Article

A Brief Roadmap into Uncertain Knowledge Representation via Probabilistic Description Logics

Rafael Peñaloza

IKR3 Lab, University of Milano-Bicocca, Italy; rafael.penaloza@unimib.it

Abstract: Logic-based knowledge representation is one of the main building blocks of (logic-based) artificial intelligence. While most successful knowledge representation languages are based on classical logic, realistic intelligent applications need to handle uncertainty in an adequate manner. Throughout the years, many different languages for representing uncertain knowledge—often extensions of classical knowledge representation languages—have been proposed. We briefly present some of the defining properties of these languages as they pertain to the family of probabilistic description logics. This limited view is intended as a way to help the interested researcher find the most adequate language for their needs, and potentially identify the gaps remaining.

Keywords: Knowledge representation; uncertainty; probabilistic reasoning; survey

Citation: Peñaloza, R. A Brief Roadmap into Uncertain Knowledge Representation. *Algorithms* **2021**, *1*, 0. https://doi.org/

Received:
Accepted:
Published:

1. Introduction

Logic-based knowledge representation [1] is one of the fundamental building blocks for (logic-based) artificial intelligence. In fact, any intelligent application has, as an unavoidable requirement, the need to represent and handle the knowledge about the domain that it works in [2]. This need has led to a plethora of knowledge representation languages targeting diverse properties and applications of the knowledge and its management. In their classical version, these languages are designed to deal with perfect knowledge, in the sense that knowledge is assumed to be precise, certain, and correct. In general, however, knowledge is not perfect, and knowledge representation and reasoning systems should be able to handle these cases as well, if they are ever to be used in practice.

One prominent case of imperfect knowledge, which arises in many natural applications including medicine and biology, but also economics and sociology is the presence of *uncertainty*. This refers to facts or situations which may hold or not, but we simply cannot know *a-priori* (without an intervention or an observation) which is the case. To deal with the uncertainty of these domains, many uncertain knowledge representation languages have been developed as well.[1] Just as in the classical case, several uncertain knowledge representation languages can be developed, depending on the desired logical, computational, and practical properties that they should have. Importantly, uncertainty adds a new dimension over which further variants can be constructed: starting from the chosen uncertainty representation, up until the source of uncertainty, passing through several additional considerations which impact not only the semantics, but also their applicability, underlying assumptions, and reasoning efficiency.

Given this large landscape of uncertain knowledge representation formalisms, it is easy for a newcomer to get lost in an attempt to understand the area, or simply to grasp the most adequate language for their needs. As a consequence, the entry cost for dealing with uncertain knowledge representation is unreasonably high, specially for users who may only be interested in *using* the formalisms, as opposed to developing or extending

[1] Importantly, *uncertain knowledge representation* refers to the representation of uncertain knowledge, not to uncertainty in the representation.

them. This is one of the largest hurdles for the adoption of uncertain representation formalisms, and draws the risk of choosing a wrong paradigm for a given application with potentially catastrophic consequences.

This paper is an attempt to decrease that entry cost, by providing a very brief roadmap into knowledge representation formalisms dealing with uncertainty. The roadmap is limited in many aspects: it focuses primarily on probabilities as uncertainty representation, and uses a well-known but by no means all-encompassing family of (classical) knowledge representation formalisms as the basis for constructing probabilistic extensions. Still, the discussion on the ideas behind the formalisms and their properties and limitations will hopefully provide enough of a background for an interested reader to understand other variants and explore the rest of the forest by themself. In particular, the pattern for extending logical formalisms with probabilities repeats almost unchanged throughout languages. What this roadmap does not provide are the tools to deal with other kinds of uncertainty representations [3,4] such as possibility theory [5,6] or evidence theory [7,8]; nor any other kinds of imperfect knowledge like vagueness [9–11], or inconsistency [12–14].

The structure of the paper is as straightforward as it can be. We first discuss the fact that *uncertainty* is a multi-faceted issue, and the need to understand which face is relevant for each specific application. In that section we also restrict our attention to probabilities, justifying our choice. Afterwards, we introduce a class of uncertain knowledge representation formalisms built as extensions of the well-known family of description logics. Although this choice limits the class of languages studied, it gives a general overview of the issues encountered, and the kinds of uncertain representation and reasoning available.

2. The Many Faces of Uncertainty

Since our goal is to formally handle uncertain knowledge, we should start to clarify what uncertainty *is*, and how it can be quantified, combined, and more generally, manipulated. Skimming over the many, and deeply interesting philosophical discussions on the topic, we consider the most etymological version of the term: uncertainty as a lack of certainty. In this respect, uncertain knowledge is *not* a lack of knowledge (which classical knowledge representation languages handle effectively by means of the so-called *open world semantics*) nor imprecise knowledge, which is the scope of fuzzy logic [11]. Instead, uncertainty refers to properties or events which either hold or not, but we cannot certainly know which is the case beforehand. Consider the typical example of a coin toss. Before the toss, we know that it will either land on *heads* or on *tails*, but we have no way of knowing *a-priori* which is the case; thus we are uncertain of the result, up until the point when the coin is tossed. Note that uncertainty of a property may be an indirect consequence of other certain or uncertain properties, some of which may not be obviously stated. For instance, if one makes a bet on the coin toss, then it is known with certainty that *if* the coin lands on heads, then they win $1, and they loose $1 otherwise. However, it is still uncertain which will be the actual case until the toss is made.

The obvious next question is how can one represent and manage such uncertainty. When first encountering this question, most people turn immediately to the notion of *probability*. Indeed, probabilities and their close friends the percentages are taught to most of us from a relatively early stage, and we encounter them almost daily in all aspects of our life; we in fact use probabilistic terminology in our daily-life interactions. This familiarity with the theory of probabilities is both a blessing and a curse. On the one hand, it greatly reduces the *entry cost* of dealing with uncertainty in a formal setting, removing the wall of introducing a new theory along with its nomenclature and notation. On the other, the heavy baggage of probabilities includes many misunderstanding and erroneous intuitions that we have grown used to accept as true. In part for this reason and in part for other issues that we will point to later on, other uncertainty representations have been proposed; most notably, possibility theory [5] and evidence theory [7,8].[2] For the scope of this paper, as in most of the literature in uncertain knowledge representation, we give more weight to the advantages of the easiness of presentation over the likelihood of misunderstandings. Thus, we consider probability theory as the basis for representing and managing uncertainty in the context of knowledge representation. However, we still need to take into account the different interpretations of probability as uncertainty.

Despite its unified name, and the use of probabilities for handling it, not all uncertainty is equal. Halpern [15] already hinted at it when describing its different types of logics of probability. Broadly speaking, Halpern's classification considers two kinds of views on uncertainty: a *statistical* one referring to a proportion of the population satisfying a property of interest, and a *subjective* one dealing with beliefs about possible worlds. The difference lies in how the uncertainty is used within a derivation or reasoning process, but mirrors existing differences from real-life use of probabilities. However, it is important to note that Halpern's classification is orthogonal to the usual distinction between frequentist and Bayesian probabilities, about which we refrain from mentioning anything further in this text.

Statistical probabilities come into play when speaking about proportionality, and a random selection of elements. Hence, when we say that a medical test has a 95% diagnostic specificity—in lay terms, that if the test is positive, then there is a 95% chance that the individual is in fact positive for the disorder under scrutiny—what we are saying is that 95 out of every 100 positive tests are correct (and the remaining 5 are wrong). Hence, if we randomly take one of these tests, it has a 95% chance of being a correct one. Note that one can only be so specific about the probability if the whole population is known.

In contrast, subjective probabilities consider unique instances which are characterised by different possibilities. The prototypical example in this direction is the weather forecast. When a meteorological model predicts a 40% chance of rain tomorrow, it cannot be read as a statistical statement saying that in 40 out of 100 *tomorrows* rain will be present. Instead, it studies different scenarios based on possible parameters like wind speed and direction, temperature, humidity, and others, to verify in which of those scenarios rain is present.[3]

Halpern's classification, however, is not fully satisfying in the context of knowledge representation, and in particular in the context of incomplete domain knowledge and expert knowledge. We use these two cases to exemplify the limitations of each of the two types of probabilities.

As mentioned already, statistical probabilities are derived as proportional observations of an event within a given population. The term *statistical* hence refers to a very basic analysis of data. The name is unfortunate, as it also evokes the use of more advanced statistical analyses, which are *not* foreseen in these logics. The most basic example is the presence of incomplete knowledge. While it is pretty straightforward to find out the exact proportion of students in a given classroom who are left-handed, the same cannot be said about e.g., COVID-19 patients who have pulmonary scars. To know this latter proportion, it is necessary to identify precisely who has been infected with the disease, and make a pulmonary plaque on all those subjects. Both of these tasks induce high economic, social, and human costs which one might not be willing to cover. Instead, it is possible to approximate this knowledge using a statistical analysis on the available data of publicly known infected individuals, and results from hospital analyses from people suspect of having lung issues arising from it. Alternatively, one can also *sample* the population to estimate these proportions. Both ideas are intended to fill the gap left by the incomplete knowledge of exactly how many people fall into each of the

2 Often referred to also as *Dempster-Shafer theory*.

3 In reality, the values provided by actual weather forecasts are more complex, as they also take into account the area of the region under consideration [16]. For the sake of the example, we do not delve deeper into these details.

categories of interest. The cost, however, is that there are (uncertain) margins of error that one needs to deal with.

Let us consider now subjective probabilities, which aim to represent beliefs about the likelihood of specific events. A common use of subjective probabilities is for modelling expert knowledge, where a (human) expert may—perhaps based on past observations—assign a probability to an event. In these cases, the numerical values underlying probabilities (and their algebraic manipulations) become more a hindrance than an advantage. Indeed, there is no-one capable of discerning a probability of 95% from one of 95.5% nor, for that matter, 60% from 70%. Importantly, even subtle differences may cause huge mismatches over a derivation process; they could even lead to inconsistency in the collected knowledge. In these cases, it is perhaps more useful to represent *comparative* statements, of the form *X is more likely than Y*. However, this requires the development of new reasoning techniques, specially in the presence of mixed statements. Moreover, it comes at the price of losing precision. On the other hand, these statements are more easily understandable by the lay person, and describable by the experts.[4]

As it can be seen from this section, representing and managing uncertain knowledge is far from trivial, even from the point of view of choosing the measure of uncertainty. The landscape of probabilistic interpretations is vast, and different applications have diverse needs for expressivity. If we attach other practical considerations like complexity of reasoning, availability of resources, or historic knowledge to use, the panorama gets even more diverse. It is important to keep this diversity in mind when studying uncertain knowledge representation languages to avoid getting lost among the variants that they induce. This is, in fact, one of the biggest obstacles faced by researchers trying to get started in the area: not knowing the differences in the probabilistic interpretations, exploring the state of the art seems a Sisyphean task.

The following section is an attempt to draw a map of the uncertain knowledge representation landscape and highlight active work and potential gaps.

3. Representing Uncertain Knowledge

Representing uncertain knowledge has a prerequisite representing knowledge, full stop. Knowledge representation, by itself, has a very long history, during which a plethora of variations, limitations, and features have been considered. A natural first step is to consider a known logic for representing knowledge; hence, one cannot avoid mentioning propositional and (first-order) predicate logic as the foundations of logic-based knowledge representation languages. However, from a practical point of view, propositional logic tends to be too inexpressive, and even the elements which can be expressed sometimes require a complex and difficult to grasp construction to handle correctly. On the other spectrum, in full predicate logic it is known that verifying the satisfiability of a formula (which in terms of knowledge representation translates to deciding whether a knowledge base is consistent) is an undecidable problem; that is, there is no algorithm which can provide a correct answer in finite time for any possible formula.

For this paper, we focus on a family of formalisms which lies mainly within these two formalisms. More specifically, most languages within this family—the family of Description Logics (DLs) [17]—are more expressive than propositional logic (thus, able to formalise more complex knowledge in a simpler manner) and at the same time less expressive than predicate logic guaranteeing decidable reasoning tasks (with consistency among them). There are a few exceptions to this statement, which only help in increasing the relevance of the family as knowledge representation formalisms. The very inexpressive DLs $\mathcal{EL}$ [18] and DL-Lite [19], which are specially targeted for tractable reasoning, do not contain the full power of propositional logic although they allow

[4] This is one of the settings where possibility theory becomes relevant: under some specific interpretations, the exact numerical values are cast aside in favour of their ordering.

for additional constructors. At the other end of the spectrum, expressive description logics like $\mathcal{SROIQ}$ [20] include constructors (like transitive closure) which cannot be directly expressed in first-order logic. These are handled in a manner that prevents undecidability of reasoning.

The semantics of description logics, which is based on interpretations akin to first-order logic—that is, with a domain representing all the relevant objects, and an interpretation function which expresses the properties of those individuals in relation to each other—is specially useful for dealing with the various interpretations of uncertainty. We will see this in detail later, but in a nutshell and using Halpern's classification, statistical probabilities are handled by adding uncertainty over the elements of the domain (i.e., the *population*) while subjective probabilities are dealt with through several potential interpretations (*possible worlds*). As mentioned before, the differences may be important.

For all these reasons, we consider description logics as a basic formalism for representing uncertain knowledge. This is meant mainly as a prototypical representation: most of the ideas that we describe apply similarly to other formalisms without major modifications. We emphasise, however, that the classical family of description logics has some limitations which we will not consider further. Most notably, it cannot handle non-monotonic [21], nor temporal knowledge [22,23] natively. Importantly, combining uncertainty with non-monotonicity and with temporal constructors is known to be specially problematic [24], both in terms of conceptual understanding and in the computational complexity of reasoning.

Without going into too many details, the basic building blocks in a description logic are *concepts* (that is, sets of individuals) and *roles*, which represent relationships between individuals; slightly more formally, concepts are unary predicates, and roles are binary predicates of first-order logic. Hence, Student is a concept that refers to all the students in the world of interest, while supervises expresses the relationship between a supervisor and their student. These symbols receive an interpretation by setting a (potentially infinite) *domain*, which contains all the objects of interest, and an *interpretation function* expressing which objects belong to which concepts, and which pairs are related via roles. What differentiates one description logic from another is the class of *constructors* used to build more complex concepts—e.g., conjunction, negation, number constraints, etc.—and how they are interpreted.

The goal of description logics is not only to express different kinds of concepts, but to actually represent the *knowledge* of a domain. This is achieved through a *knowledge base* which is a finite set of *axioms* that serve as constraints for the interpretations. That is, each axiom excludes some potential interpretations as not representing the domain knowledge. For example, an axiom could express that *"every student must have at least one supervisor."* In this case, any interpretation including a supervisor-free student will be excluded as a violation of the constraint. In general, given a knowledge base, there are still many different (actually, infinitely many) interpretations which satisfy all the constraints imposed. These so-called *models* are the only interpretations of interest in the context of the knowledge base.

When we use the term *reasoning*, we refer to the task of extracting consequences which logically follow from the knowledge expressed in the knowledge base. Recalling that the axioms within the knowledge base are simply constraints in the possible interpretations, reasoning then refers to finding other pieces of knowledge which are guaranteed by these constraints. In other words, the logical consequences of a knowledge base are those which follow in all possible models of this set of axioms. We usually say that reasoning is the task of making knowledge which is *implicitly* encoded by the knowledge base *explicit*. The motivation behind using several models for reasoning is that we consider that a knowledge base is always (necessarily) incomplete. That is, we believe that a knowledge base will always exclude some information, either because it is irrelevant, or because it is not yet known. In those cases, we want to leave open the possibility of such an assertion being true or false, until it is known. This approach is commonly known as the *open world assumption* in the literature.

Once again, a knowledge base defines a class of interpretations, each of which introduces a set of individuals. When knowledge is uncertain, we thus have two natural choices to introduce a probability distribution: it can be defined over the class of interpretations, expressing the likelihood that each of them represents the actual state of the world, or it can be defined over the individuals of the interpretation domain, differentiating the characteristics of the individuals. These two choices transfer easily to the two kinds of probabilistic logics in the classification by Halpern. This correspondence has given rise to several probabilistic description logics.

3.1. Subjective Probabilities

Consider first the case of subjective probabilities. These refer to the situation where the uncertainty is about the state of the world, and hence about the specific model under consideration. Thus, the semantics of these kinds of logics introduce a probability distribution over the class of relevant models, expressing which of them are more likely. From a syntactic point of view, it is necessary to express the likelihood of the knowledge appearing in the knowledge base. The typical approach is to associate to each axiom (that is, each constraint) a probability degree. This probability expresses the (subjective) belief that the constraint expressed by the axiom actually holds in the world [25,26]. Intuitively, if this probability is p, then the probability distribution should assign probability p to the set of all the interpretations which satisfy p and probability $1 - p$ to its complement. Unfortunately, things as not as easy as they seem at first sight, and there are many aspects to take into account.

One issue is how to relate the probabilities of different axioms with each other for the construction of the probability distribution. That is, if an axiom α has probability p and another axiom β as probability q, which probability should one assign to the class of interpretations satisfying both axioms α, β? In most cases, to simplify the language and the probability computations, it is common to assume that axioms are probabilistically independent—that is, that the truth value of one does not affect the likelihood of the other. Under this assumption, the probability of satisfying both axioms becomes pq. This assumption, however, is not always realistic in a knowledge representation application. For one, as knowledge bases tend to be big, knowledge engineers often rely on modelling guidelines, which specify how some specific kinds of knowledge should be represented in a given scenario. These guidelines commonly require simple axioms, which can only specify complex knowledge when combined with other axioms. If this complex knowledge is uncertain, it is unreasonable to assume that all the small pieces building it are probabilistically independent. The other side of the same coin are the normalisation steps often performed implicitly to aid reasoning. Again, in this case many simple axioms are generated from one complex one, but clearly they are all dependent on each other. Formally, to solve this issue one would need to specify the full probability distribution of the axioms or at the very least the joint probabilities for all relevant combinations of axioms. Unfortunately, this solution requires a complex representation and slows reasoning. Some approaches have been proposed to use only partial independence assumptions [27–29]. Another approach is to consider all possible coherent assignments of probabilities in what is known as Nilsson's semantics [30]. However, this semantics does not satisfy the axioms of probability in general, and knowledge manipulation always increases imprecision.

A second issue arises from the presence of the open world assumption. Recall that we previously said that if an axiom holds with probability p, then the interpretations that do not satisfy this axiom should have probability $1 - p$. The issue, however, is with guaranteeing that the remaining interpretations indeed do not satisfy an axiom, and deciding what precisely that (i.e., violating an axioms) means in practice. From the open world assumption, we note that knowledge which is not explicitly stated could be true or false. When dealing with uncertain axioms, we will have some interpretations where the axiom *explicitly* holds, and some—due to the open world assumption—where it may hold, but is not required to do so. This means that in reality the probability stated by the axiom is only a lower bound: the likelihood of making it true may in fact be higher. Once again: although one may conceivably construct a logic where the probabilistic value is precise, by explicitly violating the axiom in all remaining interpretations via some kind of closed-world interpretation, this can have unexpected consequences. This happens, in fact, in [28]. In this logic, the semantics guarantee a closed-world interepretation. However, this has been shown to produce some counter-intuitive behaviour, and in particular to lead to inconsistency even in simple cases.

A third issue is also closely related to the open world assumption, but arises mainly from the fact that knowledge is assumed to be incomplete within a knowledge base. Combined with the existence of *implicit* knowledge which may be extracted through reasoning, the probabilities of different axioms, and of their consequences may be in evident conflict. As a very simple example, consider a setting where knowledge is redundant, in the sense that different sets of axioms state the same knowledge. It may very well happen, due to the nature of knowledge base engineering, that the probabilities associated to these classes differ, yielding two different (conflicting) probability degrees to the same piece of knowledge. Deciding how to solve these conflicts—by computing the maximum, following a full probabilistic approach, or simply declaring inconsistency—is a design choice which impacts the accuracy, practicality, and complexity of the language and its reasoning tasks.

At this point it is perhaps worth mentioning also an approach which avoids fully specifying probabilistic degrees, but instead gives more importance to their ordering; that is, uncertainty values are specified relative to each other, rather than absolutely as probability degrees. In log-linear logics [31,32], each axiom is assigned a *weight*, which is real number not necessarily in the interval $[0,1]$. At the very basis of the interpretation of these values is an *ordering* of the probabilities of the axioms in the knowledge base: axioms with the same weight will have the same probability, and the larger the weight, the larger the probability that will be assigned to the axiom. Hence, there is also a level of *proportionality* in the sense that one can express that an axiom is much more likely than another simply by assigning a much larger weight. One should, however, be very careful when modelling uncertain knowledge through this formalism, as the relationship between weights and probabilities is not linear; in simple terms, duplicating the weight does not necessarily imply duplicating the probability. In fact, this almost never happens. The issue is further complicated by the possibility of assigning negative weights to axioms (which, however, still yield positive probabilities). For more details, see [31,33].

3.2. Statistical Logic

The second class in Halpern's classification is statistical logic, where the uncertainty is distributed over the objects of the domain, but only one "possible world" is considered at a time. That is, rather than the situation of the world, the unknown refers to the properties of specific individuals [34,35].

Syntactically, probabilistic description logics based on the statistical logic semantics often look very similar to those with subjective probabilities: each axiom is associated with a probability degree. The difference becomes apparent only in the semantics. In this case, the semantics assign a probability distribution for each property (or combination thereof) within the elements of the domain. For example, an axiom stating that a property A is a subproperty of the property B with probability p is interpreted as a *conditional probability* expressing that the probability of observing B given that A was observed is p. This in general complicates reasoning as it requires the development of new techniques for dealing with the individuals and transferring the properties among them [36].

The additional difficulties on dealing with statistical probabilistic logics become obvious from exploring the literature. Indeed, while probabilistic description logics based on subjective probabilities abound, and their properties have been deeply studied, the variants based on statistical probabilities are extremely limited. Moreover, their reasoning complexity tends to grow as well [37]. Hence, despite being very useful in many situations—indeed, providing the adequate form of uncertainty in many practical scenarios—these logics are largely unexplored.

In knowledge representation one is often interested in the practical and computational properties of the languages as a way to guarantee adequate answers within reasonable time bounds, and minimise the effort of implementing, optimising, and updating the systems. Still, the choice of the semantics is fundamental to obtain the right answers. One way to summarise the difference between subjective and statistical probabilities is that in the former, an axiom holds in all the individuals or not, depending on the world under consideration, while in the latter the properties of the axiom hold in some individuals and not in others. Consider for example the statement "a person is female with probability 0.5." Under subjective probabilities, this statement is interpreted as the knowledge that in half of all possible worlds, *every person is female*. Under statistical probabilities, instead, it is interpreted as stating that half of all persons are female. Although the difference may look very subtle at first sight, a simple reasoning question may highlight the deep differences between both: if one takes two random individuals and ask what is the probability that one is female and the other is male (assuming that there is additional knowledge about gender in the knowledge base), a statistical probabilistic approach would yield the (intuitive) answer that this probability is 0.5; a subjective probabilistic approach would instead set this probability to 0, because in its semantics either all individuals are male, or all are female (and thus, there cannot be one of one gender and one of another). This simple example can be used as a general test to understand which kind of semantics is adequate in a given application.

3.3. Other Approaches

As mentioned before, Halpern's classification does not cover the whole spectrum of uncertainty which can arise in practical applications. Indeed, one of the best known and most commonly used cases for uncertainty is not covered by this classification. In many situations with unknown relationships between properties, we gather partial information about the world through different statistical models. It is important not to confuse statistical models with statistical probabilities—the latter is the unfortunate name given to the case considered earlier in this section. One of the most common statistical models is the use of *sampling* to approximate the incidence of a property. In a nutshell, one takes a small part of the population (a sample) and queries for the property of interest. Under some reasonable assumptions (about the quality and covering of the sampling method) the proportion of the sample that satisfies the property is approximately the proportion of the full population with the same property. The quality of the approximation improves as the sample size grows, but obtaining a larger sample may be expensive and in many cases (like in the medical domain) even impossible.

Importantly, as a part of the population has not been observed, the actual incidence of the property studied is itself uncertain: it can still move in any direction although with decreasing probability as it moves farther away from the computed estimate. Dealing with the uncertainty of these approximations alongside logical properties is not an easy task. Some approaches have tried to handle it through uncertain [38,39] or imprecise probabilities [40,41]. The issue is that these approaches (despite their names) still require a precise knowledge of the probabilistic bounds, in opposition to the knowledge provided by the statistical analysis, which can provide different bounds with different degrees of certainty, even propagate the uncertainty through different reasoning steps. A preliminary approach trying to handle this information formally was presented in [42], where so-called confidence intervals appear as first-class citizens to be manipulated. However, as it is clear from that preliminary study, there remain many open gaps before these ideas can be developed into a fully-fletched uncertain knowledge representation language.

A final approach that is worth mentioning is based on the principle of *maximum entropy*. In this approach, the probabilities of the axioms define a unique probability distribution which is considered the least informative, thus preserving the idea of open-world assumption, but simplifying the reasoning process once that this so-called maximum entropy distribution has been computed. For details on how this principle is applied in description logics, see [43,44].

4. Conclusions

We have provided a very brief roadmap to the representation, managing, and handling of uncertainty in knowledge representation languages. As warned in the introduction, the roadmap is extremely limited in its scope: it considers only probabilities as uncertainty representation, and focuses on formalisms extending the well-known family of description logics. The choice of these limits necessarily leaves out a huge part of the literature on uncertain knowledge representation, and still it was a quintessential choice. On the one hand, covering the whole area would require a much larger space (see e.g., an outdated survey at [45]). On the other hand, the main features of this class of languages are already represented in the formalisms covered.

While it is true that changing the base formalism requires an additional analysis of the technical details and may deeply affect the computational properties of the resulting language, it is also the case that the main issues that should be considered, specially when trying to decide which language is best suited for a given application, are already covered in this roadmap.

It is our hope that this brief paper will be useful to newcomers trying to identify gaps in the field to work in, and to knowledge engineers trying to assess the right formalism for their needs when modelling uncertainty.

Conflicts of Interest: The author declares no conflict of interest.

References

1. van Harmelen, F.; van Harmelen, F.; Lifschitz, V.; Porter, B. *Handbook of Knowledge Representation*; Elsevier Science: San Diego, CA, USA, 2007.
2. Robinson, J.A.; Voronkov, A., Eds. *Handbook of Automated Reasoning (in 2 volumes)*; Elsevier and MIT Press, 2001.
3. Zio, E.; Pedroni, N. Literature review of methods for representing uncertainty **2013**.
4. Levitt, T.S. Choosing uncertainty representations in artificial intelligence. *International Journal of Approximate Reasoning* **1988**, *2*, 217–232. doi:https://doi.org/10.1016/0888-613X(88)90118-1.
5. Dubois, D.; Prade, H. *Possibility Theory - An Approach to Computerized Processing of Uncertainty*; Springer, 1988. doi:10.1007/978-1-4684-5287-7.
6. Dubois, D. Possibility theory and statistical reasoning. *Computational Statistics & Data Analysis* **2006**, *51*, 47–69. The Fuzzy Approach to Statistical Analysis, doi:https://doi.org/10.1016/j.csda.2006.04.015.
7. Dempster, A.P. Upper and Lower Probabilities Induced by a Multivalued Mapping. *The Annals of Mathematical Statistics* **1967**, *38*, 325 – 339. doi:10.1214/aoms/1177698950.
8. Shafer, G. *A Mathematical Theory of Evidence*; Princeton University Press, 1976.
9. Borgwardt, S. Fuzzy description logics with general concept inclusions. PhD thesis, Technische Universität Dresden, Germany, 2014.
10. Borgwardt, S.; Distel, F.; Peñaloza, R. The limits of decidability in fuzzy description logics with general concept inclusions. *Artificial Intelligence* **2015**, *218*, 23–55. doi:10.1016/j.artint.2014.09.001.
11. Hájek, P. *Metamathematics of Fuzzy Logic*; Vol. 4, *Trends in Logic*, Kluwer, 1998. doi:10.1007/978-94-011-5300-3.
12. Bertossi, L.E.; Hunter, A.; Schaub, T., Eds. *Inconsistency Tolerance*, Vol. 3300, *LNCS*. Springer, 2005. doi:10.1007/b104925.
13. Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; Savo, D.F. Inconsistency-Tolerant Semantics for Description Logics. Proc. 4th Int. Conf. on Web Reasoning and Rule Systems (RR 2010); Hitzler, P.; Lukasiewicz, T., Eds. Springer, 2010, Vol. 6333, *LNCS*, pp. 103–117. doi:10.1007/978-3-642-15918-3_9.
14. Bienvenu, M.; Bourgaux, C. Inconsistency-Tolerant Querying of Description Logic Knowledge Bases. 12th International Reasoning Web Summer School, Tutorial Lectures. Springer, 2016, Vol. 9885, *LNCS*, pp. 156–202. doi:10.1007/978-3-319-49493-7_5.

15. Halpern, J.Y. An Analysis of First-Order Logics of Probability. *Artificial Intelligence* **1990**, *46*, 311–350. doi:10.1016/0004-3702(90)90019-V.
16. Wilks, D.S. *Statistical methods in the atmospheric sciences*; Elsevier Academic Press, 2011.
17. Baader, F.; Calvanese, D.; McGuinness, D.; Nardi, D.; Patel-Schneider, P., Eds. *The Description Logic Handbook: Theory, Implementation, and Applications*, second ed.; Cambridge University Press, 2007.
18. Baader, F.; Brandt, S.; Lutz, C. Pushing the $\mathcal{EL}$ Envelope. Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI'05); Kaelbling, L.; Saffiotti, A., Eds. Professional Book Center, 2005, pp. 364–369.
19. Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; Rosati, R. Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family. *Journal of Automated Reasoning* **2007**, *39*, 385–429.
20. Horrocks, I.; Kutz, O.; Sattler, U. The Even More Irresistible $\mathcal{SROIQ}$. Proc. 10th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'06); Doherty, P.; Mylopoulos, J.; Welty, C.A., Eds. AAAI Press, 2006, pp. 57–67.
21. Brewka, G. *Nonmonotonic reasoning - logical foundations of commonsense*; Vol. 12, *Cambridge tracts in theoretical computer science*, Cambridge University Press, 1991.
22. Chomicki, J. Temporal Query Languages: A Survey. Proc. 1st Int. Conf. on Database Theory (ICDT'94); Gabbay, D.M.; Ohlbach, H.J., Eds. Springer, 1994, Vol. 827, *LNCS*, pp. 506–534.
23. Artale, A.; Franconi, E. A survey of temporal extensions of description logics. *Ann. Math. Artif. Intell.* **2000**, *30*, 171–210.
24. Maggi, F.M.; Montali, M.; Peñaloza, R. Temporal Logics Over Finite Traces with Uncertainty. Proceedings of The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI'20). AAAI Press, 2020, pp. 10218–10225.
25. Lutz, C.; Schröder, L. Probabilistic Description Logics for Subjective Uncertainty. Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning (KR 2010); Lin, F.; Sattler, U.; Truszczynski, M., Eds. AAAI Press, 2010.
26. Gutiérrez-Basulto, V.; Jung, J.C.; Lutz, C.; Schröder, L. Probabilistic Description Logics for Subjective Uncertainty. *Journal of Artificial Intelligence Research* **2017**, *58*, 1–66. doi:10.1613/jair.5222.
27. Ceylan, İ.İ.; Peñaloza, R. The Bayesian Ontology Language BEL. *Journal of Automated Reasoning* **2017**, *58*, 67–95. doi:10.1007/s10817-016-9386-0.
28. d'Amato, C.; Fanizzi, N.; Lukasiewicz, T. Tractable Reasoning with Bayesian Description Logics. Proceedings of the Second International Conference on Scalable Uncertainty Management (SUM 2008). Springer, 2008, Vol. 5291, *LNCS*, pp. 146–159. doi:10.1007/978-3-540-87993-0_13.
29. Botha, L.; Meyer, T.; Peñaloza, R. A Bayesian Extension of the Description Logic ALC. Proceedings of the 16th European Conference on Logics in Artificial Intelligence (JELIA'19); Calimeri, F.; Leone, N.; Manna, M., Eds. Springer, 2019, Vol. 11468, *LNCS*, pp. 339–354. doi:10.1007/978-3-030-19570-0_22.
30. Nilsson, N.J. Probabilistic Logic. *Artificial Intelligence* **1986**, *28*, 71–87. doi:10.1016/0004-3702(86)90031-7.
31. Niepert, M.; Noessner, J.; Stuckenschmidt, H. Log-Linear Description Logics. Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11); Walsh, T., Ed. IJCAI/AAAI, 2011, pp. 2153–2158. doi:10.5591/978-1-57735-516-8/IJCAI11-359.
32. Borgwardt, S.; Ceylan, İ.İ.; Lukasiewicz, T. Ontology-Mediated Query Answering over Log-Linear Probabilistic Data. Proceedings of The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI'19). AAAI Press, 2019, pp. 2711–2718. doi:10.1609/aaai.v33i01.33012711.
33. Koller, D.; Friedman, N. *Probabilistic Graphical Models - Principles and Techniques*; MIT Press, 2009.
34. Peñaloza, R.; Potyka, N. Towards Statistical Reasoning in Description Logics over Finite Domains. Proceedings of the 11th International Conference on Scalable Uncertainty Management SUM 2017; Moral, S.; Pivert, O.; Sánchez, D.; Marín, N., Eds. Springer, 2017, Vol. 10564, *Lecture Notes in Computer Science*, pp. 280–294. doi:10.1007/978-3-319-67582-4_20.
35. Klinov, P.; Parsia, B. Probabilistic Modeling and OWL: A User Oriented Introduction to P-SHIQ(D). Proceedings of the Fifth OWLED Workshop on OWL: Experiences and Directions, collocated with the 7th International Semantic Web Conference (ISWC-2008); Dolbear, C.; Ruttenberg, A.; Sattler, U., Eds. CEUR-WS.org, 2008, Vol. 432, *CEUR Workshop Proceedings*.
36. Klinov, P.; Parsia, B. Pronto: A Practical Probabilistic Description Logic Reasoner. Uncertainty Reasoning for the Semantic Web II, International Workshops URSW 2008-2010, Revised Selected Papers; Bobillo, F.; da Costa, P.C.G.; d'Amato, C.; Fanizzi, N.; Laskey, K.B.; Laskey, K.J.; Lukasiewicz, T.; Nickles, M.; Pool, M., Eds. Springer, 2013, Vol. 7123, *LNCS*, pp. 59–79. doi:10.1007/978-3-642-35975-0_4.
37. Baader, F.; Ecke, A. Extending the Description Logic ALC with More Expressive Cardinality Constraints on Concepts. GCAI 2017, 3rd Global Conference on Artificial Intelligence, Miami, FL, USA, 18-22 October 2017; Benzmüller, C.; Lisetti, C.L.; Theobald, M., Eds. EasyChair, 2017, Vol. 50, *EPiC Series in Computing*, pp. 6–19. doi:10.29007/f3hh.
38. Buckley, J.; Eslami, E. Uncertain probabilities I: the discrete case. *Soft Computing* **2003**, *7*, 500–505. doi:10.1007/s00500-002-0234-2.
39. Jøsang, A. A Logic For Uncertain Probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **2001**, *09*, 279–311. doi:10.1142/S0218488501000831.
40. Cano, A.; Cozman, F.G.; Lukasiewicz, T. Reasoning with imprecise probabilities. *International Journal of Approximate Reasoning* **2007**, *44*, 197–199. doi:10.1016/j.ijar.2006.09.001.
41. Cozman, F.G. Graphical models for imprecise probabilities. *International Journal of Approximate Reasoning* **2005**, *39*, 167–184. doi:10.1016/j.ijar.2004.10.003.

42. Peñaloza, R. Towards a Logic of Meta-Analysis. Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020; Calvanese, D.; Erdem, E.; Thielscher, M., Eds., 2020, pp. 672–676. doi:10.24963/kr.2020/68.
43. Baader, F.; Ecke, A.; Kern-Isberner, G.; Wilhelm, M. The Complexity of the Consistency Problem in the Probabilistic Description Logic *ALC* ^\mathsf ME. Frontiers of Combining Systems - 12th International Symposium, FroCoS 2019, London, UK, September 4-6, 2019, Proceedings; Herzig, A.; Popescu, A., Eds. Springer, 2019, Vol. 11715, *Lecture Notes in Computer Science*, pp. 167–184. doi:10.1007/978-3-030-29007-8_10.
44. Wilhelm, M.; Kern-Isberner, G.; Ecke, A.; Baader, F. Counting Strategies for the Probabilistic Description Logic *ALC*^\mathsf ME Under the Principle of Maximum Entropy. Logics in Artificial Intelligence - 16th European Conference, JELIA 2019, Rende, Italy, May 7-11, 2019, Proceedings; Calimeri, F.; Leone, N.; Manna, M., Eds. Springer, 2019, Vol. 11468, *Lecture Notes in Computer Science*, pp. 434–449. doi:10.1007/978-3-030-19570-0_28.
45. Lukasiewicz, T.; Straccia, U. Managing uncertainty and vagueness in description logics for the Semantic Web. *Journal of Web Semantics* **2008**, *6*, 291–308.