

FROM IGNORANCE TO UNCERTAINTY: A CONCEPTUAL ANALYSIS

PIETRO BARONI, GIOVANNI GUIDA AND SILVANO MUSSI

This paper aims to develop an analysis of how ignorance affects the reasoning activity and is related to the concept of uncertainty. With reference to a simple inferential reasoning step, involving a single piece of relational knowledge, we identify four types of ignorance and show how they give rise to different types of uncertainty. We then introduce the concept of reasoning attitude, as a basic choice about how reasoning should be carried out in presence of ignorance. We identify two general attitudes, analyze how they are related to different types of ignorance, and propose some general requirements about how they should affect the reasoning activity. A formalism for uncertain reasoning explicitly including the different types of uncertainty identified and satisfying the stated requirements is finally introduced and its performance is analyzed in simple examples.

1. INTRODUCTION

In artificial intelligence literature most attention has always been focused on knowledge, on the analysis of its nature, on its characterization, and on its role in intelligent reasoning process. More specifically, in the uncertain reasoning field, knowledge has been advocated as the primitive concept from which the concept of belief should be derived, as clearly and explicitly stated in [7]. According to this standpoint, uncertain reasoning is nothing else but a special case of reasoning with “certain” knowledge. Certain (logical) reasoning is assumed, in a sense, as the fundamental and perfect form of reasoning, in relation to which uncertain reasoning is perceived as an imperfect exception. The goal of this paper is to propose a different point of view about reasoning under uncertainty, with a particular attention to the role of ignorance. In particular, we emphasize the role of ignorance in determining a state of imperfect knowledge. Ignorance can be generally characterized as lack of knowledge: as we will show, different types of knowledge may be lacking and different lacks of knowledge affect the reasoning activity differently. Our fundamental claims are:

- reasoning is based on the exploitation of the knowledge available on the domain of interest and is limited by the ignorance existing about the same domain;
- the existence of ignorance makes the use of available knowledge not definitely reliable, giving rise this way to the concept of uncertainty: uncertainty is

therefore a way to take into account the ignorance that affects the knowledge we use in reasoning processes;

- since ignorance inextricably pervades any possible domain of interest (we can never be sure that our knowledge about a given subject is definitely complete), human reasoning is in its essence uncertain;
- certain reasoning is only a conventional and synthetic representation – indeed very useful in practice – to deal with those cases where uncertainty is expected to have a negligible role.

The paper is conceptually organized as follows. First, we introduce a simple schematization of the basic structure of a generic inference step, considered as the fundamental constituent of a reasoning process. We propose then a basic characterization of how ignorance may affect an inference step, showing that there exist various types of ignorance. Next we propose an analysis of how different types of ignorance affect the activity of reasoning under uncertainty, introducing the concept of reasoning attitude. The proposed ideas are framed in a (preliminary) representation and reasoning formalism, whose behavior is analyzed in some simple examples.

2. A SIMPLE INFERENCE STEP

In a very simple and basic formulation, a generic inference step, considered as the fundamental constituent of a reasoning process, involves:

- an individual, that is the specific subject of interest to which the inference step is applied;
- a piece of relational knowledge which states the relation between two properties relevant to the individual considered: the former property is called premise and the latter consequence;
- an inference mechanism which allows to infer that, if an individual has the property stated in the premise of the relation, then the individual has also the property stated in the consequence.

Using relational knowledge to reason about an individual, requires therefore the following three steps:

- verifying, exploiting the information available about an individual, whether he has the property stated in the a premise of the chunk of relational knowledge considered;
- if the individual has such a property, assuming that the chunk of relational knowledge can be applied to the individual;
- deriving then the consequence of the application of the chunk of relational knowledge to the individual and ascribing to him the property stated in the consequence.

This simple scheme holds when an ideal state of perfect knowledge is assumed (we assume that knowledge is perfect when it has all intuitively desirable properties, i. e. it is complete, consistent, error free, etc). Let us briefly examine the meaning of

this assumption. According to the schematization proposed above, an inference step involves four basic elements (namely: the individual, the premise, the consequence, and the relation between the premise and the consequence) and two activities involving them (namely: verifying the match between the individual and the premise and then adding the consequence to the set of known facts). Therefore the assumption of an ideal state of perfect knowledge implies:

- perfect knowledge about the individual, i.e. about all its properties relevant to the matching with the premise;
- perfect knowledge about the premise, i.e. about all properties which make the relation applicable to an individual;
- perfect knowledge about the relation between the premise and the consequence;
- perfect knowledge about the consequence, i.e. about all properties which are entailed by the premise.

In [7] the case of imperfect knowledge is identified with the concept of belief: a belief about a proposition φ is characterized as “the agent knows that either φ is the case or else some specific (perhaps unusual) circumstances obtain”. Therefore the concept of imperfect knowledge is somehow reduced to that of exception: knowledge degrades itself to belief when it admits exceptions. The concept of exception is also the background on which nonmonotonic reasoning theories have been built. Default rules are rules with exceptions, as stated, for instance, in [10]: “The sentence ‘Birds fly’ is not synonymous with ‘All birds fly’ because there are exceptions”.

However, the presence of exceptions is not, in general, the only possible cause of imperfection in knowledge. In particular, we emphasize the role of ignorance in determining a state of imperfect knowledge. Ignorance can be generally characterized as lack of knowledge: as we will show, different types of knowledge may be lacking and different lacks of knowledge affect the reasoning activity differently.

3. TYPES OF IGNORANCE

Let us start our analysis by supposing that we are interested in an individual, say Tom, and that our knowledge base includes just the following relations: “winged birds fly”, “liver disease DIS1 requires drug DR1”, “liver disease DIS2 requires drug DR2”.

3.1. Ignorance about the individual

In order to start any reasoning activity, we need first of all some knowledge about the individual. This knowledge – or, more generally, part of it – may however be lacking. If we simply ignore if Tom is a bird or a man, we can not infer anything about Tom. Moreover, our knowledge may also be partially lacking: for instance we might know that Tom is a bird, but ignore if it is winged or not, as well as we may know that Tom suffers from his liver, but we may ignore which is the disease that causes such problems.

3.2. Ignorance about the premise

Considering now the premise, we may partially ignore the properties which underlie the relation we are interested in. For instance, being a bird is not a sufficient premise for deducing that an individual can fly. However, we are unable to enumerate all additional conditions which would be necessary to completely specify this premise (for instance, we should mention in the premise all the cases of exceptional birds: penguins, ostriches, caged birds, etc). Similarly, we may ignore some additional conditions that make a drug ineffective in restoring health (in fact, in some cases, a drug may fail to give the desired effects).

3.3. Ignorance about the relation

As far as the relation between the premise and the consequence is concerned, let us first state that such a relation represents, in very general terms, an understanding of an aspect of a specific domain (the zoological domain for “birds fly”, or the medical domain for the other relations considered). Such understanding is based on other chunks of knowledge concerning the same domain. For instance, the fact that drug DR1 is useful for disease DIS1 is based on the fact that chemical components of DR1 contrast the negative effects DIS1 has on liver cells. However, in many cases such detailed knowledge is lacking. For instance, a statistical correlation may be detected between a (supposed) cause and an effect, but there is no clear physical understanding of the causation mechanism. On the other hand, one may be able to build a very sound physical model which explains why and how a certain cause should produce a given effect, but experimental evidence supporting the model may be lacking, either because experimental data are contradictory or simply because new, ad hoc, experiments should be carried out to verify the theory and such an experimentation campaign requires a lot of time and money. Of course other more and more complicated cases may happen; for example, given a new drug both experimental results and a suitable physical model may be available to support its use for a certain disease, but both experimental results and the physical model may be questioned by a group of scholars who can provide a different interpretation of the same data, in the frame of an alternative model. Anyway, such lack of basic domain knowledge should strongly affect our feeling about the considered relation. First of all, one might wonder whether such a piece relational knowledge, lacking sufficient support or understanding, should definitely remain in the knowledge base and be used for reasoning purposes. Moreover, one might feel that such a piece of knowledge should not be treated the same way as a well-grounded relation such as “birds fly”.

3.4. Ignorance about the consequence

Finally, also ignorance about the consequence should be considered. If one verifies that a given premise holds, he is generally interested in all relevant facts that can be derived from the premise. Of course, this interest is strongly context dependent; in the case of birds, one can simply neglect, without problems, the fact that being a bird also implies being feathered or making eggs. But in the case of drugs, one

is very interested in all side-effects they may have, that, especially for very recent drugs, are often ignored. In such cases, there is a very strong need of specifying all the consequences of a given premise, and, of course, it may happen that some of them are ignored.

4. FROM IGNORANCE TO UNCERTAINTY

After having provided the above classification of different types of ignorance, let us examine how different kinds of ignorance differently affect the reasoning process and give rise to different kinds of uncertainty. In general, we mean that a reasoner is in a condition of uncertainty when she/he is unable to answer, in a definite and not defeasible way, a given question she/he is interested in. In other words, a reasoner is in a condition of uncertainty when she/he has some reasons to believe that the assertions she/he makes will need to be revised in the future. Of course in absence of ignorance, i.e. in the ideal case of complete and perfect knowledge, every question has its right answer and all the conclusions derived from the reasoning process are certain, in the sense that they are grounded on a stable basis and there is no reason to doubt that they may be questioned in the future. In fact, since there is no ignorance, no new knowledge elements can be acquired, that could induce to retreat, modify, or extend previously derived conclusions. On the contrary, if we admit the presence of ignorance, we must also admit that our knowledge can be improved, if our ignorance can be reduced. Therefore, we are also constrained to accept that our conclusions can be retreated, modified, or extended. In a sense, our conclusions are never definitive, and, for the time being, they should be considered only as certain to a given degree, i.e. as intrinsically uncertain. Therefore, ignorance originates uncertainty. When carrying out an inference step there are different questions to be answered and different types of knowledge to be used. Therefore different kinds of ignorance should be considered, that differently affect the reasoning process and give rise to different kinds of uncertainty, as we discuss in the next sections.

4.1. Ignorance about the individual

Let us consider first the ignorance about the properties of the individual, called *I*-ignorance for short. Such properties are considered in the matching phase, in order to verify whether the relation applies to the individual. Let us assume, for the sake of simplicity, that the choice about the application of a piece of relational knowledge is two-way, i.e. either the relation is completely applicable or it is not applicable at all. Therefore, *I*-ignorance implies that the choice to apply or not to apply the relation is considered as retractable. For instance, if one just knows that Tom is an animal, he is not reasonably allowed to apply to Tom the relation “birds fly”, but if he later learns that it is a bird, the relation turns out to be applicable. However, if subsequently it turns out that Tom is a penguin, the conclusions derived from the incorrect application of the relation should be retreated.

4.2. Ignorance about the premise

Turning to ignorance about the premise, called *P*-ignorance for short, if not all the conditions which should correctly be included in the premise are known (or specified), it happens that, even if an individual matches with the properties stated in the premise, it is not guaranteed that the relation can be safely applied. For instance, if one ignores that drug DR1 is useful for disease DIS1 only in absence of disease DIS2 (perhaps because DR1 has never been experimented on patients suffering from both diseases), DR1 can be prescribed to a patient suffering from DIS1 even if it is known that he suffers from DIS2 too. Therapy results or the acquisition of more detailed knowledge will however show that the prescription is incorrect, because the relation should not be applied to such an individual.

Note that both *I*-ignorance and *P*-ignorance affect the matching between an individual and the premise of a relation, i.e. involve the applicability of the relation to the individual. In fact they affect the answer to the question: “Should the relation be applied to the individual?” and in both cases the acquisition of further knowledge may show that the relation was not correctly applied to the individual. Therefore, considering their effect on the reasoning activity, both *I*- and *P*-ignorance make uncertain the fact that a given relation should be applied to a given individual, i.e. they produce a unique type of uncertainty that concerns the applicability of a relation, called *A*-uncertainty for short.

4.3. Ignorance about the relation

Let us consider now ignorance concerning the relation between the premise and the consequence. This type of ignorance does not involve any more the individual and solely affects the relations itself. Consider, for example, the case of a drug to which a positive effect on a given disease was ascribed. If it is learned that some recoveries, initially ascribed to its chemical properties, were, in fact, due to a placebo effect, this new knowledge leads to suppress the relation between the drug and the disease from the knowledge base, as not valid. As it is clear, this suppression makes invalid all the inferences previously drawn based on the relation, independently from the individuals to which they were applied. In this case not only the application of the relation to some special individuals is questioned, but the general validity of the relation itself is challenged. Therefore we call this kind of ignorance *I*-ignorance, i.e. ignorance affecting the validity of a relation. *V*-ignorance affects the question “Is my knowledge reliable?” and produces *V*-uncertainty, i.e. it makes uncertain the fact that a given relation should be considered valid and included in the knowledge base.

4.4. Ignorance about the consequence

Finally, ignorance may concern also the consequence of a relation, *C*-ignorance for short. Let us note that this kind of ignorance concerns the fact that one may fail to identify all the relevant facts entailed by the premise. Thus, if, for instance, one learns about previously unknown side effects of a given drug, he should extend

previously drawn deductions in order to take into account the new consequences. *C*-ignorance affects the question “Am I overlooking any important consequence?” and produces *C*-uncertainty, i.e. it makes uncertain that all the relevant deductions have been drawn in the reasoning process. Therefore, differently from *A*- and *V*-uncertainty, it does not involve the potential retraction of a previous choice but just the extension of the deductions made. In other words, it directly affects the completeness of the reasoning results rather than their correctness. However, completeness may indirectly affect correctness (so requiring retractions) when the new drawn conclusions interact with the premise of a relation affected by *A*-uncertainty, i.e. when the new conclusions make inapplicable a previously applied relation.

In the above discussion we have outlined how different types of ignorance give rise to different types of uncertainty. This distinction is, in our opinion, of fundamental importance for modeling uncertain reasoning, since different types of uncertainty have different properties and, most importantly, affect the reasoning activity differently, as we will discuss in next section.

5. IGNORANCE AND REASONING UNDER UNCERTAINTY

5.1. Background

Let us now focus on how the various types of uncertainty defined above can affect the reasoning activity. Before proceeding, it is necessary to better define what we mean by reasoning activity, in fact, in presence of uncertainty, at least two basic interpretations are possible:

- the reasoning activity has the goal of ascribing truth values to propositions, however, due to the presence of uncertainty, such truth values are defeasible;
- the reasoning activity has the goal of ascribing an uncertainty quantification to pairs $\langle \text{proposition}, \text{truth value} \rangle$.

The former interpretation is the standpoint adopted by symbolic approaches, while the latter is the basic assumption of quantitative approaches.

5.2. Attitudes in reasoning in symbolic approaches

In a symbolic approach when evaluating an inference step in presence of ignorance two basic choices are possible:

- to suspend reasoning, i.e. to renounce to draw any conclusion, until new knowledge is acquired;
- to carry out inference anyway, admitting however, that it can be subsequently refuted.

In this context, we say that the former choice corresponds to a conservative attitude while the latter to an evolutive attitude. The choice between the two attitudes seems very natural in some situations. For instance, in case we just know that Tom is an animal, we adopt a conservative attitude and renounce to assume that it could be a

bird and therefore that it could fly. On the other hand, if we know that Tom is a bird, we adopt an evolutive attitude and we are easily inclined to assume that it is not an abnormal bird and, therefore, that it can fly. The choice between a conservative or an evolutive attitude might be related to the type of ignorance; however, the limited expressive capabilities of symbolic approaches do not allow a satisfactory explicit representation of these aspects. In fact, in symbolic approaches, any proposition is regarded (defeasibly) just as true or false and any relation is regarded just as valid or not (i. e., it is included in the currently used knowledge base or not). The simple criteria for choosing between conservative or evolutive attitude adopted in symbolic approaches can therefore be stated as follows:

- In case of *I*-ignorance, a conservative attitude is adopted; if the available knowledge about an individual does not allow to match it with the premise, the relation is not applied. This is intuitively justified by the fact that adopting an evolutive attitude in this case would lead to apply almost any relation to almost any individual, which is, in general, unacceptable. In fact, such choice could lead to contradictory – or at least counterintuitive – conclusions. Moreover, if the reasoning activity is costly, this choice may turn out to be fairly uneconomical.
- In case of *P*-ignorance, an evolutive attitude is adopted; even if the premise is not completely specified, if it matches with the properties of the individual and nothing explicitly prevents the deduction, the relation is applied. This is the case of typical default rules [10], where in presence of admittedly incomplete premises, defeasible conclusions are drawn anyway. This is intuitively justified by the fact that adopting a conservative attitude in this case would lead to block inference in almost any practical reasoning situation. In fact, relations for which it is actually possible to define a complete premise (i. e., for which it is possible, for instance, to enumerate all possible exceptions) are very rare in practice.
- The case of *V*-ignorance can not be dealt explicitly with within the frame of symbolic approaches. Practically, if a doubted validity relation is included in the knowledge base, an evolutive attitude is adopted; whereas if the relation is excluded from the knowledge base, this corresponds to adopt a conservative attitude. Therefore, in this case, the choice about which attitude to associate to a relation is actually committed to the person in charge of building the knowledge base, and is not explicitly dealt with at reasoning level.
- Finally, in the case of *C*-ignorance, a conservative attitude is adopted; only explicitly stated consequences are considered. Similarly to the case of *I*-ignorance, this is justified by the fact that adopting an evolutive attitude would lead to derive any consequence from a relation.

As a conclusion, from the analysis of how the four types of ignorance affect the choice of the most appropriate attitude to take in a given inference step, it can be hypothesized that a basic principle of reasoning economy underlies the choice of the attitude. In all cases, the attitude is selected that prevents either an exaggerated extension or a drastic reduction of the reasoning activity. In the case of

ignorance about validity, the choice is left outside reasoning process and delegated to an external authority, the knowledge base builder.

5.3. Attitudes in reasoning in quantitative approaches

Let us extend now our analysis to quantitative approaches. First of all let us stress that this analysis is carried out at an abstract level, without reference to any specific existing formalism. Let us consider a generic quantitative formalism, in which uncertainty quantification (say q) ranges, as it is indeed very usual, over the real interval $[0, 1]$, where 1 represents intuitively the maximum certainty and 0 the minimum (null) certainty. For the sake of generality, given a proposition P that may assume the truth values $\{true, false\}$, we assume that uncertainty quantification about P is represented by a pair of distinct (though not necessarily independent) values $[q(P, true), q(P, false)]$.

Similarly, for a relation R the complete characterization of the uncertainty about R requires two distinct quantifications:

- a quantification associated to the fact that R is applicable or not applicable to an individual I , represented by a pair $[q(applicable(R, I), true), q(applicable(R, I), false)]$;
- a quantification associated to the fact that R is valid or not, represented by a pair $[q(valid(R), true), q(valid(R), false)]$.

It has to be noted that, while validity quantification depends only on the relation R , applicability quantification depends both on R and on the individual I to which it is applied. Moreover we do not consider here quantification of C -uncertainty. Such quantification would concern the fact that the consequence of the relation is completely specified or not. In other words, it does not concern the actually drawn conclusions but an estimation about the existence of other relevant conclusions that have not been drawn, due to C -ignorance. Therefore, a quantification of C -uncertainty represents a judgement about the completeness of drawn conclusions rather than on their credibility. Even if this aspect is very interesting and may also be very important in many applications, we do not further investigate it in this paper, leaving it for future work.

The goal of a reasoning step in the frame of quantitative approaches can now be stated as follows: “Given a relation R and a fact F about an individual I , such that F matches with the premise of R , derive from the uncertainty quantifications $[q(F, true), q(F, false)]$, $[q(applicable(R, I), true), q(applicable(R, I), false)]$, and $[q(valid(R), true), q(valid(R), false)]$ the proper uncertainty quantification $[q(G, true), q(G, false)]$ about a fact G , corresponding to the consequence of the relation.” Therefore, in order to characterize the reasoning activity, the role played by six different components in determining $[q(G, true), q(G, false)]$ has to be defined. Let us examine them individually:

- The component $q(F, true)$ represents intuitively the belief degree that F holds, i.e. that I has the property stated in the premise, therefore the higher $q(F, true)$ the higher should be $q(G, true)$.

- The component $q(F, false)$ represents intuitively the belief degree that F does not hold, i.e. that I has not the property stated in the premise and, therefore, that the rule does not apply to I . A conservative attitude is appropriate in this case: if the rule does not apply to the individual because it has not the required properties, nothing should be inferred about the consequence. Therefore the higher $q(F, false)$ the lower should be $q(G, true)$ and $q(G, false)$.
- The component $q(applicable(R, I), true)$ represents intuitively the belief degree that R is applicable to a generic individual, in other words it represents the certainty that the premise is completely specified and that there are no exceptions to the relation. Of course, if the premise is completely specified and there are no exceptions, given the premise it is sure that the consequence holds. Therefore, the higher $q(applicable(R, I), true)$ the higher should be $q(G, true)$.
- The component $q(applicable(R, I), false)$ represents intuitively the belief degree that R is not applicable to a generic individual, i.e. the certainty that the premise is not completely specified and that the relation admits (several) exceptions. The role played by this component depends on our attitude towards exceptions: in a conservative attitude nothing is assumed about an exception, whereas in an evolutive attitude it is assumed that if an individual is an exception, the negation of the consequence holds. Both attitudes make sense, the choice depending mainly on the nature of the relation at hand and on context dependent conditions. Therefore, in a conservative attitude the higher $q(applicable(R, I), false)$, the lower should be both $q(G, true)$ and $q(G, false)$, whereas in an evolutive attitude the higher $q(applicable(R, I), false)$, the higher should be $q(G, false)$.
- The component $q(valid(R), true)$ represents intuitively the belief degree that the relation is valid, i.e. that it is founded on a solid understanding of the domain at hand. Therefore, the higher $q(valid(R), true)$ the higher should be $q(G, true)$.
- The component $q(valid(R), false)$ represents intuitively the belief degree that the relation is not valid, i.e. that it could be erroneous and should be canceled from the knowledge base. In this case a conservative attitude is appropriate: in absence of the relation itself, nothing can be inferred. Therefore, the higher $q(valid(R), false)$, the lower should be both $q(G, true)$ and $q(G, false)$.

Given these general requirements, we propose in the next section a preliminary proposal of a quantitative formalism which is appropriate to model uncertain reasoning taking into account the six uncertainty quantifications illustrated above.

6. A PARADIGM FOR REASONING WITH A- AND V-UNCERTAINTY

The proposal presented here has to be considered as very preliminary and aims more to substantiate some basic ideas than to introduce a new general and well-settled formalism for uncertainty management. For the sake of simplicity, we assume that facts about individuals are represented by propositions and that relations are represented in form of IF-THEN production rules.

6.1. Quantified propositions

First of all, let us introduce the concept of belief: a belief is an evidential judgement about the credibility of the truth values ($\{true, false\}$ in the case of ordinary two-valued logic) assigned to a proposition. Beliefs may assume values in an ordered set of belief degrees. For the sake of simplicity, we assume here the real interval $[0, 1]$ as the set of possible belief degrees. It is important to underline that, in our proposal, the concept of belief degree is related to the intuitive concept of “amount of evidence” supporting the credibility that a certain proposition should have a certain truth value. So, given an available body of evidence E , $bel_E(P_1, true) = 0$ means that there is null (or negligible) evidence supporting the credibility that proposition P_1 has the truth value $true$, and this is totally different from excluding that $true$ is a possible truth value for P_1 . Similarly, $bel_E(P_1, true) = 1$ means that available evidence fully supports the credibility that proposition P_1 has the truth value $true$, and this is again totally different from being absolutely certain that $true$ is the correct truth value of P_1 . If we now consider a proposition and compute the belief degrees for all its possible truth values, we obtain a global representation of the uncertainty about which truth value should be assigned to the proposition, on the basis of the available evidence. Therefore, given a proposition P_1 and a body of evidence E , the belief state of P_1 under E , denoted by $bels_E(P_1)$, is the pair $(bel_E(P_1, true), bel_E(P_1, false))$, (say (bt_{P_1}, bf_{P_1}) for short). The belief state represents therefore how much one is authorized to believe in the association between a given proposition and its possible truth values, on the basis of the available evidence. A proposition accompanied by the relevant belief state is called a quantified proposition: more formally, for any proposition P_1 , the pair $(P_1, bels_E(P_1))$ is a quantified proposition. Intuitively, if we are fully convinced, on the basis of available evidence, that a proposition is true, this will be represented by the belief state $(1, 0)$, whereas the opposite conviction will be represented by $(0, 1)$. Moreover we can represent a state of total ignorance about a proposition (due to a lack of evidence) with the belief state $(0, 0)$, which indicates the absence of evidence both supporting the value true and the value false. On the contrary, if we have, for any reason, strong evidences for both the values true and false, we can represent this contradictory situation by the belief state $(1, 1)$. Of course, all intermediate situations are possible, since the two components of a belief state are independent.

6.2. AV-quantified relations

According to the discussion presented in Section 4.3 and to the concepts introduced in Section 5.1, given a relation R represented as a production rule if P_1 then P_2 it is possible to quantify it with a pair of belief states:

- an *A*-belief state, denoted by $A-bels_E(R)$, defined as the pair:
 $(bel_E(applicable(R, I), true), bel_E(applicable(R, I), false))$
 related to the applicability of the rule (also denoted as $(btapp_R, bfapp_R)$ for short);
- a *V*-belief state, denoted by $V-bels_E(R)$, defined as the pair:
 $(bel_E(valid(R), true), bel_E(valid(R), false))$

related to the validity of the rule (also denoted as $(btval_R, bfval_R)$ for short).

The subscript E refers to the overall amount of evidence, possibly collected along years of experience, from which a general evaluation about the applicability and validity of a rule can be derived. The elicitation of these evaluations may be very critical in real knowledge acquisition tasks. We refer here to an ideal case, where all such evaluations are actually available.

The pair $(A-bels_E(R), V-bels_E(R))$ is called the AV -belief state of the relation R and denoted by $AV-bels_E(R)$. A relation accompanied by the relevant AV -belief state is called an AV -quantified relation: more formally, for any relation R , the pair $(R, AV-bels_E(R))$ is an AV -quantified relation.

6.3. A formalism for reasoning with AV -quantified relations

According to what stated in Section 4.3, given a relation R represented through a production rule if P_1 then P_2 , a basic reasoning step consists in deriving the belief state $bels_E(P_2)$ from $bels_E(P_1)$, given the $AV-bels_E(R)$. A simple way of doing this derivation, in accordance with the requirements stated in Section 3.3 and adopting an evolutive attitude as far as applicability is concerned is the following:

$$\begin{aligned} bt_{P_2} &= bt_{P_1} \bullet btapp_R \bullet btval_R \bullet (1 - bf_{P_1}) \bullet (1 - bfval_R) \\ bf_{P_2} &= bt_{P_1} \bullet bfapp_R \bullet btval_R \bullet (1 - bf_{P_1}) \bullet (1 - bfval_R) \end{aligned}$$

Let us note that the validity of these formulas strongly relies on the assumption that the two components of a belief state are completely independent. The intuitive meaning of the proposed formulas can be better appreciated through a simple example. Consider first the case of the relation “smoke causes cancer”: it admits (rare) exceptions but is fully valid, therefore assume its A -belief state is $(0.8, 0.1)$ and its V -belief state is $(1, 0)$. Suppose now it is known with certainty that Tom is a smoker, i. e. $bels(\text{“Tom is a smoker”}) = (1, 0)$. Using the above reported formulas it is then possible to derive the belief state $bels(\text{“Tom catch cancer”}) = (0.8, 0.1)$. Therefore, intuitively, we are strongly convinced that Tom will catch cancer, but we leave also a little space to the opposite hypothesis. Note also that, since P -ignorance is associated to an evolutive attitude, the sum of the components of $bels(\text{“Tom catch cancer”})$ equals that of $bels(\text{“Tom is a smoker”})$. Suppose now that evidence about Tom being a smoker is not so strong: you just have some reasonable suspects that he smokes, and, possibly, you have also some clues supporting the opposite persuasion. Therefore, the belief state about Tom being a smoker is in this case $bels(\text{“Tom is a smoker”}) = (0.7, 0.1)$. Using the above reported formulas it is then possible to derive $bels(\text{“Tom catch cancer”}) = (0.504, 0.063)$. Intuitively, in this case the presence of I -ignorance, associated to a conservative attitude, causes a reduction of both components of $bels(\text{“Tom catch cancer”})$ so that their sum is not preserved through the reasoning step. This reflects the fact that, in a conservative attitude, ignorance affects the amount of belief transferred to the consequence, whereas in an

evolutive attitude it just affects the belief distribution between the two components. One might wonder why, in this case, we have a lower belief in the fact that Tom does not catch cancer. This is coherent with our concept of belief as amount of evidence and with the conservative attitude: in fact, we have less reasons to believe anything that can be derived from the application of this relation. Consider now the (quite unlikely) case that the fact that smoke is really related with cancer is questioned by some authoritative scholar. We have therefore a new V -belief state $(1, 0.3)$, expressing the significant contradiction between the two opinions. Considering again the case where bel_s (“Tom is a smoker”) = $(1, 0)$, using the above reported formulas it can be derived that bel (“Tom catch cancer”) = $(0.56, 0.07)$. The effects of V -ignorance with conservative attitude are analogous to those of I -ignorance.

7. RELATED WORK AND DISCUSSION

In this paper we have developed an initial analysis of the different types of ignorance that may concern our knowledge and of the relations between ignorance and the uncertainty affecting inferential reasoning. Within this framework, we have then outlined a preliminary proposal of a quantitative formalism to model uncertain reasoning in presence of A - and V -uncertainty.

The main goal of the analysis developed in this paper has been to establish an ample framework, where the concept of uncertainty can be better placed and understood. In fact, from a practical perspective, it is very important that when facing a knowledge acquisition and modeling task in a specific domain, the knowledge engineer is aware of the different types of ignorance and of uncertainty he might be requested to deal with: he should be able to correctly identify and classify them, in order to avoid gross representation errors and improper knowledge use. In many challenging application domains the capability of distinguishing among different types of ignorance and of uncertainty is crucial, and a deep understanding of the different roles they play is necessary to correctly and effectively reproduce expert knowledge and reasoning.

Even if the considerations developed in this paper are not far from common intuition, they seem to defy some consolidated assumptions lying behind most of the approaches that can be found in current literature about uncertain reasoning. As we mentioned above, uncertainty about a relation is commonly associated exclusively to the presence of exceptions; also in the very rich and complete survey of Léa Sombé [6] uncertainty about a relation between propositions refers to the presence of exceptions only. Our analysis shows that uncertainty about a relation is more complex and raises the exigency of having, in the general case, a richer and more detailed representation for the uncertainty that may affect a relation than for uncertainty affecting a proposition. However, as far as we know, none of the most known quantitative approaches to uncertain reasoning meets such requirement.

In approaches based on probability theory, a probability value is associated both to propositions and to relations, expressed through conditional probabilities in Bayesian networks [9] or through logical implication relations in probabilistic logic [1, 5, 8]. Therefore, both uncertainty concerning propositions and relations is

represented the same way, actually through a single number. No explicit specification is given about the kind of uncertainty the proposed representation is intended to capture, even if it seems to be strictly related to the concept of A -uncertainty.

In possibilistic logic [4], a real number representing possibility or necessity is associated to propositions and to relations (relations are simply propositions of the form $\neg P_1 \vee P_2$). Also in this case, no explicit specification is given about the kind of uncertainty the proposed representation is intended to capture; moreover, the distinction between possibility and necessity, which is quite clear for propositions, does not seem to carry a definite and well understood meaning for relations. Consider the example presented in [4]: the rule “If John comes tomorrow, it is rather likely that Albert will come” is represented by $(\neg comes(John, m) \vee comes(Albert, m))$ ($N = 0.6$), whereas the rule “Someone will come to the meeting whose presence may (highly possibly, but not certainly at all) make the meeting not quiet” is represented by $(\neg comes(a, m) \vee \neg quiet(m))$ ($= 0.8$). In these cases the distinction between the possibility and the necessity of a relation seems to be rather a matter of subtlety in the use of words than of really different concepts, and it is easy to imagine that, asking different persons, they will express almost the same knowledge using different words such as: “If John comes tomorrow, it is possible, but not certain at all that Albert will come” or “Someone will come to the meeting whose presence is highly likely to make the meeting not quiet”. If knowledge analysis criteria are not specified, it is rather difficult to avoid the risk of an imprecise, and possibly even meaningless, use of the formalism.

In Dempster–Shafer theory (D–S theory), uncertainty quantification applies to subsets of the frame of discernment (i. e., a set of exhaustive and mutually exclusive hypotheses). Uncertainty quantification for a subset of the frame of discernment consists of a pair of real numbers, representing respectively the belief and the plausibility that the correct hypothesis belongs to the subset. Uncertainty quantification is derived from a basic belief assignment, which associates to each subset a belief mass corresponding to a given chunk of evidence. When distinct chunks of evidence are available, global uncertainty quantification is obtained through Dempster’s rule of combination. However, as it was clearly pointed out by [11], “Representing even simple patterns of generic knowledge in a D–S framework may become highly problematic”. Different ways for representing uncertain knowledge in D–S theory have been proposed, such as associating basic belief mass to implication relations [3, 14], or associating directly belief and plausibility to implication relations [11, 12]. In both cases, however, no different representation is provided for propositions and for relations.

Turning now to the concept of attitude proposed in Section 5, it should be noted that this issue has received only very limited attention in the past. In most symbolic and quantitative approaches to uncertain reasoning, the strategies underlying the reasoning mechanism adopted are just left implicit, whereas they are a key factor to verify the suitability of an approach in a given application domain. In general, it can be recognized that probability theory relies on a strongly evolutive attitude, since, in propagation, probability which is not assigned to an hypothesis is forced to be assigned to its negation. Possibilistic logic and D–S theory offer a relaxation of

this strong evolutive attitude: both allow the representation of uncommitted beliefs, since necessity (or equivalently belief) of a proposition and of its negation are not forced to sum up to 1. Therefore, a range of different attitudes could be represented in these approaches by modulating, in the propagation, the amount of uncommitted belief. However, as far as we know, this aspect has not been outlined and satisfactorily explored in the past. Consider for instance Dempster's combination rule: it implicitly adopts a conservative attitude about the belief ascribed to the whole frame of discernment, but an evolutive attitude about the belief ascribed to \emptyset which is redistributed among all the focal elements. Such evolutive attitude is no more present if one adopts an unnormalized representation [13] where it is possible to ascribe a non null belief to \emptyset . Further studies are needed in this direction: in particular the choice of normalized vs. unnormalized beliefs (and of normalizing or non-normalizing rules) should be explicitly represented and under the control of the reasoner, rather than being a design parameter embedded in the formalism. In this way, the most appropriate representation and rule could be dynamically selected and applied, depending on the current reasoner attitude. As a matter of fact, belief functions and probability theory seem more appropriate to accomodate such kind of extensions than probability theory.

Without entering a larger debate, which would be beyond the scope of the paper, it is possible to remark, in general, that the existence of different attitudes suggests the possibility of defining different propagation schemes within a single representation approach. This way, the automated reasoning mechanisms should be able to switch from a scheme to another according to the current attitude. This contrasts with the habit of defining an uncertainty representation and reasoning approach as the combination of a representation and of a unique propagation scheme, considered as generally valid for this representation.

The results discussed in this paper should be considered as preliminary achievements. Several issues will deserve specific attention in future developments. Among the most challenging and promising topics we mention:

- the analysis of the concept of ignorance, the identification of specific classes of ignorance, and the study of the relations between ignorance and knowledge;
- the study of the principles that govern the selection of the most appropriate attitudes to be adopted in an inference step;
- a closer characterization of the nature of A -uncertainty and of C -uncertainty.

(Received November 7, 1997.)

REFERENCES

-
- [1] F. Bacchus: Representing and Reasoning with Probabilistic Knowledge. A Logical Approach to Probabilities. MIT Press, Cambridge 1990.
 - [2] P. Baroni, G. Guida and S. Mussi: Modeling default reasoning through A -uncertainty. In: Proceedings IPMU' 96, International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Granada 1996, pp. 1197–1204.
 - [3] S. Benferhat, A. Saffiotti and P. Smets: Belief functions and default reasoning. In: Proc. UAI 95 11th Conference on Uncertainty in Artificial Intelligence, Montreal 1995.

- [4] D. Dubois, J. Lang and H. Prade: Automated reasoning using possibilistic logic: Semantics, belief revision, and variable certainty weights. *IEEE Trans. Knowledge Data Engineering KDE-6* (1994), 1, 64–71.
- [5] J. Y. Halpern: An analysis of first-order logics of probability. *Artificial Intelligence* 46 (1990), 311–350.
- [6] Léa Sombé Group: Reasoning under incomplete information in artificial intelligence: A comparison of formalisms using a single example. *Internat. J. Intelligent Systems* 5 (1990), 4, 323–472.
- [7] Y. Moses and Y. Shoham: Belief as defeasible knowledge. *Artificial Intelligence* 64 (1993), 299–321.
- [8] N. J. Nilsson: Probabilistic logic. *Artificial Intelligence* 28 (1986), 71–87.
- [9] J. Pearl: Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, San Mateo 1991.
- [10] R. Reiter: Nonmonotonic reasoning. *Ann. Rev. Computer Science* 2 (1987), 147–186.
- [11] A. Saffiotti: Using Dempster-Shafer theory in knowledge representation. In: *Uncertainty in Artificial Intelligence 6* (P. P. Bonissone, M. Henrion, L. N. Kanal and J. F. Lemmer, eds.), Elsevier, New York 1991, pp. 417–431.
- [12] A. Saffiotti: A belief-function logic In: *Proc. AAAI-92 10th National Conference on Artificial Intelligence*, San Jose, 1992, pp. 642–647.
- [13] P. Smets: The nature of unnormalized beliefs encountered in the transferable belief model. In: *Proc. of Uncertainty in AI92*, pp. 292–297.
- [14] P. Smets and Y. T. Hsia: Default reasoning and the transferable belief model. In: *Uncertainty in Artificial Intelligence 6* (P. P. Bonissone, M. Henrion, L. N. Kanal and J. F. Lemmer, eds.), Elsevier, New York 1991, pp. 495–504.

Dr. Pietro Baroni, Prof. Giovanni Guida, Università di Brescia – Dip. di Elettronica per l'Automazione. Italy.

e-mail: baroni@bsing.ing.unibs.it, guida@bsing.ing.unibs.it

Dr. Silvano Mussi, Consorzio Interuniv. Lombardo per l'Elaborazione Automatica. Italy.

e-mail: mussi@icil64.cilea.it