete but varying values of the bias factor $b = 0.4$ and $b = 0.7$ and the size of the frame of discernment $N = 3$ and $N = 5$
- Both S1 and S2 are assumed to be reliable (discounting factor $d_1 = d_2 = 0$) and the bias coefficient is varied for frame of discernment $N = 3$ and $N = 5$

In each case, the mass functions m_1 and m_2 are combined using the Dempster-Shafer rule of combination, and the plausibility conditioned distance is evaluated in order to determine which of the two sources of information has a more significant contribution to the final belief in A . For each experimental setting, 5000 Monte Carlo realizations are performed. The performance is evaluated by comparing the rate of identification of m_1 as the main source of contribution (the % of realizations where the plausibility conditioned distance from m_1 is smaller than that from m_2) to the identification rate of the singleton A using the maximum of pignistic probability (BetP) and maximum of plausibility (PI) criteria (see Section 2.2.4.2).

Impact of varying discounting factor In this section, we study the impact of varying the discounting factor. The simulation described above is performed for reliability values between 0 and 0.8 for $N = 3$ and between 0 and 0.5 for $N = 5$ Figures 4.3 and 4.4 show the results for

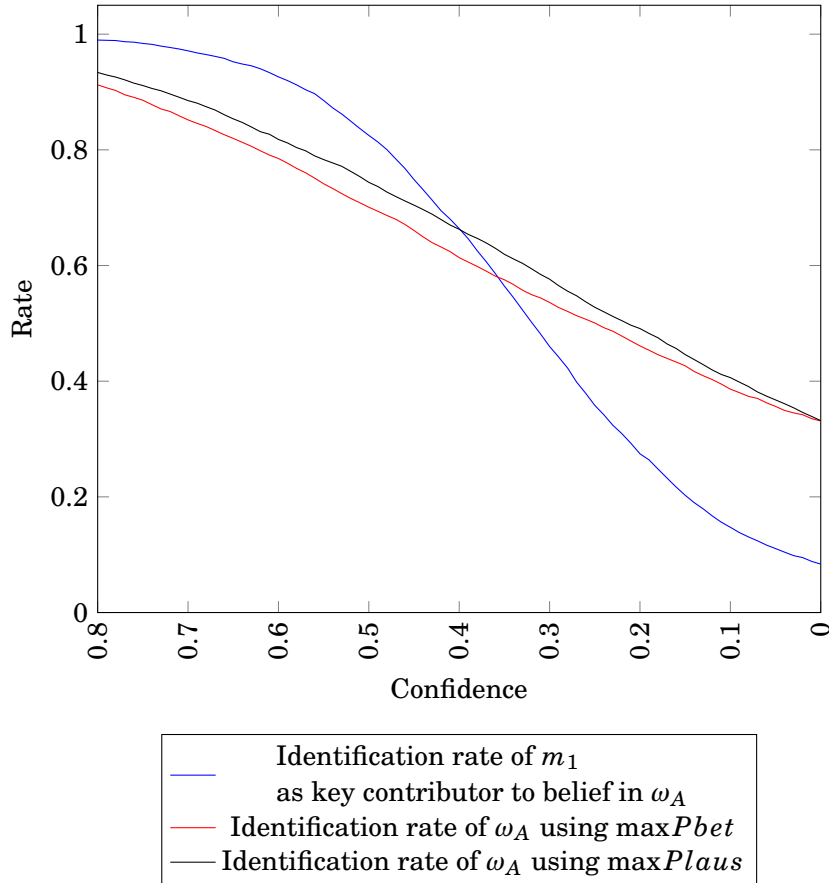


Figure 4.3: Contribution of m_1 to belief in ω_1 using $N = 3$ and $b = 0.4$

bias coefficient $b = 0.4$ and frame of discernment $N = 3$ and $N = 5$ respectively. The identification rate of m_1 as the key contributor to $Bel(A)$ is defined as the proportion of Monte Carlo realisations where the plausibility conditioned Jousselme's distance $d_J(m_A(S), m_1(S))$ is smaller than $d_J(m_A(S), m_2(S))$. Similarly the identification rate of A is the proportion of realizations where A is deemed the most likely proposition based on the maximum of pignistic probability (P_{bet}) or the maximum of plausibility (Pl) criteria. It can be seen that unless the reliability value r is very small, the rate of contribution of m_1 to A tracks the rate of identification of A relatively closely. It is important to note that for very small reliability $m_1(\Omega) \rightarrow 1$ as m_1 becomes a vacuous belief function, whereas the identification rate of A tends to $\frac{1}{N}$.

Figures 4.5 and 4.6 show the results for bias coefficient $b = 0.7$ and frame of discernment $N = 3$ and $N = 5$ respectively. Here we can observe that as the discounting factor is increased, the rate of identification of A decays slower than the rate of contribution of m_1 to A . This phenomenon can likely be explained by the dictatorial property of Dempster-Shafer fusion, where a more concentrated BBA has a larger impact on the fusion result.

It is important to note that in all of these cases, the identification rate of m_1 as the main contributor to A tracks the identification rate of A relatively closely when the reliability of m_1 is large. These experiments show that the plausibility conditioned distance on A does correlate with




Figure 4.4: Contribution of m_1 to belief in ω_1 using $N = 5$ and $b = 0.4$

the identification rate of A .

Impact of varying bias In the second setup, we study the impact of varying the bias coefficient b , while the discounting factor, or reliability, is kept constant, and both sources of information are assumed to be fully reliable. 5000 Monte Carlo realizations are performed for each value of b . Figures 4.7 and 4.8 show the results for the frame of discernment $N = 3$ and $N = 5$ respectively. Since the reliability of m_1 , $r_1 = 1$ and is kept constant, as $b \rightarrow 0$ m_1 becomes a randomly generated unbiased BBA and hence the identification rate of itself as the key contributor to the belief in A tends to 0.5. As per the previous experiment, the identification rate of A tends to $\frac{1}{N}$.

It can be observed that for both values of N , despite the different behaviour as $b \rightarrow 0$, the identification rate of m_1 as the key contributor to A tracks the identification rate of A rather closely.

Sensitivity analysis for the fusion of two BBAs In this section, we study the applicability of plausibility conditioned distance as a tool for sensitivity analysis of the final fusion outcome with respect to the discounting factor. Here we consider a fusion of two semi-reliable sources of information S_1 and S_2 generating mass assignments m_1 and m_2 listed in Table 4.1 with a




Figure 4.5: Contribution of m_1 to belief in ω_1 using $N = 3$ and $b = 0.7$

relatively small degree of conflict ($K = 0.192$).

The contribution of a source of information to the final degree of belief in a proposition is quantified using the dual of plausibility conditioned distance

$$(4.21) \quad c_i = 1 - d_J(m_A, m_i)$$

Fusion of BBAs m_1 and m_2 was performed using the Dempster-Shafer rule of combination, and c_i was computed for each of the mass assignments for reliability values in the interval $[0, 1]$ without recomputation of the fusion result. Subsequently, in order to verify the actual contribution of each of the BBAs to the fusion outcome, the fusion process was performed again while keeping the reliability of one of the mass assignments constant and varying the other in the interval $[0, 1]$. The resulting joint BBAs were then used to compute the values of $Bel(A)$. The results of this process are shown in Figure 4.9.

From Figure 4.9 we can see that the gradient of the contribution of m_1 to A around its initial reliability is significantly larger than that of the contribution of m_2 to A . Similarly the rate of change of $Bel(A)$ as the reliability of m_1 , r_1 is varied is much larger than when r_2 is varied, showing its sensitivity to the variation of r_1 . It is important to note that the computation of the contribution metric does not require recomputation of the fusion result and, as such, is much less




Figure 4.6: Contribution of m_1 to belief in ω_1 using $N = 5$ and $b = 0.7$

computationally intensive, which can have a significant impact when the frame of discernment is large or when a large number of sources of information are fused.

Table 4.1: Mass assignments and initial reliability of sources S_1 and S_2

	S_1	S_2
$m(A)$	1	0
$m(B)$	0	0.4
$m(B,C)$	0	0.4
$m(A,B,C)$	0	0.2
Initial reliability r	0.3	0.8

These results have shown that this distance-based contribution measure can be adequately used both for the identification of the key contributing source as well as for sensitivity analysis. In [95] the behaviour of this measure was further investigated in light of the decision quality indicator $q(\hat{\omega})$ as defined in Equation 2.34

Behaviour of the contribution measure in light of the quality indicator The quality indicator can be used to assess the performance of a contribution metric as follows. Given a set of

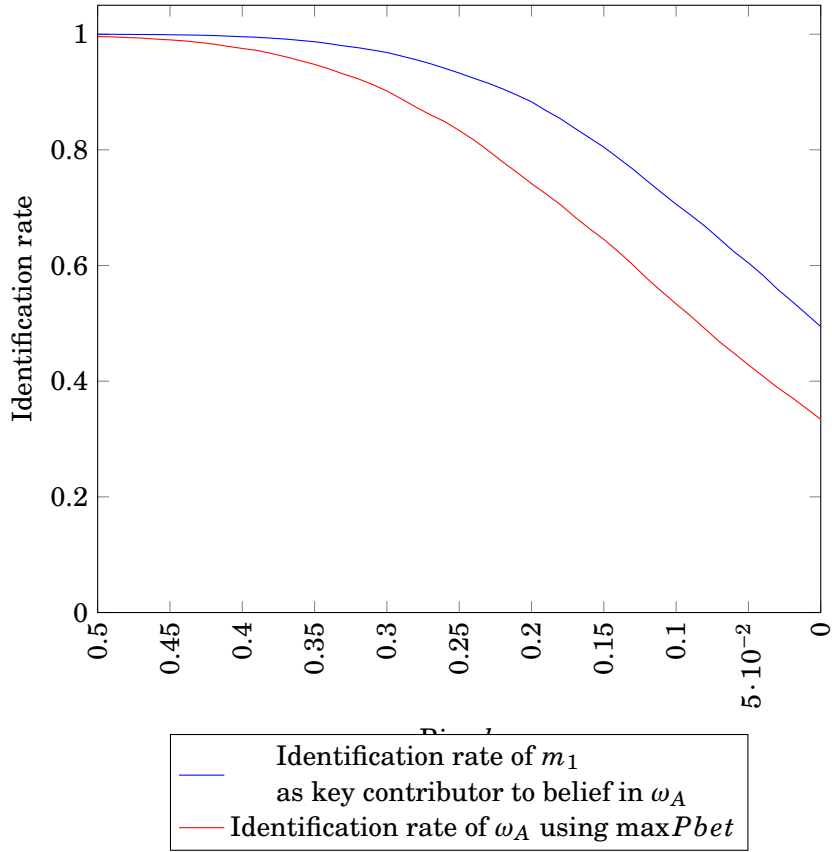


Figure 4.7: Contribution of m_1 to belief in ω_1 for different values of b using $r_1 = r_2 = 1$ and $N = 3$

contributing masses, the relative contribution of each of them can be computed. This particular BBA can then be removed from the set, the joint mass can be re-evaluated, and thus, the change in the quality of a decision can be computed. The correlation between the relative contribution of the mass and its effect on the quality of a decision can be used as a performance metric.

4.3.4 Generation of BBAs

In order to generate sets of realistic basic belief assignment, the following method was used.

1. The perturbation method

Given a seed BBA and perturbation parameter $p \in [0, 1]$ each focal element of the seed BBA is discounted by p , and the mass is randomly redistributed across all the other focal sets. This produces a family of BBAs with relatively low conflict and identical focal sets (see Algorithm 1)

2. The diverging method

Given a seed BBA and perturbation parameter $p \in [0, \frac{1}{n}]$, in order to generate n BBAs, each focal element of the seed BBA is discounted by $i \times p$ and the mass is randomly redistributed




Figure 4.8: Contribution of m_1 to belief in ω_1 for different values of b using $r_1 = r_2 = 1$ and $N = 5$

across all of its supersets, where i denotes the iteration of the algorithm. This produces a family of BBAs with decreasing informativeness (see Algorithm 2)

Conjunctive combination The behaviour of the contribution measure under conjunctive combination is investigated using various seed BBAs and perturbation coefficients in both generation methods described in the previous section. The quality indicator $q(\hat{A})$ (see Equation 2.34) is used as a benchmark. This means, that aside of investigating the impact of the contribution $\hat{C}(m_1, m|A)$ on $\Delta Q(A)$, we also look at the correlation between normalized $q_{\hat{m}_1}(A)$ and $\Delta Q(A)$, where $q_{\hat{m}_1}A = \frac{q_{m_1}(A)}{\sum_{m_i \in M \setminus m_1} q_{m_i}(A)}$ and m_1 denotes the BBA being removed from the set. This is justified as the quality indicator is calculated using only the removed BBA and as such exhibits similar behaviour to traditional measures such as the belief and plausibility measure, whilst also scaling similarly to the contribution measure. For the sake of clarity, the quality indicator is referred to as *benchmark* when used in this context throughout the rest of this chapter.

Throughout the simulations, the contribution measure exhibited a significant correlation with the change in quality indicator in almost all the cases whilst also usually outperforming the benchmark. However, the type of seed BBA used, as well as the method of generating the BBA family, would significantly affect both strength and type of the correlation.

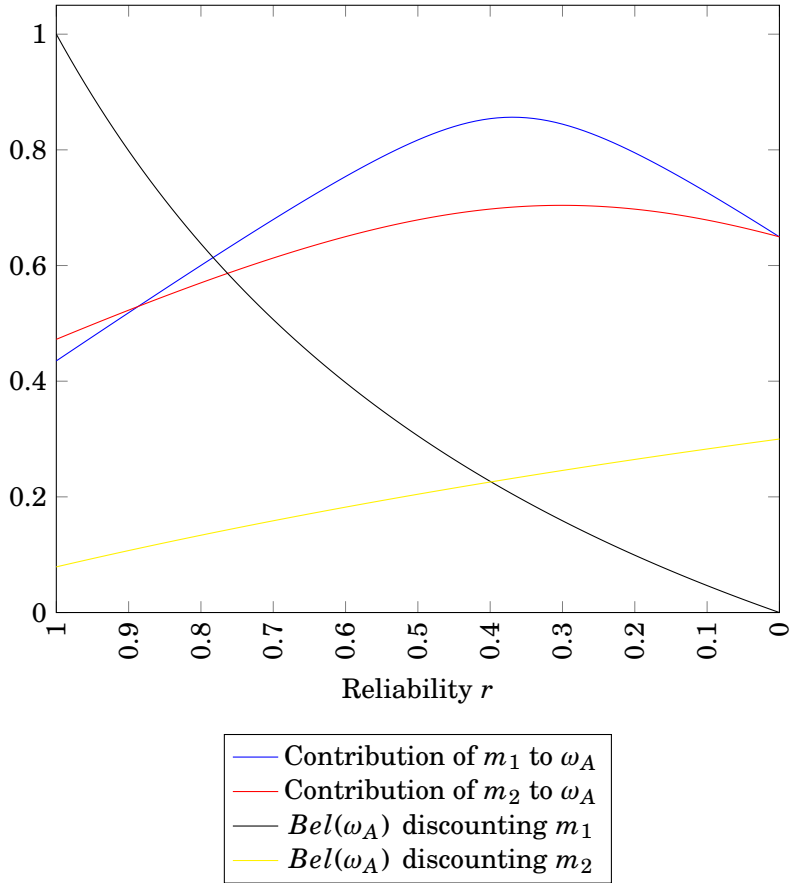


Figure 4.9: Sensitivity analysis of the final fusion outcome with respect to the reliability of information sources

```

input      :seed BBA  $m_{in}$ 
output    :perturbed BBA  $m_{out}$ 
parameter:  $p$  - perturbation rate
 $l = size(m_{in});$ 
for  $i = 1:l$  do
     $m_{temp1}(i) = m_{in}(i) \times (1 - p);$ 
    Generate an array  $r[l]$  of random numbers summing up to 1;
    for  $j = 1:l$  do
         $m_{temp2}(j) = m_{in}(i) \times p \times r(j)$ 
    end
end
 $m_{out} = m_{temp1} + m_{temp2}$ 

```

Algorithm 1: Generation of perturbed BBAs

4.3.4.1 Strongly biased seed BBA

In this section, the behaviour of a strongly focused seed BBA is discussed. An example of such a seed is displayed in Table 4.2, where Ω denotes the universal set regardless of the size of the frame of discernment. Both linear and logarithmic models are fitted to the data, and the coefficient

```

input      : seed BBA  $m_{in}$ 
output    : perturbed BBAs  $m_{out}[n]$ 
parameter:  $p$  - perturbation rate,  $n$  - number of BBAs
 $l = size(m_{in})$ ;
for  $k = 1: n$  do
  for  $i = 1: l$  do
     $m_{temp_1}(i) = m_{in}(i) \times (1 - p \times k)$ ;
     $o = \text{number of supersets of } m_{in}(i)$  ;
    Generate an array  $r[o]$  of random numbers summing up to 1;
    for  $j = 1: k$  do
      if  $m_{temp_1}(j)$  is superset of  $m_{in}(i)$  then
         $m_{temp_2}(j) = m_{in}(i) \times p \times k \times r(j)$ 
      end
    end
     $m_{out}[k] = m_{temp_1} + m_{temp_2}$ 
  end
end

```

Algorithm 2: Generation of diverging BBAs

of determination R^2 is used to assess performance. A typical response is shown in Fig 4.10, where such a BBA with frame of discernment cardinality $n = 4$ (i.e. $\Omega = \{A, B, C, D\}$) is used to generate a family of 12 BBAs under divergent generation scheme with perturbation rate = 0.06. The goodness of fit measures is displayed in Table 4.3.

Table 4.2: Highly focused seed basic belief assignment used

A	B	Ω
0.1	0.8	0.1

Table 4.3: Coefficients of determination R^2 (focused BBA, divergent mode)

Metric	Linear	Logarithmic
Contribution $\hat{C}(m_1, m A)$	0.859	0.94248
Benchmark	0.8375	0.93256

There are two key points established by these results. First of all, there exists a strong logarithmic relationship between the contribution measure and the quality indicator under this configuration. This can be explained by the general Dempster-Shafer behaviour - the more concentrated BBAs can disproportionately dominate the fusion process, whilst the impact of the more divergent BBAs is much smaller. This is in line with the well understood dictatorial property of Dempster's rule of combination.

In addition to that, our contribution measure has a slight advantage above the benchmark, albeit it is rather small.




Figure 4.10: Change in quality of decision when a single body of evidence is removed under Dempster-Shafer rule of combination combination, using highly focused seed BBA, divergent mode

Figures 4.11, 4.12 and Table 4.4 display the corresponding information for simulations performed with Algorithm 1. Here the size of the frame of discernment is irrelevant as this algorithm does not populate focal sets, which are initially set to be empty. With a low perturbation rate, the correlation is nearly linear, as every BBA has a very small impact on the combination result by itself; as the perturbation rate is increased, the correlation becomes logarithmic as less informative and more conflicting BBAs are introduced into the fusion process. In all of these cases, $\hat{C}(m_1, m|A)$ outperforms the benchmark.

Table 4.4: Coefficients of determination R^2 (focused BBA, perturbation mode)

Metric	Perturbation $p = 0.2$		$p = 0.9$	
	Linear	Log	Linear	Log
Contribution $\hat{C}(m_1, m A)$	0.983	0.991	0.465	0.971
Benchmark	0.982	0.992	0.438	0.962

4.3.4.2 Seed BBA with singletons only

The second type of seed BBA to be investigated is a BBA with singletons only, i.e. a probability distribution (see Table 4.5). Note that if the divergent mode of generation is used, the non-singleton focal sets are populated; otherwise, the entire family of BBAs generated consists of probability assignments. As per Table 4.6, both generation methods result in significant linear correlation and a slight improvement over the benchmark.

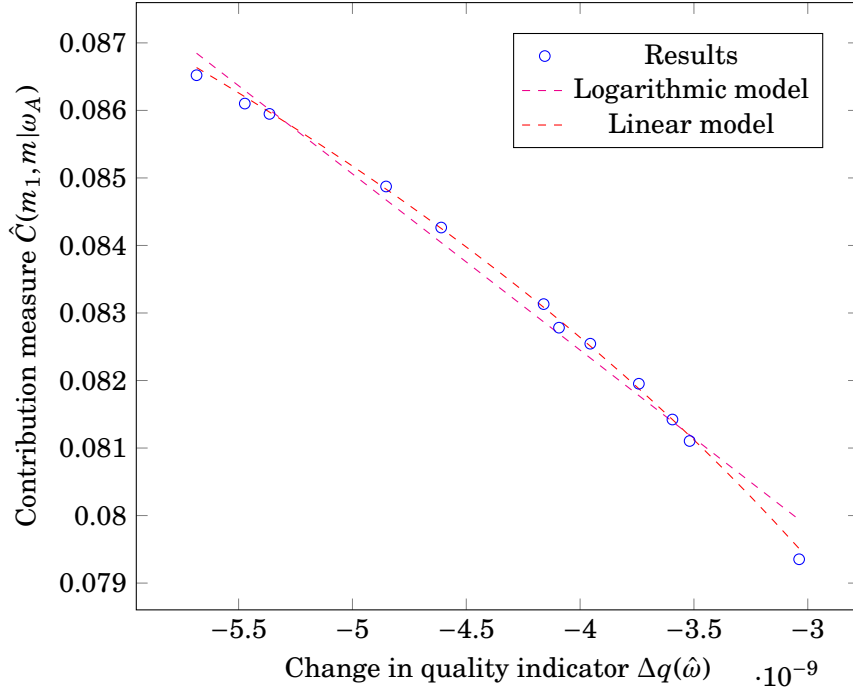


Figure 4.11: Change in quality of decision when a single body of evidence is removed under Dempster-Shafer rule of combination, using highly focused seed BBA, high perturbation

Table 4.5: Basic probability assignment used as the seed BBA

A	B	C
0.2	0.3	0.4

Table 4.6: Coefficients of correlation r and determination R^2

BBA family type	Divergent		Perturbed	
	r	R^2	r	R^2
Contribution $\hat{C}(m_1, m A)$	-0.941	0.886	-0.984	0.969
Benchmark	-0.928	0.862	-0.980	0.959

4.3.5 Randomly generated seed BBAs

In this section, we investigate the behaviour of the contribution measure in Monte Carlo simulations using randomly generated seed BBAs. The coefficient of linear correlation r is used as a performance metric in this case, where $r = -1$ corresponds to a strong correlation as per the previous section. Larger (less negative) values of the correlation coefficient imply a weaker correlation and hence worse performance of the contribution metric as a predictor for the change in the quality of the decision. Finally, a positive correlation metric means that removing a BBA with high relative contribution increases the decision quality - a counter-intuitive behaviour that should not be taking place.

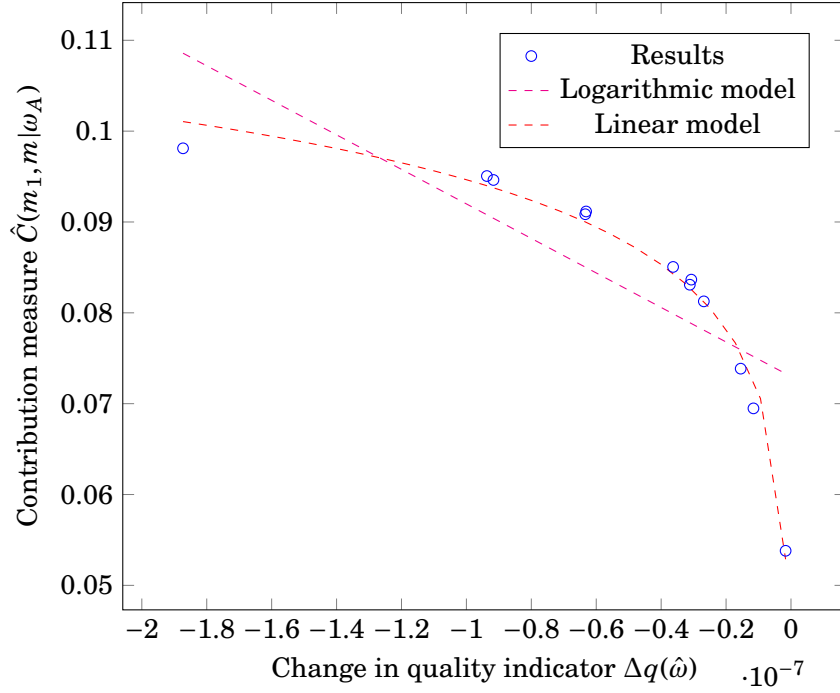


Figure 4.12: Change in quality of decision when a single body of evidence is removed under Dempster-Shafer rule of combination combination, using highly focused seed BBA, low perturbation

Table 4.7: Mean μ and standard deviation σ of r for the Monte Carlo simulation with randomly generated BBAs

	μ	σ
$\hat{C}(m_1, m A)$	-0.8652	0.1425
Benchmark	-0.4532	0.3511

Figure 4.13 - shows an example of a histogram generated with this kind of Monte Carlo simulation. 500 initial seed BBAs were generated by distributing the entire mass among 10 randomly selected focal sets in the frame of discernment of cardinality $n = 6$. Families of 6 BBAs were generated from each seed, using the perturbation generation method with $p = 0.2$. Table 4.7 shows the means and standard deviations for the correlation coefficients obtained. It can be seen that in this particular configuration, $\hat{C}(m_1, m|A)$ strongly outperforms the benchmark. However, it is important to note that these results can vary very strongly with the properties of the seed BBA.

It was also noted that there is a very strong correlation ($r > |0.95|$) between the computed coefficient of correlation r and the coefficient of determination R^2 . Thus, it can be argued that the counter-intuitive behaviour seen on the right-hand side of the histogram can be associated with a poor fit of the linear model rather than unpredicted behaviour.

These results show that in the majority of cases, there is a strong correlation between a BBA's contribution level and its impact on the quality indicator. In addition to that, in many cases, the contribution measure proposed outperforms the benchmark. The importance of the benchmark




Figure 4.13: Monte Carlo histogram using $\hat{C}(m_1, m|A)$ (top) and benchmark (bottom)

chosen stems from it being based on the contributing BBA only, similar to measures such as belief, plausibility and probability transformations. The contribution measure behaves differently, as its value is driven by the relationship between the contributing BBA and the target BBA.

It can be argued that the performance difference between the contribution measure proposed and the benchmark was relatively small. It can be, however, argued that this is due to two reasons. First of all, Jousselme's distance may not be the best dissimilarity measure for this application; its occasionally counterintuitive behaviour may negatively impact the results. Additionally, all the tests have been performed using Dempster's rule of combination. Dempster's rule has many weaknesses, and whilst experts' opinions differ on its overall usefulness, there exists a general consensus that it cannot be applied blindly to any combination of BBAs. This caveat was not properly followed whilst performing the simulations presented in this chapter, and it is very likely that it may have affected the results, in particular in the context of Monte-Carlo simulations described in subsection 4.3.5.

4.4 Information fusion as cooperative game theory - Shapley value

The problem of information fusion can be considered from the perspective of a cooperative game. In the most simple approach, the sources can be considered to form a single coalition. We consider a set \mathcal{N} of N players who cooperate. A game can be given by the characteristic function $v : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ with $v(\emptyset) = 0$. The payoff achieved by a coalition $S \subseteq \mathcal{N}$ is denoted by $v(S)$.

The Shapley value [96] is a game theoretic method of distributing the gains to the individual players or assessing each player's contribution to the overall payoff. It is given by

$$(4.22) \quad \phi_i(v) = \sum_{S \subseteq \mathcal{N} \setminus \{i\}} \frac{|S|!(N - |S| - 1)!}{N!} (v(S \cup \{i\}) - v(S)).$$

It can be interpreted as the average marginal contribution of each player to every possible coalition over each possible permutation.

The importance of Shapley value stems from the fact that it possesses a number of significant properties, some of which are of particular relevance to this problem.

Efficiency

$$\sum_i \phi_i(v) = v(\mathcal{N})$$

Symmetry If i and j are equivalent players ($v(S \cup \{i\}) = v(S \cup \{j\}) \forall S$) then:

$$\phi_i(v) = \phi_j(v)$$

.

Linearity

$$\phi_i(v + w) = \phi_i(v) + \phi_i(w)$$

Null player If player i is null (it contributes nothing to any coalition, i.e. $v(S \cup \{i\}) = v(S) \forall S$) then:

$$\phi_i(v) = 0$$

.

We may consider the individual sources contributing to an information fusion problem to be agents in a cooperative game. Several possible payoff functions may be used, such as the degree of belief in a hypothesis Bel , plausibility Pl or the pignistic probability P_{Bet} . Using Shapley value with one of these payoff function results in a provenance measure satisfying all four properties.

Consider the set of n masses \mathcal{M} . Using the belief measure we have the following provenance measure:

$$(4.23) \quad \phi_i(Bel) = \sum_{S \subseteq \mathcal{M} \setminus \{i\}} \frac{|S|!(N - |S| - 1)!}{N!} (Bel(\oplus S \cup \{i\}) - Bel(\oplus S))$$

where \oplus denotes the Dempster's rule of combination.

Non-degeneracy Follows directly from symmetry ■

Inclusion Using the same commitment definition from Definition 4.3 we consider masses m_1 and m_α such that m_α is obtained from m_1 through Shafer's discounting with discounting factor α

From the definition of Shafer's discounting and the belief measure, we know that for any A $Bel_1(A) \geq Bel_\alpha(A)$.

From Equation 4.23, $Shap_{Bel(A)}(m_1) \geq Shap_{Bel(A)}(m_\alpha)$ if $Bel_{m_1 \oplus m} \geq Bel_{m_\alpha \oplus m}$. This is always true when there is zero conflict between m and m_1 . ■

Vacuous sensitivity Directly follows from the null player property, provided that the payoff function is selected such that the null player property is satisfied (i.e. $v(m_\Omega) = 0$). ■

Usage of Shapley value for explanations in an evidential network The proposed Shapley value-based contribution measure has a significant advantage over the distance-based approaches and simple contribution measures - it can be applied directly in evidential networks. Some examples of evidential networks were discussed in the previous chapter. To recall, it is a graphical model consisting of variables and valuations (mass functions), which can either represent the relationships between variables or observations. When considering explanation generation in an evidential network, we may be interested in the contribution of observation to inference on a particular variable. This can be very difficult to do using distance-based measures or simple contribution measures. Consider the following valuation network, with variables A, B, C defined on $\Omega_A = \{true, false\}$, $\Omega_B = \{true, false\}$ and $\Omega_C = \{true, false\}$ respectively




Figure 4.14: Valuation network

Here m_1 and m_2 are observations on B and C respectively and v denotes the logical relationship between A , B and C (with the example below). Let A be the variable of interest, and consider that we are interested in contribution of m_1 to $A = true$. Neither the distance-based contribution measure, nor the simple contribution measures such as belief or plausibility measures can be applied directly, since the frame of discernment of m_1 is different to that of variable A . One intuitive workaround is to make inference on A using m_1 as the sole observation (i.e. excluding m_2):

$$m_{1A} = \left(m_1^{\uparrow A \times B \times C} \oplus v \right)^{\downarrow A}$$

And consider that mass to be the contributing mass. This will work in cases when B and C contribute independently to A , i.e. when it is possible to split $v^{A \times B \times C}$ into $v_1^{A \times B}$ and $v_2^{A \times C}$. However,

this is not the case if we consider a mass ν , which cannot be separated in such a manner. For example, the logical relation:

$$A = true \iff B = true \wedge C = true$$

induces the mass ν :

$$\nu_{A \iff B \wedge C} \left(\left\{ \begin{array}{l} A = true \times B = true \times C = true \\ A = false \times B = false \times C = \Omega_C \\ A = false \times B = \Omega_B \times C = false \end{array} \right\} \right) = 1$$

which means that inference on A obtained using only m_1 or only m_2 will always have zero mass on $A = true$ - even if both m_1 and m_2 assert both B and C to be true. Thus even with $m_1(B = true) = 1$ and $m_2(C = true) = 1$ we still have $m_{1A}(\Omega_A) = 1$.

Going back to the properties which need to be satisfied by a contribution measure this is not necessarily a problem as it does not contradict the preference property - both masses would have the same degree of contribution, as they are both necessary for the decision to be made. Regardless that does not always seem intuitively correct. In particular let us consider the scenario when we have an additional observation on C, m_3 such that $m_3(C = false) = 0.4$, $m_3(C = \Omega_B) = 0.6$. Furthermore let the two original masses be discounted by different factors - 0.8 and 0.6 respectively, such that:

$$\begin{array}{ll} m_1(B = true) = 0.8 & m_2(C = true) = 0.6 \\ m_1(\Omega_B) = 0.2 & m_2(\Omega_C) = 0.4 \end{array}$$

It is now reasonable that the contribution of the two masses should be different. In particular, m_2 is in conflict with m_3 , and its inclusion increases the plausibility of A. We can compute the Shapley value with $1 - Pl(A = true)$ as the payoff. This value is used instead of taking plausibility directly in order to satisfy the condition that the null player (a vacuous mass function) receives zero payoffs. As such, we consider the cost to maximum plausibility induced by the mass rather than the plausibility itself. With this in mind we obtain: $Shap_{Pl(A)}(m_1) = 0$, $Shap_{Pl(A)}(m_2) = 0.0633$, $Shap_{Pl(A)}(m_3) = -0.1467$, identifying the additional contribution of m_2 associated with conflict redistribution.

Another issue with such evidential networks is that it would be reasonable to assume that the contribution of the more informative mass (m_1) should be greater than that of the other, even though they are both necessary. One argument is that either of them can be considered to be a set of constraints on the possible values of C, and as such m_1 allows a greater maximum value of $m(C) = true$. Unfortunately, in this case, the contribution of the two masses is the same. Finding an appropriate method for the generation of such explanations in expert systems is still an area that needs to be researched further.

4.5 Contribution measures - the preference property

The preference property has been proposed as a desirable property of a contribution measure. It makes sense intuitively as a natural way of identifying the most significant contributor to the

decision - if removal of one body of evidence causes the decision to change and another does not, it implies that the former must have had a more significant impact on the decision than latter. However, it can be shown that for all discussed contribution measures, there exist sets of basic belief assignments for which this property does not hold.

In this section, we investigate two things. First of all, we compare the different contribution metrics with combination and decision methods to assess the proportion of cases where the preference property holds. Second of all, we investigate the properties of basic belief assignments, which may result in the preference property not being satisfied.

4.5.1 Dempster's rule of combination and maximum of belief decision making

First of all, let us narrow our considerations to the most common case - utilizing Dempster's rule of combination and the maximum of belief rule for decision making. Using the belief measure as a contribution metric should intuitively satisfy the preference property; however, in subsection 4.3.1, it was shown that there exists a set of belief functions for which this is not the case. Let us investigate these behaviours in more detail through Monte-Carlo simulations.

The procedure is performed as per Algorithm 3.

```

output      : Results log
parameter :  $n$  - number of MC iterations,  $p$  - size of frame of discernment, decision
              making formalism  $D$ , method of combination, set of contribution measures  $C$ 
 $l = \text{size}(m_{in})$ ;
for  $i = 1 : n$  do
  for  $k = 1 : 2$  do
    Generate a random number  $s$  between 1 and  $2^p$ 
    Randomly generate a BBA  $m_k$  with  $s$  focal elements
  end
  Combine  $m_1$  and  $m_2$  using the selected method of combination. Let  $m_J$  denote the
  result
  Using the selected decision making formalism  $D$ , make decisions on  $m_1, m_2, m_J$ :
   $d_1 = D(m_1), d_2 = D(m_2), d_J = D(m_J)$ 
  if ( $d_1 == d_2$  OR ( $d_1 \neq d_J$  AND  $d_2 \neq d_J$ )) then
    Go back to BBA generation step;
  else
    Note the key contributor  $m_A$  such that  $D(m_A) = D_J$ . Denote the other mass  $m_B$ 
    for every  $c$  in  $C$  do
      Compute  $c_A = c(m_A, m_J, d_J)$  and  $c_B = c(m_B, m_J, d_J)$ 
      if  $c_A \geq c_B$  then
        Log preference property satisfied for  $c$ ;
      else
        Log preference property not satisfied for  $c$ ;
      end
    end
  end
end

```

Algorithm 3: Monte-Carlo simulation for computing the proportion of cases where preference property does not hold

First of all, let us look at the relative accuracy of the different contribution measures (here the Shapley value using belief and plausibility, the distance-based contribution measure C_d (Equation 4.14), the derivative-based contribution measure d_i (Equation 4.11) and belief, plausibility and pignistic probability values themselves). It seems reasonable that as the maximum of belief is used as the contribution measure, belief-based Shapley value and the belief measure itself will be most accurate. This expectation is supported by the results. Running the simulation at 800000 iterations, we obtain the failure rates as shown in Table 4.8.

Shap(Bel)	Shap(Pl)	Distance-based	Derivative-based	Bel	Pl	BetP
32934	250084	51931	61305	32934	250044	71381
4.11%	31.3%	6.4%	7.6%	4.11%	31.3%	8.9%

Table 4.8: Proportion of cases where the preference property is not satisfied for the different contribution measures using Dempster's rule of combination and maximum of belief decision

Unsurprisingly belief-based Shapley value, as well as the belief measure itself, outperform the other measures in this scenario, closely followed by the distance-based contribution measure. Pignistic probability performance is a bit worse, and plausibility-based approaches perform very poorly.

The second question we need to answer is: what properties of the basic belief assignments result in the preference property not being satisfied? The discussion in section 4.3.1 implies that one possible cause of this behaviour can be considerable overlap between the two masses or inclusion of one within the other, relative similarity of the two belief masses, small but non-zero conflict and low specificity of one of the two masses. In the following section, we study the behaviour of the distance-based contribution measure in detail in comparison to the *Bel* measure.

The similarity of the belief mass functions can be measured using the geometrical distance d_J (Equation 2.28). The conflict between the mass functions is typically defined as the mass on the empty set of the conjunctive combination of the masses (Equation 2.18). The specificity of a belief mass, which characterizes its relative precision, is defined as follows [97]:

$$(4.24) \quad Spec(m) = \sum_{A \subseteq \Omega} \frac{m(A)}{|A|}$$

Finally, we require an approach to the assessment of relative inclusion of one belief assignment within another. In order to avoid attempting to find the "optimal" method of measuring this property, the following simple subsumption measure will be used:

$$(4.25) \quad S_{A \subseteq B} = \sum_{A \subseteq B} m_1(A)m_2(B)$$

Now in order to analyze the behaviour correctly, the data generating algorithm (Algorithm 3) needs to be modified slightly. A set of features to be logged is included as a parameter: here we use the distance $d_J(m_1, m_2)$, the conflict $m_1 \oplus m_2(\emptyset)$, the inclusion of 1 in 2 and vice-versa ($S(m_1, m_2)$ and $S(m_2, m_1)$) and the specificities $Spec(m_1)$ and $Spec(m_2)$. Furthermore, in order to

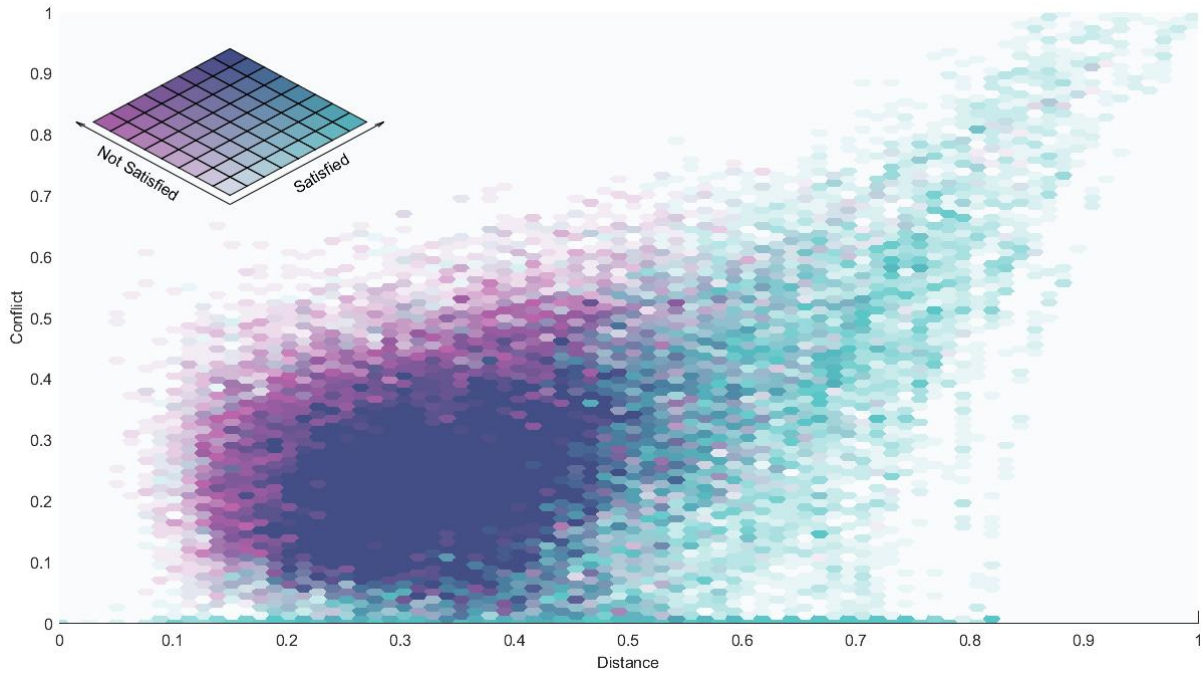


Figure 4.15: Bivariate hexagonal binning showing density of bba pairs satisfying and not satisfying the preference property for *Bel* as a function of the distance and the conflict between them, using maximum of belief decision making and Dempster's combination

keep the dataset balanced, rather than performing a Monte-Carlo simulation for a specific number of iterations, a target value n is passed, and the simulation runs until it has logged n positive (where the preference property was satisfied) and negative (where the preference property was not satisfied) instances.

As per the previous section, the best "performance", where performance refers to the rate of satisfaction of the preference property, is obtained using either the *Bel*-based Shapley value or the *Bel* function itself (as with maximum of belief decision making their performance is identical). A matrix of scatter plots showing the distribution of cases satisfying and not satisfying the preference property is shown in Figure 4.17. This can be compared with similar behaviour but using the distance-based measure as the contribution metric as per Figure 4.18.

The distance between the two BBAs and the conflict between them are some of the most significant predictors of the behaviour. Unfortunately, it is difficult to see it on a scatter plot due to the density and downsampling of the results would make the outliers more difficult to see. In order to circumvent this, bivariate honeycomb graphs are used, a technique drawing from hexagonal binning and choropleth mapping. The Figures 4.15 and 4.16 respectively show this relationship for the contribution measures *Bel* and distance-based C_d .

Conflict and distance histograms in Figures 4.17 and 4.17 can be helpful in analysis of this. In the case of *Bel* being used as a contribution measure, it can be clearly seen that the preference property is almost always satisfied when there is no conflict between the two masses and when the distance between the masses is large. In Figure 4.15, an area where the preference property is unlikely to be satisfied can be seen when the distance between the masses is small (~ 0.2)

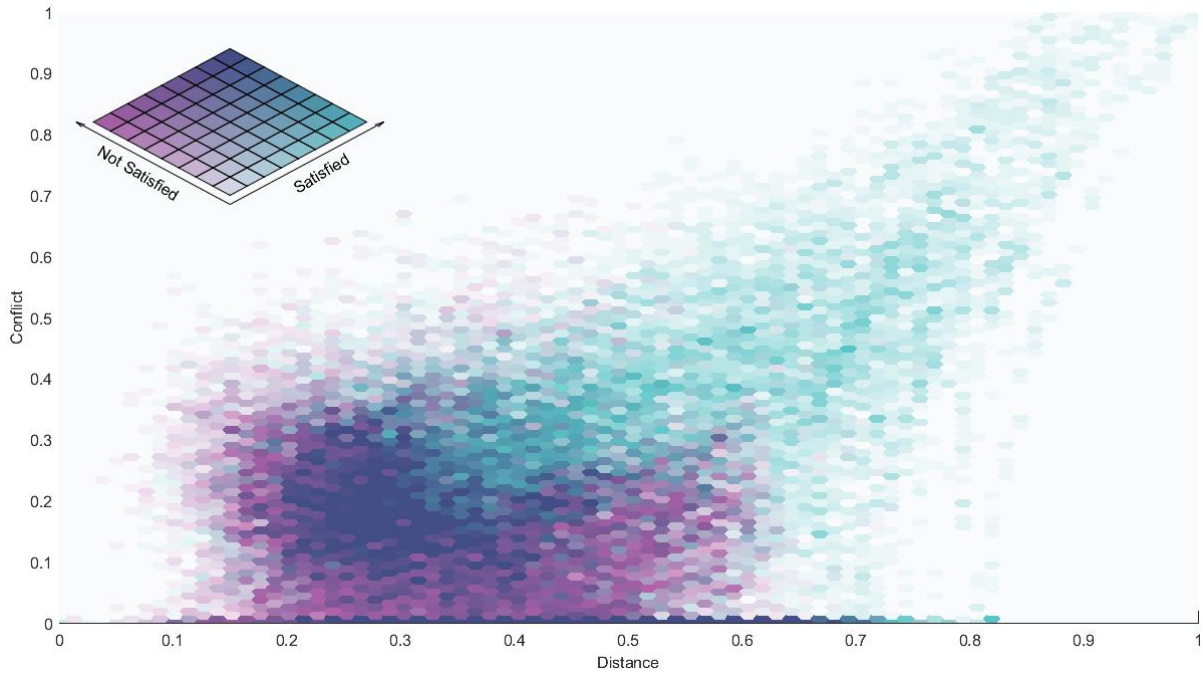


Figure 4.16: Bivariate hexagonal binning showing density of bba pairs satisfying and not satisfying the preference property for distance-based measure as a function of the distance and the conflict between them, using maximum of belief decision making and Dempster’s combination

regardless of the conflict (as long as it is non-zero) and when distance is in the $\sim 0.2 - \sim 0.5$ range and conflict is in the $\sim 0.4 - \sim 0.5$ range. When both distance and conflict are medium and similar $\sim 0.2 - \sim 0.4$, the satisfaction of the preference property is not explained by these two parameters.

A closer look at the data points within this area does not reveal much more. However, it can be seen that in this case, the preference property is less likely to be satisfied than not (satisfied in 35% of cases). A close look at the histogram of the distance values reveals that there remains a correlation between the satisfaction rate and the distance, with the satisfaction rate decreasing with distance, but, again, it is not sufficient to determine whether a BBA couple satisfies this property based on these parameters only. Finally, the histogram in Figure 4.19 shows that the relative inclusion of one BBA in the other can be used to predict with some degree of accuracy whether the preference property is satisfied regardless of the conflict and distance values.

It is interesting to note how different this behaviour is when C_d is used as the contribution measure (Figures 4.16 and 4.18). In particular, it is noteworthy that a large proportion of cases in which it is impossible to tell by distance and conflict only whether or not the preference property will be satisfied for Bel lies within the region which is very likely to be satisfied for C_d (conflict greater than ~ 0.25 and distance between ~ 0.35 and ~ 0.5). On the other hand, while the cases with low conflict virtually always satisfy the preference property for Bel , this is not the case for C_d . In fact, when the cases without conflict are disregarded, C_d performance is much closer to that of Bel as a contribution measure, even though Bel determines the decision.

Another interesting difference between the behaviour of Bel and C_d with regard to the preference property is the impact of $S_{A \subseteq B}$. This is visualized in Figures 4.20a and 4.20b. In

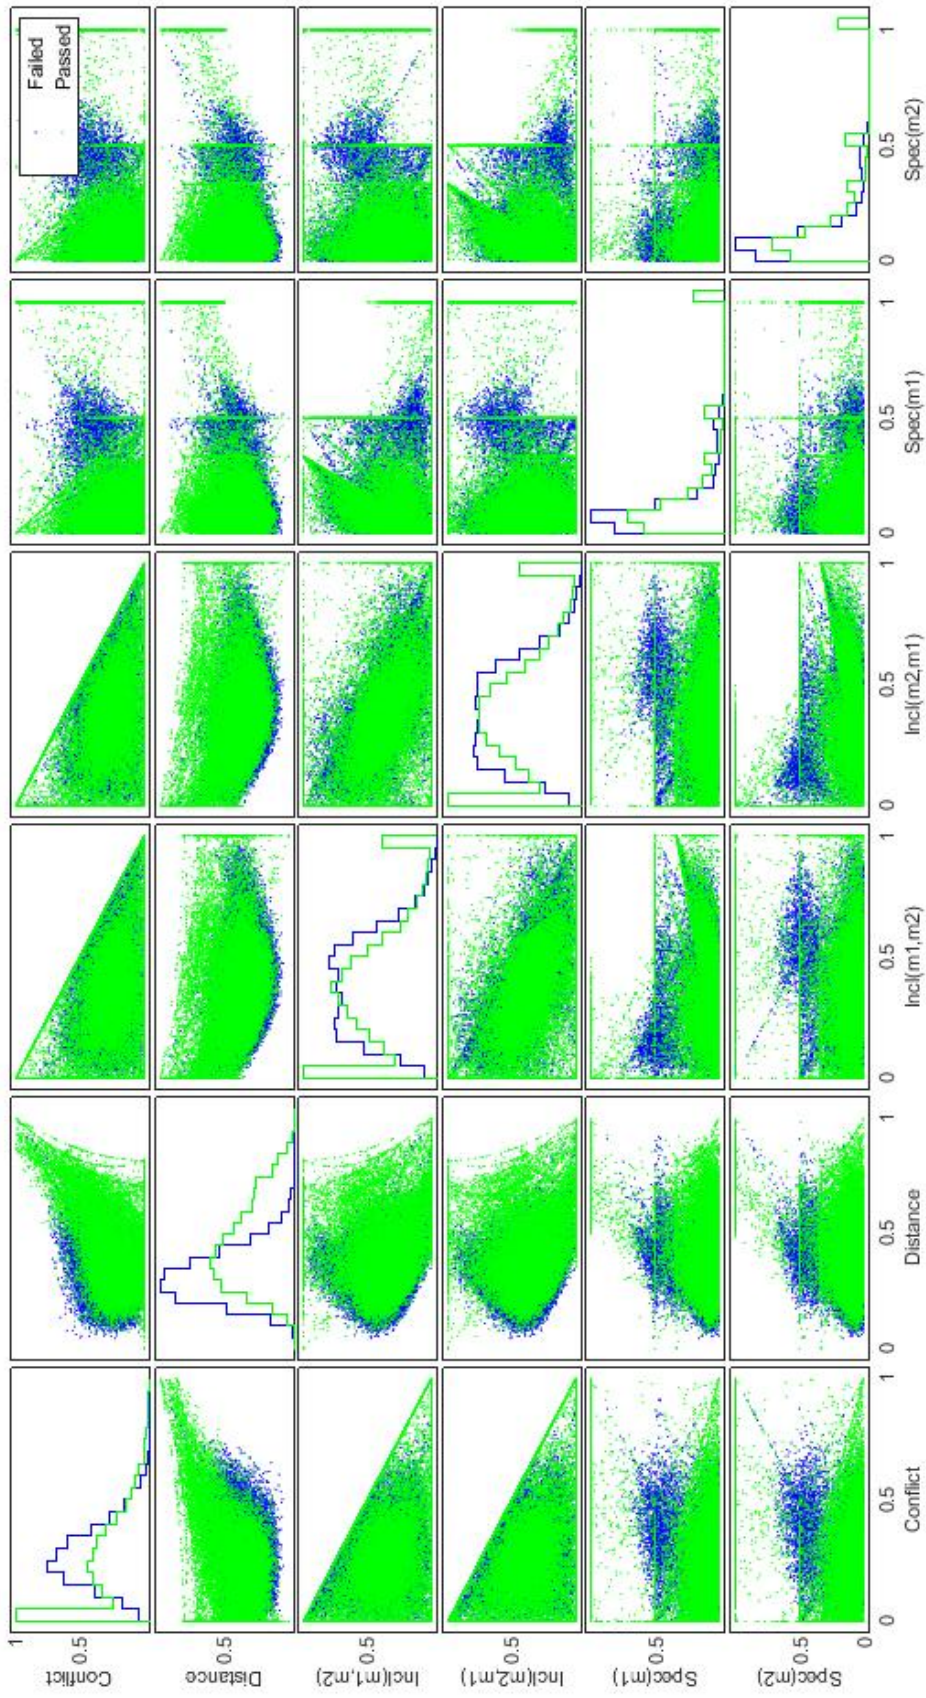


Figure 4.17: Set of scatter plots for decision making using maximum of belief, combination using Dempster's rule and the belief measure being used as the contribution measure

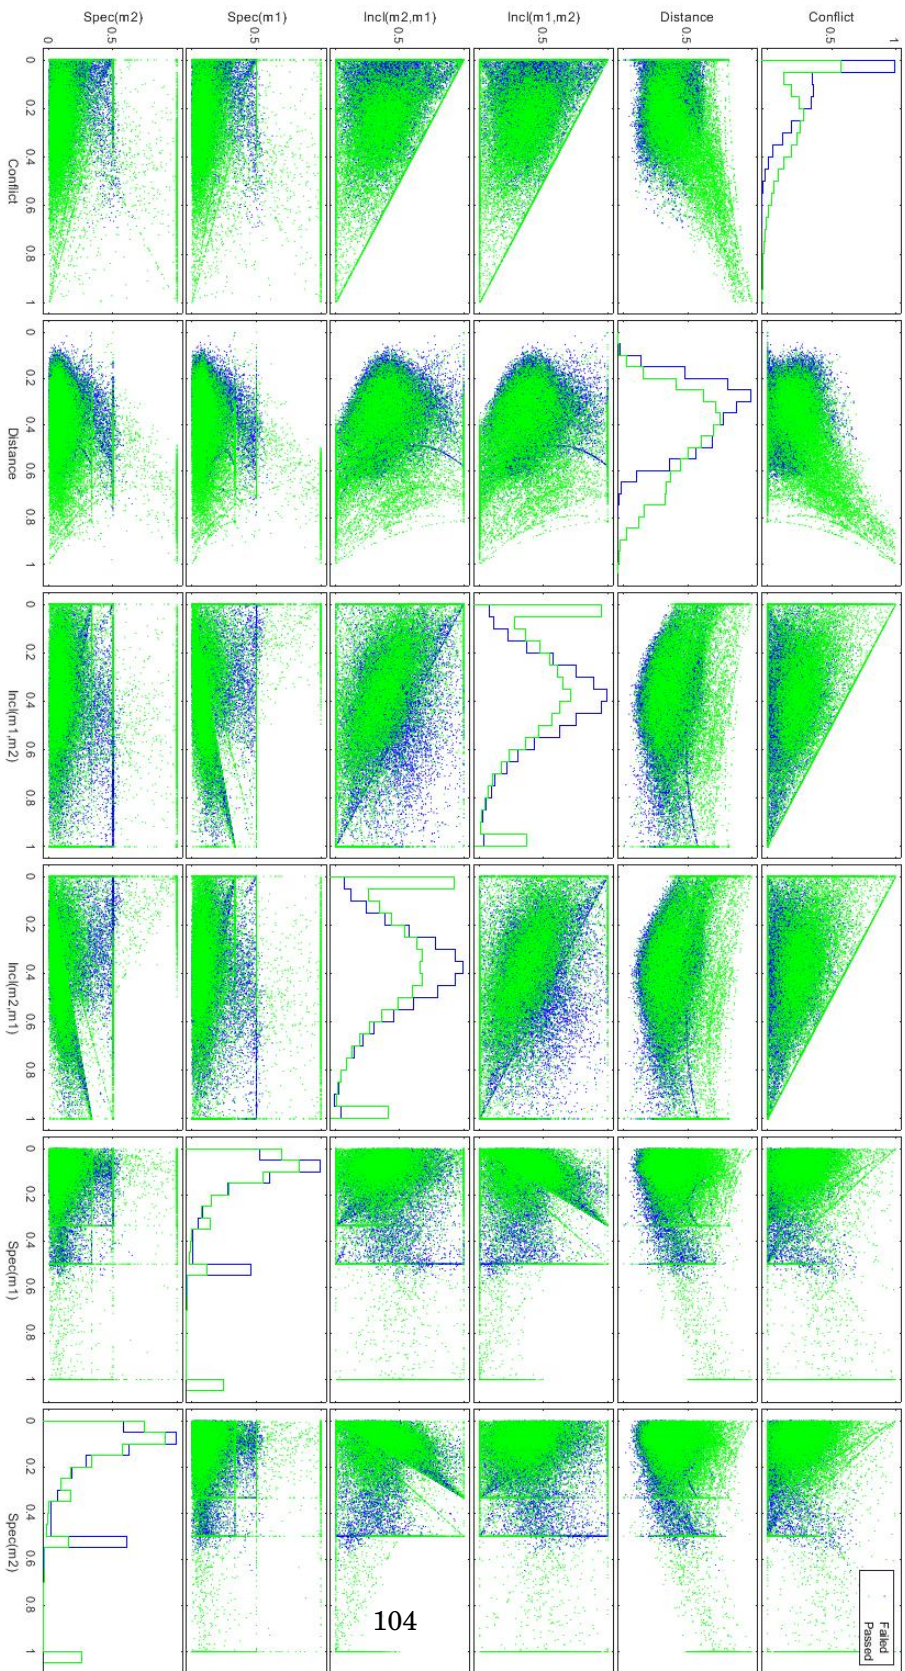


Figure 4.18: Set of scatter plots for decision making using maximum of belief, combination using Dempster's rule and the belief measure being used as the contribution measure




Figure 4.19: Histogram of bba pairs satisfying and not satisfying the preference property for Bel , where $0.2 < d_J(m_1, m_2) < 0.45$ and $0.1 < m_J(\emptyset) < 0.4$ for different values of $S_{A \subseteq B}$

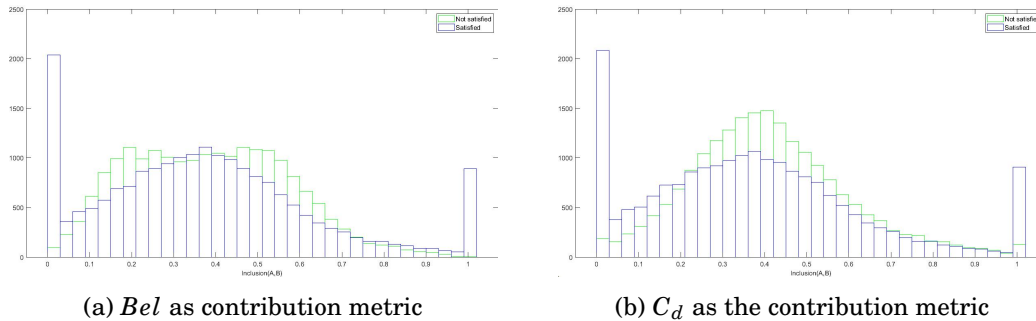


Figure 4.20: Histogram of bba pairs satisfying and not satisfying the preference property for different values of $S_{A \subseteq B}$

particular, it can be seen that the performance of C_d is affected by $S_{A \subseteq B}$ significantly more than that of Bel .

4.5.2 Dempster's rule of combination and other decision-making approaches

The next step is to analyze corresponding results when other decision-making approaches are used. It is clearly not viable to attempt to analyze every single decision-making method in as much detail as in the previous section. As such, the main focus in this section will be placed on decision making using the pignistic probability. However, some results for a maximum of plausibility and maximum of probability methods will also be presented. First of all let us compare the overall performance of different contribution measures for Bel (as presented in Table 4.8) to the corresponding performance rates for the other decision-making approaches as presented in Tables 4.9). These are obtained using an appropriately modified procedure from Algorithm 3.

There are some interesting insights to be gained from this data. First of all, as expected, the behaviour with regard to the preference property of Shapley-based measures is identical or very similar to that of the measures on which they are based. We can see that for all the decision-making methods, the distance-based contribution measure C_d is the second-best - it performs almost as

good as the "best" measure, although it is always slightly lagging behind either *Bel* or *BetP*. This suggests it may be useful in cases where a more complex decision-making method is employed, or the same explanation method needs to be applied regardless of the decision tool used.

4.5.3 Other combination rules

It is worthwhile to consider the situation where combination rules other than Dempster's are used - although the former is the most popular, it is not necessarily suitable in every case. The most likely most common alternatives are Smets' conjunctive rule and Yager's rule. The disjunctive rule of combination is occasionally used in case of highly conflicting but potentially reliable sources, but its usage in real applications is less common. The Proportional Conflict Redistribution rules proposed by Dezert and Smarandache for DsmT are also relatively well-known in particular for decision-making purposes.

Table 4.10 shows the failure rates using the different combination rules. Some interesting things to note are that the distance-based contribution measure performs very well for DSMT - it is the second-best measure regardless of the decision method used. Its behaviour is slightly worse for the conjunctive-disjunctive combination, albeit still acceptable. Finally, when the disjunctive rule of combination is used, its performance is very poor - however, *BetP* performs surprisingly well in this final case.

4.6 Contribution to conflict

The first step towards discussion of contribution to conflict has to be to define what conflict actually is in the belief function framework and how is it different to uncertainty in general. Within the belief functions community it is common to separate the conflict into two forms - the internal and external conflict, Daniel[98], Liu [99], Destercke and Burger [90] as well as Pichon et al. [100] discuss this in more detail.

The notion of conflict follows from the logical notion of consistency. Generally, a consistent theory is such that it does not entail a contradiction. This can be applied to the information provided in the form of sets. If only a single set provides information, it is inconsistent if and only

Decision	Shap(Bel)	Shap(Pl)	Distance-based C_d	Bel	Pl	BetP
Bel	4.11%	31.3%	6.4%	4.11%	31.3%	8.9%
Pl	20.9%	26.8%	19.5%	20.9%	26.7%	18.6%
BetP	17.2%	25.2%	15.3%	17.2%	25.2%	14.1%
Ignoring cases with no conflict						
Bel	4.9%	25.3%	6.1%	4.9%	25.3%	8.1%
Pl	20.7%	22.4%	18%	20.7%	22.4%	16.6%
BetP	16.3%	20.7%	13.2%	16.3%	20.7%	11.7%

Table 4.9: Proportion of cases where the preference property is not satisfied for the different contribution measures and various decision making methods using Dempster's rule of combination

if this set is the empty set \emptyset . The consistency measure for sets only has two values - a set is either fully consistent or inconsistent, with no room for fuzziness.

This is not the case with belief functions, but the notion can be naturally extended. A mass assignment modelling the empty set only ($m(\emptyset) = 1$) can be naturally associated with the total inconsistency. However, there are several possible definitions of the totally consistent mass assignment, drawing from logical and probabilistic approaches [101] [102] [103].

Definition 4.4 (Logical consistency). A mass assignment is logically consistent if and only if $\bigcap_{E \in \mathcal{F}} E \neq \emptyset$ (where \mathcal{F} is the set of focal sets)

This definition means that for a mass function to be logically consistent, its non-zero focal sets must have a non-empty intersection. It is shown in [90] that this means that such a mass considers at least one state of the world as fully plausible.

Definition 4.5 (Probabilistic consistency). A mass assignment is probabilistically consistent if and only if $m(\emptyset) = 0$

This definition is consistent with logic-based interpretation of belief function and with the original formulation as part of the Dempster-Shafer Theory. These two definitions of total consistency induce two alternative consistency measures from the set of all mass functions to $[0, 1]$ for belief functions, satisfying both the minimum and maximum values

$$(4.26) \quad \phi_{pl}(m) = \max_{\omega \in \Omega} pl(\omega)$$

and

$$(4.27) \quad \phi_m(m) = 1 - m(\emptyset)$$

DSmT PCR6						
Decision	Shap(Bel)	Shap(Pl)	Distance-based C_d	Bel	Pl	BetP
Bel	1.7%	14.9%	3.7%	1.7%	14.9%	4.2%
Pl	10.1%	13.1%	9.2%	10.1%	13.1%	9.1%
BetP	8.1%	12.4%	7.1%	8.1%	12.4%	6.9%
Dubois-Prade (Conjunctive-Disjunctive)						
Bel	2.0%	15.6%	4.4%	2.0%	15.6%	4.5%
Pl	10.3%	12.9%	10.4%	10.3%	12.9%	9.1%
BetP	8.5%	12.3%	8.5%	8.5%	12.3%	7.0%
Disjunctive rule						
Bel	8.2%	7.1%	23.6%	8.2%	7.1%	4.3%
Pl	13.9%	7.0%	22.4%	13.7%	7.0%	7.5%
BetP	12.2%	6.3%	22.7%	12.2%	6.3%	5.9%

Table 4.10: Proportion of cases where the preference property is not satisfied for the different contribution measures and various decision making methods using Dempster's, Dubois-Prade or PCR rules of combination

This can be extended to conflict between sources by considering the inconsistency resulting from information combination. Again, for sets, we either have a total conflict ($A \cap B = \emptyset$) or no conflict.

Total conflict is the straightforward case where the disjunctions of all focal elements of the two masses have an empty intersection. Let $\mathcal{D}_i = \cup_{A \in \mathcal{F}_i}$ denote this disjunction, where \mathcal{F}_i denotes the set of all focal elements with non-zero mass of the mass assignment function m_i . Then we have:

Definition 4.6 (Total conflict). m_1 and m_2 are in total conflict if $\mathcal{D}_1 \cap \mathcal{D}_2 = \emptyset$

There exist multiple measures for internal conflict of belief functions as well as that between belief functions. The traditional measure of conflict between two mass functions is the mass on an empty set following the conjunctive combination of the two masses $m_1 \oplus m_2(\emptyset)$. This was criticized by Liu [104], who pointed out the existence of masses which may be fully consistent with one another but still result in a significant mass assigned to the empty set through conjunctive combination. It was proposed to use a pair of values - the mass on the empty set and the geometric distance (Equation 2.28) to accurately assess the overall conflict between mass functions. The concept of using a pair of values was also subsequently discussed by Daniel [98] who suggested taking a minimum of several possible conflict values. Pichon et al. proposed a number of measures to quantify this external conflict [100], based on the notions of N-consistency

$$(4.28) \quad \phi_N(m) = 1 - m^N(\emptyset)$$

with m^N denoting the combination of m with itself N times

$$(4.29) \quad m^N = m^{N-1} \cap m$$

which in a way can be considered to follow from the concept of auto-conflict (1-consistency) [103]

$$(4.30) \quad m^2(\emptyset) = m \cap m(\emptyset)$$

One such measure which is relatively straightforward and of interest to us is:

$$(4.31) \quad \kappa_\pi(m_1, m_2) = 1 - \phi_\pi(m_1 \cap m_2)$$

. In [90] a consistency measure based on plausibility contour was proposed satisfying the proposed axioms along $m(\emptyset)$:

$$(4.32) \quad \phi_\pi = \max_{x \in \mathcal{X}} \pi(x)$$

where $\pi(x)$ is the plausibility function restricted to singletons. This measure is one of the many measures we can use to quantify internal conflict.

4.6.1 Conflict contribution measures

The definitions and properties of a contribution measure of conflict do not specify whether it is the internal or external conflict that is concerned. When we consider the contribution to internal

conflict, we are concerned with how the internal conflict changes when another body of evidence is included in the fusion process. This can be expected to correlate with the external conflict between the said body of evidence and the remainder of the knowledge base in the majority of cases.

In the previous section, we have discussed two main avenues of assessing the impact of a single body of evidence on the fusion process - these are the approaches based on the Shapley value and taking the derivative with respect to the discounting factor. The former was done for conflict by Ben-Abdallah, and Jousselme [79], whereas the latter was used for conflict and uncertainty through consonance and specificity in the original work by Strat and Lowrance [92]. However, this consonance measure is not a real conflict measure, and it does not satisfy the axioms from [90]. With the breadth of possible measures for the conflict, we have various possible contribution measures for quantifying conflict brought by a particular body of evidence. In general, with one of these approaches, correctly identifying the key contributor is significantly easier than in the case of analysis of the decision, as the preference property is no longer that difficult to satisfy.

In fact for Proposition 4.5 to be satisfied we have: If

$$\phi(K^*/k_1) \geq \phi(K^*/k_2)$$

then

$$\kappa(K^*, k_1) \leq \kappa(K^*, k_2)$$

Consider the method based on taking the derivative of the measure of interest with respect to discounting factor. Then we have $\kappa(K^*, k_1) = \frac{d\phi(K^*)}{d\alpha_1}$ and by definition of the discounting process $\phi(K^*, \alpha_1 = 1) = \phi(K^*/k_1)$, as such it is simple to see that the preference property is satisfied for all cases where the derivative $\frac{d\phi(K^*)}{d\alpha}$ is constant with respect to α .

Given the breadth of possible conflict measures, it is a daunting task to try to analyze all or even the majority of them. As such, we only discuss a few in greater detail. First of all, we must consider the Shapley value with auto conflict $m^2(\varnothing)$ as the payoff function, which was proposed by Ben-Abdallah²[79].

$$(4.33) \quad \text{Shap}_\kappa(m) = \sum_{S \subseteq \mathcal{N}} = \frac{|S|!(N - |S| - 1)!}{N!} (\kappa_m(S \cup U\{i\}) - \kappa_m(S))$$

with $\kappa_m(S) = \bigoplus_{j \in S} m_j(\omega)$. As per the definition of Shapley value we have $\sum_{i \in \mathcal{N}} \phi_i(\kappa) = \kappa_m(\mathcal{N})$. Clearly we can replace the payoff value with another conflict measure, such $m_J(\omega)$ or the plausibility-contour based conflict $1 - \phi_\pi$ (Equation 4.31). We can also apply all of these measures easily to the derivative-based approach $\frac{d}{d\alpha}$.

Therefore we are concerned with six measures of contribution of the individual body of evidence to the overall conflict. These are the Shapley values for auto-conflict $\text{Shap}_{m^2(\varnothing)}$, mass on the unnormalized joint empty set $\text{Shap}_{m_J(\varnothing)}$, and contour function conflict Shap_{κ_π} as payoff functions. Furthermore we have the derivatives of the three measures with respect to discounting rate of the mass of interest: $\frac{dm^2(\varnothing)}{d\alpha}$, $\frac{dm_J(\varnothing)}{d\alpha}$ and $\frac{d\kappa_\pi(m)}{d\alpha}$

With this in mind, we can have a look at the behaviour of these six measures. Consider the following scenario where a vessel type is to be identified. The frame of discernment is $\Omega =$

²At the time of writing, this was only published in a CMRE report

{*Tanker, Cargo, Fishing*}. We have four bodies of evidence with different degrees of internal inconsistency and conflict between one another

$$\begin{aligned}
 m_1(T) &= 0.6 & m_2(T, C) &= 0.7 \\
 m_1(\Omega) &= 0.4 & m_2(\Omega) &= 0.3 \\
 \\
 m_3(T) &= 0.4 & m_4(F) &= 0.7 \\
 m_3(C) &= 0.3 & m_4(\Omega) &= 0.3 \\
 m_3(\Omega) &= 0.3 & &
 \end{aligned}$$

Clearly, m_1 and m_2 are both fully internally consistent with different degrees of nonspecificity. The mass function m_3 is in partial conflict with the other two due to its own internal inconsistency or internal conflict. Finally, m_4 is in conflict with all the other masses.

Combining these four masses conjunctively results in the following joint mass function: $m(T) = 0.57$, $m(C) = 0.19$, $m(T, C) = 0.1$, $m(F) = 0.1$, $m(\Omega) = 0.04$ with $m(\emptyset) = 0.7468$ pre-normalization. Computing the conflict contribution measures we have: These results highlight the differences in

Table 4.11: Conflict contribution measures for masses m_1 to m_4

	$Shap_{m^2(\emptyset)}$	$Shap_{m_j(\emptyset)}$	$Shap_{\kappa_\pi}$	$\frac{dm^2(\emptyset)}{d\alpha}$	$\frac{dm_j(\emptyset)}{d\alpha}$	$\frac{d\kappa_\pi(m)}{d\alpha}$
m_1	0.02	0.129	-0.04	-0.193	0.110	-0.3
m_2	0.006	0.084	0	-0.191	0.059	-0.16
m_3	0.203	0.148	0.197	0.122	0.13	0.1
m_4	0.156	0.385	0.135	0.29	0.51	0.21

the behaviour of the six measures. In particular, we can see the main difference between using the Shapley index and the derivative is that the former considers the average of all the combinations (coalitions) involving the mass of interest, whereas the latter considers changing the mass of interest only and keeping the other contributors constant. This difference is particularly striking when considering the difference between m_3 and m_4 . We can see that for the auto conflict $m^2(\emptyset)$ the main contributor is identified as m_3 using the Shapley value, but m_4 taking the derivative. Similar behaviour is observed with κ_π . As discussed before, $\frac{d}{d\alpha}$ satisfies the preference property as long as it is constant with respect to α , which is the case for conjunctive combination. We can verify it for this example as $m_{\cap i=1,2,4}^2(\emptyset) = 0.3$ and $m_{\cap i=1,2,3}^2(\emptyset) = 0.265$, showing that m_4 is a greater individual contributor to conflict. However, from the Shapley value, we know that it is the greater contributor across all the possible coalitions - the latter measure is useful, despite not satisfying the preference property in this example.

4.7 Contribution to uncertainty

As discussed at length in Section 2.3.2 it is not straightforward to distinguish between the notions of "conflict" and "uncertainty". In general, for the purpose of this section, "uncertainty" means "ignorance" or "non-specificity," i.e. lack of information. As such, we are not looking at the overall uncertainty, which combines both inconsistency and second-order uncertainty, but rather attempt

to separate the two and consider them independently. Furthermore, in most information fusion applications, sources will not contribute to uncertainty (ignorance); each non-conflicting source will either reduce the overall uncertainty or keep it unchanged [1] - in other words, new sources can bring in new information, but not take it away. When dealing with belief functions, there are two main situations when we may consider a body of evidence to increase the overall uncertainty. The first case is when we consider conflicting sources and a method of combination other than Dempster's rule where this conflict is interpreted in such a manner that it increases uncertainty. Some obvious examples are the disjunctive rule of combination (where it is assumed at least one of the sources is correct and the result of the combination is the knowledge that both sources can agree on) or Yager's rule, where conflict is interpreted as uncertainty. Another example of a body of evidence increasing overall ignorance is that of contextual information, the case which has been discussed in Chapter 3 or more generally a body of evidence which is used to reason about the behaviour of another source can increase the second-order uncertainty regarding a variable of interest.

In this section, we look at some measures of nonspecificity or ignorance and discuss how they can be used to construct appropriate contribution measures using the method discussed in the previous sections. The measures are then applied to the two scenarios where individual bodies of evidence increase second-order uncertainty - non-Dempsterian combination rules and valuation networks with source correction model.

4.7.1 Measures of second-order uncertainty

A number of measures for the quantification of second-order uncertainty have been proposed. Traditionally we have the Hartley functions as a measure of uncertainty, which has been generalized by Higashi and Klir [105] who proposed the measure of U-uncertainty, also satisfying the properties of Shannon's entropy for possibility distributions and by extension to basic belief assignments describing ordered possibility distributions. This has been extended by Dubois, and Prade [102] to any basic belief assignment

$$(4.34) \quad N(m) = \sum_{A \in F} m(A) \log |A|$$

Whilst this is arguably most commonly used and most correct measure of nonspecificity we may use another one to continue considering the geometric interpretation of the belief function, i.e. the distance to the vacuous mass function $m_v(\Omega) = 1$. From 2.28 we have:

$$(4.35) \quad d(m_1, m_v) = \sqrt{\frac{1}{2}(\vec{m}_1 - \vec{m}_v)^T \mathbf{Jac}(\vec{m}_1 - \vec{m}_v)}$$

and hence we have the uncertainty measure

$$(4.36) \quad U(m) = 1 - d(m, m_v)$$

In general, the advantage of the latter is that it is bounded between 0 and 1, whereas the former is not. However, it is dependent on the size of the frame of discernment. As such, despite not being a proper nonspecificity measure, it can be useful nevertheless.

Similarly to the conflict measures, we can obtain four possible measures of the contribution of the particular body of evidence to overall nonspecificity or uncertainty, using either the derivative with respect to the discounting factor $\frac{d}{d\alpha}$ or the Shapley value. We need a payoff function for Shapley value such that it is 0 for the empty coalition; in this case, the measure used must be such that $v(m_v) = 0$, where m_v denotes the vacuous belief function. For this reason, it is easier to use measures of specificity rather than nonspecificity. Therefore we have $-Shap_{N(m_v)-N}(m)$, $-Shap_{d(m,m_v)}(m)$, $\frac{dN}{d\alpha}$ and $\frac{dU}{d\alpha}$

4.7.2 Example: combination methods facilitating nonspecificity

Recall the vessel identification example from the previous section. Now we consider combining these mass functions with Yager's combination rule or Dubois-Prade conjunctive-disjunctive rule. We do not consider the disjunctive rule itself as although it may be used for conflict handling purposes, it makes little to no sense to use it for combination of multiple mass functions. Thus we have m_Y , the result of combination using Yager's rule and m_{DP} , the result of conjunctive-disjunctive combination:

$$\begin{array}{ll}
m_{DP}(T) = 0.14 & m_Y(T) = 0.144 \\
m_{DP}(C) = 0.048 & m_Y(C) = 0.05 \\
m_{DP}(F) = 0.025 & m_Y(F) = 0.025 \\
m_{DP}(T, C) = 0.1 & m_Y(T, C) = 0.025 \\
m_{DP}(T, F) = 0.34 & m_Y(\Omega) = 0.75 \\
m_{DP}(F, C) = 0.11 & \\
m_{DP}(\Omega) = 0.238 &
\end{array}$$

Table 4.12: Specificity and nonspecificity contribution measures for masses m_1 to m_4

	Yager's Combination				Conjunctive-Disjunctive Combination			
	$-Shap_{N(m_v)-N}$	$-Shap_{d_v}$	$\frac{dN}{d\alpha}$	$\frac{dU}{d\alpha}$	$-Shap_{N(m_v)-N}$	$-Shap_{d_v}$	$\frac{dN}{d\alpha}$	$\frac{dU}{d\alpha}$
m_1	-0.267	-0.09	0.11	0.018	-0.19	-0.14	0.09	0.11
m_2	-0.018	-0.071	0.08	0.01	-0.01	-0.06	0.07	0.07
m_3	-0.35	-0.057	0.1	0.048	-0.25	-0.09	0.05	0.1
m_4	-0.19	0.072	0.7	0.32	0.11	-0.07	8 (!)	0.7

We can see significant differences between the various measures. Some of the arguments discussed with respect to conflict apply here, too, in particular regarding the difference between Shapley value and derivative and how the former takes into account the average of all the coalitions, rather than focusing just on the body of evidence of interest. However, as mentioned earlier, a major issue is that for Shapley value, we require the payoff function to be 0 for the empty coalition, which in this interpretation corresponds to the vacuous mass function. For that reason, we need to convert the nonspecificity measures into specificity measures and then re-interpret the results as contributions to nonspecificity. With both Yager's and conjunctive-disjunctive combination, each body of evidence individually increases the nonspecificity (or rather removing any individual piece of evidence reduces the overall nonspecificity, whether using the Hartley measure or distance to

the vacuous mass function). Whilst this is captured by the $\frac{d}{da}$ family of measures, it is not the case for the ones based on Shapley value. In fact, in the case of Hartley measure with Yager's combination and distance-based measure with conjunctive-disjunctive fusion, the mass m_4 , which clearly has the greatest impact on the uncertainty of this combination, is not even correctly identified as the key contributor. Whilst Shapley's value has been shown to be a good method of correctly attributing the origin of the conflict, it is not really suitable for tracking non-specificity, at least not in the way discussed here.

4.8 Discussion and summary

In this chapter, we have discussed the notions of the provenance of information, conflict and uncertainty in a framework-agnostic manner and then applied these concepts to the belief functions framework. Several propositions have been put forward to be satisfied by contribution measures, including a preference property. By this property, the contribution of a source is greatest if removing it from the knowledge base has the greatest impact - in the case of contribution to decision or knowledge, it is the only one to change the decision to be made or the state of the world assumed to be true, and in the case of conflict and uncertainty, it has the greatest impact on some uncertainty or conflict measure.

This proposition seems trivial at first glance, but a deeper investigation showed that in the case of decision-making, it is actually quite difficult to satisfy. A number of Monte-Carlo situations were performed to investigate the impact the relationship between properties of the mass functions in question has on the likelihood that the two preference property is satisfied for the decision making based off these functions for a particular contribution measure and decision-making paradigm. The rationale for this was that if some relationships can be found, it may be possible to select appropriate contribution measures given the properties of bodies of evidence and the contribution measure. Furthermore, this rate of satisfaction of preference property could be used as one of the possible benchmarks of the different contribution measures.

The existing literature on the problem was reviewed, and some existing measures based on sensitivity spaces were discussed. A novel approach based on the geometric interpretation of belief functions was proposed, discussed in detail and compared to existing solutions. The purpose of this approach was to aim towards a contribution measure independent of the method of combination used, and although it was outperformed by other measures in different settings, it showed one of the best performances across the board, something that other measures lacked.

Furthermore, several contribution measures were discussed to track the origin of conflict and nonspecificity in belief functions, based on Shapley's value or the sensitivity to discounting approach. It was shown that it is much easier for the preference property to be satisfied in this context, but also that it is not necessarily the best benchmark and that even measures not satisfying it can be very useful in some cases.

There are extensive areas where further work can be carried out. In particular, it may be worthwhile to revisit the properties the contribution measure should satisfy. While in some cases, the notion of "greatest impact" can be self-explanatory, in others, it is not. In particular, if we

consider more complex valuation networks, source correction models and such, it may be very hard to agree on what is meant by the greatest impact of a particular body of evidence. This is a major difficulty with tracking uncertainty as well. While in some cases, the uncertainty can come from conflict redistribution, in others, it may be due to source correction or source self-doubt.

SITUATION AWARENESS USING CONCEPTUAL GRAPHS

Higher-level information fusion or modelling of situations will be discussed in this section. Dealing with information concerning more than one entity is a problem at higher levels of information fusion. At the Data and Information Fusion Conference organised by The Los Alamos National Laboratory in Santa Fe, New Mexico, in August 2019, one of the main topics of discussion was the identification of the key current challenges in information fusion. It was generally agreed that these are: handling uncertainty, defining and dealing with soft information and handling high-level situations. The first problem was addressed extensively in the previous chapters; here, we focus on the latter two.

5.1 Modelling of situations

The model proposed by Joint Directors of Laboratories (JDL) has been in use since the inception of the field of information fusion. Although often criticised and hence undergoing constant refinement, it is, without doubt, the most important model employed by the information fusion community. It is widely agreed that the JDL fusion model has gained most acceptance in the information fusion community - with several revisions and modifications proposed by different authors[69][70].

The JDL model differentiates between a number of fusion "levels" thus providing distinction between different processes which relate to refinement of either "objects", "situations", "threats" and "processes". In general we are not concerned with the distinction between *knowledge*, *data*, *sensor* and *information* fusion.

The JDL data fusion model consists of 5 levels: source preprocessing (level 0), object refinement (level 1), situation refinement (level 2), threat refinement (level 3) and process refinement (level 4).

5.1.1 JDL perspective

The JDL levels are not very clearly defined - to the extent that some researchers would go as far as saying that the JDL model itself does not actually exist and rather is a very general notion that different tasks within information fusion can be split and separated into various levels. As Steinberg and Snidaro suggest [69], a Data Fusion model supporting the development of fusion systems is an engineering model, and as such, its purpose is to partition the problem space to support the development of different solutions. The previous two chapters of this thesis, in particular, Chapter 3, addressed the problem of reasoning about an individual entity or its properties. As such, although some information fusion practitioners would disagree and claim the problems addressed are complex enough to consider them on Level 2 of the JDL, I am inclined to consider the work done to this point to be within Level 1 of the JDL. For Level 2, we must consider the problem of situation awareness for which we must not only have several targets with some properties but some structured relationships between them.

5.1.2 Situation awareness

Situation awareness is concerned with making sense of the environmental elements and events. Endsley [106] defines it as the perception of these elements and events, comprehension of their meaning, projection of their future status and resolution - planning and decision-making. This is naturally similar to the processes described in Chapter 3, where rather than being concerned with the elements and events in the environment, we discussed the reasoning (comprehension step) about the reports provided by sources (perception), given context-of the problem itself as well as the contexts in which these sources operate. The key difference is that, whereas before we focused on a particular problem entity (and its context), now we operate within the greater context of some environment, which may contain multiple entities.

In this chapter, we are concerned with representing information about multiple entities and the relationships between them. Again, it is assumed that this information may be provided by partially reliable sources, furthermore as each source provides an overview of the situation regarding multiple entities, we need to consider the situation where some information provided by a source is true.

5.1.3 Situations and context

The relationship between situations and context has been briefly discussed in Chapter 3 of this thesis. Recently, Steinberg addressed this relationship by defining context as a "situation that provides information that can be used either a) to condition expectations or b) to improve the understanding of a given inference or planning/control problem" [67]. This is consistent with the notion of *context-of*, as discussed in Chapter 3. From this perspective, all information describing the situation can also characterise the context. Consistently with earlier discussions on context, a situation may be true within a certain context or as part of a greater situation. Furthermore, according to Steinberg, *context-of* and *context-for* can play a significant role in fusion at any JDL

level. Still, they are particularly relevant at higher levels, where the variables of interest are relationships or multi-entity situations, which cannot be observed directly and must be inferred.

5.2 Conceptual graphs, information graphs and ontologies

In this section we discuss in more detail the notion of a conceptual graph first introduced in section 2.4.3.

For a *specialization relation* between conceptual types $t_1 \leq t_2$ to hold the concept type t_1 must be a subtype of t_2 according to the dictionary \mathcal{V} , i.e. if t_1 is a specialization of t_2 any instance of the class denoted by t_1 is also an instance of the class denoted by t_2 .

This is extended to apply to concepts - given the order on T_C concepts defined on $T_C \times \{\text{markers} \cup *\}$ can be partially ordered by a specialization relation. Given the two concept nodes $c_1 = [T_1 : m_1]$ and $c_2 = [T_2 : m_2]$ the specialization relation is defined as follows:

$$(5.1) \quad c_1 \leq c_2 \text{ iff } \begin{cases} T_1 \leq T_2 \\ m_2 = * \quad \text{or } \text{sim}(m_1, m_2) \geq \text{thres} \end{cases}$$

where $\text{sim}(m_1, m_2)$ is a similarity function and thres is a predefined threshold.

Consider this simple example where the conceptual type `CargoShip` is more specific than `Vessel`. As such we have:

$$[\text{Vessel} : *] \geq [\text{Vessel} : \text{Rex}] \geq [\text{CargoShip} : \text{Rex}]$$

As the concepts are only partially ordered `[Vessel : Rex]` and `[CargoShip : *]` are not comparable.

Furthermore a specialization relation is defined between graphs and denoted by \sqsubseteq . Given two conceptual graphs $A = (C_A, R_A, E_A, l_A)$ and $B = (C_B, R_B, E_B, l_B)$. Let P_{AB} denote the set of graph isomorphisms between A and B . As such we have:

$$(5.2) \quad B \sqsubseteq A \Leftrightarrow \exists p \in P_{AB}, \begin{cases} p : C_A, R_A \rightarrow C_B, R_B \\ c_A, r_A \rightarrow c_B, r_B \\ \forall c_A \in C_A, \quad c_B \leq c_A \\ \forall r_A \in R_A \quad r_B = r_A \end{cases}$$

This brings us to the concept of specialisation set, an operator Π which denotes the set of all more specialised graphs of a graph so that

$$(5.3) \quad B \in \Pi(A) \Leftrightarrow B \sqsubseteq A$$

The basic information fusion model for conceptual graphs follows almost exclusively from the ideas presented in [107], as the model for uncertain information fusion presented in that paper is the foundation on which the extensions proposed throughout here are built. The methodology for soft information fusion, where semantic information is combined, utilises the notion of a *knowledge model*. This model is formally expressed as a conceptual graph with all concept nodes in

the form $[\text{ConceptType} : *]$. Throughout this chapter, such a graph is referred to as the *model* and denoted by \mathcal{M} .

Any conceptual graph A associated with observation is provided in a form consistent with the model, such that $A \sqsubset \mathcal{M}$ or $A \in \Pi(\mathcal{M})$

The information fusion process for conceptual graphs representing observed situations can be defined using fusion operator \otimes compliant with a model \mathcal{M} such that:

$$(5.4) \quad \begin{aligned} \otimes : \Pi(\mathcal{M}) \times \Pi(\mathcal{M}) &\rightarrow \Pi(\mathcal{M}) \\ (A, B) &\rightarrow A \otimes B \end{aligned}$$

Different fusion operators can implement various fusion strategies [108], and these can be specified in accordance with the requirements, context, level of importance of different items or other factors.

5.2.1 Uncertainty management

A relatively straightforward method for uncertainty management was proposed in [107], where a reliability rating (which can be interpreted as the probability that a source is fully reliable) was proposed. It is based upon the assumption that if two sources of information A and B provide their observations, four cases need to be considered: both sources are reliable (in this case, we know $A \otimes B$ about the event), one of the two is reliable but the other isn't (our knowledge is either A or B and the other one is discarded) and finally neither is reliable and our state of knowledge defaults to the model \mathcal{M}

5.2.2 Motivation

From the above, it can be seen that although some work on uncertainty representation in this context has been done, there are several areas that should be investigated further. In particular, we could argue that in semantic fusion, the assumption that the entirety of information provided by a source is either true or false is going to negatively impact the outcome of the fusion process. Given that a single source (or a statement) can provide multiple pieces of information simultaneously, it is reasonable to assume that a source may be correct regarding some of its assertions and not about the others.

This makes it necessary to discuss some methods of representing the combination of partially conflicting data as well as of modelling partial reliability of sources. Finally, we discuss the issue of dealing with modality in semantic information.

5.3 Uncertainty management for partial conflict resolution

A significant portion of the content presented in this and the next several sections has been previously published in [109]

5.3.1 Set-based approach

We propose a set-based approach to partial conflict resolution by allowing the conceptual type as well as the relation type to be labelled with a *non-empty set of non-subsumptive hypotheses*. Note that set subsumption (defined below) is different to graph subsumption. For simplification and to make this approach more consistent with existing methods of uncertainty management, we assume that no two conceptual types, which are not a specialisation or generalisation of one another, may have a common specialisation.

For the conceptual type we have a non-subsumptive set of pairs $l_G(c) = \{(type(c), marker(c))\}$, such that $type(c) \in T_C$ and $marker(c) \in markers \cup *$.

We define a non-subsumptive set as a set such that no two elements within are a generalisation or specialisation of one another. This can be formally defined as follows: A set \mathcal{S} of concept labels s_i is considered *subsumptive* if there exists a pair $\{s_1, s_2\} \in \mathcal{S}$ such that $s_1 \leq s_2$

A similar extension is proposed to aggregate conflicting information regarding relation type; however, as our approach is restricted to unordered relation types, it is unnecessary to worry about subsumption. Thus we let r be labelled by a non-empty subset of relation types $l_G(r) \subset T_r, l_G(r) \neq \emptyset$

This makes it possible to embed the most basic uncertainty model in the conceptual graph model itself. To illustrate its usefulness, let us consider this simple scenario of investigating the behaviour of a suspicious cargo vessel *Rex*. Consider the maritime type hierarchy as shown in Figure 5.1.




Figure 5.1: Excerpt from a maritime concept vocabulary

Now consider conflicting reports regarding the destination of some vessel. For instance, consider that the AIS reports state the destination of the vessel is the port in Iskenderun, but the kinematic analysis of the vessel's behaviour suggest it is heading towards Mersin. In this case, the only option of modelling this using a standard conceptual graph requires a loss of information. The representations without and with the direct embedding of uncertainty are shown in Figures 5.2 and 5.3 respectively.



Figure 5.2: Conceptual graph representation without uncertainty modelling

Similarly, let us consider the case where fewer specifics are provided - say that kinematic destination analysis only suggests that the real destination is Turkey. Meanwhile, the entry in the AIS field states that the destination is Italy. Again, using the standard conceptual graph approach,




Figure 5.3: Conceptual graph representation with uncertainty modelling

the only method of representing this is in conceptual graph form is as shown in Figure 5.4. By allowing uncertainty, we obtain a more informative representation as shown in 5.5.



Figure 5.4: Conceptual graph representation with uncertainty modelling



Figure 5.5: Conceptual graph representation with uncertainty modelling

5.3.2 Uncertain information fusion

Furthermore let us look into information fusion using conceptual graphs incorporating uncertainty. First of all let us briefly discuss the fusion strategy

The fusion strategy used throughout this section is strongly based on the method proposed in [107]. However, there are some changes required to handle uncertainty. We incorporate the concepts of conjunctive and disjunctive fusion used widely in the information fusion community [12] [1] which distinguish between two approaches to dealing with partial conflict. In general, disjunctive fusion is based on the assumption that the reliability of the sources is unknown (some of the sources may be unreliable), whereas under conjunctive fusion, we assume that all the sources are reliable. Note that this is a guideline rather than a rule, and in some cases, we may prefer to use conjunctive fusion even if we know that some of the sources may be unreliable.

The original fusion process by which the nodes and relations of the two graphs are combined to generate the fused graph has been defined as follows:

- For each conceptual type, the most generic common subtype of the two initial concept is selected, e.g. $\text{Destination:*} \otimes \text{Turkey:*} \rightarrow \text{Turkey:*}$
- For markers the markers are considered compatible if either:
 1. One of the markers is * in which case the other marker is kept
 2. For strings, one marker is included in the other (and the longest is kept)

In cases when markers are in a hard data form (e.g. timestamps, locations), the exact method of the combination needs to be proposed on the case by case basis. The markers will be

considered compatible if they are of the same type, and some similarity function is below a specified threshold. The method of combination will vary as well.

This is extended with concepts of the disjunctive and conjunctive combination of concepts.

- **Disjunction:** for each conceptual node labelled with a set of non-subsumptive conceptual types, the union of sets is selected. If the result is a subsumptive set, the conceptual types, which are specialisations of other members of the set, are removed.
- **Conjunction:** for each conceptual node, the intersection of labelling sets is selected as the new label. For this intersection, each conceptual type is considered a set of all its subtypes. This cannot be performed if the resulting intersection would be an empty set.

This can be illustrated with simple examples. Consider the following extension to the earlier type hierarchy for maritime entities.




Figure 5.6: Excerpt from a maritime concept vocabulary

1. **Simple disjunctive fusion:** $\{Trawler:*, Longliner:*\} \odot \{Longliner:*, Cargo :*\} \rightarrow \{Trawler:*, Longliner:*, Cargo ship:*\}$
2. **Simple conjunctive fusion:** $\{Trawler:*, Longliner:*\} \odot \{Longliner:*, Cargo :*\} \rightarrow Longliner:*$
3. **Conjunctive fusion with supertypes as sets of types:**
 $Longliner:* \odot \{Fishing:*, Cargo:*\} \rightarrow Longliner:*$
 since $Fishing:* \geq Longliner:*$

From this, we can define an overall fusion operator \otimes where the conceptual types of G_1 and G_2 are combined using conjunctive fusion unless it would result in an empty set, in which case disjunctive fusion is used instead.

5.3.3 Multi-source aggregation scenario using uncertain information fusion

Now, consider this simple scenario of information aggregation regarding the vessel *Marcel* properties with three conflicting sources of information. The type hierarchy used is the maritime domain type vocabulary first described in Figure 5.1 and expanded on in Figure 5.6. The basic knowledge model is defined as per Figure 5.7.

The following free sources provide information



Figure 5.7: Basic knowledge model of the topic of interest

1. A maritime enthusiast tweets "The cargo vessel *Marcel* has been spotted heading towards Turkey."
2. The AIS field for activity of *Marcel* is "Fishing" and destination is "Italy."
3. The kinematic analysis of the behaviour of *Marcel* suggests that it is a fishing vessel en route to a port in Turkey

All three sources are obviously conflicting - if we try to combine them directly, we either end up with inconsistent knowledge or converge back to the model. Let us assume that, for the time being, we do not know the reliability of any of the sources.

First, we consider only sources 1 and 2. The corresponding graphs are in Figures 5.8 and 5.9.



Figure 5.8: Graph G_1 generated from the first source



Figure 5.9: Graph G_2 generated from the second source

These two sources are in full conflict - there is no possible method of representing them on a single graph without information loss. As such, we use a disjunctive fusion approach as discussed in the previous section. The outcome of this combination process is shown in Figure 5.10.

Note that the following fusion result is questionable as it allows the existence of statements not supported by any source (what happens here is that we have two variables that are not independent but no means of modelling their interdependence in a single graph). Whether this triple is acceptable or not is largely dependent on our interpretation of the partial reliability of the original statements (i.e. whether we assume that the statement "*Marcel* is a cargo vessel en route to Turkey" can partially support a statement "*Marcel* is a cargo vessel en route to Italy"). While this makes sense in this scenario, in general, it is an issue that needs to be addressed at the fusion strategy level.



Figure 5.10: Graph $G_1 \otimes G_2$ generated using disjunctive fusion of both concepts in G_1 and G_2

If we want to make sure that every statement which can be generated from the graph is fully supported by at least one of the sources, we can instead find the most specific generalisation for

one of the elements of the triple (i.e. in this case, defaulting one of the two concepts to the type used one in the model). In this case, some information loss still occurs as we have no knowledge anymore regarding the specialisation of *Vessel* in the former case or *Destination* in the latter.

We can now combine this with G_3 . It is reasonable to use either disjunctive or conjunctive fusion. In the former case we would obtain $G_1 \otimes G_2 \otimes G_3 = G_1 \otimes G_2$. As per the fusion strategy we laid out in the earlier section we use conjunctive fusion, yielding $G_1 \otimes G_2 \otimes G_3 = G_{123}$ as shown in Figure 5.11 which is arguably the most likely outcome



Figure 5.11: Graph G_{123} , the result of combination of the three sources $G_1 \otimes G_2 \otimes G_3$ using an appropriate fusion strategy

5.4 Partial reliability ratings

As discussed earlier, a significant issue with modelling uncertain information using the approach used in literature is that every source has to be considered either entirely reliable or not, with no sliding scale.

The approach proposed in the previous section, however means that reliability cannot be simply modelled as a probability of a source of information being entirely truthful. A simple method is to extend the probabilistic approach so that if a source is reliable, some of its information can be true.

Drawing from possibility theory, we can use two numbers to describe a source's reliability, a possibility of some of its information being true, and the necessity of all its information being true. Note that although we use the same nomenclature as that used in possibility theory, the meaning of possibility and necessity is conceptually similar but inherently different. The reason why it cannot be considered a *real* possibility distribution is that there are at least three possible reliability levels of a source: it can be fully reliable (all the information is true), partially reliable (all the information could be true) or not reliable (no information provided) at all. However, by attaching possibility and necessity values, we obtain the distribution on two variables only. Consider we have a graph G provided by the source S , its reliability is modelled by the variable R . For convenience let us denote $Pos(R) \equiv Pos(S) \equiv Pos(G)$. Hence we may have a possibility distribution on $\{R, \bar{R}\}$. In order to model the three possible states, we choose to interpret $Pos(R)$ as the probability that the source is partially reliable and $Nec(R)$ as the probability that the source is fully reliable. While other axioms of possibility do hold, this interpretation is not coherent with the possibility theory.

To illustrate this, let us go back to the previous example but include possibility based reliability ratings, such that $Pos(G_1) = 0.7, Pos(G_2) = 0.6$ and $Pos(G_3) = 0.7$. The meaning of this is that we believe, with a degree of 0.7, that at least some of the information provided by the source generating the graph G_1 could be truthful. Alternatively, this can be interpreted as a belief that the source is fully unreliable with a likelihood of 0.3, which is consistent with the duality of possibility theory, i.e. $Pos(G_1) = 1 - Nec(\bar{G}_1)$.

The result (Table 5.1) is a probability distribution on possible outcomes which is similar to earlier work.

Graph G	P(G)	Concepts
$G_1 \otimes G_2 \otimes G_3$	0.294	Fishing, Turkey
$G_1 \otimes G_2$	0.126	{Fishing, Cargo} , {Turkey, Italy}
$G_1 \otimes G_3$	0.126	{Fishing, Cargo}, Turkey
$G_2 \otimes G_3$	0.196	Fishing, {Turkey, Italy}
G_1	0.084	Cargo, Turkey
G_2	0.054	Fishing, Italy
G_3	0.084	Fishing, Turkey
M	0.036	Vessel, Destination

Table 5.1: Semantic evidence combination results

Interestingly we can draw similarities between this and belief function theory [110] as now we can obtain belief and plausibility values for any statement. For example the belief in the statement that *Marcel* is en route to Turkey can be quantified as the sum of evidence supporting this belief: $Bel(Turkey) = P(G_1 \otimes G_2 \otimes G_3) + P(G_1 \otimes G_3) + P(G_1) + P(G_3) = 0.588$. Similarly the plausibility measure can be obtained by summing the evidence not contradicting this statement $Plaus(Turkey) = 1 - P(G_2) = 0.946$

Now let us revisit the concept of necessity introduced earlier. Consider we obtain a new piece of evidence G_4 , which states that *Marcel* is indeed a fishing vessel and it comes from a more reliable source so that we have $Pos(G_4) = 0.8$ but also $Nec(G_4) = 0.3$.



Figure 5.12: Graph G_4 with $Nec(G_4) = 0.3$ and $Pos(G_4) = 0.8$

To model its reliability let us define two identical graphs G_4^N and G_4^Π which have respective reliabilities 0.3 and $0.8 - 0.3 = 0.5$. For simplicity, let us only combine this new piece of evidence with G_2 .

G_4^N is now assumed to be entirely true, and as such, all information contained within has to be included in the fusion result.

As such, when performing the combination $G_4^N \otimes G_2$ only information from G_2 , which is not in conflict with information from G_4 , remains. Conversely, the two graphs are treated equally when combining $G_4^\Pi \otimes G_2$ and the same fusion rules as the earlier example are used. Both cases are shown in Figure 5.13

Finally we obtain the distribution of results displayed in Table 5.2.




Figure 5.13: Graphs resulting from combination of G_2 and G_4 assuming that G_4 is necessary (5.13a) and that G_4 is possible (5.13b)

Graph G	P(G)	Concepts
$G_4^N \otimes G_2$	0.21	Fishing, Turkey
$G_4^\Pi \otimes G_2$	0.35	{Cargo, Fishing}, Turkey
G_4	0.24	Fishing, Destination
G_2	0.14	Cargo, Turkey
M	0.06	Vessel, Destination

Table 5.2: Semantic evidence combination results including necessity ratings

5.5 Comparison of the two models

In this section, we focus on the difference between the model for dealing with uncertainty from [107] and the one proposed in the previous two sections, in particular highlighting the different approaches towards dealing with conflict and incoherence. Let us consider the fusion example provided in that paper with the following three information sources combined consecutively. For conciseness, only the model is shown in Figure 5.14 and the pieces of information are represented as a list of concepts.




Figure 5.14: Simplified model for the maritime surveillance scenario from [107], the geographical location has been removed

The following information sources are consistent with the model Ship, Flag, Activity:

1. G_1 : Ship: SEAWOLF, Flag, Piracy, reliability: 0.8
2. G_2 : Dhow, Flag: Somali, Suspicious Activity, reliability: 0.7
3. G_3 : Ship: SEAWOLF, Flag, Fishing, reliability: 0.2

where $Dhow \sqsubset Ship$; $Piracy \sqsubset Suspicious\ Activity \sqsubset Activity$; $Fishing \sqsubset Activity$

Using the approach described in Section 5.3.2 the combination results are: $G_1 \otimes G_2$: Dhow: SEAWOLF, Flag: Somali, Piracy and $G_1 \otimes G_2 \otimes G_3 = G_1 \otimes G_3 = G_2 \otimes G_3 = \diamond$, where \diamond rep-

Graph G	P(G)	Concepts
$G_1 \otimes G_2$	0.448	Dhow: SEAWOLF, Flag: Somali, Piracy
G_1	0.192	Ship: SEAWOLF, Flag, Piracy
G_2	0.112	Dhow, Flag: Somali, Suspicious Activity
G_3	0.012	Ship: SEAWOLF, Flag, Fishing
\diamond	0.188	-
M	0.048	Ship, Flag, Activity

Table 5.3: Probability distribution on outcomes of the maritime surveillance scenario using sources 1 - 3 and the original model from [107]

Graph G	P(G)	Concepts
$G_1 \otimes G_2 \otimes G_3$	0.112	Ship: SW, Flag:Som., {Piracy, Fishing}
$G_1 \otimes G_2$	0.448	Dhow: SW, Flag: Som., Piracy
$G_1 \otimes G_3$	0.028	Ship: SW, Flag, {Piracy, Fishing}
$G_2 \otimes G_3$	0.048	Dhow: SW., Flag: Som., {Susp. Act., Fish.}
G_1	0.192	Ship: SW, Flag, Piracy
G_2	0.112	Dhow, Flag:Somali, Susp. Act.
G_3	0.012	Ship: SW, Fishing
M	0.048	Ship, Flag, Activity

Table 5.4: Probability distribution on outcomes of the maritime surveillance scenario using the set-based uncertainty model for partial conflict resolution

resents the incoherent knowledge. The probability distribution on the possible outcomes ($G_1 \otimes G_2, G_1, G_2, G_3, \diamond, M$) is straightforward to compute and is shown in Table 5.3.

When the method proposed in this chapter is used instead we obtain $G_1 \otimes G_2 \otimes G_3$ Dhow: SEAWOLF, Flag:Somali, {Piracy, Fishing}, $G_1 \otimes G_3$: Ship: SEAWOLF, Flag, {Piracy, Fishing} and $G_2 \otimes G_3$: Dhow: SEAWOLF, Flag: Somali, {Suspicious Activity, Fishing}.

Treating the reliability ratings provided as possibility values, we obtain the distribution of results shown in Table 5.4

Let us now consider the degree of support for a hypothesis that the ship SEAWOLF is involved in pirate activities. Using the original method this statement is supported by graphs $G_1 \otimes G_2$ and G_1 , the probabilities of which sum to $P(\text{SEAWOLF, Piracy}) = 0.64$. However, by using the method described earlier, we can instead obtain belief and plausibility values akin to the belief function theory. The belief in this hypothesis is directly supported by graphs $G_1 \otimes G_2$ and only contradicted by G_3 , resulting in $Bel(\text{SEAWOLF, Piracy}) = 0.64$ and $Plaus(\text{SEAWOLF, Piracy}) = 0.988$. This implies that the actual likelihood of SEAWOLF being involved in piracy could be higher than that suggested by the original model.

The key difference in the outcome is that in the original model, all the incoherent knowledge is ignored or treated as a situation outside of the frame of discernment, whereas the approach proposed in this chapter allows us to extract some information from it. It is important to note that this particular example scenario is not optimal to demonstrate the performance of the model as the third source exclusively contains information that is either conflicting or already included in

the other two sources. However, it is reasonable that a source deemed incoherent and excluded from the combination process could carry valid and useful information (e.g. consider a scenario in which the ship name is not included in G_1). In that case, the original approach of disregarding information that is partially conflicting as incoherent would result in a loss of useful data.

5.6 Modal uncertainty

Another concept that we want to briefly discuss, is how to deal with semantic descriptions of uncertainty provided as part of the source (i.e. self-assessed). Here we propose a simple method of doing it in the graph itself, rather than treating it as part of reliability rating.

- We allow for every concept label, defined as a pair of the conceptual type and a marker, to have a modal parameter $M(A) \in (0, 1)$ associated with it.
- A concept A which does not have $M(A)$ associated with it is treated as if $M(A) = 1$
- Sets of concept labels $l_G(c)$ are allowed to be subsumptive of as long as the following is true: given concept labels A and $B \in l_G(c)$, if B is a specialisation of A ($A > B$) then $M(A) \geq M(B)$

This can be formalised as follows: a concept node is labelled with a set of 3-tuples: $l_G = \{(\text{type}(c), \text{marker}(c), M(c))\}$.

To illustrate this, consider the following statement, which can be modelled using the TFL station vocabulary: "The vessel X has been spotted, most likely heading towards Turkey line, possibly to Iskenderun port". Let us model the "strength" of these modal statements (*most likely* and *possibly*) as 0.8 and 0.3 respectively. Therefore the uncertain concept, regarding the destination of the vessel, can be modelled as follows: $\{[\text{Destination}:*], [\text{Turkey}:*/0.8], [\text{Turkey}:\text{Iskenderun}/0.3]\}$.

In order to develop an appropriate fusion strategy, it is reasonable to consider the concept of decomposition. It makes intuitive sense that this information should be equivalent to the following three statements

1. The vessel has some destination (possibility = 1)
2. The vessel is heading towards Turkey (possibility = 0.8)
3. The vessel is heading to the Iskenderun port in Turkey (possibility = 0.3)

These three can then be combined as per the previous section to obtain a probability distribution on the three possible concepts. Although some loss of expressiveness occurs, this shows that the two approaches are inherently compatible. Furthermore, an issue that needs to be considered here is that of double-counting of evidence. In particular, should this sentence be interpreted as a set of independent pieces of information - where the phrase "possibly to Iskenderun" supports the hypothesis that the vessel is heading to Turkey, or should it be interpreted that "The vessel (possibly) is heading to Iskenderun *given* that it is heading to Turkey"? Interpreting this information as three independent statements as above and using it to obtain a probability distribution will result

in a result where evidence may be double-counted. Alternatively, a set of probability assignments can be obtained through an α -cut based method

We consider two modal statements to be non-conflicting if all the concept labels included in one of them are either included or are specializations or generalizations of all the concept labels included in the other statement. For example the statement $\{[Destination:], [Turkey:*/0.6], [Turkey:Iskenderun/0.5]\}$ is not conflicting with the one used in the previous example, whereas $\{[Destination:], [Italy:*/0.4]\}$ is. Combination of two non-conflicting modal concepts is relatively straightforward. As they are essentially two possibility distributions an appropriate t-conorm can be used. One simple method is to simply use the maximum operator, e.g:

$$(5.5) \quad \begin{aligned} & \{[Destination:], [Turkey:*/0.8], [Turkey:Iskenderun/0.3]\} \\ \otimes & \{[Destination:], [Turkey:*/0.6], [Turkey:Iskenderun/0.5]\} \\ = & \{[Destination:], [Turkey:*/0.8], [Turkey:Iskenderun/0.5]\} \end{aligned}$$

The combination of conflicting modal sources is significantly more complex. Some existing tools for conflict resolution could be used, but it is important to bear in mind that in the context of graph-based fusion, there exists a significant risk of loss of tractability. Some of the potential methods are briefly outlined. The example $\{[Destination:], [Turkey:*/0.8], [Turkey:Iskenderun/0.3]\} \otimes \{[Destination:], [Italy:*/0.4]\}$ is used for illustration.

1. **Keep everything:** a simple method that avoids any information loss but may cause scalability or tractability issues. All the conflicting information is retained, as was done in the previous section. This would result in $\{[Destination:], [Turkey:*/0.8], [Turkey:Iskenderun/0.5], [Italy:*/0.4]\}$. This can then be used to generate possible graphs and obtain a probability distribution, as discussed earlier
2. **Retain non-conflicting information only** - a significant loss of information, but potentially justified due to the logical inconsistency. The output would simply be $[Destination:]$
3. **In case of conflict, retain information with greater modality** - the conflicting subsets are identified, and the one with the greatest modality is retained. In this case this would result in $\{[Destination:], [Turkey:*/0.8], [Turkey:Iskenderun/0.5]\}$. The behaviour is questionable when the difference in modalities is small and would be strongly affected by second-order uncertainty (i.e. how certain are we regarding modality assessments)

5.7 Fusion strategies

A major downside of the fusion strategy presented in Section 5.3.2 is that it is not insensitive to the order in which information is being combined.

As such an alternative fusion strategy is proposed, based on the notions of specificity and consistency. First of all let us recall the notion of consistency discussed at length in the previous chapter. Logically a consistent theory is one that does not entail a contradiction. Here, we can consider the degree of inconsistency or conflict between the two graphs sharing a knowledge model to be proportional to the number of incompatible concept nodes. Equation 5.6 defines an inconsistency measure as a ratio of inconsistent concept nodes to all concept nodes in the knowledge model \mathcal{M}

$$(5.6) \quad \phi(G_1, G_2) = \frac{|\{c : c \in C_1 \wedge c \circledast c_2 = \emptyset \forall c_2 \in C_2\}|}{|C_{\mathcal{M}}|}$$

Furthermore, we can consider the notion of (non)specificity of a conceptual graph with embedded uncertainty. The simplest definition is the number of possible dogmatic (non-uncertain) graphs which are consistent with the graph of interest and only include information in the graph of interest.

$$(5.7) \quad NS(G) = |\{G^* : G^* \in \Pi(G) \wedge G^* \sqsubseteq G\}|$$

With this in mind we can consider the example from Section 5.4. Table 5.5 shows all the possible fusion candidate graphs - consistent with the knowledge model and with at least one of the sources and not including any information not included in the knowledge base, and their corresponding average consistency and nonspecificity values. Average consistency is the average of consistencies of the candidate graph G_J and each member of the knowledge base G_i . The set of all graphs in the knowledge base is denoted as $\mathcal{G} = \{G_1, G_2, G_3\}$.

The purpose of the fusion procedure is to find a graph that satisfied two properties - it needs to have the least amount of conflict possible with all the elements of the knowledge base, and it needs to be as specific as possible. The degree of former can be expressed by the average value of inconsistency or conflict $\frac{1}{|\mathcal{G}|} \sum_{i \in \mathcal{G}} \phi(G_i, G_J)$, whereas the latter by the nonspecificity measure $NS(G)$.

G_J concepts	$\frac{1}{ \mathcal{G} } \sum_{i \in \mathcal{G}} \phi(G_i, G_J)$	$NS(G_J)$	$\frac{\sum_{i \in \mathcal{G}} \phi(G_i, G_J)}{ \mathcal{G} \times NS(G_J)}$
Fishing, Italy	0.5	1	0.5
Fishing, Turkey	0.66	1	0.66
Cargo, Italy	0.33	1	0.33
Cargo, Turkey	0.5	1	0.5
Fishing, {Turkey, Italy }	0.833	2	0.4166
Cargo, {Turkey, Italy }	0.66	2	0.33
{Fishing, Cargo}, Turkey	0.833	2	0.4166
{Fishing, Cargo}, Italy	0.66	2	0.33
{Fishing, Cargo}, {Turkey, Italy}	1	4	0.25

Table 5.5: Nonspecificity and average consistency values for the family of candidate fusion outcome graphs

Furthermore, we may consider some additional constraints, e.g. the graph should not be in full conflict with any element of the knowledge base, or we may want to limit ourselves only to dogmatic graphs (nonspecificity of 1). At this stage, we choose not to worry about any of these additional constraints, as they are essentially another filter on allowed fusion results and, as such trivial to implement. In order to decide on the best fusion outcome, we want to have some ranking function, maximising the consistency and minimising the nonspecificity. One possible example of such a function is the following:

$$(5.8) \quad R = \frac{\sum_{i \in \mathcal{G}} \phi(G_i, G_J)}{|\mathcal{G}| \times NS(G_J)}$$

and the graph with maximum value of R is the selected fusion outcome. Note that this is conceptually similar to the notion of Jaccard's similarity.

Based on Table 5.5, we can clearly see that the preferred solution from the previous example is selected again.

5.8 Reasoning and association

The basic conceptual graphs (BCG) formalism used throughout this chapter expresses knowledge through conjunctions of positive information, does not being suitable for general reasoning. This holds true for the uncertain graphs we have discussed in this section. Some alternatives for reasoning in knowledge bases have been proposed, such as the polarised graphs (PG) from Chein, and Mugnier [111] as part of their graph-based knowledge model. Sowa's full conceptual graphs (FG) implement negation through *boxes* and co-reference through *lines*. This allows capturing any First-Order Logic (FOL) formula with the existential qualifier, negation and conjunction.

It is noteworthy that even though the uncertain basic conceptual graphs do not capture notions such as if-then rules, they can still be expressed as FOL formulas. Having a set of possible values at a conceptual node corresponds to the disjunction operator. For example consider the graph $G_1 \otimes G_2$ from Figure 5.10. In natural language it can be interpreted as *Marcel is either a cargo vessel or a fishing vessel which has destination of either Italy or Turkey*. This statement can be written in FOL as:

$$\exists x (Fishing(Marcel) \vee Cargo(Marcel)) \wedge hasDestination(Marcel, x) (\wedge Italy(x) \vee Turkey(x))$$

Reasoning with FOL is relatively well-understood, and the above formula could be combined with a more general FOL knowledge base, containing formulas that cannot be captured in UGs.

5.9 Relation to belief function framework

There are two main ways in which conceptual graphs can be considered within the belief function framework. As mentioned in the previous section, UGs can be naturally thought of as belief functions themselves. On the other hand, we can exploit the evidential networks discussed in detail in earlier chapters to implement the fusion and reasoning and see individual graphs as joint valuations.

5.9.1 Uncertain graphs as belief functions

The link between conceptual graphs when considered in an uncertain context and the belief function framework was first brought up by Fossier *et al.* [107] when considering a probability distribution on a set of graphs. However, whilst this is similar to the evidential distribution on powerset of hypotheses, a major difference is that the combination of two graphs is considered conjunction, rather than a disjunction as in the belief function framework.

In section 5.4 it was discussed how graphs with reliability ratings could be used to obtain a probability distribution similar to earlier work. However, as discussed in the previous section, each uncertain concept in a UG (or by extension uncertain relation, although they have not been discussed) can be considered a union of several possible concept nodes (or more precisely, the label of each concept node is considered the union of several possible labels). As such, the probability distribution on the possible graphs can be considered to be a probability distribution on sets of outcomes - i.e. a belief function. This result was used in section 5.4 to obtain belief and plausibility values on hypotheses of interest. Now we look at this more generally.

An uncertain graph includes a set of concept nodes and relation nodes, labelled by sets of hypotheses. As such, each graph can be considered to contain sets (C_G, R_G) of sets (sets of labels for each conceptual node and for every relation node). Consider a biclique graph G with concept nodes C_1 and C_2 and a relation node R , labelled with sets of labels L_{C_1} , L_{C_2} and L_R respectively. Now for each node we can consider a random variable defined on the frame of the vocabulary and constrained so that its true value lies within the set described by the label. Let these be denoted by X_{C_1} , X_{C_2} and X_R . Our knowledge about these can be described by *dogmatic, focused* belief masses $m(X_{C_1} = L_{C_1}) = 1$ etc. As such the overall knowledge of labels (not the structure) of the graph can be described by the focal set $X_{C_1} = L_{C_1} \times X_{C_2} = L_{C_2} \times X_R = L_R$. If needed the random variables for context labels can be split into two, to separate the conceptual type and the marker. If it is the case, the focal set becomes $X_{C_1} = L_{C_1} \times X_{C_1}^* = L_{C_1}^* \times X_{C_2} = L_{C_2} \times X_{C_2}^* = L_{C_2}^* \times X_R = L_R$ with * superscript denoting the marker.

Now let us recall the results of the fusion of uncertain graphs. The results of the fusion process discussed earlier result in a probability distribution over a set of graphs. However as we have just shown, each graph can be considered to be a focal set. Thus this result can be seen as belief a mass distribution. Consider the results from Table 5.1. These can be rewritten as shown in Table 5.6, where C_1 and C_2 correspond to the `Vessel:Marcel` and `Destination:*` nodes respectively and **F**, **C**, **T** and **I** denote Fishing, Cargo, Turkey and Italy respectively. The set $\Omega_V = \{A \in T_C : A \subseteq \text{Vessel}\}$ and $\Omega_D = \{B \in T_C : B \subseteq \text{Destination}\}$ Finally the labels are not included as they are always "Marcel" for C_1 and * for C_2 .

This brings us to a natural extension - if we can frame a large portion of the knowledge within a conceptual graph in a belief function framework, why would we not exploit that framework for a combination of evidence, too? An obvious problem with this approach - and with representation as belief functions in general - is that of scalability - on one hand, we deal with a large number of variables, but on another with a possibly unconstrained frame of discernment for each variable. Parts of this problem can be alleviated using valuation networks and exploiting the type hierarchy

Graph G	$m(\mathcal{F})$	focal set \mathcal{F}
$G_1 \otimes G_2 \otimes G_3$	0.294	$C_1 = \mathbf{F} \times C_2 = \mathbf{T}$
$G_1 \otimes G_2$	0.126	$C_1 = \{\mathbf{F}, \mathbf{C}\} \times C_2 = \{\mathbf{T}, \mathbf{I}\}$
$G_1 \otimes G_3$	0.126	$C_1 = \{\mathbf{F}, \mathbf{C}\} \times C_2 = \mathbf{T}$
$G_2 \otimes G_3$	0.196	$C_1 = \mathbf{F} \times C_2 = \{\mathbf{T}, \mathbf{I}\}$
G_1	0.084	$C_1 = \mathbf{C} \times C_2 = \mathbf{T}$
G_2	0.054	$C_1 = \mathbf{F} \times C_2 = \mathbf{I}$
G_3	0.084	$C_1 = \mathbf{F} \times C_2 = \mathbf{T}$
M	0.036	$C_1 = \Omega_V \times C_2 = \Omega_D$

Table 5.6: Semantic evidence combination results represented as a belief function

to reduce the size of the frame of discernment. Furthermore, local computation axioms mean that the marginal over the entire frame does not necessarily need to be computed - as long as we can narrow down the set of variables of interest (the variables the values of which we need to infer).

5.9.2 Representation as evidential networks

Valuation networks, discussed at length in earlier chapters are a powerful tool for reasoning under uncertainty. One implementation is that of evidential networks, where the underlying formalism is that of belief functions. One advantage of using evidential networks in context of




Figure 5.15: Evidential network for possible implementation of trivial type hierarchy

conceptual graphs is that it allows a form of implementation of type hierarchy. It is possible to use several related variables with different frames of discernment to minimize the size of each individual frame. For a trivial example consider a type hierarchy with two top level concepts A and B , entailing sub-types $\{A_1 \dots A_n\}$ and $\{B_1 \dots B_n\}$ respectively. Now we can model this as an evidential network with three variables X , A and B , defined on frames $\{A, B\}$, $\{A \dots A_n, \neg A\}$ and $\{B \dots B_n, \neg A\}$ respectively. The logic governing the relations between these is straightforward: if and only if $X = A$ then $A = \{A_1 \dots A_n\}$ and $B = \neg B$. Conversely $X = B \iff A = \neg A \iff B = \{B_1 \dots B_n\}$. See visualization in Figure 5.15

Now let us recall the examples relating to the vessel *Marcel*. The concepts and relation in the model can be illustrated by the evidential network in Figure 5.16. Again, it is important to keep in mind that this kind of evidential network does not encode any information about the structure of the graph. In order to maintain the open world assumption an "other" element has been included in the frame of discernment for every variable. The valuation T_{c1} encodes the type hierarchy relating $X_{C1} = Fishing$ to the set $X_{C1F} = \neg(\neg Fishing)$

Let us recall the three pieces of information relating to *Marcel's* type and destination. These can be represented as focal sets \mathcal{F}_1 , \mathcal{F}_2 and \mathcal{F}_3 (including only the information not already contained

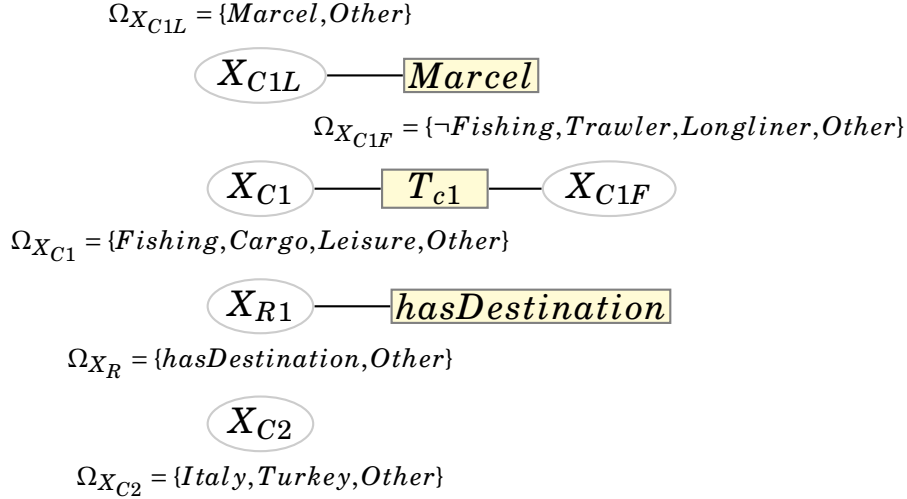


Figure 5.16: Valuation network showing the concepts and relations in the model for Marcel analysis example

in the model) and using the shorthand notation for conceptual types defined earlier:

$$G_1 : \mathcal{F}_1 = \{X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{T}\}$$

$$G_2 : \mathcal{F}_2 = \{X_{C1} = \mathbf{F} \times X_{C2} = \mathbf{I}\}$$

$$G_3 : \mathcal{F}_3 = \{X_{C1} = \mathbf{F} \times X_{C2} = \mathbf{T}\}$$

Now we can construct the valuations to be added to the evidential network based on the reliability ratings of these sources. However, recall that the reliability model we use distinguishes between the possibility and necessity of graphs - i.e. a graph is possible when any of its information is true and necessary when all its information is true. As such, bringing in the possibility values from that example: $Pos(G_1) = 0.7$, $Pos(G_2) = 0.6$, $Pos(G_3) = 0.7$ we have the set of valuations:

$$m_1 \left(\left\{ \begin{array}{l} X_{C1} = \mathbf{C} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{T} \end{array} \right\} \right) = 0.7$$

$$m_1 \left(\{ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \Omega_{X_{C2}} \} \right) = 0.3$$

$$m_2 \left(\left\{ \begin{array}{l} X_{C1} = \mathbf{F} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{I} \end{array} \right\} \right) = 0.6$$

$$m_2 \left(\{ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \Omega_{X_{C2}} \} \right) = 0.4$$

$$m_3 \left(\left\{ \begin{array}{l} X_{C1} = \mathbf{F} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{T} \end{array} \right\} \right) = 0.7$$

$$m_3 \left(\{ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \Omega_{X_{C2}} \} \right) = 0.3$$

Adding to the evidential network from Figure 5.16 and computing the joint on $X_{C1} \times X_{C2}$ yields the mass function shown in Table 5.7. By inspection, it is very similar to the results in Table 5.6. The main difference is that the results are overall less informative - whereas the assumption that "some of the information provided by the source holds true" was only taken into account in case of

conflict, here it is inherently included in the way the original focal sets are encoded, and therefore the focal sets of the joint in Table 5.6 are more specific than the ones in Table 5.7.

$m_J(\mathcal{F})$	focal set \mathcal{F}
0.294	$C_1 = \mathbf{F} \times C_2 = \mathbf{T}$
0.196	$C_1 = \Omega_V \times C_2 = \mathbf{T}$
0.126	$C_1 = \mathbf{F} \times C_2 = \Omega_D$
0.126	$C_1 = \{\mathbf{F}, \mathbf{C}\} \times C_2 = \{\mathbf{T}, \mathbf{I}\}$
0.084	$\left\{ \begin{array}{l} X_{C_1} = \mathbf{C} \times X_{C_2} = \Omega_{X_{C_2}} \\ X_{C_1} = \Omega_{X_{C_1}} \times X_{C_2} = \mathbf{T} \end{array} \right\}$
0.084	$\left\{ \begin{array}{l} X_{C_1} = \mathbf{F} \times X_{C_2} = \Omega_{X_{C_2}} \\ X_{C_1} = \Omega_{X_{C_1}} \times X_{C_2} = \mathbf{T} \end{array} \right\}$
0.054	$\left\{ \begin{array}{l} X_{C_1} = \mathbf{F} \times X_{C_2} = \Omega_{X_{C_2}} \\ X_{C_1} = \Omega_{X_{C_1}} \times X_{C_2} = \mathbf{I} \end{array} \right\}$
0.036	$C_1 = \Omega_V \times C_2 = \Omega_D$

Table 5.7: Evidential network joint distribution

To summarise this comparison, let us also consider the case where all the information within a specific graph holds true. In this case, the focal set induced by the graph is the Cartesian product of the corresponding elements only, rather than the set of all members of the joint frame resulting in a non-empty intersection. Therefore recall the second example from Section 5.4. Here we consider a slightly modified version of that example, where the graph considered provides the same information as G_1 . This source of information has $Nec(G_1) = 0.3$ and $Pos(G_1) = 0.8$, which are interpreted as the graph providing all its information with probability 0.3, some of its information with probability 0.5 and no information with probability 0.2. This corresponds to a mass function with three focal sets with corresponding masses:

$$\begin{aligned}
m_4(\{ X_{C_1} = \mathbf{C} \times X_{C_2} = \mathbf{T} \}) &= 0.3 \\
m_4\left(\left\{ \begin{array}{l} X_{C_1} = \mathbf{C} \times X_{C_2} = \Omega_{X_{C_2}} \\ X_{C_1} = \Omega_{X_{C_1}} \times X_{C_2} = \mathbf{T} \end{array} \right\}\right) &= 0.5 \\
m_4(\{ X_{C_1} = \Omega_{X_{C_1}} \times X_{C_2} = \Omega_{X_{C_2}} \}) &= 0.2
\end{aligned}$$

Combining the above with m_2 , prior to normalization we obtain the distribution shown in 5.8. A key difference between this approach and the other fusion strategies for conceptual graphs is that some probability is assigned to the inconsistent knowledge and as such needs to be normalized.

The existence of this conflict or inconsistency seems counterintuitive given that the purpose of this entire work was to avoid dismissing knowledge considered inconsistent. Furthermore, we know that this result would have been different if one of the earlier fusion strategies (not involving valuation networks) were used. For comparison, with the same sources of information, we would have obtained the results in Table 5.9. The two main differences are associated with the combination $G_4^N \otimes G_2$, which yields G_4 using conceptual graphs and \emptyset using Dempster's combination and $G_4^\Pi \otimes \neg G_2$, which should yield simply G_4 , but results in a more relaxed focal set under the belief function interpretation.

$m_J(\mathcal{F})$	focal set \mathcal{F}
0.18	\emptyset
0.3	$\left\{ \begin{array}{l} X_{C1} = \mathbf{F} \times X_{C2} = \mathbf{T} \\ X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{I} \end{array} \right\}$
0.2	$\left\{ \begin{array}{l} X_{C1} = \mathbf{C} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{T} \end{array} \right\}$
0.12	$\left\{ \begin{array}{l} X_{C1} = \mathbf{F} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{I} \end{array} \right\}$
0.12	$\left\{ X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{T} \right\}$
0.08	$C_1 = \Omega_V \times C_2 = \Omega_D$

Table 5.8: Belief distribution with one of the masses having a degree of necessity

These issues largely stem from whether we interpret conceptual graphs as a conjunction or as a disjunction. The main notion introduced in this chapter was that if necessary, we may consider only some information contained in the graph if they conflict. When we rewrite this as a belief function, we may either consider it a union of concepts or an intersection, corresponding to the notions of possibility and necessity, respectively. In other words, the notion of possibility and necessity is only relevant to conceptual graphs at the moment of combination, whereas when they are represented as valuation network, it affects the formulation of the focal sets. As such, each conceptual graph in this example can induce at least two different focal sets. This suggests that it may be more reasonable to change the approach proposed here and rather focus on a single way of converting conceptual graphs into belief functions and resolve the possibility/necessity of uncertain graphs through an appropriate combination rule.

One possible solution to the issues of conflict is the selection of another combination rule. The conjunctive-disjunctive combination rule, where inconsistent intersections are combined distinctively, is one such solution. This set of results is shown in Table 5.10. Note that the only difference between this and using Dempster's rule of combination is the focal set to which the mass of 0.18 is assigned, corresponding to combination $G_4^N \otimes G_2$.

$m_J(\mathcal{F})/P(G_J)$	G_J concepts	\mathcal{F} (conjunction-based)
0.5	$G_4 : \text{Cargo, Turkey}$	$C_1 = \mathbf{F} \times C_2 = \mathbf{T}$
0.3	$G_4^\Pi \otimes G_2 : \{ \text{Fishing, Cargo}, \text{Turkey, Italy} \}$	$\left\{ \begin{array}{l} X_{C1} = \mathbf{F} \times X_{C2} = \mathbf{I} \\ X_{C1} = \mathbf{F} \times X_{C2} = \mathbf{T} \\ X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{I} \\ X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{T} \end{array} \right\}$
0.12	$G_2 : \text{Fishing, Italy}$	$X_{C1} = \mathbf{F} \times X_{C2} = \mathbf{I}$
0.08	M	$C_1 = \Omega_V \times C_2 = \Omega_D$

Table 5.9: Belief distribution with one of the masses having a degree of necessity with fusion performed using conceptual graphs

$m_{\mathcal{J}}(\mathcal{F})$	focal set \mathcal{F}
0.18	$\left\{ \begin{array}{l} X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{T} \\ X_{C1} = \mathbf{F} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{I} \end{array} \right\}$
0.3	$\left\{ \begin{array}{l} X_{C1} = \mathbf{F} \times X_{C2} = \mathbf{T} \\ X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{I} \end{array} \right\}$
0.2	$\left\{ \begin{array}{l} X_{C1} = \mathbf{C} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{T} \end{array} \right\}$
0.12	$\left\{ \begin{array}{l} X_{C1} = \mathbf{F} \times X_{C2} = \Omega_{X_{C2}} \\ X_{C1} = \Omega_{X_{C1}} \times X_{C2} = \mathbf{I} \end{array} \right\}$
0.12	$\left\{ X_{C1} = \mathbf{C} \times X_{C2} = \mathbf{T} \right\}$
0.08	$C_1 = \Omega_V \times C_2 = \Omega_D$

Table 5.10: Belief distribution with one of the masses having a degree of necessity with fusion performed using Dubois-Prade conjunctive-disjunctive rule

5.9.3 Existential valuations

The above approach is reasonable when we consider valuation networks representation of conceptual graphs in context of the knowledge model for fusion proposed before. In other words, we are not concerned neither with the structure of the graph (which cannot be represented by valuation networks anyway) nor the existence of certain concepts and relations.

5.10 Markov logic networks

The same problem can be analysed using Markov Logic Networks. They have been discussed briefly in Chapter 2. To recall, they are a method of adding probabilistic reasoning to first-order logic - a problem we have already discussed in this chapter. Typically Markov Logic Networks handle uncertain rules with crisp evidence. Although the weights can be provided externally, in most applications they are learned. The evidence is provided in the form of a set of groundings assumed to be true. However, on a more general level, the difference between a "rule" and a "body of evidence" is largely semantic. Here we consider a situation where we treat the bodies of evidence as uncertain rules and assign weights based on log-probability. This approach can be problematic if weights for some rules are learned, which is discussed by Papai et al. [112]. In this case, however, as no learning of underlying distribution takes place, we can assume this to be suitable. The weights w_i are therefore obtained using log-odds as per Equation 5.9:

$$(5.9) \quad w_i = \ln \left(\frac{P(r_i)}{1 - P(r_i)} \right)$$

where $P(r_i)$ is the probability that the rule i with weight w_i holds. In the case the "rules" are actually provided by sources, the value of $P(r)$ is obtained through a probability transformation. For a binary frame of discernment $P(r) = 1.5m(r)$, as the weight on the universal set is redistributed equally across the two singletons.

We can define the MLN for the Marcel type-destination problem example as follows. The knowledge base consists of the following predicates: *Cargo(vessel)*, *Fishing(vessel)*, *Other(vessel)*,

$hasDestination(vessel, location!)$, where "!" operator denotes mutual exclusivity and exhaustively constraint on the variable. The domain for location is {Italy, Turkey, OtherD}. The weighted rules are shown in Table 5.11, where \cdot denotes a "hard" constraint.

#	Rule	Weight
1	$Fishing(Marcel) \cup hasDestination(Marcel, Italy)$	1.32
2	$Cargo(Marcel) \cup hasDestination(Marcel, Turkey)$	1.73
3	$Fishing(Marcel) \cup hasDestination(Marcel, Turkey)$	1.73
4	$Cargo(x) \implies (\neg Fishing(x) \cap \neg Other(x))$	\cdot
5	$Fishing(x) \implies (\neg Cargo(x) \cap \neg Other(x))$	\cdot
6	$Other(x) \implies (\neg Fishing(x) \cap \neg Cargo(x))$	\cdot

Table 5.11: Weighted rules for the Markov Logic Network

For such a simple case we can go ahead with exact inference. This yields the following set of results (Table 5.12).

P(X)	X
0.42	$Fishing(Marcel) \cap hasDestination(Marcel, Turkey)$
0.1	$Cargo(Marcel) \cap hasDestination(Marcel, Turkey)$
0.1	$Other(Marcel) \cap hasDestination(Marcel, Turkey)$
0.07	$Fishing(Marcel) \cap hasDestination(Marcel, Italy)$
0.07	$Cargo(Marcel) \cap hasDestination(Marcel, Italy)$
0.07	$Fishing(Marcel) \cap hasDestination(Marcel, OtherD)$
0.02	$Cargo(Marcel) \cap hasDestination(Marcel, OtherD)$
0.01	$Other(Marcel) \cap hasDestination(Marcel, Italy)$

Table 5.12: Markov Logic Network inference results

These results are consistent with the upper and lower bounds on the different worlds obtained either through conceptual graph fusion or valuation network-based approach. Whilst this shows how the same problem can be approached from the perspective of conceptual graphs, valuation networks and Markov logic networks, it does not really leverage the relative strengths of the different approaches. Typically a Markov logic network is used to allow uncertain reasoning with a larger knowledge base consisting of rules the relative weights of which can be learnt from data. In the example here, the rules are used to implement uncertain evidence. However, in typical MLN applications, the evidence is considered to be certain. We could argue that this example is the opposite of a typical MLN application - we have a small set of certain (hard) rules and soft evidence, whereas usual MLN approaches work with a set of uncertain rules and hard evidence.

5.11 Discussion and challenges

This chapter addressed the challenge of the fusion of information at the higher levels of the JDL framework. In particular, we consider the move from single-entity identification to more general situational awareness that may include multiple related entities. The main framework used for

knowledge representation is conceptual graphs, specifically its subset, basic conceptual graphs. However, the existing approaches to information fusion do not support the combination of uncertain and conflicting information. Since the combination of uncertain information and conflict resolution is arguably the keystone of information fusion, a solution is proposed by which uncertainty is directly embedded within the uncertain graph using an epistemic set-based approach.

As we investigate this approach further, it is reasonable to consider that there may be two levels of likelihood associated with each information source. This notion of *probability of all* versus *probability of some* is linked to the possibility theory, from which the terms necessity are borrowed. Different fusion strategies are proposed, and it is shown that the result of information aggregation can be considered a belief distribution. This link is investigated further, and a formal relationship between the uncertain graph formalism and evidential networks is made, making it possible to implement uncertain reasoning within the framework of uncertain graphs.

This approach bears similarity to Markov Logic Network, and, as such, we compare how the same problem can be solved by mapping either valuation networks or Markov Logic Networks to the uncertain graph framework. A natural area of further research is the investigation of links between valuation networks and Markov Logic Networks, as the former can be induced by a set of uncertain rules and constants - which are the elements defining a Markov Logic Network.

MARITIME SITUATION AWARENESS SCENARIOS

In this chapter, we discuss two scenarios that can be used to highlight the three key challenges in information fusion addressed in this thesis - exploitation of context, explanation of results and modelling of situations with multiple uncertainty and knowledge representation formalisms. The focus remains on JDL levels 1 and 2 with some degree of level 3 analysis through explanations. Both scenarios address realistic situations which may be encountered in maritime situation awareness, with an assumption of significant preprocessing of sources, i.e. the sources are assumed to be outputs from deep learning classifiers, analyst assessments and similar, rather than raw or barely processed data. In the first scenario, we consider the problem of illegal fishing and demonstrate the ability to reason about large scale, multi-entity situations as well as keeping track of a large number of possible worlds without compromising reasoning and explanation abilities. The second scenario focuses on a single vessel but with more complex reasoning and an in-depth analysis of the impact of source quality.

6.1 Scenario 1: multi-entity illegal fishing

The first scenario to demonstrate the notions discussed throughout this thesis is that of situation assessment involving multiple vessels with some of them involved in illegal fishing. This scenario focuses on the usage of both conceptual graphs and valuation networks as well as the exchange of information between the two. The conceptual graph is used as a knowledge base, and valuation networks are used to provide a layer for more complex uncertainty management, reasoning and source modelling.

6.1.1 Scenario overview

This scenario consists of four vessels in international waters. A marine protected area (MPA) lies within the region, where fishing is prohibited. The problem is identifying vessels that are involved

in either illegally fishing in the MPA or transporting the illegal catch.

We consider four vessels "ARAMIS", "BOHUN", "CHEKHOV" and "DOROTHY", which we will occasionally refer to as *A*, *B*, *C* and *D*. The ground truth is that vessels *A* and *B* are fishing vessels, and *A* is involved in illegal fishing. Both *C* and *D* are refrigerated cargo vessels used for transshipment, with *D* doing it legally and only receiving catch from *B*, while *C* receives both "clean" catch from *B* as well as illegal catch from *A*.

The information comes from several sources, but all of it except for the context and meta-knowledge is based on AIS transmissions. For the sake of avoiding unnecessary complexity, the vessel identities (name, type, flag) provided by AIS are assumed to be fully reliable. Other pieces of information provided include time from the last AIS transmission (is the AIS fresh or old), as well as analysis of the kinematic data provided through the AIS. This kinematic analysis includes correction and prediction based on old AIS data in case the vessel stopped transmission. From this source, we obtain an estimation of the vessel's current location and its activity with a focus on events such as fishing and rendezvous.

6.1.2 Rules and sources of information

The reasoning for this scenario is relatively straightforward. The key variables of interest (defined on a *true - false* frame) are whether a vessel is involved in illegal activity and whether the vessel performs illegal fishing. Generally, the reasoning rules we have can be divided into two main groups. On the one hand, we have certain rules which are mostly concerned with semantics (e.g. if a vessel is illegally fishing, it must be fishing). On the other, we have uncertain rules with different degrees of likelihood - e.g. if a vessel is fishing and it is involved in illegal activity, then it is likely illegally fishing - even though there is a small chance that its fishing is legal and it is actually involved in another kind of illegal activity.

As such, we have the following set of rules:

1. If a vessel is fishing and it is in an area where fishing is illegal, then it is illegally fishing
2. If a vessel is illegally fishing, then it is involved in illegal activity, and it is fishing
3. If a vessel is involved in illegal activity and it is fishing, then it is likely illegally fishing
4. If a vessel is not broadcasting its AIS, it is more likely to be involved in illegal activity
5. If a vessel is flying a flag from a country on the High-Risk list from the Paris Memorandum of Understanding on Port Control, it is slightly more likely to be involved in illegal activity
6. If a cargo vessel is performing a rendezvous with another vessel that is involved in illegal activity, the cargo vessel is likely to be involved in illegal activity

Contextual reasoning is included in this problem by modelling the impact of meteorological conditions on the likelihood that a vessel is not broadcasting its AIS, as well as the rules on the legality of fishing in different areas.

The following sources provide information describing the behaviour of the vessels:

1. *AIS transmission* - Name: *ARAMIS*, Vessel type: 30 (Fishing), Flag: TZ, recently received : false, reliability = 1
2. *AIS transmission* - Name: *BOHUN*, Vessel type: 30 (Fishing), Flag: IT, recently received : true, reliability = 1
3. *AIS transmission* - Name: *CHEKHOV*, Vessel type: 70 (Cargo), Flag: MD, recently received : true, reliability = 1
4. *AIS transmission* - Name: *DOROTHY*, Vessel type: 70 (Cargo), Flag: GB, recently received : true, reliability = 1
5. *AIS analysis* - *ARAMIS* is fishing, reliability = 0.6
6. *AIS analysis* - *ARAMIS* is in location X, reliability = 0.6
7. *AIS analysis* - *BOHUN* is fishing, reliability = 0.9
8. *AIS analysis* - *BOHUN* is in location Y, reliability = 0.9
9. *AIS analysis* - *CHEKHOV* has had a rendez-vous with *ARAMIS*, reliability = 0.75
10. *AIS analysis* - *CHEKHOV* is in location Z, reliability = 0.9
11. *AIS analysis* - *CHEKHOV* has had a rendez-vous with *BOHUN*, reliability = 0.75
12. *AIS analysis* - *DOROTHY* has had a rendez-vous with *BOHUN*, reliability = 0.75
13. *AIS analysis* - *DOROTHY* is in location Z, reliability = 0.9

For each of the four vessels, we can generate a simple conceptual graph describing the information on its identity as provided by the AIS, similar to the one in Figure 5.14 in Chapter 5. An example for *ARAMIS* is shown in Figure 6.1, which we will denote as G_{1A} . In order to keep it legible, we will avoid a full-scale graph containing all the entities at this stage. However, bear in mind that these graphs for individual vessels can be connected with the relation describing a rendezvous.




Figure 6.1: A single vessel conceptual graph shown for *ARAMIS*

6.1.3 Multi-source reasoning implementation and results

With this in mind we can consider inclusion of further sources. In order to obtain the information of interest which we want to include in the conceptual graph we will use a valuation network to reason about the information provided by the sources and combine it with some information already in the conceptual graph, as well as context and metaknowledge. In this scenario we use two valuation networks covering different areas of interest - a relatively simple one to implement the "transitivity" of illegal activity in case of a rendez-vous and a somewhat more complex one to assess the likelihood of illegal activity and illegal fishing.

This valuation network shown in Figure 6.2 implements the rules 1 to 4¹, as well as the contextual source model for AIS as proposed in Chapter 3, behaviour-based source correction model for activity and location reports and contextual information on location of MPAs. The variables of interest include IA - illegal activity, IF - illegal fishing, AIS - AIS received, IL - illegal fishing location and F - fishing, all of which are defined on the boolean frame of discernment $\Omega_{X_{IA}} = \{true, false\}$. Furthermore we have the vessel's location LOC defined on frame $\Omega_{LOC} = \{X, Y, Z\}$. The valuations m_1^R , m_2^R and m_3^R correspond to the source reports on AIS reception, vessel's activity and vessel's location respectively. The valuation m_1^Q provides the loss rate of the channel, valuations m_2^Q and m_3^Q provide the metaknowledge on sources' reliability. Finally the valuations m_1^C and m_3^C provide contextual information on weather and locations of MPA's. For simplicity we let all the uncertain implications such as $m_{IA \cap F \rightarrow IF}$ have the weight of 0.7.

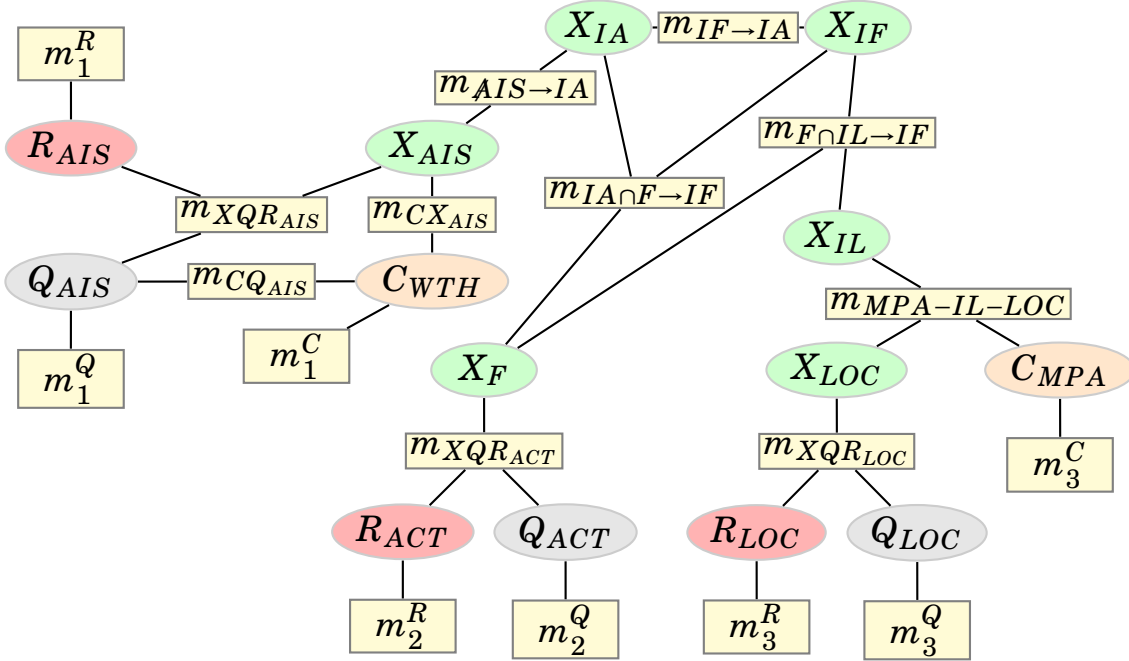


Figure 6.2: Valuation network for assessment of illegal activity and illegal fishing

In the case of Vessel A, *ARAMIS*, we have the following source reports $m_1^R(False) = 1$, $m_2^R(Fishing) = 1$ and $m_3^R(X) = 1$. The channel is assumed to be 80% lossless and 20% lossy -

¹Rule 5 is not implemented for clarity, as it has a relatively minor impact and the valuation network is sufficiently complicated

$m_1^Q(loss) = 0.2$, $m_1^Q(\neg loss) = 0.8$. The AIS analysis of both location and activity is 60% reliable - $m_2^Q(rel) = 0.6$, $m_2^Q(\neg rel) = 0.4$, $m_3^Q(rel) = 0.6$, $m_3^Q(\neg rel) = 0.4$. Finally the weather is assumed to be 60% good and 40% bad - $m_1^C(good) = 0.6$, $m_1^C(bad) = 0.4$ and the contextual information on MPA's states that X lies within an MPA with no uncertainty - $m_3^C(X) = 1$.

With this in mind, we can perform inference on our domain of interest $IA \times F \times IF$. The inference results are shown in Table 6.1. Since the results are a belief distribution, it would be difficult to convert them back to a form consistent with the conceptual graph, which we use as the main knowledge base. As such, we accept a minor loss of information and consider the product of marginals for each focal set, each of which can be used to generate a conceptual graph on its own. This belief distribution can be considered a probability distribution on a set of five possible conceptual graphs. As such, the graph in Figure 6.1 represents the true state of the world with probability corresponding to marginal Ω - $P(G_{1A}) = 0.1990$

$m_J(\mathcal{F})$	focal set \mathcal{F}	Marginals
0.44442	$IF \times F \times IA$	$IF \times F \times IA$
0.17687	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times \neg F \times IA \end{array} \right\}$	IA
0.1194	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times \neg F \times IA \\ \neg IF \times \neg F \times \neg IA \end{array} \right\}$	Ω
0.08358	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times F \times \neg IA \end{array} \right\}$	F
0.05572	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times \neg F \times IA \\ \neg IF \times F \times \neg IA \\ \neg IF \times \neg F \times \neg IA \end{array} \right\}$	Ω
0.03618	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times F \times IA \end{array} \right\}$	$F \times IA$
0.03582	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times F \times IA \\ \neg IF \times F \times \neg IA \end{array} \right\}$	F
0.02412	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times F \times IA \\ \neg IF \times \neg F \times IA \end{array} \right\}$	IA
0.02388	$\left\{ \begin{array}{l} IF \times F \times IA, \\ \neg IF \times F \times IA \\ \neg IF \times \neg F \times IA \\ \neg IF \times F \times \neg IA \\ \neg IF \times \neg F \times \neg IA \end{array} \right\}$	Ω

Table 6.1: Inference on Illegal fishing \times Fishing \times Illegal Activity for ARAMIS

For visualisation, Figure 6.3 shows the realisation of the graph corresponding to the marginals

$IF \times F \times IA$, which is the single most likely result of this combination. Note this is the only possible graph asserting IF , as it being true implies both F and IA being true.




Figure 6.3: A single vessel conceptual graph shown for *ARAMIS* obtained through inference on AIS data with $P = 0.44442$

Meanwhile, there is no evidence of illegal activity from vessel B, and as such, only two conceptual graphs are generated for *BOHUN* - the one shown in Figure 6.4 and a one more similar to Figure 6.1, where the analysis of the AIS is assumed unreliable, and hence the activity of the vessel is unknown ($P = 0.1$).




Figure 6.4: A single vessel conceptual graph shown for *BOHUN* obtained through inference on AIS data with $P = 0.9$

With this in mind we can move on to analysis of the cargo vessels and the rendez-vous. For the vessel *CHEKHOV* we have the basic overview graph as shown in Figure 6.5.




Figure 6.5: A single vessel conceptual graph shown for *CHEKHOV*

Through the combination of the information on the rendezvous, we obtain the set of graphs shown in Figure 6.6. Note that as the pieces of information are not mutually exclusive and refer to separate events, when both sources of information are assumed to be true, the two relations are

added independently, rather than using the basic knowledge model \mathcal{M} as it was done in earlier examples.

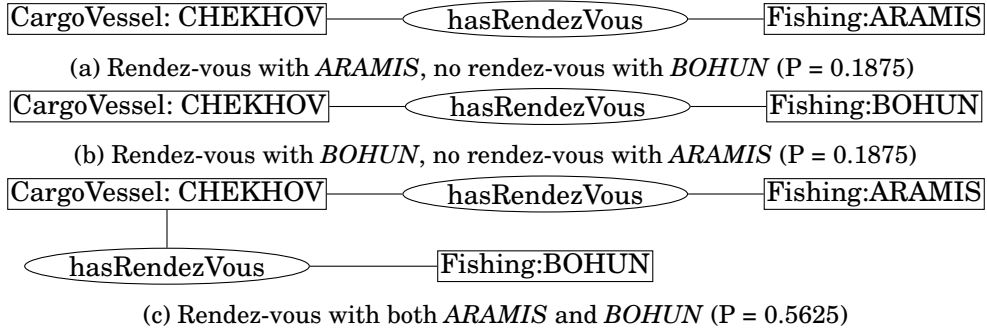


Figure 6.6: Excerpt from the greater conceptual graph, showing the different possible situations regarding rendez-vous of vessel *CHEKHOV*. The case when neither of the two sources is true is not shown here, but it corresponds to 6.5 only ($P = 0.0625$)

Note that at this stage, for clarity, all the figures shown so far depict excerpts from the full graph, which would include all four vessels.

The reasoning about whether a cargo vessel is involved in illicit activity is performed using a separate valuation network, shown in Figure 6.7 which takes the inference results from before and combines it with the implementation of rules 4 to 6. The variable $X_{RV(N)}$ represents the existence of the relation 'hasRendezVous' between the concepts corresponding to the vessel in question and the vessel N . The variable X_{HR} denotes the risk assessment of the vessel, which in this case is only assessed through its flag (although it could naturally be replaced by an approach such as the Paris Memorandum of Understanding risk assessment procedure).

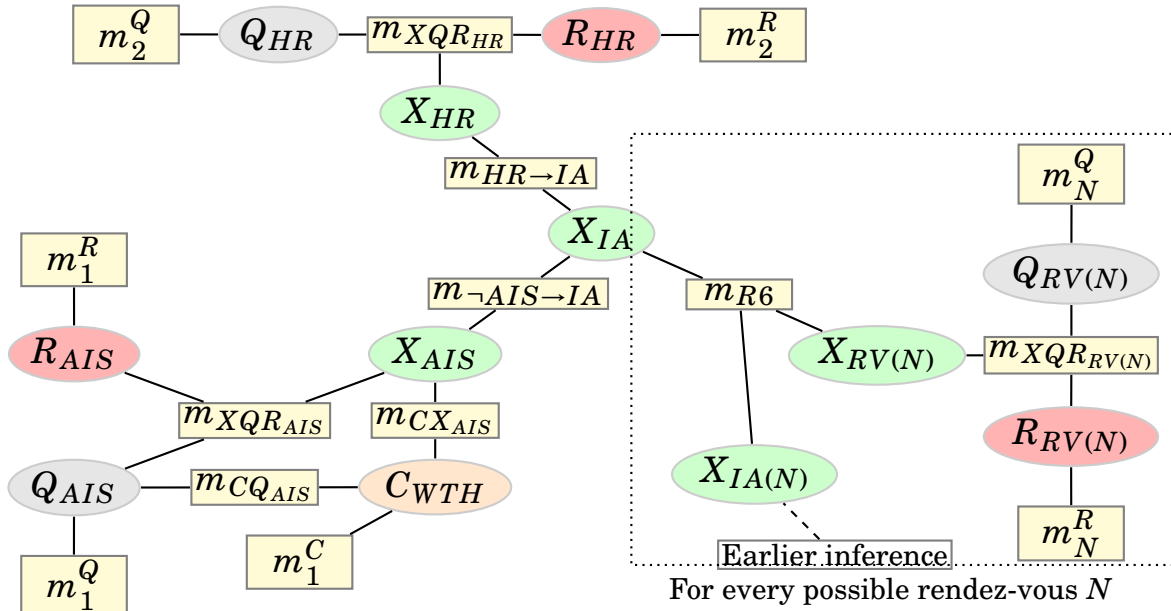


Figure 6.7: Valuation network for assessment of illegal activity for a cargo vessel

Considering the cargo vessel C, *CHEKHOV*, we have the following source reports $m_1^R(True) = 1$

and $m_2^R(Tru e) = 1$. The channel is assumed to be 80% lossless and 20% lossy - $m_1^Q(loss) = 0.2$, $m_1^Q(\neg loss) = 0.8$. The AIS broadcast of the flag is assumed to be fully reliable $m_1^Q(rel) = 1$.

For each potential rendez-vous report we obtain a separate branch. As such we have reports $m_A^R(true) = m_B^R(true) = 1$ and $m_A^Q(rel) = m_B^Q(rel) = 0.75$. However from earlier inference we have the belief that vessel A is involved in illicit activity $m(X_{IA(A)} = true) = 0.6186$, $m(X_{IA(A)} = \Omega) = 0.3814$ and no evidence supporting illegal activity by vessel B $m(X_{IA(B)} = \Omega) = 1$. This makes it possible to have an overall assessment of the belief that vessel C is involved in illegal activity - $m(X_{IA(C)} = true) = 0.55048$, $m(X_{IA(C)} = \Omega) = 0.44952$. Performing the same for vessel D, *DOROTHY* shows no evidence of illegal activity.

For the computation of probabilities of possible multi-entity worlds, we must obtain the joint inference on illicit activity by vessel C and A as well as the rendezvous between the two, similarly to the previous case where we considered the joint distribution on illicit activity and fishing. As such, we have the distribution displayed in Table 6.2

$m_J(\mathcal{F})$	focal set \mathcal{F}	Marginals
0.45496	$IA(C) \times IA(A)$	$IA(C) \times IA(A)$
0.22663	$\left\{ \begin{array}{l} IA(C) \times IA(A), \\ \neg IA(C) \times IA(A), \end{array} \right\}$	$IA(A)$
0.11702	$\left\{ \begin{array}{l} IA(C) \times IA(A), \\ IA(C) \times \neg IA(A), \\ \neg IA(C) \times IA(A) \end{array} \right\}$	Ω
0.10587	$\Omega_{IA(C)} \times \Omega_{IA(A)}$	Ω
0.09552	$\left\{ \begin{array}{l} IA(C) \times IA(A) \\ IA(C) \times \neg IA(A) \end{array} \right\}$	$IA(C)$

Table 6.2: Inference on joint belief distribution of Illegal activity (ARAMIS) \times Illegal activity (CHEKHOV)

However, to obtain the full picture, we must combine the above with the results from Table 6.1. With no evidence on the illicit activity of vessels B and D, this shows the full picture on the illicit activity of all the vessels (with the information on rendezvous between B and C, B and D, and activity of vessel B, although uncertain, being independent). This combination has a total of 51 focal sets and, as such, is not displayed here in its entirety. Instead, in Table 6.3, we display the probability distribution on the 20 marginals and by extension, the 20 conceptual graphs which can be obtained from this combination and reasoning.

At this stage, it could be combined with the results of inference on the behaviour of vessels B and D (as shown in Table 6.4. However, as the two distributions are fully independent, this combination would result in 80 possible worlds, and there is no reason to list them all.

Figure 6.8 shows the full conceptual graph corresponding to the most likely inference outcome - and also the true state of the world. Its probability is $P = 0.1333$. Although this probability may seem to be low - given that we know this to be the ground truth - it is important to bear in mind that the number of possible worlds is very large and that the other realisations do not necessarily

Marginals (Graph)	Probability
$IA(C) \times RV(A, C) \times IF(A) \times F(A) \times IA(A)$	0.26332
$IA(C) \times RV(A, C) \times IF(A) \times F(A) \times \Omega_{IA(A)}$	0.021436
$IA(C) \times RV(A, C) \times IF(A) \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.11909
$IA(C) \times RV(A, C) \times \Omega_{IF(A)} \times F(A) \times IA(A)$	0.069996
$IA(C) \times RV(A, C) \times \Omega_{IF(A)} \times F(A) \times \Omega_{IA(A)}$	0.0056982
$IA(C) \times RV(A, C) \times \Omega_{IF(A)} \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.031656
$IA(C) \times \Omega_{RV(A, C)} \times IF(A) \times F(A) \times IA(A)$	0.033331
$IA(C) \times \Omega_{RV(A, C)} \times IF(A) \times F(A) \times \Omega_{IA(A)}$	0.0027134
$IA(C) \times \Omega_{RV(A, C)} \times IF(A) \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.015074
$IA(C) \times \Omega_{RV(A, C)} \times \Omega_{IF(A)} \times F(A) \times IA(A)$	0.077773
$IA(C) \times \Omega_{RV(A, C)} \times \Omega_{IF(A)} \times F(A) \times \Omega_{IA(A)}$	0.0063313
$IA(C) \times \Omega_{RV(A, C)} \times \Omega_{IF(A)} \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.035174
$\Omega_{IA(C)} \times RV(A, C) \times IF(A) \times F(A) \times \Omega_{IA(A)}$	0.026866
$\Omega_{IA(C)} \times RV(A, C) \times IF(A) \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.044777
$\Omega_{IA(C)} \times RV(A, C) \times \Omega_{IF(A)} \times F(A) \times \Omega_{IA(A)}$	0.062687
$\Omega_{IA(C)} \times RV(A, C) \times \Omega_{IF(A)} \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.10448
$\Omega_{IA(C)} \times \Omega_{RV(A, C)} \times IF(A) \times F(A) \times \Omega_{IA(A)}$	0.0089553
$\Omega_{IA(C)} \times \Omega_{RV(A, C)} \times IF(A) \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.014926
$\Omega_{IA(C)} \times \Omega_{RV(A, C)} \times \Omega_{IF(A)} \times F(A) \times \Omega_{IA(A)}$	0.020896
$\Omega_{IA(C)} \times \Omega_{RV(A, C)} \times \Omega_{IF(A)} \times \Omega_{F(A)} \times \Omega_{IA(A)}$	0.034826

Table 6.3: Marginals obtained from the inferred joint belief distribution on $IA(C) \times RV(A, C) \times IF(A) \times F(A) \times IA(A)$

Marginals (Graph)	Probability
$RV(B, D) \times F(B)$	0.6750
$\Omega_{RV(B, D)} \times F(B)$	0.2250
$RV(B, D) \times \Omega_{F(B)}$	0.0750
$\Omega_{RV(B, D)} \times \Omega_{F(B)}$	0.0250

Table 6.4: Inference of behaviour of vessels B and D. There is no evidence supporting any implication of illegal activity for either vessel

oppose this scenario - they just do not fully support it.

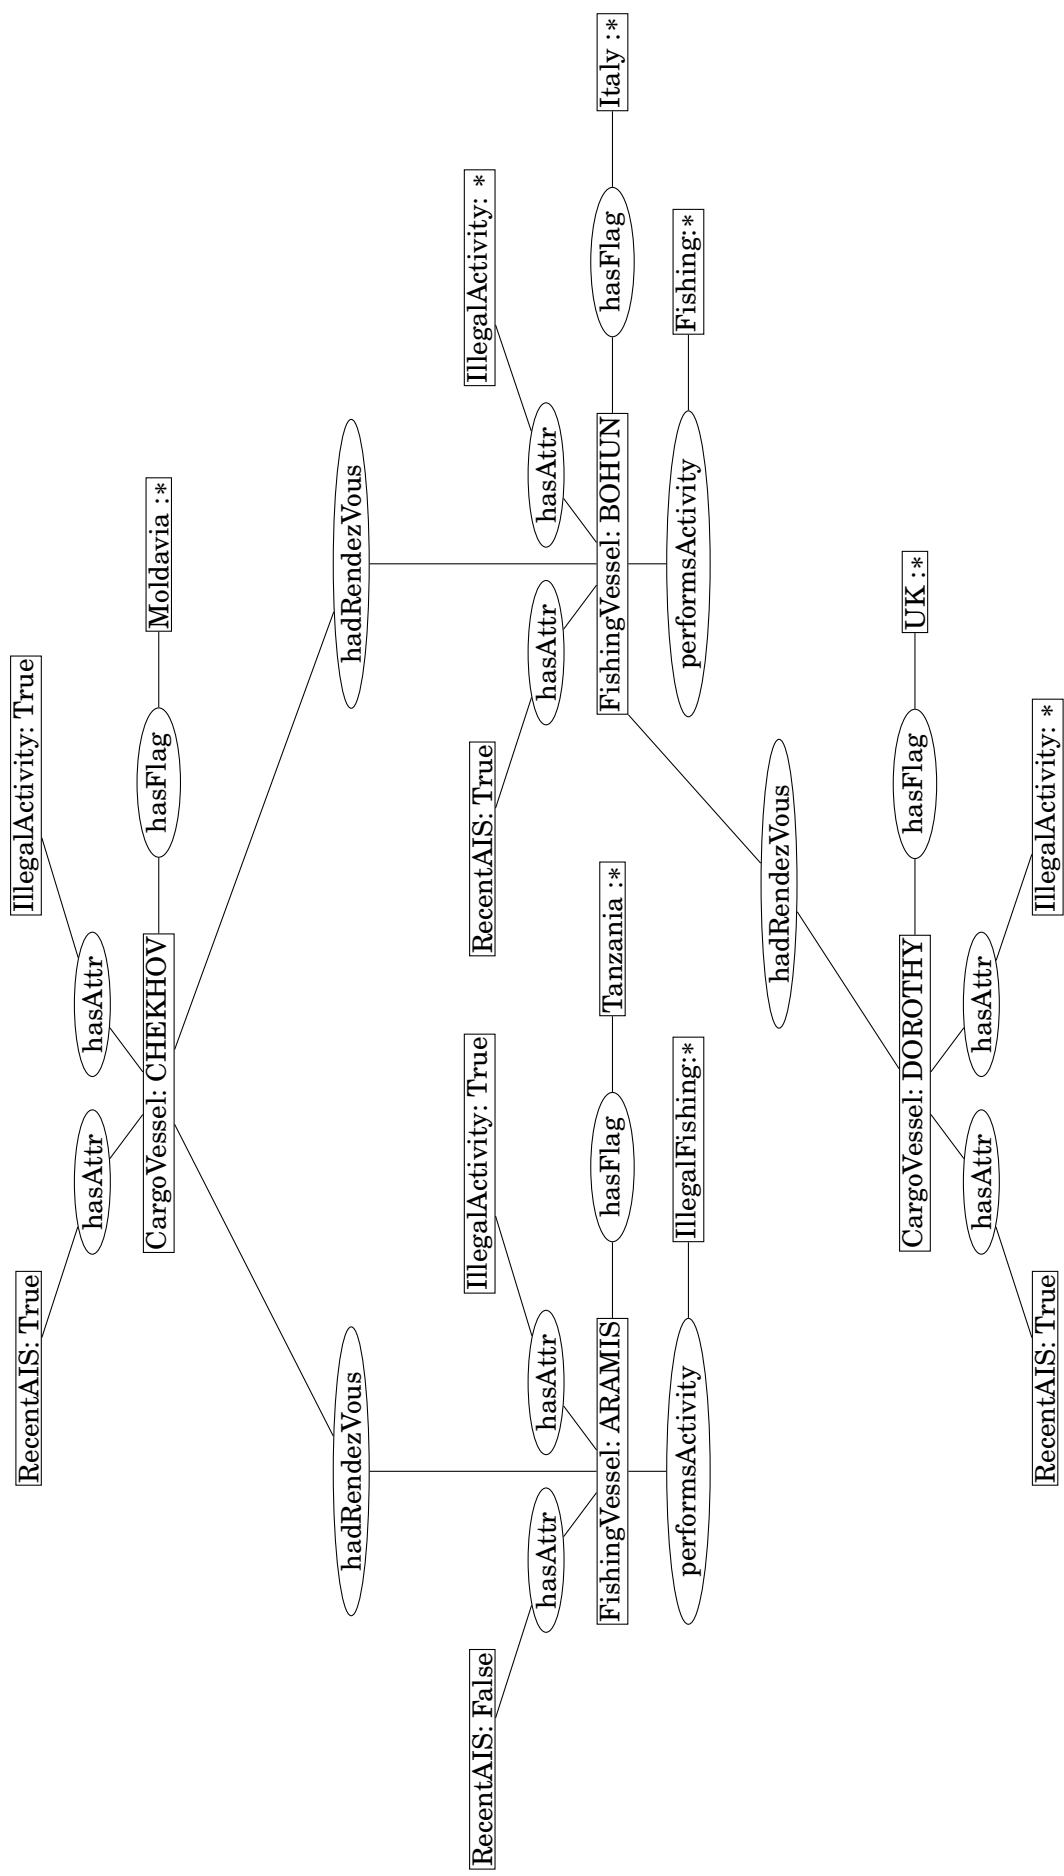


Figure 6.8: The full conceptual graph showing the most likely world including all the four vessels

As discussed in previous chapter, from distributions on graphs like this we can obtain upper and lower probabilities (or belief and plausibility values) for various queries. Some examples of interest we have here are:

1. Is there a vessel involved in illegal activity? $Bel = 0.7771, Plaus = 1$
2. Is *ARAMIS* involved in illegal fishing and *CHEKHOV* involved in illegal activity? $Bel = 0.29665, Plaus = 1$
3. Is a vessel flying a low risk flag involved in illegal activity? $Bel = 0, Plaus = 1$

It is important to bear in mind that the rules proposed here are not necessarily optimal, and changing them would have a significant effect on the inference results. For example, we could change rule 6 such that if a fishing vessel has a rendezvous with any other vessel involved in illegal activity, it, too, is involved in illegal activity. This would provide evidence for suspicious activity by *BOHUN* and, to a lesser extent, *DOROTHY*. Furthermore, note that in the reasoning system, there is no evidence opposing the notion that any vessel is involved in criminal activity - as such, the plausibility of illicit activity for any vessel is 1.

One possible way of implementing it with current sources would be to include a rule implying that a vessel flying a low-risk flag is less likely to be involved in criminal activity. This would provide evidence against the illegal activity of vessels B and D but not conflict with any other sources. Instead, let us consider a new source of information. Assume that the captain of *CHEKHOV* is asked about the behaviour of the vessel and claims not to have received any illegal cargo from *ARAMIS* and denies any wrongdoing. Clearly this kind of evidence has little weight, but we will include it anyway to see how it changes the reasoning.

Let this be denoted as `[Vessel:CHEKHOV] - (hasAttr) - [IllegalActivity:False]` with a reliability rating of 0.25. Since the illicit activity of a cargo vessel is not included as a source in any reasoning, we can simply combine it as a conceptual graph, with the set of graphs listed in Table 6.3. Using one of the fusion strategies outlined in the previous chapter, inclusion of this source of information (in cases where it is assumed reliable) would change any marginal from Table 6.3 which includes $IA(C)$ into one including $\Omega_{IA(C)}$ and similarly $\Omega_{IA(C)}$ into $\neg IA(C)$. In other words, to explain on the most likely example $IA(C) \times RV(A, C) \times IF(A) \times F(A) \times IA(A)$, with $P = 0.26332$ (used to generate the graph in Figure 6.8 when combined with information on vessels B and D) it becomes either $\Omega_{IA(C)} \times RV(A, C) \times IF(A) \times F(A) \times IA(A)$ with its new probability $P = 0.25 \times 0.26332 = 0.0658$ or remains as it was with probability $P = 0.75 \times 0.26332 = 0.1975$. The new belief marginal on variable $IA(C)$ becomes $m(IA(C)) = 0.4129, m(\neg IA(C)) = 0.1124, m(\Omega_{IA(C)}) = 0.4747$.

6.1.4 Explaining the results

The next step is to analyse how the different pieces of information affect the inference result. While the combination process is transparent by design, it is nevertheless interesting to observe the relative importance of different sources of evidence, as well as the impact of individual sources on the joint frame.

An additional degree of difficulty arises from the fact that in this scenario, we are blending different fusion approaches. Whilst most of the process is done through valuation networks, the last stage involves the combination of conceptual graphs. Note that the approach we have used is equivalent to using the conjunctive-disjunctive combination rule (with the reinterpretation of partial conflict as partial ignorance) or Yager's rule, equivalent for the combination of masses with a binary frame of discernment. As such, we can use the belief function-based approaches nevertheless.

Using the approach briefly discussed in Chapter 4, section 4.3.2.1, we visualize the impact on each body of evidence by taking the derivative of a metric of interest - belief or plausibility of a hypothesis (see Equations 4.7 and 4.8).

It is also possible to use sensitivity spaces to assess the impact of each BoE on the overall solution (as opposed to focusing on individual hypotheses). This can be done using some uncertainty measures existing within the theory of evidence. Here, as proposed in [92] specificity and consonance measures are used. The specificity of the BoE m is defined as follows:

$$(6.1) \quad Spec(m) = \sum_{A \in \Omega} \frac{m(A)}{|A|}$$

where $|\cdot|$ is the cardinality operator, while the consonance is defined as:

$$(6.2) \quad Cons_i(m) = \frac{1}{1 + Ent(m)}$$

where

$$(6.3) \quad Ent(m) = - \sum_{A \in \Omega} m(A) \log_2 Pl(A).$$

Thus the marginal impact of each BoE on the solution's specificity and consonance can be characterised as follows:

$$(6.4) \quad \widehat{Spec}_i(m) = \left. \frac{\delta Spec(m)}{\delta \alpha_i} \right|_{\alpha_i=1}$$

$$(6.5) \quad \widehat{Cons}_i(m) = \left. \frac{\delta Cons(m)}{\delta \alpha_i} \right|_{\alpha_i=1}.$$

where α_i is a discounting factor associated to m_i . Sensitivity spaces obtained by plotting sensitivities related to the focal set of interest can be used to explain to which extent each source of evidence supports or negates the particular hypothesis and similarly which sources have had the greatest impact on the decision. The sensitivities related to the properties of the entire mass function allow greater insight into how each body of evidence affects conflict and uncertainty within the inference outcome.

These sensitivity spaces for belief and plausibility are shown in Figure 6.9a and in Figure 6.9b for specificity and consonance. It can be clearly seen that the single individual source contributing to the belief that *CHEKHOV* is involved in malicious activity is the AIS analysis implying its rendezvous with *ARAMIS*, closely followed by it flying a high-risk flag. Subsequently, we have the

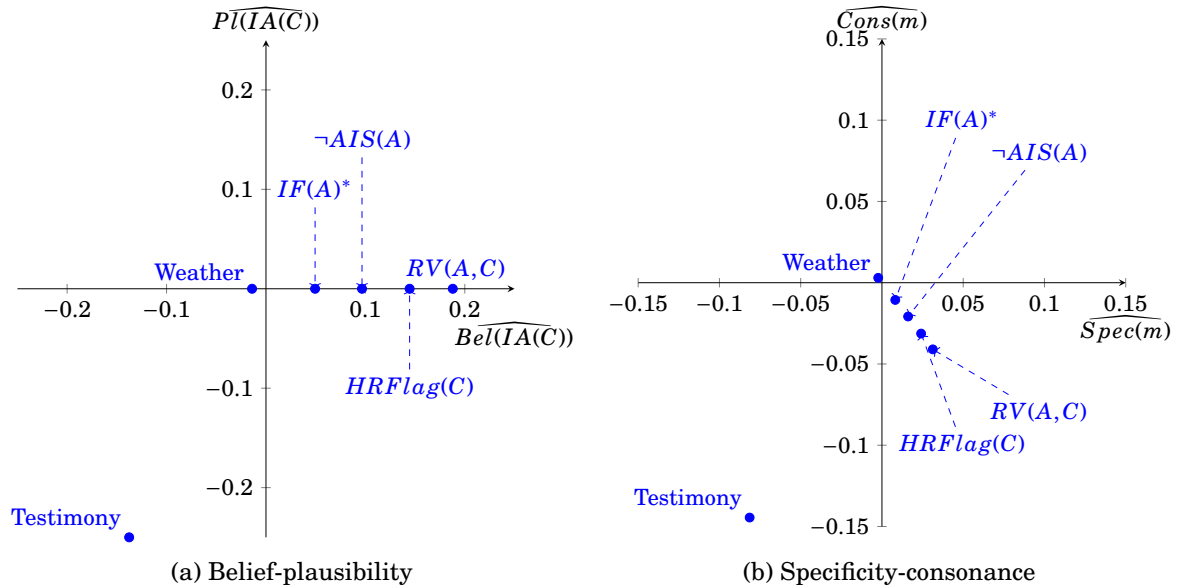


Figure 6.9: Sensitivity spaces for inference on illicit activity of *CHEKHOV*

sources implicating *ARAMIS* and thus *CHEKHOV* by association - the set of sources contributing to belief in *ARAMIS*' illegal activity (disabled AIS) and illegal fishing (AIS analysis describing its location and activity as well as contextual information on MPA's - these are jointly denoted as $IF(A)^*$). The knowledge of poor weather slightly reduces the belief in *CHEKHOV*'s involvement. On contrary, we have captain's testimony which simultaneously argues against the rest of the sources (reduces $Bel(IA(C))$), while also supporting the contrary hypothesis having a negative impact on $Pl(IA(C))$).

Another interesting aspect to consider is the difference between using the conceptual graph-based approach for a combination of evidence and simply treating this new piece of semantic evidence like any other and combining it conjunctively. We can see this yields a different sensitivity space, shown in Figures 6.10a and 6.10b for belief-plausibility and specificity-consonance respectively.

This behaviour is largely expected - with the original approach, we resolve the conflict manually, which implies that the new piece of evidence (Captain's testimony) is not directly disputed by any other sources (as seen in the belief-plausibility graph, where originally the other sources have no impact on plausibility (as they do not reduce the likelihood of the vessel's innocence), but it increases the overall degree of ignorance. When Dempster's method is used, we have greater conflict but no decrease in specificity.

Similarly, we can look at the relative contribution of each source of information to the hypotheses. Let us consider the relative contributions of the sources to some of the most likely hypotheses. As we are interested in relative overall contribution, we can use Shapley's value for this. Similarly to the method outlined in Chapter 4, we consider the fusion process to be a cooperative game of multiple sources. In this case, we consider the payoff function to be the probability of a particular solution.

As such, for the solution shown in Figure 6.8 which we consider to have the probability




Figure 6.10: Sensitivity spaces for inference on illicit activity of *CHEKHOV* using Dempster's rule of combination exclusively




Figure 6.11: Shapley values for contribution of the sources to the solution from Figure 6.8

$P = 0.1975$, we can compute the contribution of individual sources. As per the definition of Shapley's value, the total payoff (0.1975) is proportionally distributed among all the contributing sources (with the possibility of negative contributions for sources arguing against this solution). The computed Shapley values are shown in Figure 6.11.

As expected, the most significant contributions are these, which implicate both Vessels A and C. The source providing information on the location of *ARAMIS* and the contextual information on the Maritime Protected Area are equally important, as one is irrelevant to the problem without the other. The source with the lowest impact is that Vessel C is flying a flag considered high risk - this source affects only one variable of interest, is in agreement with most other sources, and its modelled impact (the probability assigned to the rule that a vessel flying a high-risk flag is engaging in criminal activity) is low. Finally, *CHEKHOV* captain's testimony is the only source of evidence that directly opposes this solution, receiving a negative value; however its impact is quite low due to its reliability.

We can use the same approach to analyse contributions to a more general hypothesis, rather than a specific conceptual graph solution.

6.2 Scenario 2: Infrastructure protection and threat assessment

The second scenario to consider is that of assessment of the threat a single vessel can pose to sensitive maritime infrastructure. This scenario has been published at the 2020 International Conference for Information Fusion ([113]) and an extended version of the paper, a portion of the content of which is included below, has been accepted for publication in the special issue of Information Fusion on Explainable Artificial Intelligence. This scenario uses only a single method of uncertainty modelling - valuation networks - but places more emphasis on an in-depth discussion of implementation and providing explanations for the impact of the source model.

Figure 6.12 displays a high-level view of some of the variables involved in this problem as well as their hierarchical relationships. We consider a vessel (i.e., the *Target*) which investigation has been triggered by an external factor. The intent of the vessel is linked to its identity (e.g. following the NATO standard identities) and will be assumed to be more likely malicious if some inconsistency is observed in the AIS signal. The opportunity of the vessel is basically linked to its current position relative to the infrastructure, while its capability depends on its classification type (and assumed corresponding equipment) while some intelligence can confirm it regardless of its type.

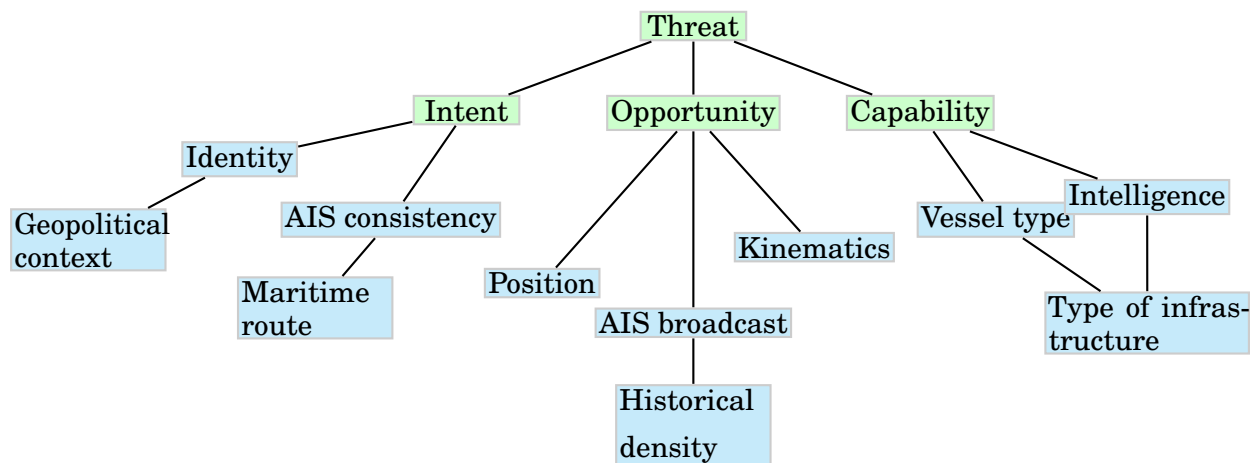


Figure 6.12: Simplified, high-level relationships between inference targets (green) and directly relevant or auxiliary information (cyan).

6.2.1 Threat assessment multi-source evidential network

This section covers implementation of a threat assessment model through an evidential network. According to the BBC model introduced in Chapter 3, variables are distinguished according to their type into focus, report and quality variables.

Table 6.5 lists the variables included in the evidential network as well as their respective frame of discernment with some comments.

For the sake of simplicity and clarity of discussion, the model used in this example is the conjunction of the three factors - intent, opportunity and capability. The validity and applicability

of this model are briefly discussed alongside the results.

Table 6.5: Variables of the threat assessment evidential network

Variable	Frame (Ω)	Variable type	Comments
THREAT	{TRUE, FALSE}	Focus	The main target variable
INTENT	{TRUE, FALSE}	Focus	Direct intelligence reports
OPPORTUNITY	{TRUE, FALSE}	Focus	can be applied at these
CAPABILITY	{TRUE, FALSE}	Focus	variables
AIS_T	{ON, OFF}	Focus	True AIS broadcast
AIST_R	{REC, NREC}	Report	Loss rate est. by correlation
AIS_Q	{LOSS, NLOSS}	Quality	of SAR and AIS contacts
AIS_INC	{TRUE, FALSE}	Focus	
AIS_INC_R	{TRUE, FALSE}	Report	Inconsistent AIS
AIS_INC_Q	{REL, NREL}	Quality	
KIN_X	{PASSING, LOITERING, STOPPED}	Focus	
KIN_R	{PASSING, LOITERING, STOPPED}	Report	Partial reliability model
KIN_Q	{PERF, INACC, NREL}	Quality	for similar behaviours
TYPE_X	{1,2,3}	Focus	
TYPE_R	{1,2,3}	Report	Vessel capability category
TYPE_Q	{OVER, UNDER, PERF., NREL}	Quality	estimated by vessel type Can be over or underestimated
ID_X	{F, N, H, A, S}	Focus	
ID_R	{F, N, H, A, S}	Report	Subset of vessel standard
ID_Q	{REL, NREL}	Quality	identities
LOCATION_X	{INFRASTRUCTURE, CLEAR}	Focus	
LOCATION_R	{INFRASTRUCTURE, CLEAR}	Report	Potential attack target
LOCATION_Q	{REL, NREL}	Quality	in vicinity

A series of rules are provided in natural language and listed in Table 6.6 to convey the general idea of the reasoning of the problem modelling described. However, the strength of the model is to handle uncertainty and ignorance.

Some corresponding logical relationships are given in Table 6.7, which are modelled using simple support belief functions, with the mass corresponding to the probability $P(R)$ that the relationship holds. The rules are parametrised using a set of parameters γ_i .

Figure 6.13 shows the constructed valuation network. Sources provide information on report variables and quality variables describe their quality, typically the relevance or reliability. The marginalisation is performed on focus variables, to summarise this information.

Sources description

The sources are modelled using a set of parameters, a combination of which corresponds to the appropriate source behaviour, as discussed in Chapter 3, and are listed in Table 6.8. The AIS reception, vessel identity, AIS consistency and location are modelled with a single parameter (loss rate for AIS reception and reliability for others), whereas kinematics and capability estimation by type are modelled with two parameters (reliability and accuracy for the former and tendency to over- and underestimate for the latter). This induces two possible behaviours for AIS reception, consistency, vessel identity and location information sources, three possible behaviours for kinematics (unreliable, reliable but inaccurate, reliable and accurate) and four possible behaviours

Table 6.6: Model of solution to threat assessment expressed in natural language rules

Threat model	
0	The vessel poses a threat if and only if it has malicious intent, capability and opportunity

Problem domain model	
1	If AIS is inconsistent then it has malicious intent
2	If AIS is disabled then it has opportunity
3	If vessel is stopped or loitering in vicinity of infrastructure then it has opportunity
4	If vessel type capability category is 1, the vessel is not capable
5	If vessel type capability category is 2 the vessel may be capable
6	If vessel type capability category is 3, the vessel is capable
7	If the vessel identity is Friend, the vessel does not have malicious intent
8	If the vessel identity is Hostile, the vessel has malicious intent
9	If the vessel identity is Suspect, the vessel likely has malicious intent
10	If the vessel identity is Assumed Friend, the vessel likely does not have malicious intent
11	If vessel is not in vicinity of the infrastructure it does not have opportunity

Sources model	
12	If AIS signal received then AIS signal broadcast
13	If AIS signal not received and channel is lossless then AIS signal not broadcast
14	If AIS is reported as inconsistent and the source is reliable then AIS inconsistent
15	If AIS is reported as consistent and the source is reliable then AIS consistent
16	If vessel reported as loitering and the source is accurate and reliable then the vessel is loitering
17	If vessel reported as stopped and the source is accurate and reliable then the vessel is stopped
18	If vessel reported as passing and the source is accurate and reliable then the vessel is passing
19	If vessel is reported as stopped and the source is partially accurate and reliable then the vessel is either loitering or stopped
20	If the vessel capability category is reported as c and the source is neither over- nor underestimating, then its capability category is c , $c \in \{1, 2, 3\}$
21	If the vessel capability category is reported as c and the source is overestimating and not underestimating, its capability category is $\leq c$, $c \in \{1, 2, 3\}$
22	If the vessel capability category is reported as c and the source is underestimating and not overestimating, its capability category is $\geq c$, $c \in \{1, 2, 3\}$

Table 6.7: Logical relationships in the threat assessment valuation network (not including source models), with parametrised probability values α_i . Values in brackets are the ones used in the current implementation

Nr.	Relationship R	$P(R)$
0	$THREAT = TRUE \leftrightarrow (OPP. = TRUE) \wedge (CAP. = TRUE) \wedge (INT. = TRUE)$	1
1	$AIS_INC = ON \rightarrow INTENT = TRUE$	γ_1 (0.7)
2	$AIS_T = OFF \rightarrow OPPORTUNITY = TRUE$	γ_2 (0.4)
3	$KIN_X = STOP. \rightarrow OPP. = TRUE$	1
3	$KIN_X = LOIT. \rightarrow OPP. = TRUE$	γ_3 (0.4)
4	$TYPE_X = 1 \rightarrow CAP. = FALSE$	1
5	$TYPE_X = 2 \rightarrow CAP. = TRUE$	γ_4 (0.3)
6	$TYPE_X = 3 \rightarrow CAP. = TRUE$	1
7	$ID_X = F \rightarrow INTENT = FALSE$	1
8	$ID_X = H \rightarrow INTENT = TRUE$	1
9	$ID_X = S \rightarrow INTENT = TRUE$	γ_5 (0.7)
10	$ID_X = A \rightarrow INTENT = FALSE$	γ_6 (0.7)
11	$LOC. = CLEAR \rightarrow OPP. = FALSE$	1

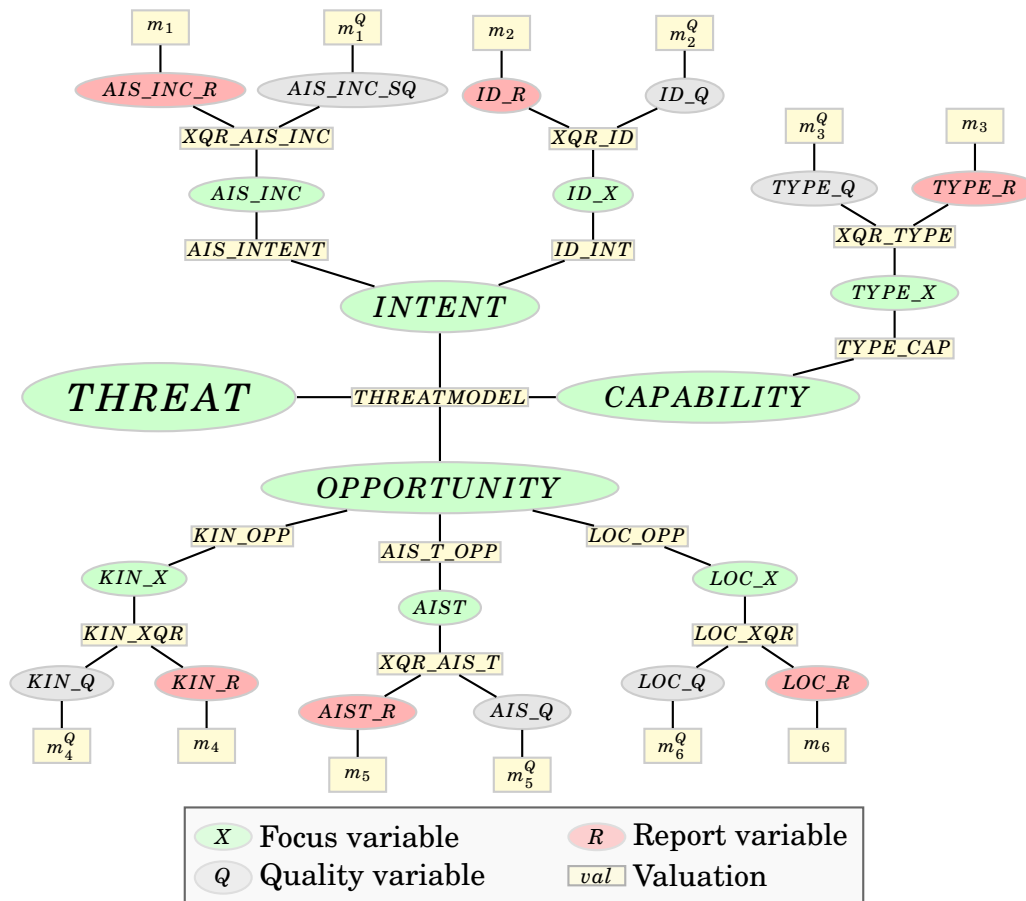


Figure 6.13: The evidential network implementing the multi-source threat assessment reasoning

for the estimation of capability by type (perfect (neither- over or underestimated), overestimated, underestimated, unreliable (either- over or underestimated)). Throughout the remainder of this

Table 6.8: List of sources

Source ID	Source	Variable	Source report
1	AIS inconsistency detector	AIS inconsistency	AIS_INC_R (m_1)
2	Intelligence	Vessel identity	ID_R (m_2)
3	Classifier	Vessel type	TYPE_R (m_3)
4	Kinematic behaviour analyser	Vessel kinematic	KIN_R (m_4)
5	AIS receiver	AIS reception	AIST_R (m_5)
6	Location prediction	Vessel position	LOC_R(m_6)

section, unless specified otherwise, the sources are assumed to behave as follows:

- The AIS is assumed to have 20% loss rate;
- The sources providing location, AIS consistency and vessel identity are all assumed to be 80% reliable;
- The estimation of kinematic behaviour of the vessel is assumed to be 70 % accurate and 90% reliable;
- The estimation of capability by vessel type has an overestimation risk of 10% and an underestimation risk of 20%.

This configuration results in the following set of quality mass functions:

$$\begin{aligned}
m_1^Q(\text{AIS_INC_SQ} = \text{REL}) &= 0.8 & m_2^Q(\text{ID_Q} = \text{REL}) &= 0.8 \\
m_1^Q(\text{AIS_INC_SQ} = \text{NREL}) &= 0.2 & m_2^Q(\text{ID_Q} = \text{NREL}) &= 0.2 \\
\\
m_3^Q(\text{TYPE_Q} = \text{PERF}) &= 0.72 & m_4^Q(\text{KIN_Q} = \text{PERF}) &= 0.63 \\
m_3^Q(\text{TYPE_Q} = \text{OVER}) &= 0.08 & m_4^Q(\text{KIN_Q} = \text{INACC}) &= 0.27 \\
m_3^Q(\text{TYPE_Q} = \text{UNDER}) &= 0.18 & m_4^Q(\text{KIN_Q} = \text{NREL}) &= 0.1 \\
m_3^Q(\text{TYPE_Q} = \text{NREL}) &= 0.02 & & \\
\\
m_5^Q(\text{AIS_Q} = \text{LOSS}) &= 0.2 & m_6^Q(\text{LOC_Q} = \text{REL}) &= 0.8 \\
m_5^Q(\text{AIS_Q} = \text{NLOSS}) &= 0.8 & m_6^Q(\text{LOC_Q} = \text{NREL}) &= 0.2
\end{aligned}$$

6.2.2 Inference results

Four different but similar scenarios involving four different types of vessels are discussed, presenting each with a different probability of threat. The mass functions m_i^j represent observation i for scenario j . In all the scenarios, no report is received regarding whether the AIS signal is consistent, and it is assumed that all the vessels satisfy the geographical constraint, i.e. they are all reported to be in the vicinity of the maritime infrastructure.

Vessel 1: Fishing vessel Here a typical behaviour of a fishing vessel is considered. Such a vessel has some, albeit limited, capability to do harm. It often disables its AIS and its kinematic

behaviour can be considered to be a mixture of passing through and loitering. In this case the information provided is on the AIS broadcast status $m_5^1(\text{AIST_R} = \text{REC}) = 1$, vessel kinematic behaviour $m_4^1(\text{KIN_R} = \text{LOITERING}) = 0.7$ and $m_4^1(\text{KIN_R} = \text{PASSING}) = 0.3$, and capability by vessel type $m_3^1(\text{TYPE_R} = 2) = 1$ with no identity report provided.

Vessel 2: High capability, possibly malicious vessel This vessel is identified as highly capable by its type. It operates with its AIS enabled, stopping and loitering within the area of interest. Given the geopolitical context, the vessel is identified as suspect. As such the observations from the sources are: $m_2^2(\text{ID_R} = \text{S}) = 1$, $m_3^2(\text{TYPE_R} = 3) = 1$, $m_4^2(\text{KIN_R} = \text{STOPPED}) = 0.4$, $m_4^2(\text{KIN_R} = \text{LOITERING}) = 0.6$ and $m_5^2(\text{AIST_R} = \text{REC}) = 1$.

Vessel 3: Assumed friendly deep-sea research vessel with infrastructure repair capabilities The vessel behaviour and capabilities are identical to Vessel 2 ($m_i^3 = m_i^2$, for $i = 3, \dots, 6$) with the only difference that this one is assumed to be friendly according to intelligence reports, $m_2^3(\text{ID_R} = \text{A}) = 1$.

Vessel 4: Leisure craft This is a case of a vessel which should pose no to little threat - its intent is unknown, opportunity and capability low. The AIS is being continuously broadcast and its kinematic behaviour and capability are typical for leisure craft. As such: capability by vessel type $m_3^4(\text{TYPE_R} = 1) = 1$, vessel kinematic behaviour $m_4^4(\text{KIN_R} = \text{STOPPED}) = 0.1$, $m_4^4(\text{KIN_R} = \text{PASSING}) = 0.9$ and AIS reception status $m_5^4(\text{AIST_R} = \text{REC}) = 1$ with no identity report provided.




Figure 6.14: Amount of evidence *supporting* (belief, lower bound), amount of evidence *consistent with* (plausibility, upper bound) and amount of evidence *against* THREAT (top graph) and its three components (bottom graph), for each of the four vessels

Inference results are shown in Figure 6.14 and can be analysed as follows. It is clear that the only vessel which poses a significant risk is Vessel 2 - it is fully capable, most likely has malicious intent, and there is significant evidence of its opportunity. Additionally, no evidence argues against malicious intent, its opportunity or its capability. Vessel 1 can be considered a plausible risk (high upper bound) due to its opportunity and degree of capability. However, there is no evidence of it having malicious intent. It is comparable to vessel 4, which intent, too, is unknown. However, the belief in the capability of the latter is much lower (as it is entirely driven by the likelihood of its capability having been underestimated), and its opportunity is reduced due to its kinematic pattern and AIS broadcast. Finally, the lowest threat is posed by Vessel 3. While it has both the capability and opportunity, just like Vessel 2, it has been positively confirmed to be a friend and, not pose a threat.

Interestingly, the model allows capturing a quite high uncertainty in the threat assessment for both Vessels 1 and 2, suggesting gathering additional information for a refined assessment. It can be argued that the threat model used underestimates the threat, as the belief in *Threat* is lower than the belief in any of the factors affecting it. However, the evidence “against” is well reflected in the threat. A more suitable model for threat assessment in this context may need to be proposed. This conservative behaviour of the model with respect to belief in threat can be considered to correspond to the actual threat level as discussed earlier and its plausibility to the potential threat level.

6.2.3 Explaining the impact of source reports

Here we consider the case of Vessel 3 as discussed earlier, but with an additional piece of information. Throughout the remainder of this section, we refer to this vessel as Vessel 3* to distinguish it from the case of Vessel 3 discussed earlier. We consider a report claiming that the vessel’s AIS is being inconsistent: $m_1^{3*}(\text{AIS_INC} = \text{TRUE}) = 1$. This increases the support in hypothesis $\text{INTENT} = \text{TRUE}$ and by extension $\text{THREAT} = \text{TRUE}$, yielding the following inference:

$$\begin{aligned}
 m^{3*}(\text{THREAT} = \text{TRUE}) &= 0.1442 \\
 m^{3*}(\text{THREAT} = \text{FALSE}) &= 0.3590 \\
 m^{3*}(\text{THREAT} = \Omega) &= 0.4968
 \end{aligned}
 \tag{6.6}$$

As discussed in the previous section, the formal explanation and analysis of these results can be done in two steps. The first one is to discuss the impact of the reports provided by the individual sources of evidence, and the second one is to consider the impact of the source model.

6.2.3.1 The contribution of the reports

We first discount the report valuations to analyse the impact of the report itself without considering the source model. The sensitivity spaces are generated as discussed previously and are laid out in Figures 6.15a and 6.15b. To consider the explanation on a higher level, we also look at the contribution of Intent, Opportunity and Capability to Threat. This is displayed in Figure 6.16. Finally, to close the loop, we may consider the contribution of the individual sources to inference on

Intent as displayed in Figure 6.17. We will not discuss the inference on Opportunity and capability in more detail as they are all driven by a single source only.

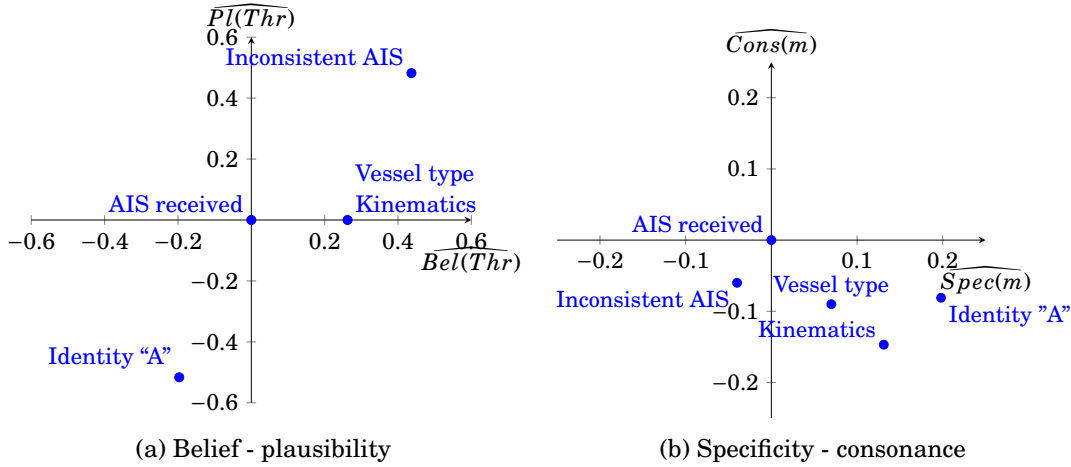


Figure 6.15: Sensitivity spaces for the impact of the individual reports on inference on THREAT

The sensitivity spaces are spanned by the measures defined in Eq. (4.7) and (4.8) in Figure 6.15a. In Figure 6.15a pieces of evidence in the top-right quadrant argue for a threat, pieces of evidence in the bottom-left quadrant argue against a threat. The top-left and bottom-right quadrant show pieces of evidence that increase and decrease ignorance, respectively. As such, it can be argued that the reception of the AIS broadcast of the vessel is inconsequential, the vessel’s type and its kinematic behaviour in the area of interest both argue for the threat. The vessel’s affiliation (Assumed friend) and its inconsistent AIS broadcast argue respectively against and for the threat while also disputing the other argument.

It is interesting to see that both mass functions related to the vessel’s kinematic behaviour and its type are equally sensitive to discounting and have no effect on the plausibility of threat, but both the vessel identity as well as the AIS inconsistency have a significant impact on $Pl(Thr)$. The overall belief and plausibility values of threat are determined by belief and plausibility levels of opportunity, capability and intent. Mass functions m_3 and m_4 solely determine capability and opportunity, respectively, since no evidence against capability or opportunity exists (they are both fully plausible), they have no effect on threat plausibility. On the contrary masses m_1 and m_2 both determine intent, however m_1 supports the hypothesis “INTENT = TRUE” whereas m_2 opposes it. The conflict between the two means that discounting one reinforces the other one, therefore, changing, increasing, or decreasing the plausibility of malicious intent and by extension of threat. This is consistent with Figure 6.16a, where opportunity and capability correspond to kinematics and vessel type. The overall sensitivity is slightly lower than when the sources are considered directly, as the reports have already been corrected by the source model. Similarly, the conflict which causes AIS inconsistency to reduce specificity in Figure 6.15b is resolved in 6.16b. Similarly, when comparing 6.17a to 6.15a, we observe similar values but scaled down due to the source correction and the threat model.

The measures 6.4 and 6.5 are shown in Figure 6.15b. Pieces of evidence at the point of origin




Figure 6.16: Sensitivity spaces for the impact of INTENT on inference on TREAT




Figure 6.17: Sensitivity spaces for the impact of the individual reports on INTENT

have no impact on the reasoning result, i.e. the reasoning result is not sensitive against the addition or removal of it. All the pieces of evidence increase the dissent between the arguments. As per the previous example, masses, m_3 and m_4 have the same behaviour. It is interesting to see that mass function m_1 has a significantly smaller effect on dissent between arguments whilst also decreasing specificity. It is expected that a decrease in specificity will also decrease disagreement between the arguments, but it is counter-intuitive that this mass should decrease the overall specificity. This can be again explained by the conflict dynamics—the two mass functions related to intent m_1 and m_2 conflict with one another. Meanwhile, the evidence supporting the hypothesis that the vessel does not have malicious intent reduces the nonspecificity of Threat more than evidence to the contrary. From Figure 6.17b it can be seen that the two mass functions have the same impact on the specificity of Intent, but it is not the case for Threat (Figure 6.15b)). The same impact of the two masses can be justified by them being in conflict with one another and only with one another. Reinforcing m_1 essentially discounts m_2 , resulting in a net decrease in specificity, as discussed in the previous paragraph.

6.2.4 Explaining the source model

The difficulty with extending this approach directly to the analysis of the impact of the source model is that the former may have multi-dimensional quality valuations. For example, whereas the approach can be applied very easily to the sources such as "Inconsistent AIS", where the valuation on quality can be discounted directly, it is more problematic for sources such as capability assessment, where the quality may change in multiple directions.

Consider the classic example proposed by Pichon et al. [82], where the same source has a behaviour that depends on the two dimensions of reliability and truthfulness. The one-dimensional description from earlier is not helpful anymore, as the only option would be to display the variation on different sensitivity spaces.

The behaviour of the source is now governed by a variable on the joint frame $\Omega_Q = \{rel \times truth, rel \times lie, \neg rel \times truth, \neg rel \times lie\}$, and defined by the mapping

$$\begin{aligned}\Gamma_{rel \times truth}(A) &= A, \quad \forall A \subseteq \Omega \\ \Gamma_{rel \times lie}(A) &= \bar{A}, \quad \forall A \subseteq \Omega \\ \Gamma_{\neg rel \times truth}(A) &= \Gamma_{\neg rel \times lie}(A) = \Omega, \quad \forall A \subseteq \Omega\end{aligned}$$

with the marginal mass functions $m(rel) = \alpha$, $m(\neg rel) = 1 - \alpha$, $m(truth) = \beta$, $m(lie) = 1 - \beta$, with α and β within $[0, 1]$ and representing respectively the probability of the source being reliable and truthful. This yields the joint mass function on Q :

$$(6.7) \quad m_Q(Q) = \begin{cases} \alpha\beta & \text{if } Q = rel \times truth \\ \alpha(1 - \beta) & \text{if } Q = rel \times lie \\ 1 - \alpha & \text{if } Q = \{\neg rel \times truth, \neg rel \times lie\} \end{cases}$$

Now we no longer can simply produce the derivative of Bel as in 4.7 with respect to a single parameter, since now $Bel(A)$ is a function of both α and β .

Hence we consider $Bel(A, \alpha, \beta)$ to be a scalar field over the space $\alpha \times \beta$, and compute the gradient derivative of the scalar field at that point:

$$(6.8) \quad \nabla Bel_i(A) = \begin{pmatrix} \frac{\delta Bel_i(A)}{\delta \alpha_i} \\ \frac{\delta Bel_i(A)}{\delta \beta_i} \end{pmatrix}$$

which corresponds to the vector of the greatest increase in gradient at that point, which magnitude is used as the sensitivity value:

$$(6.9) \quad \widehat{Bel}_i(A) = |\nabla Bel_i(A)|_{\substack{\alpha_i = \alpha_{i0} \\ \beta_i = \beta_{i0}}}$$

The directional derivatives can be used to assess the sensitivity of the computation to a specific type of variation in quality - making the source either more or less truthful, or more or less reliable. The directional derivative defines the magnitude of gradient in direction of \vec{v} :

$$(6.10) \quad \nabla_v Bel(A) = \nabla Bel(A) \cdot \hat{v}$$

and thus

$$(6.11) \quad \widehat{Bel}_i(A)_v = \nabla_v Bel_i|_{\substack{\alpha_i=\alpha_{i0} \\ \beta_i=\beta_{i0}}} = \nabla Bel_i(A)|_{\substack{\alpha_i=\alpha_{i0} \\ \beta_i=\beta_{i0}}} \cdot \hat{v}$$

where \hat{v} denotes the unit vector in direction \vec{v} .




Figure 6.18: Source behaviour as variation of α and β , with directions corresponding to maximum positive changes of belief, plausibility, specificity and consonance

Figure 6.18 displays the different directions in which quality can vary in the $\alpha - \beta$ space. Some directions of particular interest are these where the source directly tends towards one of the unique source behavioural modes. – in this case, the fully reliable and truthful source (“Perfect”), a reliable liar (“Liar”) or an unreliable source (“Unreliable”), as well as the directions of maximum changes of $Bel(X = True)$, $Pl(X = True)$, $Cons(m)$ and $Spec(m)$. Such direction is the unit vector from the initial source behaviour configuration to the point corresponding to the behaviour of interest:

$$(6.12) \quad \vec{v} = \frac{\begin{pmatrix} \alpha_1 \\ \beta_1 \end{pmatrix} - \begin{pmatrix} \alpha_0 \\ \beta_0 \end{pmatrix}}{\left\| \begin{pmatrix} \alpha_1 \\ \beta_1 \end{pmatrix} - \begin{pmatrix} \alpha_0 \\ \beta_0 \end{pmatrix} \right\|}$$

where the point (α_1, β_1) corresponds to the source configuration of interest: $(1, 1)$ for the perfect source, $(1, -1)$ for the liar and $(-1, 0)$ for the unreliable source².

By extension we consider all the sensitivity loci, which trace an ellipse:

$$(6.13) \quad \left(\widehat{Bel}_i(A) \right)_{(\theta)} = \begin{pmatrix} \nabla Bel_i(A) \cdot \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} \\ \nabla Pl_i(A) \cdot \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} \end{pmatrix} = \begin{pmatrix} \frac{\delta Bel_i(A)}{\delta \alpha_i} \cos \theta + \frac{\delta Bel_i(A)}{\delta \beta_i} \sin \theta \\ \frac{\delta Pl_i(A)}{\delta \alpha_i} \cos \theta + \frac{\delta Pl_i(A)}{\delta \beta_i} \sin \theta \end{pmatrix}$$

²It can be argued that the unreliable source should correspond to the line $\alpha = 0$, but we choose to ignore this for the sake of clarity

where θ denotes the direction of change in quality on $\alpha \times \beta$. The resulting sensitivity spaces with $\alpha_0 = 0.6$ and $\beta_0 = 0.6$ are shown in Figures 6.19a and 6.19b. It can be clearly seen that, as expected, if the conflicting source is treated as perfect it has a significant negative effect on both belief and plausibility of $X = True$, a minor increase in specificity and a very large decrease in consonance. Treating the source as a Liar increases the belief of $X = True$ as well as the specificity, with marginal negative impact on consonance. Simply reducing the relevance of the report increases the plausibility as well as the consonance, while marginally decreasing the belief and specificity. It is interesting to note that treating the three main quality directions are not symmetric, i.e. they all have distinct effects.

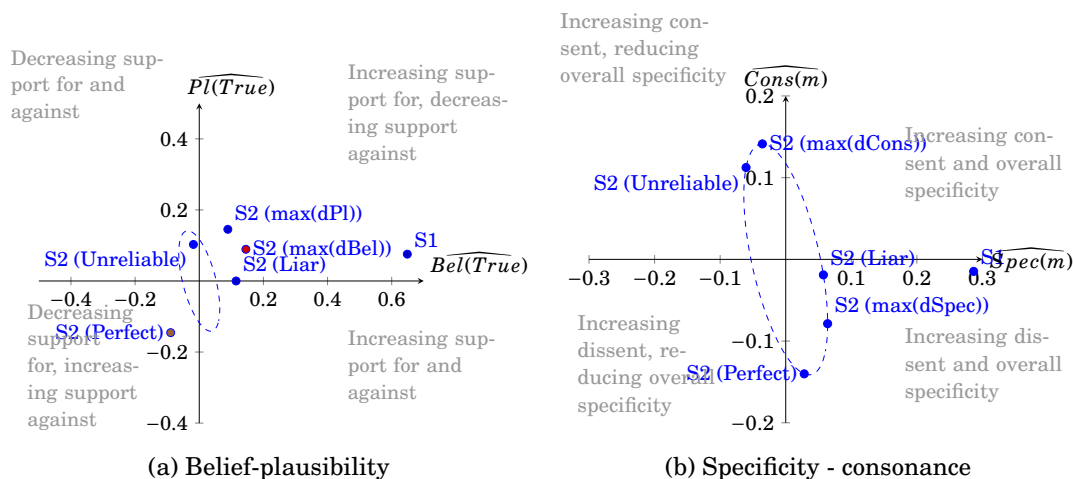


Figure 6.19: Sensitivity spaces for belief and plausibility with two-dimensional quality valuation and different directions of source uncertainty shift for source 2

The approach described here can be naturally applied to the more general BBC model with a set of arbitrary source behaviours. For this, each behaviour should be parametrised with a set of n binary variables $Z = \{Z_1 \dots Z_n\}$, each being true with probability p_i and false with probability $1 - p_i$ such that there exists a truth table relating each configuration of Z to a possible source behaviour.

This is the way we have chosen to parametrise the behaviours of both capability estimation (with the two parameters being the probability of underestimation and overestimation), as well as the kinematic behaviour analysis (the parameters being accuracy and reliability).

We apply this to the example in question. Unlike in the example discussed earlier, the lack of external conflict between many of the sources means that there is the little possible impact of changing the quality valuations on the plausibility of threat as well as a very strong correlation between change in consonance and specificity. This means that the short radius of the ellipse is so small that the shape traced is essentially a straight line segment. Therefore, depending on the direction of change of quality, the impact on inference result can lie anywhere on this line segment. These sensitivity spaces are shown in Figure 6.20.

By inspection and comparison of Figures 6.20a and 6.20b to Figures 6.15a and 6.15b respectively it can be seen that both AIS inconsistency as well as vessel identity's impact are similar

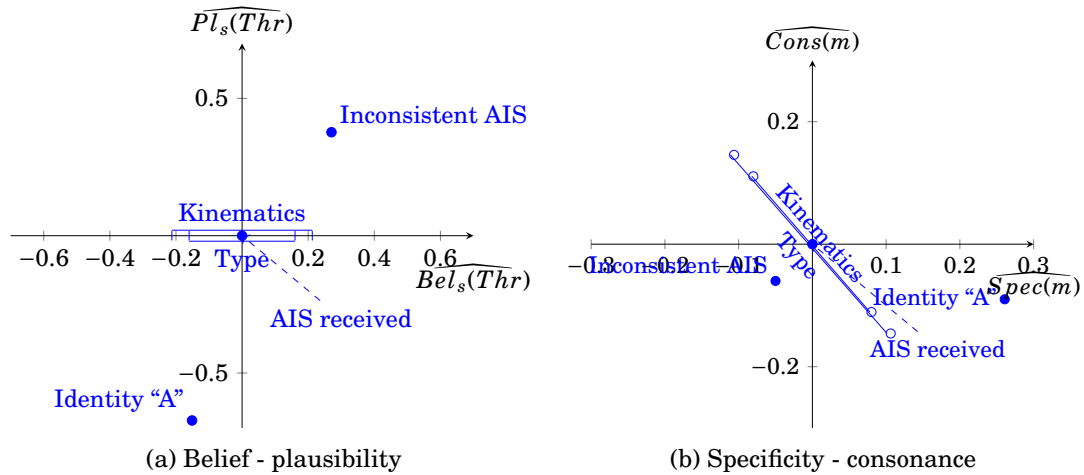


Figure 6.20: Sensitivity spaces for inference on Threat where the parameters describing source behaviour are varied

but shifted down (more negative and less positive) for plausibility and its magnitude is reduced for belief. In general, the inference gradient with respect to reliability is expected to be slightly smaller than discounting the report directly since it takes into account the initial conditions. Due to the way conflict is redistributed, as discussed before, this makes the impact on plausibility more negative or less positive.

The impact of changing the quality estimate of kinematics or type can lie anywhere on the line segment marked on the sensitivity space. Note that the sensitivity values from Fig. 6.15a can be made negative if the source is discounted rather than reinforced. The corresponding positive and negative values bound the line segment for kinematics and vessel type from Figure 6.20a. As such, the maximum sensitivity to the source model is less than the sensitivity to the reports. Similar arguments follow when considering the specificity-consonance sensitivity in Figure 6.20b. However, it is interesting to note the strong correlation between impact on consonance and specificity for this kind of conflict-free source. The real added value of this type of analysis comes from the ability to assess the impact an incorrect source assessment may have on the inference result. Considering a finite set of resources, which could be used to improve some of the sources, it can be shown that the ones the quality of which may have the greatest impact should be the ones to be prioritised. A downside of this approach is that it is impossible to distinguish the different directions of quality shift only through inspection of the sensitivity space.

Finally, we consider the overall impact of the source model. Whereas the previous figure shows the impact of source quality assessment, Figure 6.21 enables understanding of the overall impact of the model. It shows the difference in sensitivity of the inference process to reports when the sources are considered to be perfect (red) and when the source model is used (blue). A perfect source is one with no correction due to the source model, i.e. its report is assumed to accurately represent the knowledge we can obtain from it. Thus it can be seen that employing the source model generally dramatically reduces the sensitivity to individual reports while increasing the impact on the overall consonance for the two conflicting sources (Identity and AIS inconsistency).

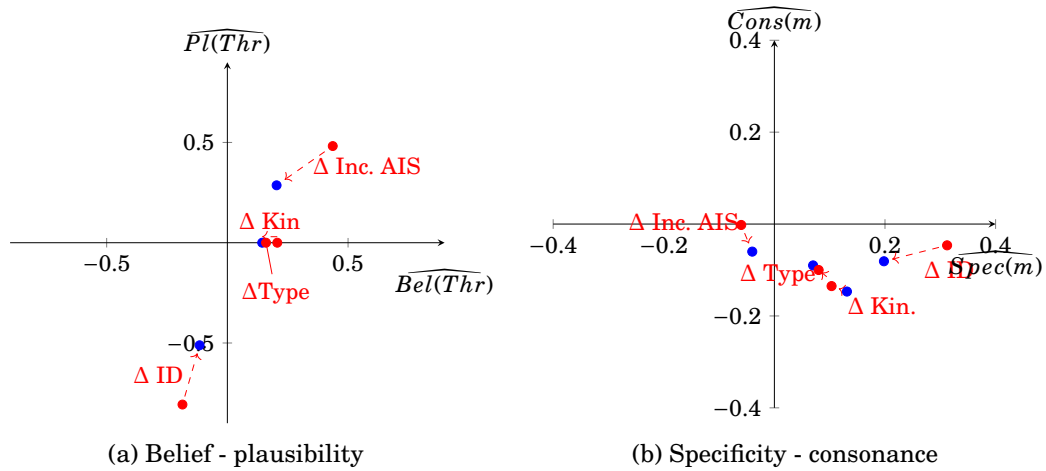


Figure 6.21: The comparison of sensitivity of the inference result to the report with (blue) and without (red) the source model included

This also shows that the impact of the source model is relatively low for non-conflicting sources (vessel type and kinematics).

6.3 Summary

In this section, two scenarios for the fusion of uncertain information within a maritime context were presented. The first was concerned with situational awareness and sense-making on information regarding multiple entities and the relationships between them. The framework of uncertain conceptual graphs was used for this purpose. The reasoning capabilities were provided using the link between uncertain conceptual graphs and valuation networks. These tools were used to obtain a distribution on possible worlds showing the relationships between all the entities, each of them described by a conceptual graph. This has shown the applicability of the approach in a relatively realistic scenario. These results were subsequently analysed using some explainability tools. Specifically sensitivity spaces were produced to describe the impact of individual sources on the belief and plausibility of hypotheses of interest, as well as their impact on conflict and epistemic uncertainty in the process. Shapley value was used to assess the impact of individual sources to the overall solution. The explanations produced were consistent with the expectations.

Typically, we would like to assess the processes used throughout both scenarios with respect to the ground truth. In both cases presented we can see that the inference result is consistent with the ground truth, however since the scenarios themselves are synthetic and the sensing process is not modelled it is insufficient to use this to assess the reasoning process. Furthermore, the purpose of both of these scenarios is to showcase methods and tools. If an inconsistency existed it would be difficult to address whether it is a weakness of the tools themselves - the valuation networks or the conceptual graphs - or rather the problem model. For example, in Scenario 2, the threat model used is likely under-sensitive and could be modified if deployed. However, this weakness of the problem model does not imply a problem with the tools and methods.

The second scenario is the one presented in the conference paper at 2020 International Conference on Information Fusion [113]. It is concerned with assessment of the threat a vessel may pose to sensitive underwater infrastructure, such as the subsea communications cables. The focus of this scenario is on modelling the partial reliability of sources, as well as its impact on the reasoning. Results for four different vessels are presented, and one vessel of particular interest is analysed in greater detail. A novel method for assessment of the latter is proposed, by extending the notion of sensitivity space to multi-dimensional source quality models. This makes it possible to assess what would the impact of the source be if the assessment of its quality would change. As such, it gives us an estimation of the impact of the second-order uncertainty, or the uncertainty regarding the quality of the sources. This can be useful in a real-life scenario, where we may have a finite amount of resources and it is necessary to prioritize which sources should be improved. The method proposed makes it possible to assess how much the source impact would change in this case.

DISCUSSION

This thesis addressed the challenges in the field of uncertain information fusion through three key lenses - that of context exploitation for partially reliable sources, explainability and situation awareness.

7.1 Thesis summary and next steps

The notion of context has been discussed in Chapter 3 and aligned with relevant works in the field of belief functions. The contextual reasoning model used currently in literature has been framed as *context-of* the problem, as opposed to *context-for* the problem. A novel method of inclusion of *context-for* in evidential reasoning has been proposed. Furthermore, we have partitioned the universe of discourse into the context of the problem and the context of the sources. Individual pieces of contextual information can exist within either context while acting as *context-for* either. These two behaviours correspond to whether a piece of contextual information makes it possible to reason about the source quality or about the variable of interest itself. This approach supports interpretability of evidential reasoning, improves the reasoning results and can even increase the expressiveness of contributing sources.

Furthermore, in this chapter, other issues related to source modelling were addressed. This includes dealing with sources that are not independent. If such a situation occurs and is not addressed, double counting of information takes place. While solutions for a combination of sources with known dependence relation exist, alongside ones that can be used when this relation is unknown, an alternative solution has been proposed. An interpretation of the Behaviour-Based Correction model (or as referred to throughout this thesis, its variant, the XQR model) was proposed to allow for uncertain dependency relations between contributing sources.

Several exciting areas of work remain that need to be omitted. One of them is the link between contextual reasoning and source dependency. A natural extension is to consider that a source dependency relation is conditional given some context. A trivial example is a set of data available

to an analyst. It is reasonable to consider this shared data to be the contextual information for the analyst. Consider that there are multiple analysts, and they all operate on the same information. In this case, their reports are not going to be independent. This is of particular interest in a multi-agency collaboration scenario, where the analysts may be unaware that they use the same information.

Smets' canonical decomposition of belief functions has been used for interpretation of results. It is a great tool to demonstrate the added expressiveness, but there is room for greater exploitation of it. One way inclusion of context increases the expressiveness of the belief mass has been shown (using the canonical decomposition), but there are other types of pseudo-BBA's that can be induced by contextual information. An example is the one discussed in the last section of Chapter 3.

Chapter 4 focuses on the second challenge for multi-source information fusion, that of provenance or explainability. The purpose of addressing this is that while graphical evidential reasoning is transparent by design, the process of information aggregation itself can be opaque, especially with multiple sources and large frames of discernment. If this process is used to make a decision, especially in safety-critical applications or where large stakes are associated with the outcome of such a decision, it is necessary to be able to understand which of the sources are driving this decision. The contribution measure based on a geometric interpretation of belief masses was proposed. As this was a novel approach, this chapter is largely exploratory, and several possible evaluation methods are proposed. One which turned out to be of particular importance is the preference property. Originally its formalisation suggested that it should be satisfied by any contribution measure. However, a more in-depth analysis showed that not only it is not satisfied by default, but for any contribution measure considered, there exists a proportion of cases where this property is not satisfied at all. As such, this property has been reinterpreted to be treated as a performance evaluation tool, where the lower proportion of cases where it does not hold, the better the performance of the particular contribution measure. Interestingly, the distance-based contribution measure proposed performed relatively well regardless of the decision-making method and the combination rule used. While in all cases it has been outperformed by some other contribution measure, it always showed at least satisfactory performance, as opposed to other possible measures which tended to perform very well in optimal conditions but significantly worse in others.

With this chapter largely treading new ground, there is a great breadth of possible further work. The postulates for the properties the contribution measure should satisfy should be revisited. While in some cases, the notion of "greatest impact" can be self-explanatory, in others, it is not. In particular, if we consider more complex valuation networks, source correction models and similar methods, it may be very hard to agree on what is meant by the greatest impact of a particular body of evidence. An example of such a case is in Chapter 7, where the threat model based on the conjunction of three variables is used. Clearly, all three are required for inference on the threat, and arguably the one with the greatest impact is the least-committed one, as it reduces the plausibility of threat. This is a major difficulty with tracking uncertainty as well. While in some cases, the uncertainty can come from conflict redistribution, in others, it may be due to source correction or source self-doubt. In order to track this correctly, this separation between lack of

information (source self-doubt) and low source quality should be addressed correctly.

Chapter 5 addresses the problem of multi-entity situation modelling. It draws very heavily from the work of Laudy et al., but with a much greater focus on the modelling of uncertainty itself. An extension of the framework of conceptual graphs is proposed to allow for embedding uncertainty in the graphs themselves. Two reliability levels are included, which make it possible to consider an alternative interpretation of the notion of partial reliability. Since a source provides multiple interconnected pieces of information, a partially reliable source can provide some valid information, whilst some can be false. This approach contrasts with previous work, where the sources were given a reliability rating, but it corresponded to the likelihood that all the information provided is true. Based on this, two methods of combining information from such uncertain graphs are proposed. The second half of that chapter addresses the links between the uncertain conceptual graphs and other uncertainty formalisms, in particular valuation networks. This makes it possible to use the conceptual graphs framework for information storage and some information fusion and the valuation networks for uncertain reasoning on the same problem. Furthermore, the uncertain conceptual graph framework is compared to Markov Logic Networks, as the two are intuitively linked. It is shown that the same problem can be solved with MLN's as well. It can be argued that the example discussed does not leverage the relative strengths of different approaches. However, it does show that the different approaches are compatible.

An interesting area of work that has been omitted in that chapter is the relationship between valuation networks and Markov Logic Networks. An MLN induces a Markov network, and the set of rules for an MLN can be used to generate a set of valuations over the same variables which exist in the Markov network. This suggests that it is possible to generate an equivalent valuation network for any Markov Logic Network. This is something that should be addressed in the future to assess the feasibility and usefulness of such an approach. Another area of further work is that of extending the notions from the field of multi-target association to the fusion of conceptual graphs. In particular, the models used in Chapter 5 do not sufficiently address the question of ensuring that the identity of targets matches, as well as targets having more than one property of the same time. Since the purpose of using conceptual graphs is to deal with multi-entity fusion, it is important to ensure that these "semantic tracks" are correctly associated with the corresponding targets.

These three key areas of interest - context, explainability and situation awareness - are brought together in Chapter 6. There two scenarios from the maritime situation awareness domain are discussed. One is a multi-entity situation awareness scenario involving several vessels involved in illegal fishing. This demonstrates the ability to reason about several entities simultaneously, exploit context and provide some level of explanations. The second scenario is focused on a single vessel that is likely to pose a threat to sensitive maritime infrastructure. In this case, the main focus is on more complex reasoning, explainability and linking explanation abilities with the partially reliable source model. The notion of sensitivity spaces is extended to allow quantification of the impact of the source model itself. Both of these scenarios show that these three areas of research can be used together, and, in fact, a significant synergy exists between them. While the examples have been designed to be realistic, they are nevertheless synthetic and suffer from

several oversimplifications. For example, the threat model itself is clearly overly conservative. The multi-entity illegal fishing scenario, too, operates on some assumptions which would not necessarily be adequate in a real-life situation. Nevertheless, the two examples show that the tools proposed throughout this thesis have the ability to solve a certain type of problem.

7.2 Further research and relation to other areas within Artificial Intelligence

There are several areas that could be investigated next. In this thesis, it was originally intended to place a greater focus on source heterogeneity. Although it was addressed to some extent in Chapter 3, it could have been investigated in more detail. In particular, the fusion of semantic information was only briefly mentioned in Chapter 5, in the light of situational awareness. In recent years, the general field of natural language processing has been advanced significantly, in particular through deep learning techniques. These can be applied to information extraction from text much more easily than at the time when work on this thesis began. Therefore a natural next step is to re-address this work and to align the deep neural network-based relation extraction and ontology population with the tools for the fusion of semantic information outlined in Chapter 5.

The utilisation of deep learning approaches was briefly discussed in Chapter 3, but only as an example of a partially reliable source. With greater work on explainable and interpretable deep learning models, as well as the more recent work on evidential deep learning, new challenges for integrating these with evidential networks arise. This applies to contextual information, too, as some works in integration context in neural networks have been undertaken throughout recent years. It is important to note that it is a well-known issue with deep learning models (and neural networks in general) that misclassification can take place due to imbalanced contextual information (such as time of the day or season in computer vision tasks). The selection of ensembles of classifiers by context is another example of context utilisation in machine learning, the use of which is strongly related to the work presented in Chapter 3.

Decision-making in the theory of evidence was very briefly outlined in Chapter 2 and then recalled in Chapter 4. Very clearly, there is much more to that. With a possible advent of evidential deep learning, it is likely that evidential decision-making will be an area of interest in the near future. A particular topic of interest is that of information value. In a Bayesian approach, information value is associated with the expected increase in utility when a source is sampled. However, this problem becomes more complicated in the evidential framework. If we consider a partially reliable source that can provide information with a degree of self-doubt, we may have a situation where the information provided by the source decreases the expected utility but reduces the interval between maximum and minimum utility. If we consider that ambiguity aversion is an important aspect of decision making in human agents, it stands to reason that the reduction in epistemic uncertainty associated with a decision should in itself be considered a form of utility. This becomes particularly important in light of the discussion in Chapter 6 - where we are concerned with the possible cost of improving a source. Another area related to the aforementioned problem

is confirmation bias. This is something that can be positioned between the discussions in Chapter 3 (on source behaviour and partially independent sources) and Chapter 4 (explainability). In a multi-intelligence scenario, it is possible that the sources which are independent are more likely to report on a particular area in the frame of discernment. This ties in with the information value problem, as well as with context, which can shed light as to which areas of the frame of discernment are likely to be over-represented in the reports.

This thesis began with a very open-ended research question: "How can sensemaking in a multi-intelligence scenario, with uncertain and possibly unreliable sources, be improved?". I believe that the three key areas addressed in this thesis: context, explainability and situational awareness, framed in the light of uncertainty modelling, form a major part of the answer to this question. The challenge of using multiple sources to analyse multiple targets, as described in the introduction, is essentially the definition of situational awareness. Context exploitation fits perfectly into the research question, too, as, on the one hand, it can be used to reduce the complexity of the problem and, on the other, to better interpret the information provided by unreliable sources. Finally, a degree of explainability is necessary to make these solutions applicable in a scenario where misinformation may have tangible costs. The three concepts are seemingly unrelated but given a more in-depth investigation, they turn out to be significantly intertwined, and it becomes difficult to discuss one without addressing the others, and they are all equally important for multi-intelligence sensemaking.

BIBLIOGRAPHY

- [1] D. Dubois, W. Liu, J. Ma, and H. Prade, “The basic principles of uncertain information fusion. An organised review of merging rules in different representation frameworks,” *Information Fusion*, 2016, ISSN: 15662535. DOI: 10.1016/j.inffus.2016.02.006.
- [2] G. L. S. Shackle, *Decision order and time in human affairs*. Cambridge University Press, 2010.
- [3] F. Cuzzolin, *The geometry of uncertainty : the geometry of imprecise probabilities*. Springer, 2021, ISBN: 978-3-030-63152-9.
- [4] K.-D. Sommer, O. Kühn, F. P. León, and B. R. Siebert, “A Bayesian approach to information fusion for evaluating the measurement uncertainty,” *Robotics and Autonomous Systems*, vol. 57, no. 3, pp. 339–344, 2009. DOI: 10.1109/MFI.2006.265657.
- [5] R. E. Kass and L. Wasserman, “The selection of prior distributions by formal rules,” *Journal of the American Statistical Association*, vol. 91, no. 435, pp. 1343–1370, 1996. DOI: 10.2307/2291752.
- [6] J. Sander and J. Beyerer, “Bayesian fusion: Modeling and application,” in *Sensor Data Fusion: Trends, Solutions, Applications (SDF), 2013 Workshop on*, IEEE, 2013, pp. 1–6. DOI: 10.1109/SDF.2013.6698254.
- [7] B. Chen and P. K. Varshney, “A bayesian sampling approach to decision fusion using hierarchical models,” *IEEE Transactions on signal processing*, vol. 50, no. 8, pp. 1809–1818, 2002. DOI: 10.1109/TSP.2002.800419.
- [8] W. R. Gilks, “Markov chain monte carlo,” in *Encyclopedia of Biostatistics*. American Cancer Society, 2005, ISBN: 9780470011812. DOI: <https://doi.org/10.1002/0470011815.b2a14021>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/0470011815.b2a14021>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/0470011815.b2a14021>.
- [9] S. Blackburn, *The Oxford Dictionary of Philosophy*. Oxford: Oxford University Press, 2008, ISBN: 9780199541430.
- [10] B. Khaleghi, A. Khamis, F. O. Karray, and S. N. Razavi, “Multisensor data fusion: A review of the state-of-the-art,” *Information Fusion*, vol. 14, no. 1, pp. 28–44, 2013, ISSN: 15662535. DOI: 10.1016/j.inffus.2011.08.001.

- [11] D. Dubois and H. Prade, "Possibility theory and data fusion in poorly informed environments," *Control Engineering Practice*, vol. 2, no. 5, pp. 811–823, 1994. DOI: 10.1016/0967-0661(94)90346-8.
- [12] —, "Possibility theory in information fusion," in *Information Fusion, 2000. FUSION 2000. Proceedings of the Third International Conference on*, IEEE, vol. 1, 2000, PS6–P19. DOI: 10.1109/IFIC.2000.862412.
- [13] H. Borotschnig, L. Paletta, and A. Pinz, "A comparison of probabilistic, possibilistic and evidence theoretic fusion schemes for active object recognition," *Computing*, vol. 62, no. 4, pp. 293–319, Jun. 1999, ISSN: 0010-485X. DOI: 10.1007/s006070050026. [Online]. Available: <http://dx.doi.org/10.1007/s006070050026>.
- [14] M. Oussalah, H. Maaref, and C. Barret, "New fusion methodology approach and application to mobile robotics: Investigation in the framework of possibility theory," *Information Fusion*, vol. 2, no. 1, pp. 31–48, 2001. DOI: 10.1016/S1566-2535(01)00021-5.
- [15] S. Benferhat and C. Sossai, "Reasoning with multiple-source information in a possibilistic logic framework," *Information Fusion*, vol. 7, no. 1, pp. 80–96, 2006. DOI: 10.1016/j.inffus.2005.01.006.
- [16] L. Liu, R. Hao, and S. Cheng, "A possibilistic analysis approach for assessing environmental risks from drinking groundwater at petroleum-contaminated sites," *Journal of Environmental Informatics*, vol. 2, no. 1, pp. 31–37, 2003. DOI: 10.3808/jei.200300015.
- [17] G. Shafer *et al.*, *A mathematical theory of evidence*. Princeton University Press, 1976, vol. 1.
- [18] G. Shafer, "Dempster-shafer theory," *Encyclopedia of artificial intelligence*, pp. 330–331, 1992.
- [19] D. Mercier, B. Quost, and T. Denœux, "Contextual discounting of belief functions," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2005, ISBN: 3540273263. DOI: 10.1007/11518655_47.
- [20] A. P. Dempster, "Upper and Lower Probabilities Induced by a Multivalued Mapping," *The Annals of Mathematical Statistics*, vol. 38, no. 2, pp. 325–339, 1967. DOI: 10.1214/aoms/1177698950.
- [21] P. Smets, "Belief Functions: The Disjunctive Rule of Combination and the Generalized Bayesian Theorem," *International Journal of Approximate Reasoning*, 1993, ISSN: 0888613X. DOI: 10.1016/0888-613X(93)90005-X.
- [22] R. R. Yager, "On the dempster-shafer framework and new combination rules," *Information Sciences*, 1987, ISSN: 00200255. DOI: 10.1016/0020-0255(87)90007-7.
- [23] D. Dubois and H. Prade, "On the unicity of Dempster rule of combination," *International Journal of Intelligent Systems*, vol. 1, no. 2, pp. 133–142, 1986. DOI: 10.1002/int.4550010204.

- [24] T. Denœux, “Conjunctive and disjunctive combination of belief functions induced by nondistinct bodies of evidence,” *Artificial Intelligence*, 2008, ISSN: 00043702. DOI: 10.1016/j.artint.2007.05.008.
- [25] P. Smets, “The Canonical Decomposition of a Weighted Belief,” in *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 2*, ser. IJCAI’95, Montreal, Quebec, Canada: Morgan Kaufmann Publishers Inc., 1995, pp. 1896–1901, ISBN: 1558603638.
- [26] A. L. Jousselme and P. Maupin, “Distances in evidence theory: Comprehensive survey and generalizations,” in *International Journal of Approximate Reasoning*, 2012. DOI: 10.1016/j.ijar.2011.07.006.
- [27] F. Cuzzolin, “A geometric approach to the theory of evidence,” *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 38, no. 4, pp. 522–534, 2008.
- [28] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, Inc., 2007, ISBN: 1596930926, 9781596930926.
- [29] A. L. Jousselme, D. Grenier, and É. Bossé, “A new distance between two bodies of evidence,” *Information Fusion*, 2001, ISSN: 15662535. DOI: 10.1016/S1566-2535(01)00026-4.
- [30] M. Bouchard, A. L. Jousselme, and P. E. Doré, “A proof for the positive definiteness of the Jaccard index matrix,” *International Journal of Approximate Reasoning*, 2013, ISSN: 0888613X. DOI: 10.1016/j.ijar.2013.01.006.
- [31] P. Smets, Y. T. Hsia, A. Saffiotti, R. Kennes, H. Xu, and E. Umkehrer, “The transferable belief model,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 1991, ISBN: 9783540546597. DOI: 10.1007/3-540-54659-6_72.
- [32] B. R. Cobb and P. P. Shenoy, “On the plausibility transformation method for translating belief function models to probability models,” *International Journal of Approximate Reasoning*, vol. 41, no. 3, pp. 314–330, 2006, ISSN: 0888-613X. DOI: <https://doi.org/10.1016/j.ijar.2005.06.008>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0888613X05000472>.
- [33] J. Dezert, D. Han, J.-M. Tacnet, S. Carladous, and Y. Yang, “Decision-making with belief interval distance,” in *Belief Functions: Theory and Applications*, Springer International Publishing, 2016, pp. 66–74, ISBN: 978-3-319-45559-4.
- [34] *STANAG 2511 (January 2003, Intelligence reports, NATO Unclassified)*.
- [35] E. Blasch, K. B. Laskey, A.-L. Jousselme, V. Dragos, P. C. G. Costa, and J. Dezert, “URREF Reliability versus Credibility in Information Fusion (STANAG 2511),” in *16th International Conference on Information Fusion*, Istanbul, 2013.
- [36] G. J. Klir and B. Yuan, *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc, 1995, ISBN: 9780131011717.

- [37] G. J. Klir, “Generalized information theory,” *Fuzzy Sets and Systems*, 1991, ISSN: 01650114. DOI: 10.1016/0165-0114(91)90049-V.
- [38] P. Smets, “Imperfect information: Imprecision and uncertainty,” in *Uncertainty management in information systems*, Boston: Springer, 1997, pp. 225–254. DOI: 10.1007/978-1-4615-6245-0_8.
- [39] P. Bosc and H. Prade, “An Introduction to the Fuzzy Set and Possibility Theory-Based Treatment of Flexible Queries and Uncertain or Imprecise Databases,” in *Uncertainty Management in Information Systems*, 1997. DOI: 10.1007/978-1-4615-6245-0_10.
- [40] D. Dubois and H. Prade, *Possibility Theory*. Boston, MA: Springer US, 1988, ISBN: 978-1-4684-5289-1. DOI: 10.1007/978-1-4684-5287-7. [Online]. Available: <http://link.springer.com/10.1007/978-1-4684-5287-7>.
- [41] R. V. Hartley, “Transmission of Information,” *Bell System Technical Journal*, 1928, ISSN: 15387305. DOI: 10.1002/j.1538-7305.1928.tb01236.x.
- [42] G. J. Klir and M. Mariano, “On the uniqueness of possibilistic measure of uncertainty and information,” *Fuzzy Sets and Systems*, 1987, ISSN: 01650114. DOI: 10.1016/0165-0114(87)90090-X.
- [43] L. A. Zadeh, “Fuzzy sets as a basis for theory of possibility,” *Fuzzy Sets and Systems*, vol. 100, no. 1, pp. 9–34, 1999. DOI: 10.1016/S0165-0114(99)80004-9.
- [44] D. Dubois and H. Prade, *Fuzzy set and possibility theory-based methods in artificial intelligence*, 2003. DOI: 10.1016/S0004-3702(03)00118-8.
- [45] M. Wierman, “Measuring conflict in evidence theory,” 2002. DOI: 10.1109/nafips.2001.943815.
- [46] J. F. Sowa, “Generating language from conceptual graphs,” *Computers and Mathematics with Applications*, 1983, ISSN: 08981221. DOI: 10.1016/0898-1221(83)90005-6.
- [47] F. van Harmelen, F. van Harmelen, V. Lifschitz, and B. Porter, “Handbook of knowledge representation,” in San Diego, USA: Elsevier Science, 2007, pp. 823–824, ISBN: 0444522115, 9780444522115.
- [48] G. Pilato, A. Augello, M. Missikoff, and F. Taglino, “Integration of ontologies and Bayesian networks for maritime situation awareness,” in *IEEE 6th International Conference on Semantic Computing, ICSC 2012*, IEEE, 2012, pp. 170–177, ISBN: 9780769548593. DOI: 10.1109/ICSC.2012.25.
- [49] K. B. Laskey, P. C. Costa, and T. Janssen, “Probabilistic ontologies for knowledge fusion,” in *Information Fusion, 2008 11th International Conference on*, IEEE, 2008, pp. 1–8.
- [50] M. Richardson and P. Domingos, “Markov logic networks,” *Machine Learning*, vol. 62, no. 1-2 SPEC. ISS. Pp. 107–136, 2006, ISSN: 08856125. DOI: 10.1007/s10994-006-5833-1.
- [51] P. P. Shenoy, “Using Dempster-Shafer’s belief function theory in expert systems,” in *Advances in the Dempster-Shafer Theory of Evidence*, M. Fedrizzi, J. Kacprzyk, and R. R. Yager, Eds., New York, NY: John Wiley & Sons, 1994, pp. 395–414.

- [52] A. N. Steinberg and C. L. Bowman, "Adaptive context exploitation," B. D. Broome, D. L. Hall, and J. Llinas, Eds., May 2013, p. 875 804. DOI: 10.1117/12.2015623. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2015623>.
- [53] A.-L. Jousselme, "Context-based reasoning for Maritime Situation Awareness," NATO STO Centre for Maritime Research and Experimentation, La Spezia, Tech. Rep., 2015.
- [54] M. A. Razzaque, S. Dobson, and P. Nixon, "Classification and modeling of the quality of contextual information," *Journal of Autonomic and Trusted Computing*, 2006.
- [55] L. Snidaro, I. Visentini, K. Bryan, and G. L. Foresti, "Markov Logic Networks for context integration and situation assessment in maritime domain," in *15th International Conference on Information Fusion, FUSION 2012*, 2012, ISBN: 9780982443859.
- [56] D. Poole and N. L. Zhang, *Exploiting contextual independence in probabilistic inference*, 2003. DOI: 10.1613/jair.1122.
- [57] O. E. Gundersen, "Situational Awareness in Context," in *Modeling and Using Context*, Springer Berlin Heidelberg, 2013, pp. 274–287.
- [58] L. Snidaro, J. Garcia, J. Llinas, and E. Blasch, "Recent Trends in Context Exploitation for Information Fusion and AI," *AI Magazine*, vol. 40, no. 3, pp. 14–27, 2019.
- [59] M. Weiser, "The computer for the 21st Century," *IEEE Pervasive Computing*, vol. 1, no. 1, pp. 19–25, Oct. 2002, ISSN: 1536-1268. DOI: 10.1109/mprv.2002.993141.
- [60] M. Baldauf, S. Dustdar, and F. Rosenberg, "A survey on context-aware systems," *Int. J. Ad Hoc Ubiquitous Comput.*, vol. 2, no. 4, pp. 263–277, Jun. 2007, ISSN: 1743-8225. DOI: 10.1504/IJAHUC.2007.014070. [Online]. Available: <https://doi.org/10.1504/IJAHUC.2007.014070>.
- [61] A. K. Dey, "Understanding and Using Context," *Personal and ubiquitous computing*, vol. 5, no. 1, pp. 4–7, 2011. DOI: 10.1007/s007790170019.
- [62] P. Brézillon, "Context in Problem Solving: A Survey," *The Knowledge Engineering Review*, vol. 14, no. 1, pp. 47–80, 1999. DOI: 10.1017/S0269888999141018.
- [63] T. Winograd, "Architectures for Context," *Human-Computer Interaction*, vol. 16, no. 2-4, pp. 401–419, 2001. DOI: 10.1207/S15327051HCI16234_18.
- [64] McCarthy, John, "Notes on Formalising Context," in *IJCAI'93 Proceedings of the 13th international joint conference on Artificial intelligence*, 1993. DOI: 10.5555/1624025.1624103.
- [65] Leiguang Gong, "Contextual Modeling and Applications," in *IEEE International Conference on Systems, Man and Cybernetics*, 2006. DOI: 10.1109/icsmc.2005.1571176.
- [66] A. N. Steinberg and G. L. Rogova, "System-Level Use of Contextual Information," *Context-Enhanced Information*, 2016. DOI: 10.1007/978-3-319-28971-7_26.

- [67] A. N. Steinberg, "Situations and Contexts," *Perspectives on Information Fusion*, vol. 1, no. 1, pp. 16–24, 2016.
- [68] J.-C. Pomerol and P. Brézillon, "About Some Relationships between Knowledge and Context," in, Nov. 2007, pp. 461–464. DOI: 10.1007/3-540-44607-9_44.
- [69] A. N. Steinberg and L. Snidaro, "Levels?" In *18th International Conference on Information Fusion (Fusion)*, Washington, 2015, pp. 1985–1992.
- [70] A. N. Steinberg, C. L. Bowman, and F. E. White, "Revisions to the JDL Data Fusion Model," *Sensor Fusion: Architectures, Algorithms, and Applications*, vol. 3719, no. III, 1999.
- [71] L. Snidaro, J. García, and J. Llinas, "Context-based Information Fusion: A survey and discussion," *Information Fusion*, 2015, ISSN: 15662535. DOI: 10.1016/j.inffus.2015.01.002.
- [72] F. Caron, E. Duflos, D. Pomorski, and P. Vanheeghe, "GPS/IMU data fusion using multisensor Kalman filtering: Introduction of contextual aspects," *Information Fusion*, 2006, ISSN: 15662535. DOI: 10.1016/j.inffus.2004.07.002.
- [73] M. Vespe, M. Sciotti, F. Burro, G. Battistello, and S. Sorge, "Maritime multi-sensor data association based on geographic and navigational knowledge," in *2008 IEEE Radar Conference, RADAR 2008*, 2008, ISBN: 9781424415397. DOI: 10.1109/RADAR.2008.4720782.
- [74] G. Battistello, M. Ulmke, F. Papi, M. Podt, and Y. Boers, "Assessment of vessel route information use in Bayesian non-linear filtering," in *2012 15th International Conference on Information Fusion*, Singapore, 2012, pp. 447–454.
- [75] L. Snidaro, I. Visentini, and K. Bryan, "Fusing uncertain knowledge and evidence for maritime situational awareness via Markov Logic Networks," *Information Fusion*, vol. 21, no. 1, pp. 159–172, 2015, ISSN: 15662535. DOI: 10.1016/j.inffus.2013.03.004.
- [76] A. Jouan, Y. Allard, and Y. Marcoz, "Coastal Activity Monitoring with Evidential Fusion of Contextual Attributes from Multi-Pass RADARSAT-1 Data," in *7th International Command and Control Research and Technology Symposium*, Quebec City, 2002.
- [77] J. Gómez-Romero, M. A. Serrano, J. García, J. M. Molina, and G. Rogova, "Context-based multi-level information fusion for harbor surveillance," *Information Fusion*, 2015, ISSN: 15662535. DOI: 10.1016/j.inffus.2014.01.011.
- [78] G. Rogova, M. Hadzagic, M. O. St-Hilaire, M. C. Florea, and P. Valin, "Context-based information quality for sequential decision making," in *2013 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support, CogSIMA 2013*, 2013, ISBN: 9781467324373. DOI: 10.1109/CogSIMA.2013.6523818.
- [79] N. Ben Abdallah and A.-L. Joussemme, "On reasoning with partially reliable and conflicting sources, CMRE–FR-2019-004," NATO STO CMRE, La Spezia, Tech. Rep., 2019.
- [80] L. Bovens and S. Hartmann, *Bayesian Epistemology*. 2005, ISBN: 0199269750. DOI: 10.1093/0199269750.001.0001.

- [81] R. Haenni and S. Hartmann, “Modeling Partially Reliable Information Sources: a General Approach based on Dempster-Shafer Theory,” *Information Fusion*, vol. 7, no. 4, pp. 361–379, 2006. DOI: 10.1016/j.inffus.2005.06.005.
- [82] F. Pichon, D. Dubois, and T. Denoeux, “Relevance and truthfulness in information correction and fusion,” *International Journal of Approximate Reasoning*, vol. 53, no. 2, pp. 159–175, 2012. DOI: 10.1016/j.ijar.2011.02.006. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-00656775>.
- [83] F. Pichon, D. Mercier, F. Delmotte, and É. Lefèvre, “Truthfulness in contextual information correction,” in *3rd International Conference on Belief Functions, BELIEF 2014*, F. Cuzzolin, Ed., Oxford, UK, Sep. 2014, pp. 11–20.
- [84] D. Mercier, B. Quost, and T. Denoeux, “Contextual discounting of belief functions,” in *ECSQARU 2005: Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, ser. Lecture Notes in Computer Science, vol. 3571, Springer, 2005, pp. 552–562, ISBN: 978-3-540-31888-0.
- [85] D. Dubois, F. Faux, and H. Prade, “Prejudiced information fusion using belief functions,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11069 LNAI, Springer Verlag, 2018, pp. 77–85, ISBN: 9783319993829. DOI: 10.1007/978-3-319-99383-6_11.
- [86] N. Ben Abdallah, A. L. Jousselme, and F. Pichon, “An ordered family of consistency measures of belief functions,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11069 LNAI, Springer Verlag, 2018, pp. 199–207, ISBN: 9783319993829. DOI: 10.1007/978-3-319-99383-6_25.
- [87] F. Pichon, T. Denœux, and D. Dubois, “Relevance and truthfulness in information correction and fusion,” *International Journal of Approximate Reasoning*, vol. 53, pp. 159–175, 2012.
- [88] R. Haenni, “Uncover Dempster’s Rule Where It Is Hidden,” in *2006 9th International Conference on Information Fusion*, 2006, pp. 1–8. DOI: 10.1109/ICIF.2006.301675.
- [89] A. Hunter and S. Konieczny, “Measuring Inconsistency through Minimal Inconsistent Sets,” in *Eleventh International Conference on Principles of Knowledge Representation and Reasoning*, 2008. DOI: 10.5555/3031661.3031704.
- [90] S. Destercke and T. Burger, “Toward an axiomatic definition of conflict between belief functions,” *IEEE transactions on Systems, Man, and Cybernetics*, 2012. DOI: 10.1109/TSMCB.2012.2212703.
- [91] T. M. Strat, “Decision analysis using belief functions,” *International Journal of Approximate Reasoning*, 1990, ISSN: 0888613X. DOI: 10.1016/0888-613X(90)90014-S.
- [92] T. M. Strat and J. D. Lowrance, “Explaining evidential analyses,” *Int. Journal of Approximate Reasoning*, vol. 3, pp. 299–353, 1989. DOI: 10.1016/0888-613X(89)90020-0.

- [93] H. Xu and P. Smets, "Some strategies for explanations in evidential reasoning," *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*, vol. 26, no. 5, pp. 599–607, 1996, ISSN: 10834427. DOI: 10.1109/3468.531907.
- [94] P. Kowalski and T. Martin, "Assessment of the contribution of a source of information to the final belief in a proposition," in *ICNC-FSKD 2017 - 13th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery*, 2017, ISBN: 9781538621653. DOI: 10.1109/FSKD.2017.8392970.
- [95] —, "Provenance Across Evidence Combination in Theory of Belief Functions," in *2018 21st International Conference on Information Fusion, FUSION 2018*, 2018, ISBN: 9780996452762. DOI: 10.23919/ICIF.2018.8455764.
- [96] L. S. Shapley, "17. A Value for n-Person Games," in *Contributions to the Theory of Games (AM-28), Volume II*, 2016. DOI: 10.1515/9781400881970-018.
- [97] R. R. Yager, *Entropy and specificity in a mathematical theory of evidence*, 2008. DOI: 10.1007/978-3-540-44792-4_11.
- [98] M. Daniel, "Conflicts within and between belief functions," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6178 LNAI, pp. 696–705, 2010, ISSN: 03029743. DOI: 10.1007/978-3-642-14049-5_71.
- [99] W. Liu, "Analyzing the degree of conflict among belief functions," *Artificial Intelligence*, vol. 170, no. 11, pp. 909–924, 2006, ISSN: 00043702. DOI: 10.1016/j.artint.2006.05.002.
- [100] F. Pichon, A. L. Jousselme, and N. Ben Abdallah, "Several shades of conflict," *Fuzzy Sets and Systems*, vol. 366, pp. 63–84, 2019, ISSN: 01650114. DOI: 10.1016/j.fss.2019.01.014. [Online]. Available: <https://doi.org/10.1016/j.fss.2019.01.014>.
- [101] F. Cuzzolin, "On consistent approximations of belief functions in the mass space," in *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Springer, 2011, pp. 287–298.
- [102] D. Dubois and H. Prade, "A note on measures of specificity for fuzzy sets," *International Journal of General Systems*, vol. 10, no. 4, pp. 279–283, 1985, ISSN: 15635104. DOI: 10.1080/03081078508934893.
- [103] A. Martin, "Conflict management in information fusion with belief functions," in *Information Quality in Information Fusion and Decision Making*, Springer, 2019, pp. 79–97.
- [104] B. Liu, Q. Yang, X. Wu, Y. Guo, and S. Fang, "An Efficient Method to Measure Evidence Conflict," Atlantis Press, 2017. DOI: 10.2991/caai-17.2017.109.
- [105] M. Higashi and G. J. Klir, "Measures of uncertainty and information based on possibility distributions," in *Readings in Fuzzy Sets for Intelligent Systems*, 1993. DOI: 10.1016/b978-1-4832-1450-4.50022-5.
- [106] M. R. Endsley, "Towards a theory of situational awareness in dynamic systems," *Human Factors*, vol. 37, no. 1, 1995. DOI: 10.1518/001872095779049543.

- [107] S. Fossier, C. Laudy, and F. Pichon, “Managing Uncertainty in Conceptual Graph-Based Soft Information Fusion,” in *16th International Conference on Information Fusion*, International Society of Information Fusion Institute of Electrical and Electronics Engineers, 2013, ISBN: 9786058631113.
- [108] C. Laudy, “Semantic Knowledge Representations for Soft Data Fusion,” in *Efficient Decision Support Systems - Practice and Challenges From Current to Future*, 2011. DOI: 10.5772/17762.
- [109] P. Kowalski and T. Martin, “Embedding Uncertainty in Conceptual Graphs for Semantic Information Fusion,” in *Proceedings - 2019 IEEE Conference on Cognitive and Computational Aspects of Situation Management, CogSIMA 2019*, Institute of Electrical and Electronics Engineers Inc., Apr. 2019, pp. 1–8, ISBN: 9781538695999. DOI: 10.1109/COGSIMA.2019.8724358.
- [110] G. Shafer, “A Mathematical Theory of Evidence turns 40,” *International Journal of Approximate Reasoning*, vol. 79, 2016, ISSN: 0888613X. DOI: 10.1016/j.ijar.2016.07.009.
- [111] M. Chein and M.-L. Mugnier, *Graph-based Knowledge Representation*. Springer London, 2008. DOI: 10.1007/978-1-84800-286-9.
- [112] T. Pápai, S. Ghosh, and H. Kautz, “Combining subjective probabilities and data in training Markov logic networks,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 7523 LNAI, no. PART 1, pp. 90–105, 2012, ISSN: 03029743. DOI: 10.1007/978-3-642-33460-3_11.
- [113] P. Kowalski, M. Zocholl, and A.-L. Joussetme, “Explainability in threat assessment with evidential networks and sensitivity spaces,” in *Proceedings of 2020 23rd International Conference on Information Fusion, FUSION 2020*, 2020, ISBN: 9780578647098. DOI: 10.23919/FUSION45008.2020.9190475.

