Modeling default reasoning through A-uncertainty

Pietro Baroni, Giovanni Guida

Università di Brescia
Dip. di Elettronica per l'Automazione
Via Branze 38, 25123 Brescia, Italy
e-mail: {baroni, guida}@bsing.ing.unibs.it

Silvano Mussi

CILEA (Consorzio Interuniversitario Lombardo per la Elaborazione Automatica)
Via R. Sanzio 4, 20090 Segrate, Italy
e-mail: mussi@icil64.cilea.it

Abstract

In this paper we present a novel approach to default reasoning, based on the concept of A-uncertainty, namely uncertainty concerning rule applicability. The paper first presents a focused analysis, based on some simple default reasoning examples, of the limitations of some well-known approaches to default reasoning. Then, the novel approach, based on the explicit representation of A-uncertainty is introduced. It is based on a concept of A-uncertainty intended as a property concerning both an inference rule and the individual to which it is applied, and it is shown to be appropriate to overcome most of the limitations found in classical approaches in a natural and effective way. A general default reasoning scheme is then proposed, which exploits the advantages of the introduced representation and sets up a promising background for efficient implementations of automated reasoning .

1. INTRODUCTION

The capability of drawing defeasible conclusions in presence of partial information is a crucial factor of intelligent behavior. To achieve this capability, human beings resort to a particular kind of knowledge, called *default knowledge*. The most significant property of default knowledge is that it can be exploited in the reasoning process even if there is only partial information about the satisfaction of the preconditions which allow its application, on condition that there is no reason to believe that such preconditions are not satisfied. If new information becomes available from which the falsity of such preconditions can be deduced, the conclusions derived from the application of default knowledge have to be retracted. This particular form of reasoning, involving the use of default knowledge, is called *default reasoning*.

In order to build automated reasoning systems including default reasoning capabilities, many extensions of classical logic have been proposed as models of default reasoning. Among the most notable and classic proposals in this field we mention default logic [10] and nonmonotonic logic [8]. The aim of this paper is to propose a new approach to modeling of default reasoning, relying on a different conceptual background with respect to all previous ones. The proposed approach is grounded on the concepts of A- and V-uncertainty we have introduced in a previous paper [2] and allows overcoming in a natural and effective way some conceptual and practical limitations of previous logic-based approaches.

The paper is organized as follows. In section 2 we develop a conceptual analysis of default reasoning, also introducing some default reasoning cases which show significantly different features from the default examples normally found in previous literature. In section 3 we briefly review some approaches to default reasoning and we show that they are inappropriate to deal with the default reasoning cases presented in section 2. In section 4 we shortly recall the concepts of A- and V-uncertainty. In section 5 we introduce our approach to default reasoning, grounded on the concept of A-uncertainty, and we show that it allows a natural treatment of the cases presented in section 2, thus overcoming some limitations of other classical approaches. A final discussion and some conclusions are presented in section 6.

2. DEFAULT REASONING: A CONCEPTUAL ANALYSIS

In a very simple and basic formulation, an inference step of default reasoning involves the two following entities:

- a chunk of default knowledge, and
- an individual to which such knowledge can be applied.

Default knowledge concerns properties of the individual to which it can be applied and, in very general terms, it can be characterized as a form of *relational knowledge*, in the sense that it states a relation between two properties, the former called *premise* and the latter *consequence*. If an individual has the property stated in the premise, the chunk of default knowledge allows inferring that the individual has also the property stated in the consequence. An example of common relational knowledge is "the property of being human entails the property of being mortal". Using relational knowledge to reason about an individual, requires the following three steps:

1. verify, exploiting the information available about an individual, whether he has the property stated in the premise of the chunk of knowledge considered;
2. if the individual has such a property, assume that the chunk of relational knowledge can be applied to the individual;
3. derive the consequence of the application of the chunk of relational knowledge to the individual and ascribe the property stated in the consequence to the individual.

Let us examine now the case of default knowledge. An example of default relational knowledge is "the property of being a bird entails the property of flying". Default knowledge has not a unique semantics; to the above chunk of knowledge several different meanings can be associated: "most birds fly", "normal birds fly", "when we talk about birds, we generally assume they flies", etc. What is common to all these different meanings is that the property of being a bird is a good reason but not definitely sufficient for deducing the property of flying. This introduces a particular kind of uncertainty in the reasoning scheme presented above: uncertainty concerns the fact that the premise is incomplete and, therefore, even if we are totally sure that Tom is a bird (i.e., even if there is no uncertainty about the validity of the premise), we still are not sure whether the available chunk of default knowledge applies to it. Therefore, the reasoning scheme for this type of knowledge can be restated as follows:

1. verify, exploiting the information available about an individual, whether he has the property stated in the premise of the chunk of knowledge considered;
2. if the individual has such a property, there are good reasons to assume that the chunk of relational knowledge can be applied to the individual with a certain belief degree;
3. derive the consequence of the application of the chunk of relational knowledge to the individual and ascribe the property stated in the consequence to the individual with an appropriate quantification of its belief degree.

Of course, if new information becomes available which contradicts the belief that the available chunk of relational knowledge applies to the individual, the belief degree of the consequence has to be revised. A well-known example used to illustrate this form of reasoning concerns penguins (or ostriches or other non-flying birds) and the scheme adopted for managing default reasoning is as follows: given the default knowledge that birds fly, when one knows that an individual is a bird, he is authorized to believe that it flies, however if, from more specific information, it has to be deduced that the individual can not fly, the previous deduction has to be retracted.

In order to make the discussion a little bit more general, let us introduce two more examples which falls, in our opinion, within the area of default reasoning but show significantly different features with respect to the common example of birds and penguins.

The first example concerns truly flying birds, such as canaries, seagulls, condors, etc. Suppose your initial information state is "Tom is a bird": then you have reasonably good reasons to believe it flies. Now suppose you learn that Tom is a seagull: you have now better reasons to believe it flies. In fact, even if also non-flying seagulls exist, this information should corroborate the belief that default knowledge applies to Tom and, therefore, should also corroborate the belief in the consequence. This belief corroboration is dual with respect to belief defeat that would take place if you learned Tom is a penguin. In a sense, a principle of symmetry should be applied: as more specific information contradicting the default is taken into account for retreating the consequence of the default, so more specific information confirming the default should be taken into account for corroborating its consequence. Even if this corroboration effect may look not very important when reasoning about birds, it may be dramatically important in other more realistic application contexts. Suppose there is a medical treatment which is normally very useful but has undesirable strong side effects in rare hypersensitive patients. If you are a physician, by default, you should prescribe the treatment to a patient, unless you know he is hypersensitive. However if you know also that he is not hypersensitive, you may prescribe it with more confidence or in more massive doses.

It is our opinion that both corroboration and defeat should be considered as two facets of default reasoning, which is however normally presented just as the art of retreating.

The second example concerns bird natural enemies: cats. Normally cats have tail. However cats of certain abnormal breeds have no tail. Consider now a crossbred cat, whose parents were a "tailed" and an "untailed" cat. Unless there are fixed genetic rules which apply to such crossbreed, it is impossible to state, in general, if such kind of cat is a tailed or an untailed one. Anyway, it should be considered as an abnormal cat, to which default knowledge does not apply.

Stated in other terms, this case represents a different type of exception with respect to those presented in the birds example. In fact, in those examples, the concept of exception is associated to individuals for which it is known that the property that would be implied by default knowledge is false. This way the world of birds is divided into two classes: normal birds, which fly, and abnormal birds, which do not fly. In the cats example, it is evidenced that the concept of exception is more general, and should be extended to individuals to which default knowledge can not be applied, independently from the fact that they have or have not the property that would be implied by default knowledge. Again, this distinction may look not very important when reasoning about cats, but it may be remarkable in other contexts.

Consider, for instance, the following example, taken from the development of a real application concerning preventive diagnosis of power transformers [1]. Normally, in power transformers, the presence of a significant concentration of C_2H_2 inside insulating oil is a straightforward symptom of a serious internal problem, namely the presence of arcs. Some special transformers include an additional internal component, called OLTC (On Load Tap Changer), whose operation produces C_2H_2. If there is some (unobservable) leakage from the OLTC towards the insulating oil, presence of C_2H_2 can be identified in the oil and, if the general rule is applied, a false alarm is generated, with substantial economical loss. However, if OLTC has no leakage, C_2H_2 presence suggests the presence of arcs and if no intervention are undertaken the transformer may explode. Therefore, transformers with OLTC should be considered abnormal transformers, to which default knowledge about C_2H_2 should not be applied, not because the converse holds, but simply because it is totally uncertain if such knowledge should be applied or not.

3. DEFAULT REASONING: A FOCUSED REVIEW

This section presents a short review of some approaches to default reasoning. Since there is a huge number of papers and different approaches about this topic, we focus on some of the most well-known ones. The specific goal of this review is to analyze the unsuitability of these approaches to model the different aspects of default reasoning discussed in the conceptual analysis carried out in previous section. Since reasons of this unsuitability lie essentially in the conceptual foundations of these approaches, more than in their technical details and variants, we believe that our observations may be extended almost straightforward to nearly all extended and modified versions of the few basic approaches examined here.

3.1 DEFAULT LOGIC

In default logic [10], default knowledge is represented through specific inference rules called *defaults*. A default is an expression of the form: p(x) : j(x)/ c(x), where p(x) is the prerequisite, j(x) is the justification, and c(x) is the consequent of the default. The meaning of the formula is: if p(x) is known and if j(x) is consistent with what is known, then it is possible to deduce c(x).

A default is called normal if j(x) = c(x).

A typical example of normal default is: bird(x) : fly(x) / fly(x), which means: if x is a bird and it is consistent with other available information to assume that x flies, then infer that x flies.

Without entering in more details about default logic, it is important to note that, in default logic perspective, default reasoning is modeled by means of a property of propositions, namely the property of consistency. More precisely, default reasoning consists of an assumption of consistency concerning the consequence of the default rule. An inference step is retreated if the consequence of the inference step is no more consistent with some new information acquired, in other words, it is retreated if the negation of the consequence is entailed by such new information.

Modeling default reasoning activity as a matter of a property of the consequent of default rules is, however, not adequate with respect to the conceptual analysis we carried out above and makes impossible to correctly model some aspects of default reasoning we evidenced in the analysis. In fact, default rules are rules where the premise is only partially stated, so that the matching between the premise of the rule and an individual is in turn partial and uncertain.

Therefore, from a more general and conceptually correct point of view, one should retreat the result of the application of a default rule, because the new information affects the matching between the premise of the rule and the individual, not because it affects the consequent.

In default logic perspective, there is an incorrect modeling superimposition and confusion between two conceptually different facts: the fact that a new information affects the matching between the premise of a default rule and an individual and the fact that, in addition, this new information implies the negation of the consequence of the rule.

This superimposition is not evident and does not give rise to problems in the bird example, because it is implicitly assumed that every bird to which the default rule does not apply is also a non-flying bird. However the problem becomes evident considering the cat example. If you learn that the cat Tom, you has never seen before, is a crossbreed between a tailed and an untailed cat, you might have no preference at all about the fact that Tom has or has not tail. Therefore you might want to suspend your judgment and retreat the assumption that it has tail, derived when you only knew it is a cat.

However this is not possible in Reiter's default logic, because in order to retreat a previous statement you must explicitly state the opposite statement, which is not natural and incorrect in the cat example and in other similar cases.

Besides this conceptual flaw, it has also to be stressed that the need for a consistency check in order to apply default rules is per se a significant disadvantage.

First of all, from the theoretical point of view, consistency checking is generally undecidable.

Furthermore, practical realisations of consistency checking may give rise to unacceptable computational burdens in the application of the theory.

Among the many variants of default logic, we mention here two approaches that seem to partially capture the conceptual aspects we have raised to attention.

Besnard [3] proposes the use of defaults without prerequisites, having the form $:p(x) \rightarrow c(x)/p(x) \rightarrow c(x)$.

The meaning of such defaults is: if it is consistent to assume that the premise implies the consequence, then the premise implies the consequence. In this approach what is subjected to revision is the implication relation, rather than the consequence proposition. This change in the focus of attention should be regarded as a promising step toward a more correct modelization of default reasoning activity, in the sense presented before.

However checking consistency of the implication relation $p(x) \rightarrow c(x)$, still means verifying if the converse holds, that is verifying if $p(x) \wedge \neg c(x)$ holds. Therefore, also in Besnard's approach, the explicit negation of the consequence is a necessary condition for retreating an assumption.

The use of assertion predicates [4] is very close to capture the spirit of our remarks. Supposing all the default rules are identified by an ordering number i, the assertion predicate Ri(x) means that i-th default rule may be applied to the individual x.

Default rules Di have the following form:

$Di = p(x) : c(x) \wedge Ri(x) / c(x)$.

It is important to note that, in order to retreat the default conclusion, it is not necessary to state the negation of the conclusion c(x), because it is sufficient to state the negation of Ri(x), i.e. to state that the rule can not be applied to x.

This approach suffers however, as the previous ones, from the difficulties intrinsic to consistency checking activity. Moreover it is unable, as well as all others, to model corroboration aspects, which will be discussed in more detail in section 3.3.

3.2 NON MONOTONIC LOGIC AND AUTOEPISTEMIC LOGIC

In nonmonotonic logic, [8], the concept of "conceivable", is represented through a modal operator M. The formula Mp means that p is conceivable, that is equivalent to state that $\neg p$ is not provable.

Default knowledge is represented by means of implication relations of the form:

$p(x) \wedge Mc(x) \rightarrow c(x)$.

Therefore also non monotonic logic relies on the same modeling superimposition evidenced for default logic: attention is focused on the consequence, consistency checking is required and the cat example can not be correctly modeled.

Autoepistemic logic [9] is a derivation of nonmonotonic logic, which moves the attention from reasoning about the conceivability of propositions to reasoning about what is believed about propositions.

In autoepistemic approach, the reason why we infer that Tom can fly from the fact it is a bird is that if Tom could not fly, we would know it. So, in absence of more specific information, we assume we can reason as if no further interesting information could come, because we believe that if it could, we knew it. The modal operator of nonmonotonic logic, which will be denoted here as Lp, assumes therefore the meaning of "p is believed".

Default knowledge can then be represented as:

$p(x) \wedge \neg c(x) \rightarrow L \neg c(x)$,

which is the formal translation of " if the premise hold and the consequence does not, I would know it".

Focusing attention on what we believe, rather than on the abstract properties of consistency or conceivability is a step which can be fully subscribed, as we will explain later when presenting our proposal. However, in Moore's proposal, belief concerns again the consequence proposition, so that it is prone to the same modeling superimposition evidenced for other proposals.

Moreover also in autoepistemic logic, only belief retreatment is envisaged, so neglecting belief corroboration.

3.3 INHERITANCE HIERARCHIES

Inheritance hierarchies [12] are directed acyclic graphs used to represent subsumption relations among classes of objects. In the graph two type of links are included: positive links which assert that one class is a subclass of another, e.g. birds are flying things, and negative links which assert that one class is a subclass of the complement of another, e.g. penguins are not flying things. If a class inherits contrasting properties from their superclasses, the most specific property, which roughly corresponds to the shortest path in the graph, is preferred.

Skipping other details, we just remark that also inheritance hierarchies rely on the negation of the property of interest, e.g. flying, and are therefore prone to the same limitations discussed above for other approaches.

3.4 CIRCUMSCRIPTION

Circumscription [6] is an alternative technique which aims to formally represent the intuitive tendency to minimize the assumptions concerning some "abnormal" notion, restricting abnormalities only to cases where they are manifest. Circumscription can be used to model default reasoning [7] by introducing a predicate ab, denoting abnormality, which has to be circumscribed, and representing default knowledge by implication rules having the form:

$p(x) \wedge \neg ab(x) \rightarrow c(x)$.

Since abnormality is not necessarily related to the negation of the consequence, with circumscription it is possible to correctly model situations like those presented in the cat example, because it is possible to classify as abnormal the crossbred cats.

However, also in circumscription it is not possible to model the corroboration effect, since attention is paid only to cases of manifest abnormality, whereas cases of manifest normality are neglected.

4. AV-UNCERTAINTY

We review in this section the concepts of A- and V-uncertainty we introduced in a previous paper [2].

The meaning of term uncertainty is not univocal when dealing with uncertainty affecting relational knowledge, that is knowledge which can be expressed by means of relations between the truth values of propositions. In fact, uncertainty about a relation can carry two different meanings, depending on whether it affects the applicability or the validity of the relation.

Default rules are the most celebrated example of uncertainty about applicability: here uncertainty arises from the fact that there are exceptions to these rules and that it is practically impossible to enumerate and explicitly represent all the exceptions in the premise of the rule. Therefore, the premise of the rule is inherently ill-stated and, even if the properties of an individual match the premise of the rule, it is not sure that the rule can be applied to the individual. In other words, we are simply unable to articulate all the conditions (that indeed exist) that make the rule applicable to a specific individual, either because they are too many and too intricate or because they are (partially) unknown. We call this type of uncertainty that affects the applicability of a relation *A-uncertainty*.

In other cases, uncertainty may affect not the applicability, but the validity of relational knowledge in the sense that it is not certain that the stated relation between premise and consequence really exists. This is due to the fact that in many domains it is uncertain if a given phenomenon is the cause of another (for example, it is uncertain if the smoke causes atherosclerosis, or if cholesterol is really dangerous for heart). Also in this case, if the properties of an individual match the premise of a rule, you are not sure that the consequence hold for this individual, but the reason of uncertainty is quite different. While in the former case you were not sure about the applicability of the rule, assumed to be generally valid, to the individual, in this case you are not sure whether the rule is valid in general, i.e. if a relation exists between the premise and the consequence or if they are totally unrelated, independently of the specific individual considered. We call this type of uncertainty that affects the validity of a relation, but not its applicability, *V-uncertainty*.

Note that A-uncertainty concerns the applicability of a certain chunk of knowledge to an individual, therefore it may be considered as a property of the pair (knowledge, individual). V-uncertainty concerns a chunk of knowledge, independently from individuals, therefore it is a property of the knowledge only. Since A-uncertainty is a property of the pair (knowledge, individual), it depends both on the features of knowledge and of individuals, so that it is possible to imagine, in principle, a different A-uncertainty assessment for each individual to which a given chunk of knowledge has to be applied. Moreover, as long as new information about the individual are acquired, it is possible to adjust the assessment of A-uncertainty relevant to the individual. On the other hand, V-uncertainty is a more stable property of knowledge: it has to be assessed once for all for a given chunk of knowledge and can be adjusted only as long as progress in the considered domain gives new reasons to trust or mistrust such chunk of knowledge.

5. DEFAULT REASONING THROUGH A-UNCERTAINTY

Ignorance (i.e. lack of knowledge) is the source of uncertainty. A statement about the world is either true or false, however if one has incomplete knowledge about the world then he/she is unable to decide if the statement is true or false, i.e. he/she is uncertain about either the truth or the falsity of the statement.

The principle that uncertainty is generated from ignorance (incomplete knowledge about the world) on one hand, and the existence of two kinds of uncertainty about relations (i.e. A-uncertainty and V-uncertainty) on the other hand, prompt us to identify two kinds of incompleteness (of the knowledge about the world): A-ignorance and V-ignorance respectively generating A-uncertainty and V-uncertainty.

A-ignorance is ignorance about the individual the rule is applied to. In other words, A-ignorance may be considered as incompleteness in rule premise.

For example, let us consider the rule: “inveterate smokers catch lung cancer”. The cause-effect relation between smoke and cancer is well known, i.e. the rule is considered to be fully valid. However it has been proved that some inhibitors (inside specific individuals) which prevent smoke from causing lung cancer, have a 15% statistical probability of being active. As a consequence, incomplete knowledge (i.e. A-ignorance) about an individual who is an inveterate smoker will result in considering that the belief in the applicability of the rule is 0.85 (1 means that the rule is fully believed to be applicable).

V-ignorance is ignorance about domain knowledge. In other words, V-ignorance may be considered as incompleteness in knowledge about physical laws the existence of the rule is based on.

For example, the rule “inveterate smokers catch atherosclerosis” is not considered fully valid just because of incomplete knowledge (i.e. V-ignorance) in the medical domain: 60% of scholars indicate smoke as a cause of atherosclerosis, whereas other 40% exclude this hypothesis. As a consequence such an ignorance will result in considering that the belief in the validity of the rule is 0.6 (1 means the rule is fully believed to be valid).

Such considerations show that default reasoning involves both A-uncertainty and V-uncertainty. In fact if new knowledge about the individual is acquired (i.e. A-ignorance diminishes) then the belief in the applicability of the rule may either diminish, carrying out this way a virtual partial retraction (virtual total retraction if the belief diminishes up to 0), or increase (corroboration effect). Similarly, if new domain knowledge is acquired (i.e. V-ignorance diminishes) the belief in the validity of the rule may either diminish (retraction effect) or increase (corroboration effect).

In this paper we analyze only the case where new knowledge about the individual is acquired, which is simpler (concerning only the role played by A-uncertainty) and closer to the classical interpretation of default reasoning. The examination of the case where new domain knowledge is acquired and, therefore, of the role played by V-uncertainty, requires a generalization of the concept of default reasoning, with respect to the meaning normally ascribed to it in previous literature. This issue will be the subject of future research work.

In order to explain how to model default reasoning by means of A-uncertainty, let us briefly survey the approach to uncertainty representation proposed in [2] and adopted here. This proposal, has to be considered as preliminary and aims more to give a sketch of some basic ideas than to introduce a new general and well-settled formalism for uncertainty management. For the sake of simplicity, we assume here that default knowledge is represented by means of production rules.

In our approach, uncertainty about a proposition, say the proposition A, in presence of a chunk of available evidence E, is represented by means of a pair ($\underline{\text{bel}}_E$(A, true), $\underline{\text{bel}}_E$(A, false)), say (bt_A, bf_A) for short. Such pair is called *belief state*, bels for short. The belief state represents 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. Each one of the two elements of a belief state is a *belief degree*. 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, $\underline{\text{bel}}_E$(A, true)=0 means that there is null (or negligible) evidence supporting the credibility that proposition A has the truth value true (note that this is totally different from excluding that true is a possible truth value for A). Belief degrees may assume values in an ordered set (even infinite) of symbols. For instance, we can assume the real interval [0, 1] as the set of possible belief degrees or a discrete set such as {NULL, WEAK, STRONG, FULL}.

The concept of belief state is extended to production rules as follows (we introduce here a variation with respect to [2] where applicability was described as static). Given a production rule R, an individual x, and a body of evidence E, the *AV-belief state* of R with respect to x under E, also denoted by $\underline{\text{AV-bels}}_E$(R, x), is the pair ($\underline{\text{bel}}_E$(applicable(R, x), true), $\underline{\text{bel}}_E$(valid(R), true)), say (ba_R, bv_R) for short.

In intuitive terms, the AV-belief state of a production rule R has the following semantics:

- $\underline{\text{bel}}_E$(applicable(R, x), true) provides a measure of how much one is authorized to believe that the rule is applicable to a given individual x;
- $\underline{\text{bel}}_E$(valid(R), true) provides a measure of how much one is authorized to believe that the rule is fully valid.

According to what stated in section 4, $\underline{\text{bel}}_E$(applicable(R, x), true), namely the belief degree concerning the applicability of a rule to a given individual, should be regarded as a dynamic entity, whose value is subject to change as far as new information is acquired.

Consider the rule R_{fly} = IF bird(x) THEN fly (x). It is possible to state a general value of $\underline{\text{bel}}_E$(applicable(R_{fly}, x), true), which applies when what is known (the available evidence E) is that x is a bird. For instance we could state

$\underline{\text{bel}}_{x=bird}$(applicable($R_{fly}$, x), true) = 0.95 or
$\underline{\text{bel}}_{x=bird}$(applicable($R_{fly}$, x), true) = STRONG, or even
$\underline{\text{bel}}_{x=bird}$(applicable($R_{fly}$, x), true) = FULL.

If more specific information is available, it is possible to explicitly envisage adjustments of the belief concerning applicability, such as:

$\underline{\text{bel}}_{x=canary}$(applicable($R_{fly}$, x), true) = 0.99, or
$\underline{\text{bel}}_{x=canary}$(applicable($R_{fly}$, x), true) = FULL or
$\underline{\text{bel}}_{x=penguin}$(applicable($R_{fly}$, x), true) = 0, or
$\underline{\text{bel}}_{x=penguin}$(applicable($R_{fly}$, x), true) = NULL.

As stated in [2], the cases in which a rule does not apply or is not valid may be interpreted by assuming either a *conservative* attitude or an *evolutive* attitude. For example, recall the previous rules: (R1) “inveterate smokers catch lung cancer” (for which $\underline{\text{bel}}_{x=smoker}$(applicable(R1, x), true) = 0.85), and (R2) “inveterate smokers catch atherosclerosis” (for which $\underline{\text{bel}}$(valid(R1), true) = 0.6). Let us wonder: what happens in the cases where R1 does not apply and in the cases where R2 is not valid ?

Two answers are possible:

- according to a conservative attitude: we know nothing about what might happen to inveterate smokers,
- according to an evolutive attitude: inveterate smokers do not catch lung cancer (considering R1), inveterate smokers do not catch atherosclerosis (considering R2).

In cases of V-ignorance it seems more appropriate adopting a conservative attitude. In fact ignorance about nature laws (V-ignorance) prompts us to be cautious about stating what might happen if the rule is not valid, or, in other words, prompts us to state that we do not know what might happen if the rule is not valid (conservative attitude).

Conversely, in cases of A-ignorance it seems more appropriate adopting an evolutive attitude. For example, given the rule R1, ignorance about a specific individual (A-ignorance) prompts us to implicitly expect that the individual has 15% probability of not catching cancer. In other words, we implicitly state that in the cases in which

rule R1 does not apply, the inveterate smoker does not catch lung cancer (evolutive attitude).
In conclusion, uncertainty propagation is carried out according to an evolutive attitude in rule R1 (since in R1 uncertainty originates from A-ignorance), according to a conservative attitude in R2 (since in R2 uncertainty originates from V-ignorance).
Similarly, in the case of crossbreed cats uncertainty originates from lack of domain knowledge (V-ignorance), i.e. incomplete knowledge about genetics, while in the case of a generic cat uncertainty originates from lack of knowledge about the individual (A-ignorance), i.e. incomplete knowledge about the specific cat. As a consequence, uncertainty propagation through the rule: R_{tail} : IF cat(x) THEN tailed (x), will be carried out according to a conservative attitude in case x is a crossbreed cat, and according to an evolutive attitude in case x is a simple cat.
Since the explicit representation of A- and V-ignorance is beyond the scope of this paper, the additional knowledge concerning attitude, can simply be represented as a dynamic property $\underline{att}_E(R, x)$ which can assume two values, E (evolutive) or C (conservative). For each rule, a general value of $\underline{att}_E(R,x)$ can be stated, which has to be modified if new specific information about the individual is acquired.
Two simple propagation schemes for computing the belief state of the consequence B(x), from the belief state of the premise A(x) and the AV-belief state of the rule R= IF A(x) THEN B(x) can be proposed:

In the evolutive attitude:

$$bt_B = bt_A \cdot ba_R \cdot (1 - bf_A) \cdot bv_R$$
$$bf_B = bt_A \cdot (1 - ba_R) \cdot (1 - bf_A) \cdot bv_R$$

In the conservative attitude:

$$bt_B = bt_A \cdot ba_R \cdot (1 - bf_A) \cdot bv_R$$
$$bf_B = 0.$$

Let us now consider some examples.
Suppose you know for certain that Hugh is a cat: $\underline{bels}$(cat(Hugh)) = (1, 0). Given the rule R_{tail} : IF cat(x) THEN tailed (x), the AV-belief state $\underline{\text{AV-bels}}_{x=cat}(R_{tail}, x)$ = (0.95, 1), and the attitude $\underline{att}_{x=cat}(R_{tail}, x)$ = E, we can derive: $\underline{bels}$(tailed(Hugh)) = (0.95, 0.05), expressing an almost full belief that Hugh has a tail.
If subsequently we learn that Hugh is a Man Island cat (i.e. an untailed cat), we have that $\underline{\text{AV-bels}}_{x=Man}(R_{tail}, x)$ = (0, 1) and $\underline{att}_{x=cat}(R_{tail}, x)$ = E; therefore, we can derive: $\underline{bels}$(tailed(Hugh)) = (0, 1), expressing a full belief that Hugh has not tail. This new belief state, is supposed to overwrite the previous one, so giving rise to a non monotonic belief revision.
On the other hand, if we learn that Hugh is a tabby, we have that $\underline{\text{AV-bels}}_{x=tabby}(R_{tail}, x)$ = (1, 1)
and $\underline{att}_{x=cat}(R_{tail}, x)$ = E;
therefore, we can derive: $\underline{bels}$(tailed(Hugh)) = (1, 0), expressing a full belief that Hugh has tail and obtaining a corroboration effect.
Finally if we learn that Hugh is a crossbreed between a Man Island cat and a tabby cat, we have that
$\underline{\text{AV-bels}}_{x=crossbreed}(R_{tail}, x)$ = (0, 1) and
$\underline{att}_{x=crossbreed}(R_{tail}, x)$ = C;
therefore, we can derive: $\underline{bels}$(tailed(Hugh)) = (0, 0), correctly expressing a state of total ignorance about the question if Hugh has tail.
As shown in the subscripts, it is not necessary to explicitly represent all possible cases of dependency of $\underline{\text{AV-bels}}$ and $\underline{att}$: only significant cases have to be asserted in the knowledge base. So for instance, $\underline{att}_{x=tabby}(R_{tail}, x)$ needs not to be explicitly represented: it is assumed that, if not differently stated, $\underline{att}_{x=tabby}(R_{tail}, x) = \underline{att}_{x=cat}(R_{tail}, x)$.Therefore the general properties are used unless more specific information is acquired. When a rule is applied to an individual, the known properties of the individual are considered and it is verified if they affect the values of $\underline{\text{AV-bels}}$ and $\underline{att}$, then the proper propagation mechanism is applied.
The approach briefly outlined above seems to offer more correct interpretation of the concept of default reasoning than those advocated by previous proposals. For example, in our approach, cases of interacting defaults such as the celebrated "Nixon diamond" [11] are simply and naturally dealt with. Consider in fact the two default rules, "Quakers are pacifists", "Republicans are not pacifists", and the fact "Nixon is a Quaker and a Republican". In our approach, it is possible to state that being a Quaker makes inapplicable the rule about pacifism concerning Republicans, and, vice versa, being a Republican makes inapplicable the rule about pacifism concerning Quakers. In this way, neither rule is applied and the judgment about Nixon's pacifism is, as it seems natural, suspended, instead of reaching a contradiction state, which can be solved by forcing a (possibly unwanted) preference order between defaults.

6. DISCUSSION AND CONCLUSIONS

6.1 A NOTE ON CONDITIONAL PROBABILITIES

The use of symbols such as $\underline{\text{AV-bels}}_{x=cat}$ and $\underline{\text{AV-bels}}_{x=crossbreed}$ recalls the use of conditional probabilities, such as P(tail(x) | cat(x)) and P(tail(x) | cat(x), crossbreed(x)), so that one might wonder if all our requirements could be satisfied adopting a suitable conditional probability representation. The answer is, in our opinion, negative, for two reasons.

The first point concerns the representation of V-uncertainty. As we pointed out in [2], standard conditional probability representation is unable to capture the concept of V-uncertainty. The concept of V-uncertainty could however be represented as a particular case of imprecise probability, as pointed out by Jensen [5]. In fact if one of the extremes of an imprecise probability assessment P(B|A) coincides with the value of P(B), we obtain a representation of the fact that B is completely unrelated with A. However, this proposal has both conceptual and practical weak points. From the conceptual point of view, it mixes imprecision of a probability assessment and uncertainty about the existence of a relation, two very different concepts, in the frame of a unique representation, which comes out to be somehow forced and unnatural. From the practical point of view, it requires that the value of P(B) is known. However, one of the proclaimed advantages of conditional probability approaches is just that they do not need such a priori values, which are normally very difficult or impossible to assess.

The second point concerns the ability to represent conservative and evolutive attitude. Probability calculus rules are fixed and correspond to what we have defined as evolutive attitude: the conservative attitude, which corresponds to suspend the judgment about the negation of the consequence, can not be included in the frame of a probabilistic approach, where there is the constraint that $p(\neg A) = 1 - p(A)$.

6.2 REPRESENTATION AND COMPUTATION ADVANTAGES

With respect to other approaches to default reasoning, our proposal requires the explicit representation of the influence that the acquisition of more specific information has on the AV-belief state (as far as applicability is concerned) and on the attitude of a rule. Therefore, a more rich and complex knowledge base has to be built and maintained. This is indeed the price one has to pay for a more articulate and precise representation.

The study of a complete and efficient uncertainty propagation mechanism and of its computational properties is beyond the scope of this paper and will be the subject of future research work. It has to be noted, however, that our proposal presents two main advantages from the computational point of view with respect to most previous approaches. First of all, it does not rely on any consistency checking: a rule is applied to an individual according to current belief in its applicability and attitude, without the need of other verifications. Secondly, if new information becomes available, a unidirectional propagation process, starting from modifications of the applicability and of the attitude of some rules, is sufficient to carry out all belief revisions required, without the need of detecting contradictions, removing them (if possible), and then backpropagating the effect of these modifications.

References

[1] P. Baroni, F. Cremonesi, G. Guida, S. Mussi, and S. Yakov, ASTRA: a knowledge based system for state assessment, preventive diagnosis, and intervention planning of power transformers, *Proc. ISAP 94 Int. Conf. on Intelligent System Application to Power System*, Montpellier, F, 1994

[2] P. Baroni, G. Guida, and S. Mussi, Modeling uncertain relational knowledge: the AV-production rules approach, *Proc. ECSQARU 95, 3rd European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Fribourg, CH, 1995, 18-27

[3] P. Besnard, *An introduction to default logic*, Springer-Verlag, Berlin, 1989

[4] C. Froidevaux and D. Kayser, Inheritance in semantic networks and default logic, in P. Smets, A. Mamdani, D. Dubois, and H. Prade (eds.) *Non-standard logics for automated reasoning*, Academic Press, London, 1988, 179-212

[5] F.V. Jensen, Personal Communication at *ECSQARU 95, 3rd European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Fribourg, CH

[6] J. McCarthy, Circumscription - A form of non-nonotonic reasoning, *Artificial Intelligence* 13, 1980, 27-39

[7] J. McCarthy, Applications of circumscription to formalizing commonsense knowledge, *Artificial Intelligence* 28, 1986, 89-116

[8] D. McDermott and J. Doyle, Non-monotonic logic I, *Artificial Intelligence* 13, 1980, 41-72

[9] R. Moore, Autoepistemic logic, in P. Smets, A. Mamdani, D. Dubois, and H. Prade (eds.) *Non-standard logics for automated reasoning*, Academic Press, London, 1988, 105-136

[10] R. Reiter, A logic for default reasoning, *Artificial Intelligence* 13, 1980, 81-132.

[11] R. Reiter and G. Criscuolo, On interacting defaults, *Proceedings of 7th IJCAI, International Joint Conference on Artificial Intelligence*, Vancouver, B.C., 1981

[12] D. Touretzky, *The mathematics of inheritance systems*, Pitman, London, 1986