CONDITIONING IN A FULLY NONMONOTONIC CONTEXT

Pietro Baroni*, Giovanni Guida*, Silvano Mussi°*

* Università di Brescia, Dipartimento di Elettronica per l'Automazione, Via Branze 38, 25123 Brescia, Italy
tel: +39 30 3715469 fax: +39 30 380014 - e-mail: {baroni, guida}@bsing.ing.unibs.it

° Consorzio Interuniversitario Lombardo per l'Elaborazione Automatica, Via Sanzio 4, 20090 Segrate (MI), Italy
tel: +39 2 26922434 fax: +39 2 2135520 - e-mail: mussi@icil64.cilea.it

Abstract

In this paper we explore how the concept of full nonmonotonicity, introduced by the authors in a recent paper, can be related to the classical and widely accepted view of conditioning and belief revision which has been axiomatically set up by Gardenfors [Gardenfors 88].
After recalling the concept of full nonmonotonicity and its importance for a correct modeling of many practical reasoning situations, we show that the classical view of conditioning is unable to encompass all the situations that can be covered by full nonmonotonicity. Then we argue that a correct modeling of fully nonmonotonic reasoning requires a substantial extension of the conceptual model of conditioning. In particular we remark the importance of an explicit representation for the concepts of uncertainty about the applicability of a piece of knowledge and of reasoning attitude. Finally, we introduce a preliminary reasoning formalism featuring the full nomonotonicity property.

1. Introduction

In the frame of reasoning under uncertainty, conditioning can be understood, in very general terms, as the activity of modifying in some way one or more of the beliefs currently held by the reasoner and, then, of updating accordingly (all) the other current beliefs.
In symbolic approaches to uncertain reasoning, the current belief state of the reasoner is represented by symbolic values (e.g., truth values of propositions in propositional logic) and conditioning consists of modifying such symbolic values. On the other hand, in quantitative approaches, the belief state is represented by uncertainty quantifications (e.g., probability values associated to events in probability theory) and conditioning consists of modifying such quantifications. The basic conceptual scheme of conditioning is anyway the same: a local modification of the current belief state is introduced and then its effects are propagated to the overall belief state.

From a general point of view, Gardenfors [Gardenfors 88] has introduced a systematic classification of different types of conditioning, related to the different types of local modification that can originate it. He distinguishes three cases, namely:

- expansion, when a new belief is introduced which does not contradict the previously held ones;
- revision, when a new belief is introduced which contradicts the previously held ones;
- contraction, when one of the previously held beliefs is retracted.

For each one of the three cases defined above [Gardenfors 88] provides a set of postulates that should be respected by a reasoner in order to realize such kind of conditioning in a correct way.
Such postulates are defined for a generic symbolic approach, however their generality has been widely recognized and allows their application also to quantitative approaches. The relation with probabilistic

models is directly discussed in chapter 5 of Gardenfors' book, but these postulates have also been used as a reference for other quantitative theories. For instance, in possibility theory, the definition of possibilistic revision is in accordance with Gardenfors' postulates (see the very recent [Dubois and Prade 97]).
Katsuno and Mendelzon [Katsuno and Mendelzon 91] have extended Gardenfors' classification, by taking into account the reasons why the modification is introduced. They distinguish two cases, namely:

- the case where a modification is introduced because a new information has been acquired about a world which is basically static;
- the case where a modification is introduced because a change has occurred (and has been noticed) in a dynamic world.

The former case can be correctly represented using Gardenfors' postulates about revision and contraction, whereas for the latter one Katsuno and Mendelzon introduce two new sets of postulates and two new kinds of conditioning, called updating and erasure respectively.

Another classification of conditioning is proposed in [Dubois et al. 96]. They refer only to the case of acquisition of a new information (not considering retraction) and divide the beliefs held by a reasoner in two classes, namely: factual evidence, that is information concerning the specific case at hand, and generic knowledge, pertaining to the whole class of possible situations. On the basis of this classification, they distinguish two types of conditioning:

- focusing, that is conditioning the generic knowledge by the factual evidence; for instance, changing the reference class for an individual in the light of the evidences collected about it;
- revision, that is extending either factual evidence or generic knowledge (F-revision or G-revision, respectively) by enforcing a new piece of information.

The aim of this paper is to further explore the concept of conditioning, by exploiting the concept of full nonmonotonicity recently introduced by the authors [Baroni et al. 97a]. After reviewing and discussing the concept of full nonmonotonicity, it is shown how, in a fully nonmonotonic reasoning context, the concept of conditioning has to be generalized and extended with respect to existing classifications. In particular, in a fully nonmonotonic context, conditioning may be related to a modification of the adopted reasoning attitude [Baroni et al. 97b] rather than to a direct modification of the currently held beliefs. A preliminary version of a formalism for uncertain reasoning which can correctly model the activity of conditioning in a fully nonmonotonic context is then introduced and its application to some simple examples is discussed.

2. The concept of full nonmonotonicity

In general, in any formal reasoning context, some general rules are given that define how expressions (which are the basic way to represent facts about the world) should be formed (for example, rules defining Well Formed Formulas in a logic formalism). These rules allow to generate the (possibly infinite) set E of all *expressions* one can consider when reasoning about a given fragment of the world.

At any stage of the reasoning activity, the set E can be partitioned in two classes: the class AE of expressions for which a truth value (namely, *true* or *false* in a two-valued logic) has been explicitly asserted (AE stands for *asserted expressions*) and the class UE of expressions to which a truth value has not been explicitly ascribed (UE stands for *unasserted expressions*). The set UE includes both the expressions to which it is actually impossible to ascribe a truth value, for instance due to a lack of information, and the (very numerous) expressions which are not considered in the reasoning activity, though their truth value could be derived given the available information (for instance, if a proposition P is asserted

to be *true*, also any expression of the type $p \vee q$, could be explicitly asserted to be *true*, but, of course, the bothersome and useless activity of deriving all these assertions is normally avoided).
Moreover, the class *AE* can be further partitioned in two subclasses, the class *TE* of expressions whose asserted truth value is *true*, and the class *FE* of expressions whose asserted truth value is *false*.

Each step of the reasoning activity can be seen, in general, as a modification of the partitions defined above, which moves one or more expressions from a class to another one. In monotonic reasoning, this modification is always and only the shift of expressions from *UE* to *AE*.
In fact, classical logical formalisms satisfy a property called *monotonicity*, which, in a restricted form, can be defined as follows:
given three different sets A, B, C of expressions, if A $\vdash$ B, then (A $\cup$ C) $\vdash$ B.
The monotonicity property implies: (i) that asserted expressions can not revert to unasserted and (ii) that the truth value ascribed to an expression can not be changed. So the monotonicity property implies both that the set *AE* is monotonically growing and that there are no cross shifts between *TE* and *FE*.
That being stated, an approach can be defined *nonmonotonic*, in general terms, if it fails to satisfy either one of the properties (i) and (ii) entailed by monotonicity. Therefore, in a nonmonotonic reasoning context, we can define two types of nonmonotonicity (denoted as AU- and TF-respectively):

- a reasoning formalism is said *AU-nonmonotonic*, if it allows an expression to move from *AE* to *UE*;
- a reasoning formalism is said *TF-nonmonotonic*, if it allows cross movements of expressions between *TE* and *FE*.

It is then possible to define the property of *full nonmonotonicity* as follows: a reasoning formalism is said *fully nonmonotonic*, if it is both *AU-nonmonotonic* and *TF-nonmonotonic*.
As it has been shown in [Baroni et al. 97a], most approaches to defeasible reasoning presented in literature are only TF-nonmonotonic. In fact in such approaches, including default logic [Reiter 80], nonmonotonic logic [McDermott and Doyle 80] and autoepistemic logic [Moore 88], the modification of a previously asserted expression requires necessarily that the opposite truth value is explicitly asserted.
However, in many cases, common sense reasoning features the full nonmonotonicity property, so that TF-nonmonotonicity comes out to be a partial and insufficient model of practical reasoning activity.
In general it can be stated that:

- TF-nonmonotonicity captures the situations where the acquisition of a new piece of information allows us to rectify an erroneous assertion, by modifying the asserted truth value;
- AU-nonmonotonicity captures the situations where the acquisition of a new piece of information shows that we are more ignorant than we believed to be and therefore urges us to unassert a previously asserted expression.

In [Baroni et al. 97a] we present a somewhat picturesque example (called "the ill secretary example") where we show how, in a common sense reasoning context, both TF- and AU-nonmonotonicity are present. In a nutshell, the example runs as follows: by default you believe that your secretary will come at work today, so you include its coming among the expressions asserted as true and arrange your plans under the hypothesis that she will actually come. If you learn that yesterday night she had an accident and broke her leg, you realize that she will not come (TF-nonmonotonicity), and rearrange your plans in order to take into account her absence. On the other hand, if you learn that yesterday she was feverish you are now less inclined to believe that she will come but you are also not sure that she will not. In this case you are completely uncertain about her coming (and therefore also about the best arrangement of your plans): this is a case of AU-nonmonotonicity and can be correctly modeled only by making unasserted the expression concerning secretary's coming.

This example concerns a common, trivial situation, however it is not difficult to devise the need of AU-nonmonotonicity also in more serious reasoning contexts. For instance, suppose that a physician has founded its diagnosis about a patient on a the interpretation of a given set of symptoms, and that then a new symptom emerges that was not expected on the basis of the previous diagnosis. In this case the previously asserted diagnosis is questioned, but also excluding it on the basis of the new symptom is not correct. Simply, the new symptom shows that previous diagnosis was not so certain as it seemed to be (unfortunately, this happens very often in practice) and that such diagnosis should therefore be unasserted (which is very different from denying it) until a new interpretation has been formulated coherent with all symptoms collected.

Many other examples of AU-nonmonotonicity can be easily formulated covering many application domains (just think about finance or about military strategies) and sharing the same main feature: sometimes the new information we acquire may reduce the set of expressions we are able to assert.
In other words, there are cases where we retract one of our previous conclusions not because it has to be replaced by a more reliable one, but because we are unable to draw any conclusion.
It may seem a little bit paradoxical that the acquisition of new information reduces the scope of our deduction capability. However, the ill secretary example shows that this case can be found even in very simple examples of common sense reasoning. Moreover, the paradox reveals to be only apparent if one considers that, also in this case, the acquisition of new information increases our knowledge. Simply, in this case the new piece of knowledge we learn is our inability to make deductions. In fact, learning that we are ignorant about something, definitely increases our global state of knowledge, that includes both what we know and what we know to ignore.

3 Conditioning and full nonmonotonicity

After having introduced the concept of full nonmonotonicity, we explore in this section how it can be related to the concept of conditioning and which new requirements it introduces with respect to current conditioning methods.
In order to limit the scope of the discussion, we focus here on the case where conditioning is related to the acquisition of a new piece of factual evidence about a static world situation (called revision in the terminology of [Gardenfors 88] and F-revision in the terminology of [Dubois et al. 96]).

3.1 A fundamental limitation

In the traditional view of conditioning, presented by [Gardenfors 88] (and inherited from a consolidated research trend [Harper 77] [Levi 77]), the revision of a belief set *BS* by the addition of the belief in a proposition *A*, is stated to be equivalent to the sequence of a contraction and an expansion, namely the contraction of *BS* with respect to $\neg A$, and then the expansion of the resulting belief set by the assertion of *A*. This is fully coherent with the classical view of TF-nonmonotonicity. In fact, it is possible to show that such equivalence holds only when the new piece of information entails the negation of a previously held belief.
Consider the contraction operator as characterized by [Gardenfors 88]: the contraction with respect to a proposition *P* consists of retracting (unasserting) the proposition *P* itself along with all those propositions that entail *P*. This definition is justified by the fact that if one would not unassert the propositions entailing *P*, *P* would be asserted again, due to deductive closure.

It should be noted that contraction is a nonmonotonic operator *per se*, since the "unassertion" of a proposition violates the monotonicity property. Moreover, contraction is an AU-nonmonotonic operator, since the effect is to make unasserted some previously asserted propositions. Therefore if a reasoner decides spontaneously to apply a contraction to its belief state, such operator realizes a form of AU-nonmonotonicity.

However the cases of AU-nonmonotonicity we have considered in the previous section are a little bit different. It is the acquisition of a new piece of information, not just a spontaneous decision, that prompts the reasoner to retract some of its previously held beliefs.
A correct modeling of this kind of situations should involve therefore the execution of AU-nonmonotonic operations in the context of belief revision, namely in the context of the execution of the operator that is devoted to revise previous beliefs when a new one is acquired.
Revision is, by definition, a nonmonotonic operator. However, as we will show, in the definition provided by [Gardenfors 88] the revision operator is only TF-nonmonotonic and is therefore unable to satisfy the need for AU-nonmonotonicity expressed above.
In other words, it might be stated that this traditional definition of revision limits (and, in a sense, wastes) the modeling capabilities provided by contraction operator, which is AU- nonmonotonic.

Before demonstrating this assertion, it is necessary to recall the main properties of the third operator introduced by [Gardenfors 88], namely expansion.
Expansion of a belief set *BS* with respect to a proposition P is, by definition, a monotonic operator which consists of deriving the deductive closure of $P \wedge BS$. By definition, in the case of expansion, P does not contradict *BS*, so the operator is guaranteed to be monotonic.

We are now ready to recall the fact that in a traditional view [Harper 77] [Levi 77] [Gardenfors 88], revision operator is stated to be equivalent to a sequence of a contraction and an expansion.
More formally, given a belief set *BS* and a proposition P to be added to *BS* as asserted, it is stated that:
$revision(BS, P) = expansion(contraction(BS, \neg P), P)$.

In the light of the considerations presented above, it appears that in the application of the revision operator, a first step involving AU-nonmonotonicity is provided by the contraction operator, which takes care of unasserting a given set of propositions.
The subsequent expansion operator is monotonic with respect to the results of contraction but is TF-nonmonotonic with respect to the original belief set *BS*, since the expansion of the contracted set may lead to assert the truth of some propositions that were asserted as false in *BS*.
So, at a first glance, it seems that both AU- and TF-nonmonotonicity are encompassed by the classical definition of revision. However, a more accurate examination shows that only TF-nonmonotonicity is actually captured by this approach.
In fact the explicit goal of the operation $revision(BS, P)$ is asserting P. So the contraction of $\neg P$ is bound to be followed by the assertion of P: if $\neg P$ was not asserted, we have simply a case of monotonic expansion of P, if instead $\neg P$ was previously asserted, we have a case of TF-nonmonotonicity since P becomes true from false.
Therefore, only TF-nonmonotonicity can be achieved in the classical view of revision as far as P is concerned.
What can be said now about other propositions retracted during the operation of $contraction(BS, \neg P)$?
These propositions are those which entail $\neg P$ and, of course, are retracted only if they were previously asserted. When the subsequent expansion is performed, P is asserted and its deductive closure is derived.

Reconsider now that the propositions retracted in the contraction phase was those entailing $\neg P$. Now, from the assertion of P and from the entailment relation, we can derive, by modus tollens, that the previously retracted propositions must be false. Therefore, also for these propositions, only TF-nonmonotonicity is possible in a classical revision context.

In conclusion, the discussion above shows that the classical view of revision is strictly bound to the concept of TF-nonmonotonicity and is unable to capture any form of AU-nonmonotonicity.
In other words, in such view the acquisition of a new piece of information is bound to cause either a monotonic expansion or a TF-nonmonotonic revision of previously held beliefs.
However, this is not the most general case, since it may happen that the newly acquired piece of information implies that a previously held belief becomes unknown, rather than being denied, as we have discussed in section 2. In this case, the equivalence:
"*revision = contraction + expansion*"
which is one of the cornerstones of the classical theory of conditioning, does not hold any more.
This means that the classical theory of conditioning should be modified and extended in order to cover also the case of fully nonmonotonic reasoning.

3.2 Requirements for a theory of fully nonmonotonic conditioning

We claim that such modification requires a substantial extension of the conceptual background underlying the theory, rather than an adjustment concerning merely technical details.
The key point consists of providing a suitable representation for the fact that a new information implies that a previously held belief becomes unknown. This seems very difficult to accomodate in a classical knowledge representation framework, where implication between propositions is the only way to represent relations between the truth values of different propositions.
The assertion of an implication relation $A \rightarrow B$ introduces some constraints between the truth values A and B may assume. In practice, if A is *true*, B is constrained to be *true*, whereas if B is *false*, A is constrained to be *false*.
Now one could be tempted to introduce a similar representation in order to encompass the fact that the assertion of a previously unasserted proposition constrains another proposition to become unasserted. In this way, AU-nonmonotonicity could be explicitly encompassed by a properly defined and non standard implication relation, where this relation introduces at least the constraint that if A is *true* then B is *unknown*.
However the introduction in a classical framework of such a new type of implication stating that "*A=true implies B=unknown*" is technically difficult, practically disadvantageous, and, most importantly, conceptually incorrect.

First of all it has to be noted that, rigorously speaking, *true* and *false* are *truth values*, whereas *asserted* and *unasserted*(*unknown*) are not truth values but *assertion states*, which represent the opinion of a reasoner about a proposition.
Recalling that a proposition is intended to represent a fact about a reference world (either the "real" world or an abstract one, such as the world of numbers) *true* and *false* represent the two possible states allowed for the considered fact in the reference world. If this fact may have more than two states, we have to adopt a multi-valued logic rather than a binary one.
Implication relations (and any other form of relational knowledge) represent constraints holding among truth values (that is among states of facts) in the reference world. In a sense, such relations limits the arrangements of truth values that may occur in the reference world. So a relation "*A implies B*" represents a constraint excluding the worlds where A is *true* and B is *false*.

In absence of any knowledge and information, any world should be considered possible and any world could be the actual one (of course, at least in "real" world, there exists exactly one state of affairs at a given time instant). A reasoner, however, is normally interested to have at least a reasonably accurate idea about what the actual world is. The reasoner has therefore to acquire and exploit both generic knowledge and specific information concerning the actual world in order to restrict the set of worlds to which the actual one belongs.

Both generic knowledge and specific information are a private asset of the reasoner and represent what the reasoner believes about the world: we will call therefore the set of knowledge and information possessed by the reasoner the *belief state* of the reasoner. Of course different reasoners, sharing the same reference world, may have very different knowledge and information about it (and this is indeed the case in many practical situations).

It has to be remarked that both knowledge and information are not a property of the world but of the reasoner. It is the reasoner that takes the responsibility of the correctness of the knowledge and of the information he uses.

In a simple setting not including learning capabilities, we can assume that knowledge is an a priori asset of the reasoner, whereas pieces of specific information are dynamically asserted (as true or false) and unasserted within its belief state. Therefore, within the belief state of the reasoner, a piece of specific information (let say, a proposition) may have two *assertion states*, namely *asserted* and *unasserted,* that are really different from the *truth values*, namely *true* or *false* it may have in the actual world. The difference can be further remarked by the fact that the state *asserted* admits a further specification since a proposition may be asserted as *true* or as *false*.

In this context, the a priori knowledge, whose basic meaning is to represent a constraint existing in the real world, also constrains the assertions a reasoner may make (at least as far as the reasoner does not want to contradict its knowledge). In fact, if a reasoner knows that, in the real world, *A* implies *B* and, at the same time, asserts the truth of *A* in its belief state, it is also constrained to assert that *B* is true in its belief state, because the believed world(s) must be coherent with its knowledge or, in other words, the set of believed world(s) must be included in the set of the possible ones. Therefore, these constraints drive the reasoner's reasoning activity and in particular the activity of conditioning: in fact, once the reasoner introduces a local modification in its belief state, it is forced to globally modify it in order to respect these constraint and at the same time the general principles expressed by conditioning postulates. However, it should be clear that, even though the constraints provided by knowledge also affect reasoner's assertions, their basic nature is that of representing constraints holding among truth values in the real world, not constraints among assertions in the believed world.

After this discussion, it should also be clear that the state of unasserted (or unknown) of a proposition concerns the believed world and that, strictly speaking, it is not a truth value. Therefore, a relation such as "*A= true implies B=unknown*" is clearly improper and conceptually incorrect since it introduces an undue constraint between two very different objects: a truth value and an assertion state.

In order to encompass full nonmonotonicity, one could then hypothesize to introduce a different kind of constraint, involving assertion states only, such as "*A=asserted(true) implies B=unknown*".

Technical difficulties related to the proper definition and use of such kind of constraint can easily be imagined. However, most importantly, this kind of constraint would not solve the problem anyway.

Reconsider the ill secretary example: in this case we want to unassert the fact that she will come, because we learned that yesterday she was feverish. However, if we see her coming in, wrapped up in a wool shawl, we realize that she was and is still feverish but that she has decided to come anyway: in this case the fact that she was (and is) feverish remains asserted, but the question concerning her coming is not unknown anymore. Of course, this can not be captured by the constraint presented above. Moreover, we might also learn that the secretary is feverish from herself, after she arrives regularly at her default hour. Also in this case the new information does not make her coming unknown, since she is already in.

These simple examples show that the reason why a new information *NI* (in the example, the fact that the secretary was feverish) induces to unassert a previously held belief *PB* (in the example, the belief in the presence at work of the secretary) does not lie in a direct relation between *NI* and *PB*: in some cases *NI* does actually affect *PB* and in other cases it does not. As the examples show, this depends on the reason why *PB* is held: if *PB* was based on default knowledge, *NI* affects it, whereas if it based on other justifications (for instance direct evidence), *NI* does not affect it anymore.
Therefore, the impact of *NI* on the previously held beliefs should be related to the effect *NI* has on the knowledge used to deduce *PB* and, in particular on the applicability of such knowledge to the situation at hand. In other words, the acquisition of *NI* may reveal to the reasoner that the knowledge used to derive *PB* is no more applicable to the case at hand. So, when *PB* was derived from default knowledge, (the fact that normally the secretary comes everyday), *NI* reveals that we are in a situation of exception, where default knowledge does not apply. On the other hand, if *PB* was derived from direct evidence (or from a piece of knowledge that does not admit exceptions) *NI* has not a significant impact on it.

3.3. A conceptual model for full nonmonotonicity

Such kind of reasoning mechanism may be correctly modeled by introducing an explicit representation for two concepts:

- the *uncertainty about the applicability* of a given piece of generic knowledge to a specific individual;
- the *reasoning attitude*, either *evolutive* or *conservative*, that may be adopted in the use of a piece of uncertain knowledge.

Detailed discussions about these concepts and their relation with nonmonotonic reasoning can be found in a sequence of papers [Baroni et al. 95] [Baroni et al. 96] [Baroni et al. 97b] and are quickly summarized here for the convenience of the reader.

Uncertainty about the applicability of a piece of knowledge (let us say, of a rule for the sake of brevity) arises from the fact that there are exceptions to this rule 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 certain 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.*
Since A-uncertainty concerns the applicability of a given chunk of knowledge to an individual, it may be considered as a property of the pair (knowledge, individual). Therefore 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 is acquired, it is possible to adjust the assessments of A-uncertainty relevant to the rules applicable to the individual.

The concept of A-uncertainty can then be related to the one of reasoning attitude. In fact, the cases in which a rule does not apply may be regarded by assuming either a *conservative* or an *evolutive* attitude. In fact given a rule *R*, an individual x and the fact that we may be uncertain either the rule applies to x or not, we may wonder: what should we believe about the cases where *R* does not apply to x ?

Two answers are possible:

- according to a conservative attitude: we know nothing about x,
- according to an evolutive attitude: we know that the consequent of the rule is false for x.

The relation between A-uncertainty, reasoning attitudes and full nonmonotonicity is now straigthforward.
In fact, in presence of a rule *R* affected by A-uncertainty and of an individual x whose properties match the premise of the rule, the reasoner may draw some conclusion about x which are defeasible, in the sense that the acquisition of a new piece of information may reveal that the rule was not actually applicable to x and that therefore the consequences of the application of the rule have to be properly revised. Therefore the presence of A-uncertainty causes the need for nonmonotonicity.
In fact the acquisition of new information about the individual x may reveal that there is an exception to the rule and therefore the applicability of *R* to x is affected .
Then the application of the rule has to be reassessed, selecting the most suitable reasoning attitude (either evolutive or conservative) about this exceptional situation.
The attitude selection allows to encompass both types of nonmonotonicity. In fact, the evolutive attitude corresponds to TF-nonmonotonicity, since the consequent, previously asserted as true, is now asserted as false. On the other hand, the conservative attitude corresponds to AU-nonmonotonicity, since the previously asserted consequent is then regarded as unknown.
The selection of the proper attitude depends, of course, on the nature of the exceptional situation the reasoner is actually facing and has to be driven by domain dependent criteria. Such criteria have to be explicitly encoded within the reasoner's knowledge base in order to properly manage exceptional situations.

Referring to the ill secretary example, the default rule that she will come is applied with an evolutive attitude in absence of more specific information.
If we learn that she has a broken leg, we realize that we are dealing with an exceptional situation, namely our belief in the applicability of the rule suddenly decreases, and we adopt an evolutive attitude with respect to this exception: the secretary will not come.
On the other hand if we learn that she was feverish, again we realize that we are dealing with an exceptional situation and our belief in the applicability of the rule decreases, but this time a conservative attitude is more suitable: secretary's coming is now an open question and so remains unasserted.

The use of the new acquired information to select the proper reasoning attitude introduces a new viewpoint about the concept of conditioning. In fact, in our proposal, conditioning does not affect only the results of a reasoning process, but, instead, it may also modify some parameters affecting the way the reasoning process is performed.
In particular, the acquisition of new information about an individual x may affect:

- the belief in the applicability to x of the rules that were previously applied to it;
- the reasoning attitude concerning the same rules.

It is in a subsequent step, when the modifications of these parameters are put at work, that we obtain also a modification of the propositions concerning x and previously derived using these rules.
So conditioning is regarded, in our view, as a two steps activity: the first step involves the readjustment of some reasoning parameters, the second one involves the production of modified results in accordance with the adjusted parameters. This view contrasts with the classical one, where conditioning is regarded as the activity of modifying reasoning results, in the light of new information, using a fixed reasoning scheme.
Our approach is therefore more flexible and offers several advantages. Among the most significant ones we mention that:

- it allows a conceptually richer and cognitively plausible modelization of the overall activity of reasoning under uncertainty. A more articulated knowledge representation is provided, including A-uncertainty and reasoning attitudes, which seems very suitable to capture some common, but often neglected, aspects of reasoning under uncertainty.

- it is able to provide a straightforward representation for the concept of full nonmonotonicity which instead can hardly be included within other conditioning paradigms.

A detailed comparison of our approach with other conditioning paradigms is beyond the scope of the present paper and will be the subject of future works.

4 A formalism for conditioning in a fully nonmonotonic context

In this section we introduce a quantitative formalism for representing conditioning activity in a fully nonmonotonic context. The formalism can support all the conceptual aspects discussed in section 3 and constitutes a substantial advancement with respect to the preliminary formulation presented in previous papers (see for instance [Baroni et al. 97a] [Baroni et al. 97b]).
Anyway, the proposal presented here has still to be considered as preliminary and aims more to substantiate some basic ideas than to definitely 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 rules.
The formalism is based on the definition of a proper representation of four basic elements, namely uncertainty about propositions, uncertainty about the applicability of a relation to an individual, reasoning attitudes, and uncertainty propagation mechanism.

4.1 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, if $bel_E(P1, true)$ is zero, this means that there is null (or negligible) evidence supporting the credibility that proposition $P1$ has the truth value *true*, and this is totally different from excluding that *true* is a possible truth value for $P1$. Similarly, $bel_E(P1, true)=1$ means that available evidence fully supports the credibility that proposition $P1$ has the truth value *true*, and this is again totally different from being absolutely certain that *true* is the correct truth value of $P1$.

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 $P1$ and a body of evidence E, the *belief state* of $P1$ under E, denoted by $bels_E(P1)$, is the pair:
$(bel_E(P1, true), bel_E(P1, false)$, (say (bt_{P1}, bf_{P1}) 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. A proposition accompanied by the relevant belief state is called a *quantified proposition*: more formally, for any proposition $P1$, the pair ($P1$, $bels_E(P1)$) is a quantified proposition.

We remark that the two components of the belief state are completely independent: all couples of values are possible. This choice contrasts with many uncertainty theories ranging from probability theory, to belief functions theory, to possibility theory, where uncertainty quantifications relevant to the truth and to the falsity of the same proposition are not independent, since they are subjected to more or less restrictive mutual constraints. However our choice allows a richer expressiveness and is especially apt to represent uncertain situations where multiple evidences, arising from independent sources, have to be taken into account, possibly giving rise to strong contradictions (a similar representation has been proposed in [Benferhat et al. 95]).

For a better understanding of the concept of belief state, we can consider the concept of *belief plane*. The belief plane is the Cartesian plane of all possible belief states of a generic proposition P: the x-axis represents the possible values of bt_P, and the y-axis the possible values of bf_P. Both the x-axis and the y-axis range on the interval [0..1] and a belief state $bels_E(P)$ is represented by a point in the belief plane. Four points in the belief plane are worth special attention, since they convey a clear intuitive semantics; namely:

the *B* point (1, 0) means	"totally believed"
the *D* point (0, 1) means	"totally disbelieved"
the *U* point (0, 0) means	"totally unknown"
the *C* point (1, 1) means	"totally contradictory".

In fact, the belief state (1, 0) means that the available evidence fully supports the truth of the proposition and that there is no evidence supporting its falsity: therefore we should be fully convinced, on the basis of the available evidence, that the proposition is *true*. Similarly, the opposite conviction will be represented by the point (0, 1). Moreover the belief state (0, 0), which indicates the absence of evidence both supporting the value *true* and the value *false*, represents a state of total ignorance about a proposition (due to a lack of evidence). On the contrary, the belief state (1, 1) means that there are, for any reason, strong evidences for both the values *true* and *false*, and the situation is contradictory. Of course, all intermediate situations are possible, since the two components of a belief state are independent.

4.2 Uncertainty about the applicability

The concept of belief state can be extended to the property of applicability of a rule to an individual. Given a production rule R, an individual x, we define a predicate *applicable*(R, x) which express the fact that x is included in the set of individuals to which R correctly applies. Given a body of evidence E, the reasoner may believe at a given extent that R has to be applied to x.
Therefore we can model this situation through a belief state $bels_E(applicable(R, x))$, that is the pair ($bel_E(applicable(R, x), true)$, $bel_E(applicable(R, x), false)$), say ($bt_{app(R,x)}$, $bf_{app(R,x)}$) for short.

Of course, if the body of evidence changes and we acquire new information about x, also the belief about the applicability of the rule to x may change. Therefore $bels_E(applicable(R, x))$, should be regarded as a dynamic entity, whose value is subject to change as far as new information is acquired during the reasoning process.

4.3 Reasoning attitude

As explained above, also an explicit representation of reasoning attitudes is required. In the simplest way, the attitude concerning the application of a rule *R* to an individual x can simply be represented as a dynamic property $att_E(R, x)$ which can assume two values, *E* (evolutive) or *C* (conservative). The initial reasoning attitude, whose value may be assigned by default, can be modified if new specific information about the individual is acquired.

4.4 Uncertainty propagation and conditioning

We concentrate here on the simple case of a single proposition whose belief state has to be dynamically readjusted according to the acquisition of new information. Following the definition presented in section 4.1, the belief state of a proposition is conceived as a mobile point on the belief plane.
In absence of any relevant information, the belief state of a proposition *P* lies on the *U* point, representing a state of full ignorance.
When a new piece of information arrives, affecting the belief state of *P*, the position of *P* has to be modified accordingly within the belief plane.
The basic mechanism for this modification is inspired to a spring model: the resulting belief state is conceived as an equilibrium point between different points of attraction, each one connected to the equilibrium point by an elastic spring. As we will explain in the following, the position of the attraction points and the elastic constant of the spring are determined by the different uncertainty quantifications involved in the reasoning step and by the adopted reasoning attitude.

Now given an individual x and a rule R = IF *prop(x)* THEN *cons(x)*, we have to define how the belief state $bels_E(cons(x))$ can be determined starting from the initial belief states $bels_E(prop(x))$, $bels_E(applicable(R, x))$, and from $att_E(R, x)$.
It is important to note that we do not face here the very important problem of how such initial belief states are derived from the available evidence. Evidence interpretation is a very important and complex task, which is however far beyond the limits of this paper. In a very simplified setting, we can assume that the reasoner is endowed with domain-specific knowledge that allows it to directly map collected evidence into these initial belief states.

We will now discuss how these belief states are individually translated into our spring model by examining the role played by their components in determining the belief state of the consequence.
Let us consider first of all the premise of the rule: intuitively the higher the belief in the truth of the premise is, the higher the belief in the truth of the consequence should be. To say it in other words, the belief in the truth of the consequence should be equal to the belief in the truth of the premise, if there are no reasons to decrease it.
On the other hand, the belief in the falsity of the premise attracts the belief in the consequence towards the *U* point, since, if the premise is false, the rule is not relevant for the individual and therefore nothing can be said about the consequence.

We can model therefore the effect of $bt_{prop(x)}$, $bf_{prop(x)}$ in the following way. An attractor is positioned on the belief plane in correspondence with the point $A1=(bt_{prop(x)},0)$ and another attractor is positioned in correspondence with the point $U=(0, 0)$. The value of the elastic constant of the spring fixed in *A1* is $bt_{prop(x)}$, whereas the value of the elastic constant of the spring fixed in *U* is $bf_{prop(x)}$. In this way we obtain a first point of equilibrium: $E1 = (bt^2_{prop(x)} / bt_{prop(x)}*bf_{prop(x)}, 0)$.
E1 is in an intermediate position between *A1* and *U* and is closer to either of the two point depending on the ratio between $bt_{prop(x)}$, and $bf_{prop(x)}$.

In the following we will use the short notation (x_E, y_E) to denote the values of the coordinates of a point E.

Now we have to start from E1 and take into account $bels_E(applicable(R, x))$ and $att_E(R, x)$ in order to determine the final equilibrium point for $bels_E(cons(x))$.
Intuitively if there is a full belief in the applicability of the rule (there is full belief that x is a normal individual rather than an exception) the resulting point should be *E1* itself. Otherwise, if the belief in the applicability of the rule is not full (that is, $bt_{app(R,x)} < 1$) our belief in the truth of the consequence should decrease. This fact can be represented by a decrease of the first coordinate of point *E1* proportional to the lack of belief in the truth of $bt_{app(R,x)}$: we obtain therefore a point
$E2 = (x_E * bt_{app(R,x)}, 0)$.

The *E2* point would coincide with the final equilibrium point *EF* if we had no reason to believe that x is actually an exception, that is if $bf_{app(R,x)}$ were 0. Otherwise, the belief in the falsity of the applicability determines the presence of a further attractor. The position (and therefore the role) of this attractor depends on the actual value of the reasoning attitude. If the reasoning attitude is conservative we don't want to make any assertion about the consequence in the case x is an exception. Therefore, the attractor in this case is positioned in the point U=(0, 0). On the other hand, if the reasoning attitude is evolutive, we want to explicitly state that, in the case x is an exception, the consequent is false. Therefore the attractor should be in some position $A2$=(0, y_{A2}) with $y_{A2} > 0$. An intuitive consideration leads to state that $y_{A2} = bt_{prop(x)} * bf_{app(R,x)}$. In fact, as the belief in the truth of the consequence should not exceed the belief in the truth of the premise when the individual is normal, in the same way the belief in the falsity of the consequence when the individual is an exception should not exceed the belief in the truth of the premise (recall that if the premise is not believed, nothing should be said). Moreover, if $bf_{app(R,x)}$ is not full, the belief in the falsity of the consequence should be decreased accordingly.
Therefore, finally a spring is fixed to *E2*, with elastic constant $bt_{app(R,x)}$, and another spring, with elastic constant $bf_{app(R,x)}$, is fixed either in U or in $A2$ depending on the actual reasoning attitude.
The equilibrium point *EF* between these two springs represents the final belief state of the consequence of the rule.
When a new piece of information is acquired, it may affect any of the parameters used to determine *EF* (including the reasoning attitude). In this case after the parameters have been adjusted, a new equilibrium point can be calculated: the displacement of this point represents the impact of the new piece of information on previous belief.

This formalism is able to support full nonmonotonicity in a natural way.
Recalling the secretary example, suppose you have a default rule *R1* = "IF today is a working day THEN the secretary will come". Suppose you are sure that today is a working day for your secretary (she is not on paid holidays, otherwise you would know it) and that you have an almost full belief that the rule is applicable. Therefore you have *bels*(*workingday*(*today*)) = (1, 0) and *bels*(*applicable*(*R1*, *today*)) = (0.99, 0).
The value of *bels*(*applicable*(*R1*, *today*)) can be explained as follows: there is no evidence supporting the fact that today should be an exception (so the belief in the falsity of the applicability is 0), however your belief in the truth of the applicability can not be full, since you can not be definitely sure that nothing prevents the secretary to come.
In this case the first belief state you obtain about secretary's coming is (0.99, 0), expressing a very strong conviction that she will arrive.
After a while, you may learn that she has a broken leg: the belief in the truth of the applicability goes to zero as well as the belief in the falsity of the applicability raises to 1, and your attitude in this case is evolutive (of

course the relation between a broken leg and these parameter modifications should be properly encoded in the reasoner's knowledge base). In this case the resulting belief state is (0, 1), expressing the definite conviction that the secretary will not come: this is clearly a form of TF-nonmonotonicity.
On the other hand, if you learn that the secretary was feverish, again you notice that the situation is exceptional, so that the belief in the truth of the applicability goes to zero whereas the belief in the falsity of the applicability raises to 1, but your reasoning attitude is conservative. In this case the resulting belief state is (0, 0), expressing total ignorance about secretary's coming: this is clearly a form of AU-nonmonotonicity.
It is worth to note that this is not the only possible way to model the impact of this information on previous beliefs. For instance, one could prefer a less drastic decrease of the belief in the truth of the applicability, possibly associated to an increase of the belief in the falsity of the applicability.
This would correspond to an alternative view, where the new information is conceived as a source of contradictory belief rather than of ignorance. We remark however that also this different view can be easily encompassed in our formalism, still yielding cognitively plausible results.

The simple example presented above concerns only some extreme reasoning cases. A detailed illustration of the behavior of the formalism referring to a sufficiently ample set of nonmonotonic reasoning cases is beyond the scope of this paper. We remark, however, that one of the most important features of the proposed formalism is its rich expressiveness, which allows the explicit and tailored representation of a very ample range of reasoning situations, which could be hardly captured by most existing formalisms.

References

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

[Baroni et al. 96] P. Baroni, G. Guida, S. Mussi, Modeling default reasoning through A-uncertainy, *Proc. IPMU 96, International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, Granada, E, 1996, 1197-1204

[Baroni et al. 97a] P. Baroni, G. Guida, S. Mussi, Full nonmonotonicity: a new perspective in defeasible reasoning, *Proc. of ESIT 97, European Symposium on Intelligent Techniques*, Bari, Italy, 1997, 58-62

[Baroni et al. 97b] P. Baroni, G. Guida, S. Mussi, From ignorance to uncertainty: a conceptual analysis, *Proc. of 4th WUPES 97, Workshop on Uncertainty Processing*, Prague, 1997, 1-10

[Benferhat et al. 95] S. Benferhat, D. Dubois and H. Prade, A local approach to reasoning under inconsistency in stratified knowledge bases, *Proc. ECSQARU 95 3rd European Conf. on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Fribourg, CH, 1995, 36-43

[Dubois et al. 96] D. Dubois, H. Prade and P. Smets, Representing partial ignorance, *IEEE Trans. on Systems, Man and Cybernetics*, vol. 26, n. 3, 1996, 361-377

[Dubois and Prade 97] D. Dubois and H. Prade, Qualitative possibility theory and its applications to reasoning and decision under uncertainty, *Proc. of ESIT 97, European Symposium on Intelligent Techniques*, Bari, Italy, 1997, 301-308

[Gardenfors 88] P. Gardenfors, Knowledge in Flux - Modeling the dynamics of epistemic states, Cambridge, MA, MIT Press, 1988

[Katsuno and Mendelzon 91] H. Katsuno and A. O. Mendelzon, On the difference between updating a knowledge base and revising it, *Proc. of 2nd International Conference on Principles of Knowledge Representation and Reasoning*, Cambridge, MA, 1991, 387-394

[Harper 77] W. L. Harper, Rational conceptual change, *PSA 1976*, Philosophy of Science Association, vol 2, 462-494

[Levi 77] I. Levi, Direct inference, *The journal of philosophy*, 74: 5-29

[McDermott and Doyle 80] D. McDermott and J. Doyle, Non-monotonic logic I, *Artificial Intelligence* 13, 1980, 41-72

[Moore 88] 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

[Reiter 80] R. Reiter, A logic for default reasoning, *Artificial Intelligence* 13, 1980, 81-132.