Uncertainty Representation and Reasoning Within the Threat Intelligence Domain

Ritten Roothaert[1]

[1]Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV, Amsterdam, The Netherlands

Abstract

National intelligence agencies face the challenge of data overload as they analyse vast amounts of information to address national security threats. To help analysts manage this complexity and ensure decision-making compliance, the Threat Intelligence Decision Ontology (TIDO) was developed. However, a major limitation of TIDO is the absence of a formal representation for the inherent uncertainties in threat intelligence (TI) data, which stems from imperfect sources, biases, and incomplete information. This ongoing PhD research addresses this critical gap by investigating how to integrate a suitable uncertainty representation into the TIDO framework.

We present the preliminary results of a focused analysis of three quantitative uncertainty representations: general probability theory, interval probabilities, and belief function theory. To evaluate their suitability for the TI domain, we address four key research questions: (1) What are the characteristics of data used in TI? (2) What are the benefits and limitations of the selected uncertainty representations? (3) How can the data characteristics of TI be captured by these representations? And (4) what tooling is available for each representation?

Our approach combines a literature review with insights from a focus group of domain experts to answer these questions. By comparing the needs of the TI domain with the capabilities of each uncertainty representation, we aim to determine the most suitable method for integrating uncertainty into the TIDO ontology. The ultimate goal is to provide a more robust and accurate decision-making framework that accounts for the limitations of real-world threat intelligence.

Keywords

uncertainty representation and reasoning, threat intelligence, interval probabilities, belief function theory

1. Introduction

Motivation. One of the main tasks of national intelligence agencies is to investigate and act against organizations or persons that could pose a threat to the national security. During such investigations, analysts must interpret and analyse incoming data and act accordingly. However, similar to other data-driven domains, technological advances have created new sources of information that can be used during investigations, increasing the amount of information available for decision making, and in turn increasing the risk of data-overload [1]. This data-overload not only complicates the ability of TI analysts to asses the situation and make their decisions, but also the analysis of those decisions and whether or not those decisions are compliant with the applicable legislation and internal policies. To mitigate this challenge of data-overload, the Threat Intelligence Decision Ontology (TIDO) was developed[1]. It is designed to support the analyst *during* the situation assessment and action selection processes of the TI analysis while simultaneously capturing the decision trace that can be used for further post-analysis and compliance checking.

A major limitation of the TIDO ontology is in its current form, is that there is no specific representation for the uncertainties involved in the decision making process. AI solutions and human experts will always have biases or imperfect models, sources might be poor, or misleading, and data-sharing might be limited for operational, strategic or legal reasons. Consequently, at each step of Threat

RuleML+RR'25: Companion Proceedings of the 9th International Joint Conference on Rules and Reasoning, September 22–24, 2025, Istanbul, Turkiye

h.m.roothaert@vu.nl (R. Roothaert)
https://ritten11.github.io/ (R. Roothaert)
0009-0008-7843-6513 (R. Roothaert)

[1]Available at https://w3id.org/tido. The accompanying paper is currently under review at K-CAP 2025

Intelligence Hybrid Workflow (TIHW), the analyst must revise this limited and uncertain information and recommend actions [2]. To do so, a suitable representation of the uncertainty associated with the information and considerations used within the TI decision process is needed.

Uncertainty representation is a broad and complicated field, ranging from quantitative representations, often expressed using probabilities, to qualitative representations, often expressed in natural language, and ordinal representations, which are used to express that one statement is more likely to be true that another. For expressing the uncertainties in the TIDO ontology, we focus on the quantitative uncertainty representations as we aim to propagate uncertainties throughout the TIHW using a formalized and sound theoretical background, capable of expressing *how much* more likely one statement is over another. Here, a pure probabilistic-based perspective would work well when high quality and complete information is available. However, this assumption generally doesn't hold in the domain of TI [3, 4]. This concern has sparked several investigations into alternative methods to represent uncertainty, which can be categorized into various groups. According to [3], there are five main groups: (1) general probability theory, (2) interval probabilities, (3) probability bound analysis, (4) belief function theory, and (5) possibility theory. Although this is not an exhaustive list, it provides a manageable number of widely used approaches with which to start our study. Each of these groups comes with different prerequisites, assumptions, and uncertainty interpretations, and thus the decision to opt for one representation over the other requires a careful consideration. To keep the number of comparisons manageable, we decided to focus our attention on general probability theory, interval probabilities, and belief function theory. Probability bound analysis was excluded as it can be considered a combination of general probability theory and interval probabilities, and possibility theory was excluded as it is considered to be subsumed by belief function theory [3].

Ethical considerations The author is aware of the sensitivity of the subject and the potential implications with respect to the privacy of citizens. The motivation, though, is not to collect new data but to provide a vocabulary to enrich and structure data already obtained by the intelligence services. This enables the decision makers to make their assumptions and considerations explicit, increasing their accountability and improving the testability of their decisions with respect to compliance with legislation and internal policies. The proposed research does not provide intelligence agencies with additional capabilities, nor justifications, to collect more data, and therefore, does not impact the privacy of citizens.

Approach and goals To determine which uncertainty representation is most suitable for the TI domain, we split the problem into four research questions: (**RQ1**) What are the characteristics of the data used for decision making within the domain of TI? (**RQ2**) What are the benefits and limitations of the selected uncertainty representations? (**RQ3**) How can the data characteristics identified in **RQ1** be captured by the representations analysed in **RQ2**? And finally, (**RQ4**) what tooling is available for the analysed uncertainty representations? **RQ1** is assessed through a literature analysis, in combination with findings derived from a focus group with domain experts. **RQ2** will be assessed solely though a literature analysis, and **RQ3** will be addressed by means of a comparison between the answers of **RQ1** and **RQ2**. The final question, **RQ4**, is briefly touched upon in Section 3.4, but we hope to obtain additional suggestions during the discussions at the RuleML+RR doctoral consortium.

Organisation Section 2 introduces the TIDO ontology, and explains the core functionalities of the ontology using an example of a common decision made within the TI domain. Section 3 provides an analysis of the characteristics of the data used within the TI domain, along with an in indication on how these characteristics manifest themself within the TIDO ontology, or how a future iteration of TIDO could capture these characteristics. Section 4 provides an overview of the current progress and individual contributions of the student attending this doctoral consortium, as well as brief description of the potential impact of the overall project.

2. The Threat Intelligence Decision Ontology

The Threat Intelligence Decision Ontology (TIDO) was developed to provide a vocabulary that can be used to describe the decision processes within the TI domain, together with the information that was used to come to those decisions. This ontology is an ongoing work, and at the moment of submission of this doctoral consortium paper, the paper accompanying the TIDO ontology is under review at the Thirteenth International Conference on Knowledge Capture (K-CAP 2025). Therefore, this introduction for the TIDO ontology will not focus on the development of the ontology, but instead focus on the main ideas and provide an intuitive explanation on how the ontology can be used to describe decision process in the TI domain. In section 4, this example will also be used to describe our current ideas on how to extend the TIDO ontology to allow for the expression of uncertainty.

Figure 1 depicts the TIDO ontology. This ontology is an extension of the PROV ontology [6], and directly inherits the relationship between entities, in the TI context interpreted as *pieces of information*, activities and agents. The TIDO ontology adds four additional modules to the PROV ontology:

TIDO-Process: The process module provides a vocabulary to describe different steps in the procedural element of decision making. These steps are derived from Bales' phases in group decision making [7], and while decision process generally follow the patter of an `:Orientation`-step, followed by a `:Evaluation`-step and finished with a `:Resolution`-step, real decision processes may deviate from this pattern and therefore the TIDO ontology makes no assumptions on the ordering of these steps.

TIDO-Sense: The sense-making module provides a vocabulary to describe the relationship between in-

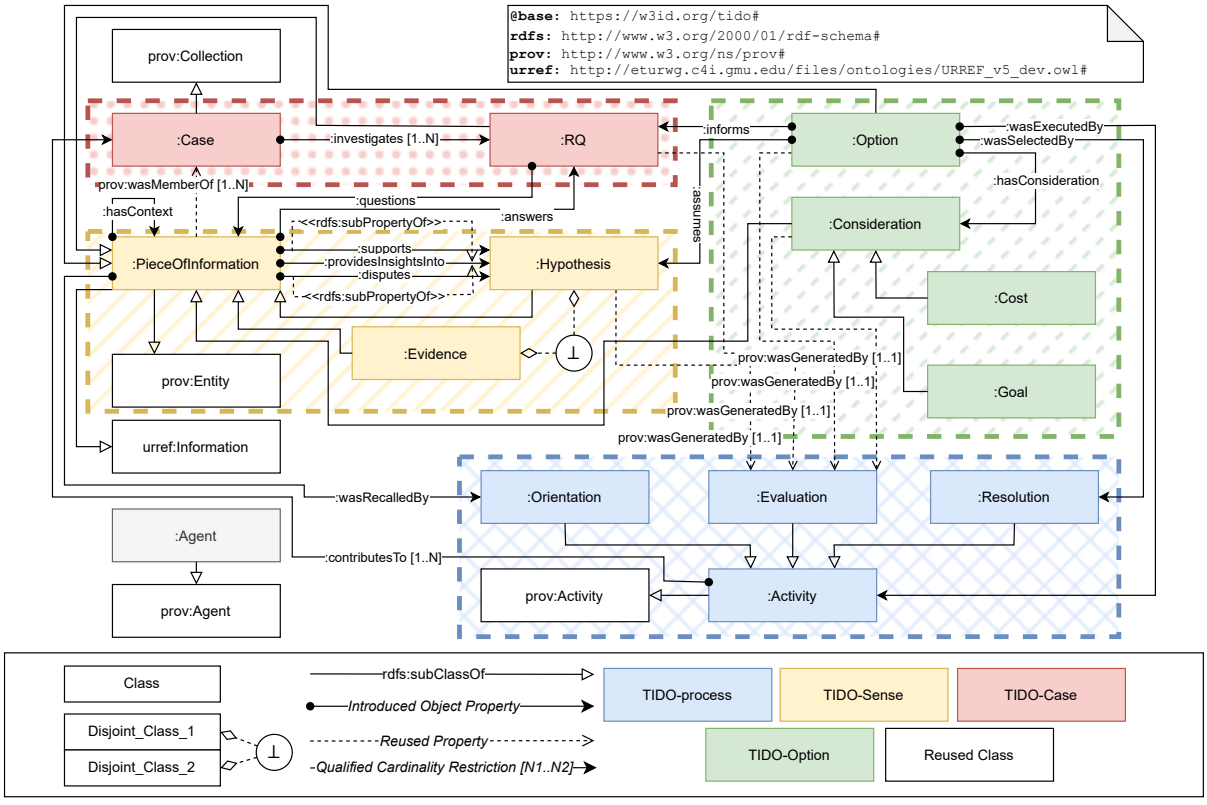

Figure 1: An overview of the TIDO ontology, depicted using the CHOWLK [5] ontology conceptualisation notation. For explanatory purposes, TIDO is divided into four conceptual modules: TIDO-Process describes the procedural element of decision-making, TIDO-Sense describes the sense-making element of decision-making, TIDO-Case describes the case and what it is investigating, and TIDO-Option describes the available options, along with their considerations.

formation, evidence and hypotheses. Here, the class `:PieceOfInformation` adopts the very broad interpretation of information in the URREF ontology [8], meaning it can range from anything to a sensory measurement, a fact, common sense, or an uncertainty statement. However, for practical purposes, the remainder of this paper will interpret the `:PieceOfInformation` class as anything that can be expressed in a natural language sentence. Both the `:Evidence` and `:Hypothesis` classes are sub-classes of `:PieceOfInformation`. Here, the `:Evidence` class describes testimonies, observations and sensory measurements, whereas `:Hypothesis` provides an interpretation of these testimonies, observations and sensory measurements. Additionally, other pieces of information can be indicated to provide insights into instantiations of the `:Hypothesis` class, with either a positive `:supports`, a negative `:disputes`, or a neutral `:providesInsightsInto` relation. This is not possible for instantiations of the `:Evidence` class.

TIDO-Option: The option module provides the vocabular to describe which options in the decision process were considered, which considerations are associated with each option, and which option was selected during a `:Resolution` step. Note that both the `:Option` and `:Consideration` classes are subclasses from the `:PieceOfInformation` class, meaning that instances of either class can also be used to provide insights into instances of the `:Hypothesis` class.

TIDO-Case: The case module is used do describe what is being investigated. Here, the `:Case` class, modelled as a sub-class of `prov:Collection`, is a collection which contains all the pieces of information that bear some degree of relevance to the investigation. The `:RQ` class contains all the *research questions* that are investigated in the decision process.

Figure 2 provides an example of how the TIDO ontology can be used to describe a decision process in the TI domain. This example concerns the start of an investigation by the Dutch Civil Intelligence and Security Service (CISS) into a potential threat and is derived from the first episode of a podcast that was co-produced by the CISS [9]. In the first step ($step_1$) depicted in this example, agent a_2 recalls a set of evidence E that has been produced by agent a_1 at an earlier point in time. Afterwards in $step_2$, after

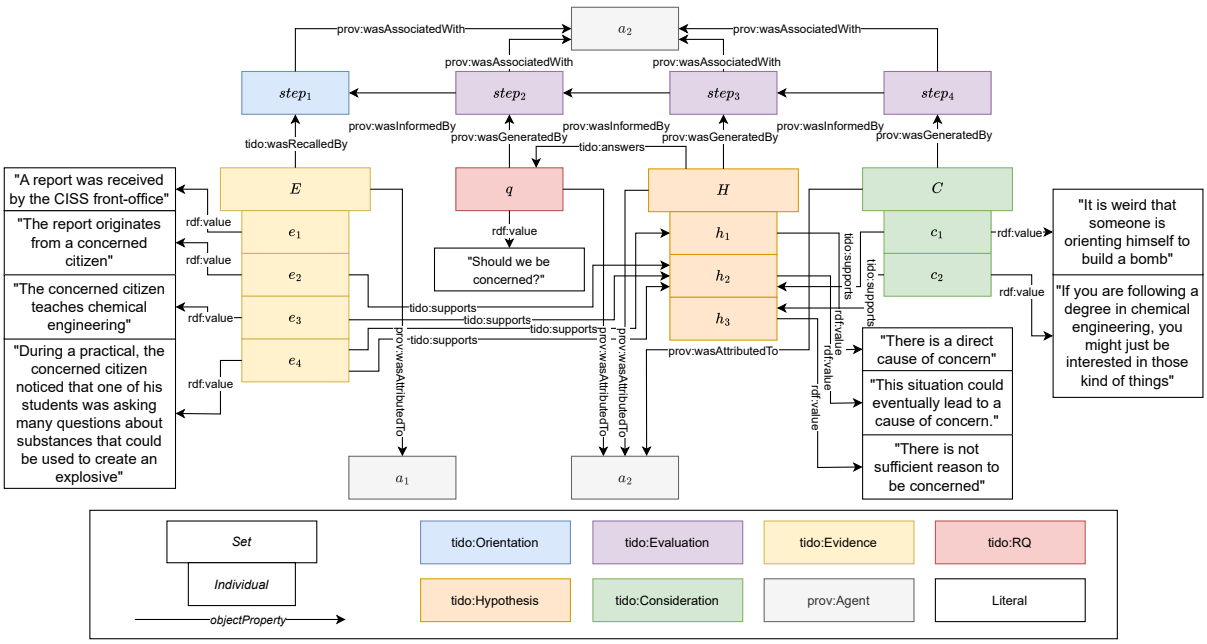

Figure 2: An example of part of a decision process in the TI domain that is annotated according to the TIDO vocabulary. Here, the boxes represent instances of the classes that are indicated by the colour of the box. Also note that, for the sake of readability, the concept of a *Set* is introduced. In order to make the notation compatible with the TIDO vocabulary, simply apply all incoming and outgoing object properties connected to a set to all individuals contained within this set.

having read the presented evidence, agent a_2 drafts the question q that he or she will try to answer. In $step_3$, a set of potential hypotheses H is drafted that could answer q, and either supporting or disputing connections are made between the pieces of evidence in set E and the hypotheses in set H. In the final step ($step_4$), agent a_2 uses his or her common sense and expert knowledge to draft a set of considerations (C) that also either support or dispute the hypotheses in H.

3. Uncertainty Representation (for Threat Intelligence)

The TIDO ontology provides a vocabulary to represent the pieces of information, activities and agents that are part of a decision process in the TI domain. However, as mentioned in Section 1, the TIDO ontology does not have a dedicated vocabulary to expressing uncertainties associated with the decisions made within a TIHW. In order to determine which uncertainty representation is most appropriate, Section 3.1 first address **RQ1** and characterize the data used within the TI domain. This analysis was primarily derived from Miller [4], van Gerwen [1], and the Joint Dutch Doctrine for Intelligence from the Dutch Ministry of Defence [10]. Afterwards, Section 3.2 addresses **RQ2** and introduces three potential uncertainty representations, where the notation and the main benefits and limitations are primarily adopted from Zio and Pedroni [3]. Section 3.3 provides preliminary ideas on how the uncertainty representations from 3.2 and be used to represent the data described in 3.1, providing a beginning point for answering **RQ3**. As this is ongoing work, a prioritisation is made on which components will be addressed first. Finally, Section 3.4 provides a brief overview of RDF compatible tools that would enable the reasoning capabilities of the to-be-selected uncertainty representation.

3.1. Data Characterisation

Miller states that 'intelligence' in the TI domain is an epistemic product of some process of analysis and evaluation of information that is done with respect to different criteria, including the likelihood that it is true, its importance, and relatedly, the reliability of the source. These evaluation criteria are less concrete and inherently contain some degree of subjectivity, complicating their formalisation. Starting with the 'likelihood of information being true', a distinction should be made according to the multiple ways a 'likelihood' can be interpreted. This is important as the meaning of likelihood in our specific context determines how it should be represented in a model. Using an inappropriate categorization may result in a under- or overestimation of the risk or threat [11]. For the remainder of this manuscript, we will use the term *certainty* to describe the belief in favour of a proposition and the term *uncertainty* the describe the lack of a belief for a proposition.

First, a distinction should be made in the *nature* of the uncertainty, where *aleatoric* uncertainty is the uncertainty that originates from a random process (e.g. a coin toss), and *epistemic* uncertainty is the uncertainty that originates from a lack of knowledge (e.g. limited sensor observability). In the TI domain, both types occur, with human behaviour begin a good example for a source of aleatoric uncertainty, and missing information concerning a target's current location being a good example for a source of epistemic uncertainty. Here, the uncertainty regarding the location of the target can be resolved by gathering more intelligence regarding the location of the target, whereas the behaviour of the target will always contain some element of randomness, regardless of the quality of the analyst's prediction.

Secondly, there are different ways in how uncertainty is *derived*. The *objective* approach uses models to calculate the uncertainty, which are generally grounded by data from past events. Generally, these models focus on representing the aleatoric component of uncertainty, but they are also capable of integrating epistemic uncertainty, such as some imperfections of sensors, as long as this uncertainty can be quantified. The *subjective* approach is, as the name already suggests, a less well-defined way of quantifying the uncertainty and is usually performed by domain experts, leaving room for multiple interpretations. Generally, this type of uncertainty quantification is associated with epistemic uncertainty as such quantifications are generally made when no sufficiently accurate model exists, or when data from past events is missing. However, the subjective approach can also be used to express forms of

aleatoric uncertainty, such as the previously mentioned example of predicting human behaviour. Again, in the TI domain, both types of uncertainty quantification are used. For instance, the uncertainty of an image classification algorithm may be expressed using a objective approach, whereas the motives of a target are better suited to be represented using a subjective approach.

Third, there is the way the uncertainty is *interpreted*. The *probabilistic* interpretation focusses on representing uncertainty as the probability of a specific event occurring. Here, probability is defined as '*the fraction of times an event A occurs if the situation considered were repeated an infinite number of times*' [3]. An example here would be the classical coin flip, which will land on its head every 1 out of 2 times on average when flipped x times, or the expected number of employees that will click on a phishing link if distributed among y employees. The alternative, the *non-probabilistic* interpretation, represents uncertainty as 'that what is unknown'. This interpretation suggests that we can only believe things that are true based on the evidence that supports it. For everything else we don't have evidence for, we don't know whether it is true of false. An example of a non-probabilistic uncertainty is the assignment of an individual being considered a threat or not according to an assessor. Here, a 0.6 could indicate that the assessor has a 0.6 degree of confidence in the individual being a threat and 0.4 expresses the assessor's epistemic uncertainty about the claim. This interpretation aligns with the open-world assumption: unknown facts may exist, so a lack of evidence for a hypothesis does not imply evidence against it.

The fourth dimension is whether the uncertainty is *ergodic* or *non-ergodic*. Here, ergodicity refers to whether the true (aleatoric) uncertainty of the system can be measured by taking the average deviation of the mean value over time. In an ergodic system, the value of system parameters (and their variance) does not vary over time or the phase of the system. This implies that the uncertainty derived from the analysis of past observations also generalizes to the current state of the system. Within the domain of TI, most uncertainties are expected to be to some extent non-ergodic. For example, the behavioural patterns of individuals or organisations might change over time, for instance by the discovery of a new strategy or the rise of new technologies. The previously mentioned image classifier would be an example of an ergodic uncertainty, provided that the target variable remains constant[2].

Besides the 'likelihood of being true', Miller mentions that information should also be evaluated according to its *importance*. Here, it is helpful to make a distinction between *importance* and *relevance*. The difference between *intelligence* and *information* is that intelligence is considered to be 'information or data (expressible as a statement or, more likely, structured set of statements) that are acquired for various institutional purposes'[4]. Within the domain of TI, this purpose is to assess whether a particular individual or organisation should be considered a threat to the national security. The example given by Miller describing *relevance* is that if an arbitrary individual breaks a leg, this information is not relevant for (most) case investigations, but if this individual is the same as the target that is investigated, this information suddenly becomes very relevant. Importance however, we interpret to be more closely associated with the question that an analyst is trying to answer. Here, the information that the target has broken its leg is less important when creating an overview of the social network of the target, whereas it is very important when assessing the capabilities or the expected location of the target. *Importance* in our interpretation therefore carries a higher degree of context dependence. From a formal point of view, we can say that we use the term *relevance* to refer to the relevance of a piece of information with respect to the case, and the term *importance* to refer to the relevance of a piece of information with respect to a research question.

The final evaluation criteria mentioned by Miller is reliability of the source. Again, no operationalized description of the term 'reliability' is given, but we interpret it as *the degree or probability with which can be assumed that the information received from this source is correct.* Referring back to the example given in Figure 2, the set of evidence E is provided to agent a_2 by agent a_1. If a_2 has past experiences or background information that could indicate that evidence provided by a_1 could have been manipulated or misrepresented, agent a_2 could adjust a *reliability* score of the agent a_1. In turn, this reliability score

[2]If an image classifier is used to classify pictures containing cars, the interpretation of the concept 'car' should not change. If at some point in the future flying cars are invented, changing our interpretation of what a car should look like, the uncertainty associated with our car-classifier is no longer ergodic.

could be used as a discount or correction factor to the degree of uncertainty associated with the evidence set E.

The previous formalisations focussed on representing the characteristics of the data used within the TI domain. However, what is missing is which data is *useful for experts* in the TI domain. To address this element, we re-analysed the competency questions (CQs) of the TIDO ontology [12] that were produced by domain experts during a focus group conducted for the development of the TIDO ontology, which was conducted using standard ontology development practices[3]. Out of the 14 in-scope CQs that are not or only partially answerable by TIDO, 4 questions concern the identification of missing information. This highlights that the uncertainty representation should not only focus on depicting the information that is known, but also which information is *not* known. In other terms, it is important to quantify the *ignorance* associated with both the questions and the hypotheses considered in the investigation. For the questions it is important to know whether the set of considered hypothesis covers the full set of plausible answers, and and for the hypothesis it is important to know whether these is missing information that could help in assessing the likelihood of the hypothesis being true.

3.2. Description uncertainty representations

Broadly speaking, there are two main 'schools' in uncertainty representation, diverting in how uncertainty is interpreted. There are those that adopt the 'probability' perspective and aim to map every type of uncertainty to the well-defined 'frequentist' interpretation of uncertainty. This has the benefit of being compatible with the powerful Bayesian uncertainty aggregation/reasoning paradigm, but the downside that every expression of uncertainty *must* be quantified in some form of a frequency distribution which is often difficult and sometimes impossible. The other school adopts a 'non-probabilistic' perspective, which has a less rigid representation of uncertainty and can therefore be considered more expressive. However, the downside here is that a custom uncertainty reasoning/aggregation paradigm is needed, depending on the expressivity required for the application. Each of these 'schools' have different flavours and providing a complete overview is well beyond the scope of this paper. Instead, we will focus on the main characteristics of some variants that are particularly relevant for the task of risk analysis. Additionally, we will focus only on uncertainty representations where the possible values of a random value of interest can be described using a discrete set of elements. For a more detailed description of the discussed uncertainty representations, along with how they might be applied to continuous variables, we refer to [3].

3.2.1. Probability Theory

Probability theory adopts the frequentist perspective on uncertainty and assumes that all variables within the system can be described by a *probability distribution function* (PDF) $d_Y(y) : \Omega \rightarrow [0, 1]$, where Ω is the sample space of *all* the values that a discrete random variable Y can assume. Here, $\sum_{y \in \Omega} d_Y(y) = 1$, meaning that the sum of of the PDF over all possible values in Ω is equal to 1. The probability $P(A)$ of any measurable subset A of Ω, called an event, can then be described by

$$P(A) = \sum_{y \in A} d_Y(y) \tag{1}$$

.

Example. *—- Unfortunately, an example for a potential extension to TIDO using general probability theory is still ongoing work —-*

Benefits:

- It is relatively simple to implement and explain [3].
- It can use information about correlations between variables using conditional probability distribution functions and Bayes' rule [3].

[3]A specific description of these methodologies is outside the scope of this work, but is discussed in other work: w3id.org/tido

Limitations:

- The uncertainty assessments from the expert must be numerical, even when the expert is not able to provide an exact value. This results in the analysis being forced to make subjective and often unjustified assumptions and guesses [3].
- Propagating/aggregating the uncertainty requires complex knowledge concerning the correlations between variables. Determining these correlations can be a time and resource intensive endeavour [3].
- It confounds ignorance with variability, making it difficult to distinguish between aleatoric and epistemic uncertainty [3].

3.2.2. Interval Probabilities

Interval probabilities, a form of imprecise probabilities, address the previously mentioned limitation that probabilities need to be assigned the to possible values $y \in \Omega$ of discrete random variable Y as a single numerical value. Instead, the likelihood of a subset $A \subseteq \Omega$ being true is represented using a lower probability $\underline{P}(A)$ and an upper probability $\overline{P}(A)$, creating a probability interval $[\underline{P}(A), \overline{P}(A)]$ where $0 \leq \underline{P}(A) \leq \overline{P}(A) \leq 1$. The difference

$$\Delta P(A) = \overline{P}(A) - \underline{P}(A)$$

is referred to as the *imprecision* in the representation of A. The single valued probabilities are the same as the special case when $\underline{P}(A) = \overline{P}(A)$ such that $\Delta P(A) = 0$.

Example. *—- Unfortunately, an example for a potential extension to TIDO using interval probabilities is still ongoing work —-*

Benefits:

- Imprecise probabilities are still relatively simple to implement and explain [3].
- No matter what the distribution of the data, or the correlations within it, imprecise probabilities will put sure bounds on the upper and lower limit of the uncertainties of each variable [3].
- It is not necessary to make assumptions on the probability distributions of random variables [13].
- Imprecise probabilities, and by extension interval probabilities, are completely based on classical probability theory and can be regarded as its generalization [13].

Limitations:

- The imprecision, i.e. the difference between the upper and lower bound on the uncertainty, can grow very quickly, creating an increasingly more conservative estimate as arithmetic operations are applied to the imprecise probabilities [3].
- While the absence of assumptions on the distribution between the upper and lower limit of the uncertainty is considered a benefit, a consequence of this is that the basic form of imprecise probabilities is not able to utilize this distribution if the data is available. Even when there are clear intuitions on what the true uncertainty value is, this intuition cannot be captured by only representing the upper and lower limit of the uncertainty [3].

3.2.3. Belief Function Theory

Belief Function Theory (BFT), also known as evidence theory or Dempster-Shafer theory [14, 15], is intended for situations where there is more information available than what can be used for interval probabilities, but not enough information to a specific probability distribution function [3]. Instead, BFT allows for the incorporation and representation of incomplete information.

In BFT, uncertainty is represented using a *basic probability assignment* (BPA) in the form of a mass distribution $m(A)$ over all sets A in the power set $\mathscr{P}(S)$, where S is the set of possible values a variable s might take. This representation of uncertainty satisfies the following requirements:

$$m : \mathscr{P}(S) \to [0,1], \quad m(\emptyset) = 0, \quad \sum_{A \in \mathscr{P}(S)} m(A) = 1 \tag{2}$$

Using this representation, the belief, plausibility and commonality measures are defined by

$$Bel(A) = \sum_{B \subseteq A} m(B), \quad Pl(A) = \sum_{B \cap A \neq \emptyset} m(B), \quad Q(A) = \sum_{B \supseteq A} m(B), \quad \forall A \in \mathscr{P}(S) \tag{3}$$

Example. To give an explanation on how to interpret these measures, let us have another look at the example given in Figure 2. Let us assume that in this example, the *answer to q* is the variable that we wish to evaluate, and H is the set of possible values the answer to q might take. Then, we can define a mass function $m_q : \mathscr{P}(H) \to [0,1]$. Potential knowledge elicitation methods to obtain this mass function from an expert are discussed in Section 3.3, but for now, let us make an educated guess of what such a mass function could look like. Taking a closer look at the `tido:support` relations in Figure 2, we see that the pieces of evidence e_2 and e_3 and consideration c_1 support that h_2 is the correct answer to q. Considering that the biggest set of available information only supports h_2, we can assign a relatively large mass to the set $\{h_2\}$, so suppose this value is $m_q(\{h_2\}) = 0.4$. Piece of evidence e_4 supports both h_1 and h_2. Looking at the content of e_4, we could argue that this is the most important piece of information when evaluating the possible answers to q, hence we can also assign a relatively high mass to the set $\{h_1, h_2\}$, so suppose the mass is $m_q(\{h_1, h_2\}) = 0.3$. Consideration c_2 supports h_3, but arguably, this piece of information is less important than e_4. Therefore, we can assign a smaller mass to the set $\{h_3\}$, so suppose that $m_q(\{h_3\}) = 0.1$. Finally, we can also model that there is still some degree of ignorance over the complete set of hypothesis H, where the analyst cannot distinguish between any potential answer to q. Here, suppose that the mass to this ignorance set is $m_q(H) = 0.2$. As there is no evidence to support any other set from the power set $\mathscr{P}(H)$, the mass for all other sets is equal to 0.

To summarize, the mass function $m_q(\cdot)$ would take the following values:

$$m_q(A) = \begin{cases} 0.2, & \text{if } A = H \\ 0.3, & \text{if } A = \{h_1, h_2\} \\ 0.4, & \text{if } A = \{h_2\} \\ 0.1, & \text{if } A = \{h_3\} \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

Now, let us revisit the belief, plausibility and commonality measures defined in 3. In the current example, *belief* ($Bel(\cdot)$) is interpreted as the degree of certainty to which the analyst believes that, based on the available evidence in set E and his considerations in set C, the true answer to q is assigned to the set A or a subset of A. The *plausibility* measure ($Pl(\cdot)$) can be interpreted as the degree of certainty to which evidence E and considerations C assign the answer for q to any set B in $\mathscr{P}(S)$ that overlaps with A. Finally, the *commonality* measure ($Q(\cdot)$) interpreted as the degree of certainty to which evidence E and considerations C assign the answer for q to A or a superset of A.

Combining 3 with 4, we get the following values for the belief, plausibility and commonality functions:

$$\begin{aligned}
Bel(H) &= 1, & Pl(H) &= 1 & Q(H) &= 0.2 \\
Bel(\{h_1, h_2\}) &= 0.9, & Pl(\{h_1, h_2\}) &= 0.9 & Q(\{h_1, h_2\}) &= 0.5 \\
Bel(\{h_1, h_3\}) &= 0.1, & Pl(\{h_1, h_3\}) &= 0.6 & Q(\{h_1, h_3\}) &= 0.2 \\
Bel(\{h_2, h_3\}) &= 0.4, & Pl(\{h_2, h_3\}) &= 1 & Q(\{h_2, h_3\}) &= 0.2 \\
Bel(\{h_1\}) &= 0.0, & Pl(\{h_1\}) &= 0.5 & Q(\{h_1\}) &= 0.5 \\
Bel(\{h_2\}) &= 0.4, & Pl(\{h_2\}) &= 0.9 & Q(\{h_2\}) &= 0.9 \\
Bel(\{h_3\}) &= 0.1, & Pl(\{h_3\}) &= 0.3 & Q(\{h_3\}) &= 0.3
\end{aligned} \tag{5}$$

Looking at the values in 5, we see that out of the singleton values h_1, h_2 and h_3, the belief, plausibility and commonality measures for h_2 are the highest, indicating that h_2 would most likely be the best answer to q. However, there is still a gap between $Bel(\{h_2\})$ and $Pl(\{h_2\})$, indicating that there is still a relatively high degree of epistemic uncertainty associated with h_2. Therefore, an analyst might decide to first investigate the case further to lower his or her degree of epistemic uncertainty before making a definitive decision on what he or she considers to be the best answer to q.

Benefits:

- The BPA can be constructed with almost any kind of data. For instance, regular probabilities can be assigned to the singleton values in S, and complete ignorance can be modelled by setting $m(S) = 1$ [3].
- BFT encompassed probability theory when the BPA is only defined over singletons or disjoint sets [3].
- It produces bounds that get narrower with better empirical information [3].
- Again, it is relatively simple to implement [3].

Limitations:

- When combining multiple BPAs, not only should you be cautious of dependencies between the used evidence, the shape of the mass functions should also be considered when selecting an appropriate combination rule. For instance, one of the most popular rules to do so, Dempster's rule, produces unintuitive results when faced with contradicting mass functions [3, 16].
- The output measures of belief, plausibility and commonality tend to be more difficult to translate to specific decision points than the standard probability over the singletons. A translation is needed that can transform the beliefs from a credal level to a pignistic level [17]. Smets introduces such a translation in [17], but other translations exist as well [18].

3.3. Future Steps

The presented work in the previous section provides a preliminary idea of how a potential uncertainty representation might be integrated into the TIDO ontology, focussing on how the uncertainty of hypotheses derived within a TIHW could be represented. We consider this to be the first step answering **RQ3**, but other steps are still to be taken. We have roughly divided the required extensions to TIDO into five steps, as depicted in Figure 3, and we plan on addressing these steps one at the time.

Step 1: Uncertainty Representation over the Hypotheses Set

We believe that the expression over the uncertainty over the available hypotheses should be the first point addressed in the uncertainty representation of TIDO as this represents how the analyst has interpreted the available information and provides the foundation for follow-up actions. Being certain about a hypothesis, or a group of hypotheses, being true could be used as an argument for mitigating actions, where as a high degree of uncertainty could be used as an argument to investigate further. The example provided in Section 3.2.3 provides an indication of what an uncertainty representation for the hypothesis set might look like. However, the following points should also be addressed before step 1 can be considered finished:

Finish examples: As of yet, only the example for BFT has been worked out. Before a proper comparison can be made between BFT and basic probability theory and interval probabilities, the examples for these cases will need to be worked out as well. Alternatively, we could focus our attention on the *probability bound analysis* (PBA) paradigm, which combines basic probability theory with interval probabilities [3] and already has some implementations in the field of threat assessment [19, 20].

Expert elicitation: For the example in Section 3.2.3, educated guesses were made on what a *possible* mass function could look like given the data depicted in Figure 2. However, if the TIDO ontology is to be used by domain experts, a quantified degree of belief over the available hypothesis must be derived from the experts' input. For the elicitation of mass functions, many methodologies exist, such as the audit risk model [21], the Analystical Hierachy Process (AHP) [22], a pair-wise ranking based procedure [23], and a methodology using Likert-scales [24]. Expert elicitation techniques for obtaining subjective probabilities or probability intervals still need to be investigated. The feasibility of implementing a suitable elicitation methodology should also play an important role in the selection of the uncertainty representation for TIDO.

Representation in a knowledge graph: In order to capture the data, it must somehow be stored in a database. As the TIDO ontology is based on the Web Ontology Language (OWL), it is important to think about how the uncertainty can represented such that it is compatible with the OWL syntax, or more specifically the Resource Description Framework (RDF). For example, directly implementing BFT would require mass values associated not only with the singletons in the set of hypotheses H, but also with the elements in the powerset of H. Worst case scenario, every element in the powerset of H would need a unique node in the graph, resulting in an exponential growth in the amount of nodes in the KG with respect to the size of H. Not only could this substantially increase the required storage capacity, it might also make it more difficult to retrieve this information from the graph using SPARQL queries. Some OWL ontologies already exist to represent uncertainties, and before designing a custom extension to TIDO, potential mappings from TIDO to these ontologies should be investigated. For probability theory, PR-OWL [25] could be an option and for BFT, BeliefOWL [26] could be an option. No OWL ontology was found for probability intervals.

Step 2: Conditional Probabilities or Belief Functions

Once step 1 is completed, there is an uncertainty representation that allows a TI expert to represent their belief about the uncertainties associated with the hypotheses in H, which, ideally, the expert derives

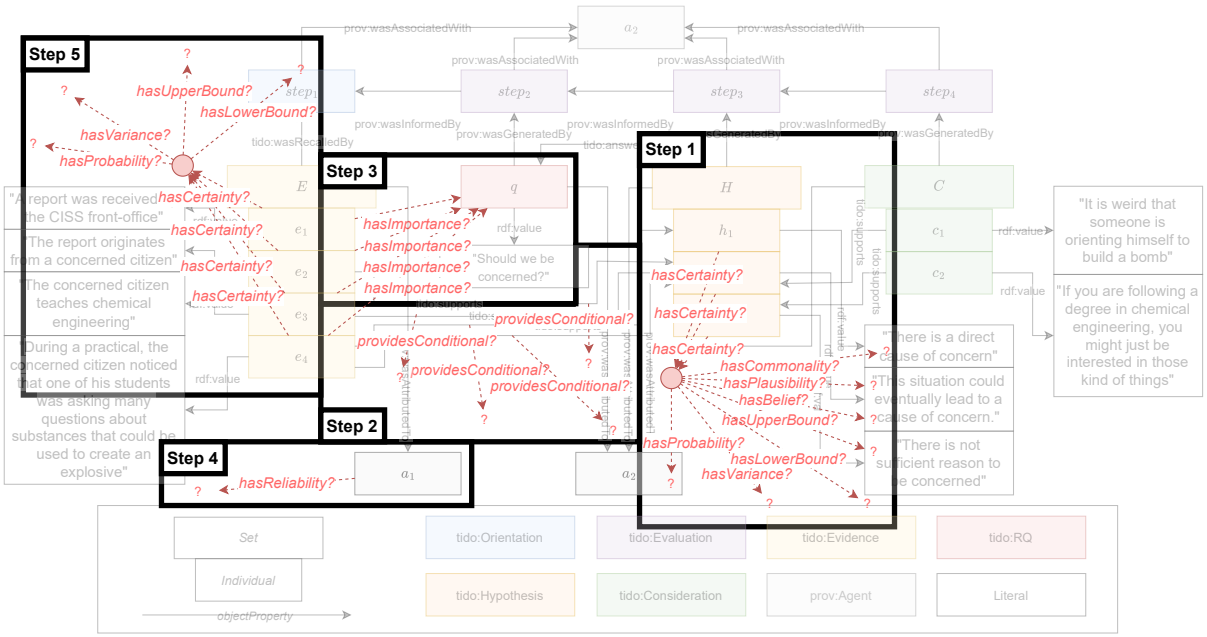

Figure 3: An indication on how the characteristics of the data used within the TI domain, as identified in Section 3.1, could be integrated into a representation of uncertainty applied to the example provided in Figure 2. Here, the steps indicate which extension to the TIDO ontology is prioritized.

from the pieces of evidence in E and the considerations in C. However, there is no vocabulary for the expert to precisely express *how* these pieces of evidence and considerations influence the uncertainty associated with the elements in H. This information could be very useful when determining to what degree each piece of evidence or consideration contributed to the uncertainty distribution over the hypothesis, and consequently, the final decision. The TIDO ontology does provide the opportunity for the expert to express whether pieces of information support (using `tido:supports`) or dispute (using `tido:disputes`) a hypothesis, but the effect of using these relations on the uncertainty associated with the hypotheses in H is not formalized. To formalize this relationship, some form of a conditional probability distributions ($P_q(A|\cdot)$) or conditional belief masses ($m_q(A|\cdot)$) would be needed, where $A \in \mathscr{P}(H)$ and $\cdot$ would be a piece of information relevant for answering research question q.

To implement this step, the same points need to be addressed as for step 1: What would a working example look like for each of the selected uncertainty representations? How would we elicit the needed numerical values from the experts? And how could this be represented in a KG? As of yet, these points have not been investigated, although we expect to encounter challenges when addressing the dependencies that exist within the available information. For example, the pieces of evidence e_3 and e_4 depicted in figure 2 show a clear dependency as e_3 provides an indication that 'concerned citizen' in e_4 is knowledgeable about explosive substances, increasing the degree to which e_4 provides an indication for hypothesis h_1.

Step 3: Importance with Respect to the Research Question

In Section 3.1 we discussed that we use the term *relevance* for refer to the relevance of pieces of information to the case investigation, and the term *importance* to describe the relevance of a piece of information with respect to the research question. Therefore, we can use formal representations of *relevance* to model our notion of *importance*. For example, we can capture this context dependence by introducing a distance measure, as it is commonly done in frameworks that need to deal with different levels of relevance [27, 28]. In our case, we could define a distance between the pieces of information used within the analysis and the research question.

Step 4: Reliability of the Source

Several options exist for accounting for the reliability of the source in a potential uncertainty representation for TIDO. The author of [13] suggests to use the formalisms used to model imprecise probabilities to model the reliability of a source instead. Alternatively, the authors of [29] describe how a combination of the reliability of the source and the degree of disagreement within evidence can be used to compose a discount factor for a mass function, similarly to how [30] uses the truthfulness of the source as a correction factor.

Step 5: Uncertainty Representation over information from External Sources

The final step that warrants investigation is the representation of uncertainty over the evidence set E. Both the hypotheses set H and the consideration set C are generated by agent a_2, implying that agent a_2 is also responsible for providing the uncertainties associated with the elements in H and C. However, the elements in E are provided by a different agent a_1. In theory, this agent could represent the uncertainties associated with the pieces of information completely different from how a_1 would represent these uncertainties. For example, a_1 could use interval probabilities to express its uncertainties, whereas a_2 could us BFT. If a_2 is to use the uncertainties provided by a_1 to make assessments about the hypotheses in H, the uncertainties from a_1 would need to be translated to a representation compatible with the uncertainty representation of a_2. Following this line of reasoning, it would make sense for a_2 to use the most expressive form of uncertainty representation, but the possibility utilizing some form of a translation function should be investigated as well.

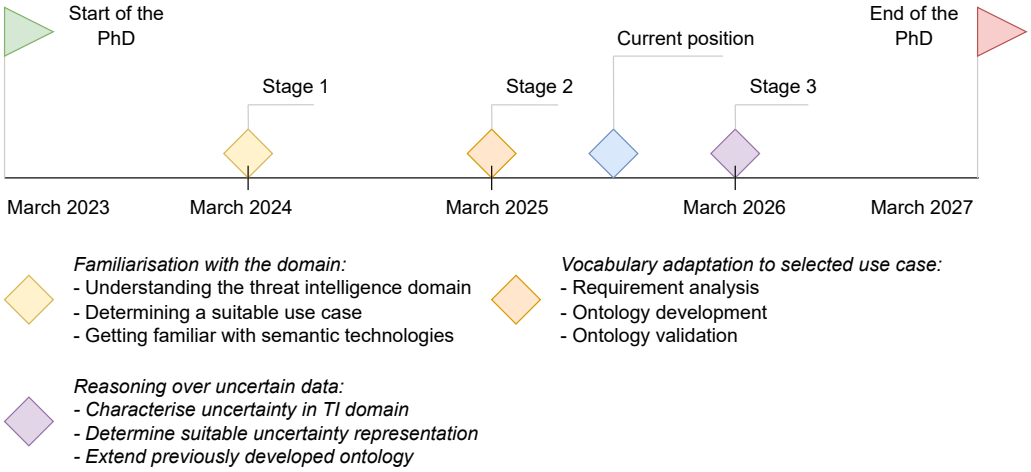

Figure 4: Timeline of the authors PhD plan

3.4. Available Tooling

As mentioned in Section 1, our overview of the available tooling to represent and reason over uncertainties is limited, and we hope to obtain interesting leads during the doctoral consortium. So far, with respect to modelling probabilistic ontologies, we have found PR-OWL [25] and BayesOwl [31], which both represent the uncertainties in the form of Bayesian networks. While either of these frameworks appear to be suitable if we opt to represent the uncertainties in TIDO using standard probability theory, these frameworks do not appear to be compatible with interval probabilities or BFT. With respect to BFT, we only found a single position paper introducing the core concepts of BeliefOWL [26], but a complete implementation of this ontology appears to be absent. For interval probabilities, we were not able to find an published OWL-based ontology. Instead, we could investigate the usage of the `pba-for-python` Python package [32] if we do not wish to implement such an ontology ourselves.

4. Approach, Progress and Conclusions

The aim of the presented research is to develop a vocabulary capable of representing a Threat Intelligence Hybrid Workflow (TIHW). The contribution of the student to this work is summarized in three stages: (1) The start-up of the project, where the main goal is to get familiar with the domain, (2) the adaptation of available vocabularies to a use-case identified during the first phase, and (3) the integration of an uncertainty reasoning paradigm into the ontology developed in the second phase. A timeline of these phases is shown in Figure 4.

The first two phases can be considered as completed with the development of the TIDO ontology described in Section 2. The remainder of the work presented in this consortium paper represents the first steps towards the third phase. In this phase, the student extended the domain analysis of the TI domain with a characterisation of the data used within this domain, as presented in Section 3.1. Afterwards, three different uncertainty representations, general probability theory, interval probabilities and belief function theory, are analysed in Section 3.2. Based on this analysis, a first step is made in the integration of an uncertainty representation into the TIDO ontology in Section 3.3.

Up until now, our focus has mainly been on belief function theory, but the upcoming time will be spend on drafting similar examples for the general probability theory and interval probability implementations. Afterwards, we will investigate how to integrate conditional probabilities and belief functions, our notion of importance with respect to the research question, the reliability of the source, and the uncertainty representation for information originating from external sources. Additionally, we still wish to investigate what tooling is available that allows for reasoning over uncertain information using the formalisms we previously analysed.

As for the potential achievements of this research, we believe that incorporating a suitable uncertainty representation into the TIDO ontology would substantially improve its ability to accurately capture the TIHW. As the decision processes within the TI domain are greatly influenced by the uncertainties associated with the information used within these processes, an accurate representation of these uncertainties would not only allow an analyst to make a better assessment of the situation, but also allow for the reconstruction of the perceived uncertainties during an audit of such a TIHW.

Acknowledgments

This PhD project is supervised by Professor Fabio Massacci, and co-supervised by Associate Professor Stefan Schlobach and Assistant Professor Lise Stork. Additionally, the author would like to thank Daira Pinto Pietro for her discussions on how to possible integrations of belief function theory into the TIDO ontology. Finally, we gratefully acknowledge the funding support by the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) under the KIC HEWSTI Project with grant no. KIC1.VE01.20.004.

Declaration on Generative AI

During the preparation of this work, the authors used Gemini 2.5 Flash in order to: Paraphrase and reword, Grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] S. Van Gerwen, J. Constantino, R. Roothaert, B. Weerheijm, B. Wagner, G. Pavlin, B. Klievink, S. Schlobach, K. Tuma, F. Massacci, To Know What You Do Not Know: Challenges for Explainable AI for Security and Threat Intelligence, in: Artificial Intelligence for Security, Springer Nature Switzerland, Cham, 2024, pp. 55–83. doi:10.1007/978-3-031-57452-8_4.

[2] F. Massacci, A. Papotti, B. Weerheijm, S. A. M. van Gerwen, E. Mezzi, J. Constantino Torres, R. Roothaert, Hybrid Explainable Workflows for Security and Threat Intelligence, 2023. URL: https://www.nwo.nl/en/projects/kich1ve0120004.

[3] E. Zio, N. Pedroni, Literature Review of Methods for Representing Uncertainty, Technical Report, Fondation pour une culture de sécurité industrielle, 2013.

[4] S. Miller, On National Security Intelligence, 1 ed., Routledge, London, 2024, pp. 33–50. doi:10.4324/9781003106449-3.

[5] S. Chávez-Feria, R. García-Castro, M. Poveda-Villalón, Chowlk: From UML-Based Ontology Conceptualizations to OWL, in: P. Groth, M.-E. Vidal, F. Suchanek, P. Szekley, P. Kapanipathi, C. Pesquita, H. Skaf-Molli, M. Tamper (Eds.), The Semantic Web, volume 13261, Springer International Publishing, Cham, 2022, pp. 338–352. doi:10.1007/978-3-031-06981-9_20.

[6] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, PROV-O: The PROV Ontology, 2013.

[7] R. F. Bales, F. L. Strodtbeck, Phases in group problem-solving. 46 (1951) 485–495.

[8] P. Costa, A.-L. Jousselme, K. B. Laskey, E. Blasch, V. Dragos, J. Ziegler, URREF: Uncertainty representation and reasoning evaluation framework for information fusion, Journal of Advances in Information Fusion 13 (2018).

[9] L. Rasker, De Dienst Podcast, 2021. URL: https://podcastluisteren.nl/pod/De-Dienst.

[10] Dhr. P.A. Brouwer, Maj M. Scholten, Joint Doctrine Publicate 2: Inlichtingen, 2012.

[11] A. D. Kiureghian, O. Ditlevsen, Aleatory or epistemic? Does it matter?, Structural Safety 31 (2009) 105–112. doi:10.1016/j.strusafe.2008.06.020.

[12] R. Roothaert, Supplementary materials used for the development of the TIDO ontology, 2025. doi:10.5281/ZENODO.15400486.

[13] L. V. Utkin, IMPRECISE RELIABILITY: AN INTRODUCTORY REVIEW, 2003.

[14] A. P. Dempster, Upper and Lower Probabilities Induced by a Multivalued Mapping, The Annals of Mathematical Statistics 38 (1967) 325–339. doi:10.1214/aoms/1177698950.
[15] G. Shafer, A Mathematical theory of Evidence, Princeton University Press, 1976.
[16] J. A. Barnett, Computational Methods for A Mathematical Theory of Evidence, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 197–216. doi:10.1007/978-3-540-44792-4_8.
[17] P. Smets, R. Kennes, The transferable belief model, Artificial Intelligence 66 (1994) 191–234. doi:10.1016/0004-3702(94)90026-4.
[18] T. Denœux, Decision-making with belief functions: A review, International Journal of Approximate Reasoning 109 (2019) 87–110. doi:10.1016/j.ijar.2019.03.009.
[19] S. Ferson, Probability Bounds Analysis Solves the Problem of Incomplete Specification in Probabilistic Risk and Safety Assessments, in: Risk-Based Decisionmaking in Water Resources IX, American Society of Civil Engineers, 2001, pp. 173–188. doi:10.1061/40577(306)16.
[20] R. Flage, T. Aven, C. L. Berner, A comparison between a probability bounds analysis and a subjective probability approach to express epistemic uncertainties in a risk assessment context – A simple illustrative example, Reliability Engineering & System Safety 169 (2018) 1–10. doi:10.1016/j.ress.2017.07.016.
[21] R. P. Srivastava, L. Liu, Applications of Belief Functions in Business Decisions: A Review, Information Systems Frontiers 5 (2003) 359–378. doi:10.1023/b:isfi.0000005651.93751.4b.
[22] A. Ennaceur, Z. Elouedi, E. Lefevre, Belief AHP Method — AHP Method with the Belief Function Framework, International Journal of Information Technology & Decision Making 15 (2016) 553–573. doi:10.1142/S0219622016500139.
[23] L. V. Utkin, Ranking procedures by pairwise comparison using random sets and the imprecise Dirichlet model, Applied Mathematics and Computation 183 (2006) 394–408. doi:10.1016/j.amc.2006.05.069.
[24] T. N. Hoang, T. T. Dao, M.-C. Ho Ba Tho, A Method for Uncertainty Elicitation of Experts Using Belief Function, in: Studies in Computational Intelligence, Springer International Publishing, Cham, 2018, pp. 39–49. doi:10.1007/978-3-319-76081-0_4.
[25] R. N. Carvalho, K. B. Laskey, P. C. Costa, PR-OWL – a language for defining probabilistic ontologies, International Journal of Approximate Reasoning (2017) 56–79. doi:10.1016/j.ijar.2017.08.011.
[26] A. Essaid, B. B. Yaghlane, BeliefOWL: An Evidential Representation in OWL Ontology, in: Proc. of the Fifth International Workshop on Uncertainty Reasoning for the Semantic Web, 2009.
[27] L. Chen, L. Diao, J. Sang, Weighted Evidence Combination Rule Based on Evidence Distance and Uncertainty Measure: An Application in Fault Diagnosis, Mathematical Problems in Engineering 2018 (2018) 1–10. doi:10.1155/2018/5858272.
[28] D. D. In, H. Kim, Distance-Based Relevance Function for Imbalanced Regression, Stats 8 (2025) 53. doi:10.3390/stats8030053.
[29] Y. Yang, D. Han, C. Han, Discounted combination of unreliable evidence using degree of disagreement, International Journal of Approximate Reasoning 54 (2013) 1197–1216. doi:10.1016/j.ijar.2013.04.002.
[30] F. Pichon, D. Mercier, F. Delmotte, E. Lefèvre, Truthfulness in Contextual Information Correction, in: Lecture Notes in Computer Science, Springer International Publishing, Cham, 2014, pp. 11–20. doi:10.1007/978-3-319-11191-9_2.
[31] Zhongli Ding, Yun Peng, A probabilistic extension to ontology language OWL, in: Proceedings of The 37th Annual Hawaii International Conference on System Sciences, IEEE, Big Island, HI, USA, 2004, p. 10 pp. doi:10.1109/HICSS.2004.1265290.
[32] N. Gray, S. Ferson, M. De Angelis, A. Gray, F. Baumont De Oliveira, Probability bounds analysis for Python, Software Impacts 12 (2022) 100246. doi:10.1016/j.simpa.2022.100246.