Possibilistic Handling of Uncertain Default Rules with Applications to Persistence Modeling and Fuzzy Default Reasoning

F. Dupin de Saint-Cyr and **H. Prade**
florence.bannay@irit.fr henri.prade@irit.fr
IRIT,
Université Paul Sabatier,
Toulouse (France)

Abstract

Default rules express concise pieces of knowledge having implicit exceptions, which is appropriate for reasoning under incomplete information. Specific rules that explicitly refer to exceptions of more general rules can then be handled in this non-monotonic setting. However, there is no assessment of the certainty with which the conclusion of a default rule holds when it applies. We propose a formalism in which uncertain default rules can be expressed, but still preserving the distinction between the defeasibility and uncertainty semantics in a two steps processing. Possibility theory is used for representing both uncertainty and defeasibility. The approach is illustrated in persistence modeling and in fuzzy default reasoning problems.

Introduction

Reasoning under incomplete information by means of rules having exceptions, and reasoning under uncertainty are two important types of reasoning that artificial intelligence has studied at length and formalized in different ways in order to design inference systems able to draw conclusions from available information as it is. However, the joint handling of exceptions and uncertainty has received little attention. This is the topic of this paper.

Default rules are useful in order to express general behaviors concisely, without referring to exceptional cases. Moreover they only require general information to be fired, which agrees with the situations of incomplete information. In practice, reasoning from a set of exception-tolerant default rules in presence of incomplete knowledge first amounts to select default rules. The selected set of rules should focus on the current context describing the particular incomplete information situation that is considered, and then this set of rules can be applied to this information situation in order to draw plausible conclusions. When new information is available on the current situation, these conclusions may be revised at the light of more appropriate default rules. The selection problem is solved in practice by rank-ordering the default rules in such a way that the most specific rules whose conclusion may conflict with the conclusion of more general defaults, receive a higher level of priority (Pearl 1990), following the idea first proposed by (Touretzky 1984). Clearly, the level of (relative) priority of a particular rule depends on the whole set of default rules which are considered.

However, conclusions that we want to privilege in a given context may themselves be pervaded with uncertainty. Indeed, when a rule of the form "if p then q generally" is stated, no estimate of the certainty of having q true in context p is provided, even roughly. The status of being a default rule, is just a proviso for possible exceptional situations to which other rules in the knowledge base may refer. The priority level of a default rule in a set of such rules cannot be regarded as a kind of qualitative certainty level. In fact, it may happen that a specific rule provides default conclusions that are less certain than more general rules, or on the contrary strengthens the certainty of its conclusion. For instance, the rule "birds with large wings fly" is more certain than "birds fly", while one may consider that the rule "Antarctic birds fly" is less certain than "birds fly", assuming that there are many penguins (that do not fly) together with some more flying sea birds in Antarctic. But, even if it is less certain, the specific rule that fits the particular context of incomplete information at hand, is the right one to use. More generally, the uncertainty attached to a rule is not necessarily related to its specificity level.

As already said, reasoning with default rules and under uncertainty are two important research trends that have been developed quite independently from each other in AI. They indeed address two distinct problems, respectively using symbolic and numerical approaches in general. Default rules are concise pieces of knowledge (by omitting some propositional variables only appropriate for describing exceptional situations), which are especially useful in case of incomplete information. Reasoning under uncertainty normally requires the complete specification of all relevant variables. However, handling uncertainty, at least qualitatively, in a given incomplete information context is a need in various situations. For example, high level descriptions of dynamical systems often requires both the use of default rules expressing persistence for the sake of concise representation and the processing of uncertainty due to the limitation of the available information.

This paper outlines a joint handling of defaults and uncertainty in qualitative possibility theory, where there already exists separate treatments for them (although other uncertainty representation settings could be considered). Separate refreshers on the possibilistic handling of uncertainty and defaults are given in Annex A and B while the problem raised by their joint processing is first discussed. Then three methods for default reasoning are presented before integrating uncertainty in these methods. The approach is illustrated on two different problems: persistence handling in dynamical environments (persistence rules are by nature default rules), and reasoning with fuzzy defaults such as "young birds cannot fly" understood as "the younger the bird, the more certain it cannot fly". Lastly links with related works are discussed.

Uncertain default rules

We assume a representation language $\mathscr{L}$ built on a set of propositional variables $\mathscr{V}$. The set of interpretations associated with this language is denoted by Ω. An interpretation $\omega \in \Omega$ represents a state of the system under study. In order to have a more expressive representation formalism, we now introduce the notion of uncertain default rule.

Definition 1 *An* uncertain default rule *is a pair* $(a \rightsquigarrow b, \quad \alpha)$ *where* a *and* b *are propositional formulas of* $\mathscr{L}$*, and* α *is the certainty level of the rule, the symbol* $\rightsquigarrow$ *is a non classical connective encoding a non-monotonic consequence relation between* a *and* b*.*

In the following, for simplicity, we use for certainty levels the real interval scale $[0, 1]$. However a qualitative scale could be used, since only the complete preorder between the levels is meaningful. The intuitive meaning of $(a \rightsquigarrow b, \quad \alpha)$ is "by default" if a is true then b has a certainty level at least equal to α. For instance, let b, f, y stand for "bird", "fly", "young". Then $(b \rightsquigarrow f, \alpha_1)$ means that "a bird generally flies" with certainty α_1. It is a default rule since it admits exceptions mentioned in other rules: for instance, $(b \wedge y \rightsquigarrow \neg f, \alpha_2)$: "young birds generally do not fly". But it is also an uncertain rule since when all we know is that we are in presence of a bird, the certainty level α_1 is attached to the provisional conclusion that it flies. Thus, the α's provide an additional information with respect to the default rule. Moreover, the more specific rule about young birds is again an uncertain default rule since some ones may fly. Note that, in general, as suggested by the above example where there is no clear inequality constraint between α_1 and α_2, there is no relation between the certainty level associated with a default rule and the certainty level associated with a more specific rule. In particular, it would be wrong to assume that the more specific rule always provides a more certain conclusion.

The core of the treatment of uncertain default rules proposed in this paper is based on the idea of translating them into a set of uncertain (non defeasible) rules. This can be done in different ways, depending on how default rules are handled and what kind of uncertainty representation framework is used. In the following, uncertainty is modeled in the qualitative setting of possibility theory (Dubois & Prade 1988; 1998) and possibilistic logic (see Annex A). Indeed, this agrees with the qualitative nature of default rules. We present several approaches for dealing with default rules.

Roughly speaking, default reasoning amounts to apply a set of default rules Δ to a factual propositional base FC describing a context at hand.

- A first idea is then to select the subset of the rules of Δ that is appropriate for the factual context FC under consideration and remove the other rules, and to turn the selected rules into classical propositional rules. As we shall see, this idea is not entirely satisfactory, because many information are lost (due to a drowning effect that leads to a problem of inheritance blocking).
- A method that copes with this difficulty, still relying on the context, named contextual entailment, has been proposed in (Benferhat & Dupin de Saint-Cyr 1996). This method may be too cautious and has no known efficient algorithmic counterpart. Based on this idea, we propose a contextual rational entailment that is less cautious. The problem is that the context should be given at once before each deduction so for each change of context a compilation of the default base must be done.
- Another approach that we also explore further in the following is to rewrite each default rule into a propositional rule by making its condition part more precise (by explicitly naming the exceptions mentioned in the default base). This approach is more satisfactory with respect to the problems encountered by the previous methods. However, to be able to deal with incomplete information, this set of rewritten rules should be augmented with an additional set of rules that depends on the context and states in what respect this context is not exceptional. These additional rules aim at completing the factual context in order to be able to apply the rewritten rules.

In the next section, we discuss in detail the three above alternatives for handling default rules before presenting the treatment of *uncertain* default rules in a new section.

Handling default rules

A normative approach to default reasoning is provided by System P (Kraus, Lehmann, & Magidor 1990) that defines a "preferential" inference relation obeying one axiom:

Reflexivity: $a \mid\!\sim a$

and five inference postulates:

Left logical equivalence if $\vdash a \leftrightarrow b$ and $a \mid\!\sim c$ then $b \mid\!\sim c$
Right weakening: if $a \vdash b$ and $c \mid\!\sim a$ then $c \mid\!\sim b$
Cut: if $a \wedge b \mid\!\sim c$ and $a \mid\!\sim b$ then $a \mid\!\sim c$
Cautious monotony: if $a \mid\!\sim b$ and $a \mid\!\sim c$ then $a \wedge b \mid\!\sim c$
Or: if $a \mid\!\sim c$ and $b \mid\!\sim c$ then $a \vee b \mid\!\sim c$

The set of conclusions that one can obtain by using a "preferential" entailment is usually regarded as the minimal set of conclusions that any reasonable non-monotonic consequence relation for default reasoning should generate.

Lehmann and Magidor (Lehmann & Magidor 1992) have defined a more adventurous consequence relation (which allows to draw more conclusions), named "rational closure entailment", which is a "preferential" relation that also obeys a Rational Monotony rule.

Rational monotony: if $a \mid\!\sim b$ and $a \mid\!\not\sim \neg c$ then $a \wedge c \mid\!\sim b$

Another landmark work in the treatment of default rules is the system Z (Pearl 1990) for stratifying a set of default rules according to their specificity (see Annex B). Given a set of default rules Δ, System Z stratification partitions it into subsets $\Delta_0, \ldots, \Delta_n$, where rules in Δ_i have priority over the ones in Δ_j if $i > j$. These priorities reflect specificity levels since specific rules get higher priority. System Z is a rational closure entailment. Besides rational closure entailment and System Z entailment have been shown to be equivalent to a possibilistic treatment of default rules briefly recalled in Annex B (Benferhat, Dubois, & Prade 1997).

In the following, we consider a set Δ of default rules, together with a propositional factual base FC describing all the available information about the context. Threw methods for drawing plausible conclusions from FC using Δ are presented below.

The factual base FC is supposed to be consistent. Moreover, we also assume that the set Δ is consistent. This means that we cannot encounter a situation where it is not possible to compute the specificity levels of Δ. This consistency condition is equivalent to the existence of a possibility measure Π satisfying the set of constraints $\Pi(a \wedge b) > \Pi(a \wedge \neg b)$ associated with each default in the base Δ, leading to a possibilistic logic handling of the specificity levels (see Annex B and A). This is the basis of the first method presented now.

Method 1: Possibilistic selection of the rules in a given context

Given a set Δ of default rules and a factual base FC, the possibilistic approach proceeds in two main steps:

• Associate to each default rule $r = a \rightsquigarrow b \in \Delta$ its specificity level $d(r) = \frac{Z(r)+1}{n+2}$, where $Z(r)$ is the rank of the stratum of the rule r once the system Z procedure has been applied (see Annex B).

Let D_π be the possibilistic knowledge base s.t. $D_\pi = \{(a_i \rightarrow b_i, d(a_i \rightsquigarrow b_i)) | a_i \rightsquigarrow b_i \in \Delta\}$ where $\rightarrow$ is the classical material implication.

Besides, each proposition φ in FC is encoded in a possibilistic format: $(\varphi, 1)$, which amounts to consider the factual information as totally certain.

Then compute the inconsistency level $Inc(D_\pi \cup FC)$ (see Annex B).

• Applying default rules in Δ to FC amounts to reason with the formulas in $D_\pi \cup FC$ that are above $Inc(D_\pi \cup FC)$. Hence, remove each formula $(a_i \rightarrow b_i, \sigma_i)$ from D_π such that $\sigma_i \leq Inc(D_\pi \cup FC)$.

The remaining rule base is: $D = \{a_i \rightarrow b_i | a_i \rightsquigarrow b_i \in \Delta$ and $d(a_i \rightsquigarrow b_i) > Inc(D_\pi \cup FC)\}$.

Definition 2 ***(rational closure entailment and possibilistic consequence)***
A formula ψ is said to be a rational closure consequence of Δ given a factual context FC, denoted by $FC \mid\!\sim_{RC,\Delta} \psi$, if and only if ψ is a classical consequence of $FC \cup D$, where $D = \{a_i \rightarrow b_i | a_i \rightsquigarrow b_i \in \Delta$ and $d(a_i \rightsquigarrow b_i) > Inc(D_\pi \cup FC)\}$:

$$FC \mid\!\sim_{RC,\Delta} \psi \quad \textit{iff} \quad FC \cup D \vdash \psi$$

Example 1 *We consider the following default base, describing the fact that birds generally fly and young birds generally do not fly:*
$\varphi_1 : \quad b \rightsquigarrow f$
$\varphi_2 : \quad b \wedge y \rightsquigarrow \neg f$
System Z gives: $\Delta_0 = \{\varphi_1\}$, $\Delta_1 = \{\varphi_2\}$. The specificity levels associated to the rules of Δ_0 and Δ_1 are 1/3 and 2/3 respectively. Let D_π be the possibilistic knowledge base associated to Δ :
φ_1*:* $\quad b \rightarrow f, \quad$ *1/3*
φ_2*:* $\quad b \wedge y \rightarrow \neg f, \quad$ *2/3*
Let $FC = \{(b \wedge y, 1)\}$, meaning that we are considering a young bird. Then $Inc(D_\pi \cup FC) = 1/3$ since $D_\pi \cup FC \vdash_\pi (F, 1/3)$ from rule φ_1, we have also $D_\pi \cup FC \vdash_\pi (\neg f, 2/3)$ from rule φ_2, hence $D_\pi \cup FC \vdash_\pi (\bot, 1/3)$ (applying the resolution rule of possibilistic logic, where $\vdash_\pi$ denotes the possibilistic entailment, see Annex A). So, the final base D only contains the formula $(b \wedge y \rightarrow \neg f)$.
So $FC \cup D_\pi \vdash_\pi (\neg f, 2/3)$. One concludes that a young bird is unable to fly.

However, this method suffers from the "drowning effect". For instance, if we had the rule "birds generally have legs (l)" then it will not be possible to conclude that "young birds generally have legs" since the rule $b \rightsquigarrow l$ will have 1/3 as specificity level.

Method 2: Contextual rational entailment

The previous approach is not entirely satisfactory, since it faces the "blocking property inheritance" problem. A new approach, based on an idea presented in (Benferhat & Dupin de Saint-Cyr 1996), may remedy this drawback. In this work, the authors studied under which conditions they can infer b from $a \wedge c$, given a rule "generally, a's are b's". Classical logic always answers that $a \wedge c$ infers b (monotony property). Default reasoning should answer like classical logic except when the c's are exceptions of the rule. Hence, it is important to check if $a \wedge c$ is an exception of the rule "generally, a's are b's".

Benferhat *et al.* used System P in order to answer this latter question since System P never draws undesirable conclusions. Here, Benferhat *et al.* approach is extended by using "rational closure" inference relations instead of "preferential" inference relations. It is based on the identification of rules having exceptions in a given context (the definition is similar to the one given by Benferhat *et al.*, but uses rational closure instead of preferential closure):

Definition 3 *Let $FC \in \mathscr{L}$ be a classical propositional formula considered as the current context and Δ be a set of default rules. A default rule $a_i \rightsquigarrow b_i$ of Δ* has an exception *in FC if and only if one of the two following conditions is satisfied:*

1. *$a_i \wedge FC \wedge b_i$ is inconsistent, or*
2. *$a_i \wedge FC \mid\!\sim_{RC,\Delta} \neg b_i$*

where $\mid\!\sim_{RC,\Delta}$ *is an inference relation based on the rational closure of the set obtained by interpreting each default* $a_i \rightsquigarrow b_i$ *of* Δ *as* $a_i \mid\!\sim b_i$.

For each rule $a_i \rightsquigarrow b_i$ of Δ we can check if it is exceptional or not in the given context. If not, we change it into a strict rule $a_i \rightarrow b_i$, else we delete it. Let $\Sigma_{FC} = \{a_i \rightarrow b_i | a_i \rightsquigarrow b_i \in \Delta$ has no exception in $FC\}$.

Definition 4 (contextual rational entailment) *A formula* ψ *is said to be a CRE-consequence (C for context and R for rational) of* Δ *given a factual context FC, denoted by* $FC \mid\!\sim_{CRE,\Delta} \psi$*, if and only if* ψ *is a classical consequence of* $\Sigma_{FC} \cup \{FC\}$*:*

$$FC \mid\!\sim_{CRE,\Delta} \psi \quad \textit{iff} \quad \Sigma_{FC} \cup \{FC\} \vdash \psi$$

Using the same reasoning as Benferhat *et al.*, we can argue that $\mid\!\sim_{CRE,\Delta}$ is non-monotonic, since increasing the context reduces the set of rules that have no exception.

Proposition 1 *If* $FC \vdash FC'$ *then* $\Sigma_{FC} \subseteq \Sigma_{FC'}$.

It can be established that $\mid\!\sim_{CRE,\Delta}$ is "rational". It means that the conclusions obtained by the first method of this paper can be obtained by contextual rational entailment as well.

Proposition 2 $\forall\Delta$, $\mid\!\sim_{RC,\Delta} \subseteq \mid\!\sim_{CRE,\Delta}$.

Proof : Indeed, if a rule $a_i \rightsquigarrow b_i$ has exceptions in a given context FC, then it means that $\{a_i\} \cup FC \mid\!\sim_{RC,\Delta} \neg b_i$. So this rule has a specificity level smaller or equal to the level of inconsistency of $D_\pi \cup FC$ (where D_π is the possibilistic knowledge base associated to Δ, $D_\pi = \{(a_i \rightarrow b_i, d(a_i \rightsquigarrow b_i)) | a_i \rightsquigarrow b_i \in \Delta\}$). Hence, a rule having exception in a given context cannot be used by $\mid\!\sim_{RC,\Delta}$. Since we translate every default rule that has no exception into a material implication, and use classical entailment on the set obtained, we use at least all rules that are kept by $\mid\!\sim_{RC,\Delta}$. So, this system can at least draw every conclusion obtained by $\mid\!\sim_{RC,\Delta}$. $\square$

Corollary 2.1 $\mid\!\sim_{CRE,\Delta}$ *verifies* Reflexivity*,* Left logical equivalence*,* Right weakening*,*Or*,* Cautious monotony*,* Cut *and* Rational monotony*.*

Proof : This property is a direct consequence of the previous property. Since $\mid\!\sim_{RC,\Delta}$ is a rational entailment relation, so does $\mid\!\sim_{CRE}$. $\square$

Moreover, contextual rational entailment can obtain more conclusions than rational entailment, namely it does not suffer from the drowning effect:

Example 2 *Let us consider the following default base* Δ*:*
$\varphi_1 : \quad b \rightsquigarrow f$
$\varphi_2 : \quad b \wedge y \rightsquigarrow \neg f$
$\varphi_3 : \quad b \rightsquigarrow l$
We have $\Sigma_{b \wedge y} = \{b \wedge y \rightarrow \neg f, b \rightarrow l\}$*, so* $b \wedge y \mid\!\sim_{CRE,\Delta} l$.

Note that some scholars have criticized "rational closure" as allowing to deduce undesirable results concerning examples where ambiguity should be preserved. Namely, let us consider the following variant of "Nixon Diamond" example:

Example 3 *Let* Δ *be a default base representing that "Quakers normally are pacifists", "Quakers are generally Americans", "Americans normally like base-ball", "Quakers generally do not like base-ball" and "Republicans are generally not pacificists":*
$\varphi_1 : \quad Q \rightsquigarrow p$
$\varphi_2 : \quad Q \rightsquigarrow A$
$\varphi_3 : \quad A \rightsquigarrow b$
$\varphi_4 : \quad Q \rightsquigarrow \neg b$
$\varphi_5 : \quad R \rightsquigarrow \neg p$
Then $Q \wedge R \mid\!\sim_{CRE,\Delta} p \wedge A \wedge \neg b$*,*
since $\Sigma_{Q \wedge R} = \{Q \rightarrow p, Q \rightarrow A, Q \rightarrow \neg b\}$.

The result "pacifist" can be debatable (note that the two other conclusions are desirable). One can argue that it would be better to not conclude anything about the plausibility of having p true or false. In our opinion, it is not the fault of "rational closure" but, it is rather due to the ambiguity of the example. In this example, there is only one piece of information about "Republicans". Indeed, here, "Republican" can be considered as a general property, as general as "American". So its specificity level is as low as the American property. Meanwhile, if we learn that Republicans are Americans that have a given particularity (if they were only Americans, then the two words would be synonymous) then the conclusions would change. Hence as discussed in (Benferhat, Dubois, & Prade 1998), it is not rational monotony that leads to undesirable conclusions, but it is rather a lack of information in the knowledge base. A too adventurous conclusion is only caused by missing pieces of knowledge that the system cannot guess on its own, and these pieces can be always added to the default base (without leading to inconsistency) in order to get the desirable conclusion as prooved in (Benferhat, Dubois, & Prade 1998).

To conclude on this approach, it gives better results than the first one, but it has a drawback: the computation depends on the context, it means that a computation of the set of rules having no exception should be done before any new contextual deduction.

Method 3: Rewriting the rules by expliciting their exceptions

The first method handles default reasoning by deleting all the rules under a level of inconsistency in a given context. It has the "drowning effect" as a drawback: rules that are not directly involved in the inconsistency may be deleted, while the second method correctly addresses this problem. The problem of the second method is that the computation depends on the context: before each deduction a computation of the rules that are kept must be done. Indeed, this computation may be heavy since the whole set of default rules Δ should be examined with respect to any new context. Hence, we propose another method that somewhat handles these drawbacks. The idea is to transform the default rules independently of any context into a set of non-defeasible rules. The idea is to generate automatically from Δ a set of non-defeasible rules D in which the condition parts explicitly state that we are not in an exceptional context to which other default rules refer. In the same time, strict rules called "completion rules" stating that we are not in an exceptional situation are added to a new set CR. The use of these completion rules is motivated by the need of reasoning in presence of incomplete information: the completion allows us to still be able to apply the modified rules which now have a more precise condition part.

Note that the rules in CR will only be used if they are consistent with the context described in FC (takind D into account). Hence, it only requires to do a consistency test each time the context FC is changed.

Definition 5 (Explicit Rule and Completion Rule) *Let $\Delta = \{a_i \leadsto b_i\}_{i=1..k}$ be a set of default rules, and let us consider D its associated set of strict rules:*
$D = \{a_i \to b_i | a_i \leadsto b_i \in \Delta\}$.
For any given default rule $r = a \leadsto b$, we define:
The set of exceptions *in Δ to the rule r:*

$$E(a \leadsto b, \Delta) = \left\{a_i \;\middle|\; \begin{array}{l} a_i \leadsto b_i \in \Delta \textit{ and} \\ \{a_i \wedge a\} \cup D \textit{ consistent and} \\ \{b_i \wedge b\} \cup D \textit{ inconsistent} \end{array} \right\}$$

The explicit rule *associated with r is*

$$a \wedge \bigwedge_{x \in E(r,\Delta)} \neg x \to b$$

A completion rule *associated with r is of the form*

$$a \to \neg x \textit{ where } x \in E(r, \Delta)$$

We consider a set of default rules Δ which has been stratified by System Z into subsets $\Delta_0, \ldots, \Delta_n$, where rules of Δ_i are more specific than rules of Δ_j if $i > j$.

Rewriting Algorithm

{ Let i be the rank of the current stratum. Let Δ' and Δ'_i be the set of all rules already rewritten from Δ and Δ_i respectively. Let CR be the set of completion rules.}

$i := n - 1; CR := \varnothing$
while $i \geq 0$ **do repeat**
  $\Delta'_i := \varnothing$;
  while $\Delta_i \neq \varnothing$ **do repeat**
    for each rule $r = a \leadsto b \in \Delta_i$ **do**:
      remove r from Δ_i {r is being considered }
      $E(r, \Delta') := \varnothing$;
      while $\exists a' \leadsto b' \in \Delta'$
      such that $\left\{ \begin{array}{l} \{a \wedge a'\} \cup D \text{ consistent} \\ \text{and } \{b' \wedge b\} \cup D \text{ inconsistent} \end{array} \right.$
      where $D = \{a_i \to b_i | a_i \leadsto b_i \in \Delta'\}$.
      do add a' to $E(r, \Delta')$;
        add $a \to \neg a'$ to CR
      {r has no more exception}
      add $a \wedge \bigwedge_{x \in E(r,\Delta')} \neg x \to b$ to Δ'_i
  add Δ'_i to Δ';
  $i := i - 1$ {examine the previous stratum }

Note that the rules of the last stratum n do not admit exceptions with respect to the knowledge base Δ since they are the most specific ones. This is why they are directly transformed into strict classical rules. Then the algorithm begins with the rules of the stratum $n - 1$. The stratum $n - 1$ contains rules that admit exceptions only because of rules in the last stratum. More generally, a stratum i contains rules that admit exceptions only because of rules in strata with rank greater or equal to $i + 1$. More precisely for each rule in a given strata, all its exceptions (coming from strata with a greater rank) are computed in order to rewrite this rule by explicitly stating that the exceptional situations are excluded in its condition part. Moreover, completion rules are added for each exceptional case found; as already said, completion rules are useful to state in what respect the current context is not exceptional. For instance, if b is the only exception to the rule $a \leadsto c$, then the rule is modified into $a \wedge \neg b \to c$, and the completion rule, associated with it, has the form $a \to \neg b$. This completion rule will only be used if it is consistent with the current context and the set of rewritten strict rules.

Proposition 3 *This algorithm terminates.*

Proof : The algorithm examines each rule of each stratum. For a rule of a stratum Δ_i, the algorithm executes at most two consistency tests with each rule of strata of rank greater or equal to $i + 1$. Since each stratum is finite, the algorithm terminates. $\square$

Proposition 4 *The set $D = \{a_i \to b_i | a_i \leadsto b_i \in \Delta'\}$ of strict rules associated to the set Δ' of modified rules obtained by this algorithm is consistent.*

Proof : At the beginning D is consistent since it is built on the set Δ_n of rules tolerated by the set $\Delta \backslash (\Delta_0 \cup \ldots \Delta_{n-1}) = \Delta_n$. It means that it exists $\omega_0 \models a_{n1} \wedge b_{n1}$ where $a_{n1} \leadsto b_{n1}$ is the first rule of Δ_n and satisfying every other rules of Δ_n. Hence $\omega_0 \models a_{n1} \wedge b_{n1} \wedge \{(\neg a_{ni} \vee b_{ni}) \mid a_{ni} \leadsto b_{ni} \in \Delta_n\}$.

At each step, a rule is added to D only if its conclusion is consistent with every conclusion of a rule of D. For a rule $r = a \leadsto b$ from a stratum Δ_i, if it exists a rule $a' \to b'$ in D such that $b' \wedge b \wedge D \vdash \bot$, then r is replaced by $a \wedge \neg a' \leadsto b$. Note that $a \wedge \neg a'$ is consistent since, by construction, every rule of Δ_{i+1} is tolerated by r, it means that it exists $\omega \models a \wedge b \wedge D \wedge (\neg a' \vee b')$, i.e, $\omega \models a \wedge \neg a' \wedge b$. r modified by specifying all its exceptions is added to D only when there is no more rule in Δ_{i+1} whose conclusion is inconsistent with b. So D remains consistent. $\square$

Note that each rule of the initial default knowledge base is present, modified or not, in the resulting rule base. So, there is no loss of information as with the previous method. Moreover the addition of rules $a \leadsto \neg a'$ and $a \wedge \neg a' \leadsto b$ in situations such that $a \leadsto b$ and $a \wedge a' \leadsto \neg b$ hold, is in full agreement with postulates of rational closure (Lehmann & Magidor 1992). Indeed, from $a \wedge c \mathrel{|\!\sim} \neg b$, we have by consistency, $a \wedge c \mathrel{|\!\not\sim} b$. Then from $a \mathrel{|\!\sim} b$ and $a \wedge c \mathrel{|\!\not\sim} b$, we get $a \mathrel{|\!\sim} \neg c$ by rational monotony. Moreover from this result and $a \mathrel{|\!\sim} b$ we obtain $a \wedge \neg c \mathrel{|\!\sim} b$ by cautious monotony.

Definition 6 (Rewriting entailment) *A formula ψ is said to be a $RW-consequence$ (RW for rewriting) of Δ given a factual context FC, denoted by $FC \mathrel{|\!\sim}_{RW,\Delta} \psi$, if and only if ψ is a classical consequence of D and CR which are respectively the sets of strict rules and completion rules obtained after the rewriting algorithm:*

$$FC \mathrel{|\!\sim}_{RW,\Delta} \psi \quad \textit{iff} \quad FC \cup D \cup CR' \vdash \psi$$

where CR' is a subset of CR consistent with $FC \cup D$.

Proposition 5 $\forall\Delta, \quad \mid\!\sim_{RC,\Delta} \subseteq \mid\!\sim_{RW,\Delta}$

Proof : As previously noticed, the addition of rules $a \wedge \neg a' \rightsquigarrow b$ in situations such that $a \rightsquigarrow b$ and $a \wedge a' \rightsquigarrow \neg b$ hold, is in full agreement with postulates of rational closure. Moreover the consistency of D computed from Δ (prop 4) allows us to transform $\rightsquigarrow$ into $\rightarrow$. More formally, it gives: $FC \mid\!\sim_{RC,\Delta} D$.

The same reasoning can be done for the completion rules: $a \rightsquigarrow \neg a'$. It leads to $FC \mid\!\sim_{RC,\Delta} CR$. Hence, $FC \mid\!\sim_{RC,\Delta} D \cup CR$, by right weakening, we get $FC \mid\!\sim_{RC,\Delta} D \cup CR'$ where $CR' \subseteq CR$.

So, if $FC \mid\!\sim_{RC,\Delta} \psi$ then, by cautious monotony, $FC \cup D \cup CR' \mid\!\sim_{RC,\Delta} \psi$, i.e., $FC \mid\!\sim_{RW,\Delta} \psi$. □

Corollary 5.1 $\mid\!\sim_{RW,\Delta}$ *verifies* Reflexivity*,* Left logical equivalence*,* Right weakening*,*Or*,* Cautious monotony*,* Cut *and* Rational monotony*.*

Example 4 *Now we can rewrite the rule of example 2 by describing explicitly their exceptions starting from the last stratum. It gives the following knowledge base D:*

$\varphi_2: \quad b \wedge y \rightarrow \neg f$
$\varphi_1: \quad b \wedge \neg y \rightarrow f$
$\varphi_3: \quad b \rightarrow l$

There is only one completion rule: $CR = \{b \rightarrow \neg y\}$*, hence, in the context* $FC = \{b\}$*, the completion rule is consistent, so it allows us to deduce* $f \wedge l$*. In the context* $FC = \{b \wedge y\}$ *we cannot add the completion rule since it is inconsistent with* FC *so we can conclude* $\neg f \wedge l$*.*

It is now interesting to check if method 3 retrieves all the conclusions of method 2. We can establish that it is the case.

Proposition 6 $\forall\Delta$, $\mid\!\sim_{CRE,\Delta} \subseteq \mid\!\sim_{RW,\Delta}$.

Proof : $\mid\!\sim_{CRE,\Delta}$ is based on the use of classical entailment from the set $\Sigma_{FC} \cup FC$ in a given context FC, meanwhile $\mid\!\sim_{RW,\Delta}$ uses classical entailment from the set $D \cup CR'$ where D is the set of rewritten rules from Δ and CR' is a subset of completion rules that is consistent with $FC \cup D$. Hence, in order to compare the two entailments it is enough to compare the two sets Σ_{FC} and $D \cup CR'$.

Let us consider a given rule $a_i \rightsquigarrow b_i$ of the initial default base Δ. Let $E(a_i \rightsquigarrow b_i, \Delta)$ its set of exceptions in Δ.

- if $\{a_i \wedge b_i\} \cup FC$ is consistent then
  - if $\{a_i\} \cup FC \not\mid\!\sim_{RC,\Delta} \neg b_i$ then $a_i \rightarrow b_i$ will be present in Σ_{FC}. Moreover it means that for any exception a' of the initial rule, $FC \cup \{a_i\} \not\vdash a'$. Indeed assume that $FC \cup \{a_i\} \vdash a'$ and a' being an exception, we have $\{a' \wedge a_i\} \cup D \vdash \neg b_i$. That would imply that $FC \cup \{a_i\} \cup D \vdash \neg b_i$, which is in contradiction with our starting hypothesis. Hence, finally, $FC \cup \{a_i\}$ is consistent with every completion rule associated to $a_i \rightsquigarrow b_i$, so also consistent with the rewritten condition part of this rule. Hence, the conclusion b_i can also be drawn by method 3.
  - else $\{a_i\} \cup FC \mid\!\sim_{RC,\Delta} \neg b_i$ so $a_i \rightarrow b_i \notin \Sigma_{FC}$. Note that it implies that $\exists a' \in E(a_i \rightsquigarrow b_i, \Delta)$ such that $FC \cup \{a_i\} \vdash a'$ (by a reasoning similar to the above one). Hence there is a completion rule, namely, $a_i \rightarrow \neg a'$, belonging to the set of completion rules associated to $a_i \rightsquigarrow b_i$ that is not consistent with FC. Hence the initial rule $a_i \rightsquigarrow b_i$ whose condition part has been rewritten will neither fired in method 3.
- else $\{a_i \wedge b_i\} \cup FC$ is inconsistent. In this case, for method 2, $a_i \rightarrow b_i$ will not be present in Σ_{FC}. For method 3, there are two cases
  - either $\{\bigwedge_{x \in E(a_i \rightsquigarrow b_i, \Delta)} \neg x\} \cup FC$ is inconsistent. It means that the explicit rule $a_i \wedge \bigwedge_{x \in E(a_i \rightsquigarrow b_i, \Delta)} \neg x \rightarrow b_i$ could not be used, leading to the same result as in method 2.
  - or $\bigwedge_{x \in E(a_i \rightsquigarrow b_i, \Delta)} \neg x$ is consistent with FC, it means that the rule $a_i \wedge \bigwedge_{x \in E(a_i \rightsquigarrow b_i, \Delta)} \neg x \rightarrow b_i$ is inconsistent with FC. Then the third method will face an inconsistency in $FC \cup D$, hence, every proposition and its negation will belong to the set of possible conclusions.

□

The last part of the proof has also pointed out that the method 3 based on the rewriting of the default rules is only protected against existing exceptions that can be discovered by compiling the default base. In case the context FC corresponds to a new exception to which Δ does not refer, method 3 cannot conclude anything meaningful (as it is the case of method 1), while method 2 would lead to non trivial conclusions. However, we may assume that the default rule base refers to any exception that can be encountered in practice. Otherwise, it would mean that there is some missing information in Δ.

Handling uncertain default rules

Let $U\Delta$ be a set of uncertain default rules of the form $(a \rightsquigarrow b, \alpha)$. In this paper, two types of levels are involved: namely levels encoding specificity and levels of certainty. Although in the first approach specificity levels are handled by possibilistic logic in the same manner as the certainty levels will be processed in this section, the two types of levels should not be confused and the inference process uses the two scales separately. In fact in each of the three above methods for handling default rules, specificity is used to determine which rules are appropriate in the current context. Then in the resulting base containing strict rules only (except for the second method), the certainty levels should be taken into account in agreement with possibility theory in order to draw plausible conclusions with their certainty levels.

Using the first method, an uncertain default rule $(a \rightsquigarrow b, \alpha)$ is considered under the form $(a \rightsquigarrow b)$ and on the basis of its specificity level is selected or not with respect to the current context. If the rule is selected, it is then rewritten into the form $(a \rightarrow b, \alpha)$.

Using the second method, an uncertain default rule $(a \rightsquigarrow b, \alpha)$ is also considered under the form $(a \rightsquigarrow b)$. If it is not exceptional in the given context according to rational closure then it is changed into a strict rule as in the previous method. Otherwise it is deleted. If the rule is selected, it is then rewritten into the form $(a \rightarrow b, \alpha)$.

For the third method, an uncertain default rule $(a \rightsquigarrow b, \alpha)$ is considered under the form $(a \rightsquigarrow b)$ and on the basis of its specificity, its set of exceptions is computed, say $a'_1, \ldots, a'_k$. Then this rule is rewritten into the form $(a \wedge \neg a'_1 \wedge \ldots \wedge \neg a'_k \rightarrow b, \gamma)$. Moreover, k completion rules are created and added to the set of completion rules CR, namely, $(a \rightarrow \neg a'_1, \delta_1), \ldots (a \rightarrow \neg a'_k, \delta_k)$. Remind that each rule in CR is used only if it is consistent with the context and the set of rewritten rules.

This can be justified in the following way. On the one hand, as already said, the addition of a rules $a \rightsquigarrow \neg a'$ and $a \wedge \neg a' \rightsquigarrow b$ in situations such that $a \rightsquigarrow b$ and $a \wedge a' \rightsquigarrow \neg b$ hold, is in full agreement with postulates of rational closure (Lehmann & Magidor 1992). On the other hand, we have to assess the certainty levels γ and $\delta_1, \ldots, \delta_k$ associated with the added default rules. This can be done easily by interpreting the certainty levels of the default rules we start with, as lower bounds of conditional necessity, namely $N(b|a) > \alpha$ and $N(\neg b|a \wedge a'_i) > \beta_i$, and noticing[1] that when the bounds are strictly positive, they coincide with the necessity of the corresponding material implication. Then from $N(\neg a \vee b) > \alpha$ and $N(\neg a \vee \neg a'_i \vee \neg b) > \beta_i$, applying possibilistic resolution rule (see Annex A), we get $N(\neg a \vee \neg a'_i) > min(\alpha, \beta_i)$. Then we can take $\delta_i = min(\alpha, \beta_i)$. Moreover, the rule $a \wedge \neg a' \rightarrow b$ is at least as certain as $a \rightarrow b$ by monotonicity of necessity measure (see Annex A), so we can take $\gamma = \alpha$.

Example 5 *If we consider the following uncertain default base $U\Delta$, describing the fact that birds generally fly with certainty α_1 and young birds generally do not fly with certainty α_2:*

φ_1 :	$b \rightsquigarrow f,$	α_1
φ_2 :	$b \wedge y \rightsquigarrow \neg f,$	α_2
φ_3 :	$b \rightsquigarrow l,$	α_3

Then the possibilistic knowledge base associated with $U\Delta$ by the first method is the following (at this step, the ignored certainty levels are kept between parentheses) :

φ_1:	$b \rightarrow f,$	1/3	(α_1)
φ_2:	$b \wedge y \rightarrow \neg f,$	2/3	(α_2)
φ_3:	$b \rightarrow l,$	1/3	(α_3)

Let $FC = \{(b \wedge y, 1)\}$, meaning that we are considering a young bird. As previously computed, $Inc(D_\pi \cup FC) = 1/3$. Hence the final uncertain base UD_π contains only the uncertain formula $(b \wedge y \rightarrow \neg f, \alpha_2)$. So $FC \cup UD_\pi \vdash_\pi (\neg f, \alpha_2)$. It means that it is certain at level α_2 that a young bird is unable to fly, but we cannot conclude anything about its legs.

The second method rejects the rule φ_1 since it admits exceptions in the given context $b \wedge y$, leading to the resulting base:

φ_2 :	$(b \wedge y \rightarrow \neg f,$	$\alpha_2)$
φ_3 :	$(b \rightarrow l,$	$\alpha_3)$

It means that it is certain at level α_2 that a young bird is unable to fly, and at α_3 that it has legs.

The third method gives the following knowledge base D:

φ_2 :	$(b \wedge y \rightarrow \neg f,$	$\alpha_2)$
φ_1 :	$(b \wedge \neg y \rightarrow f,$	$\alpha_1)$
φ_3 :	$(b \rightarrow l,$	$\alpha_3)$

together with the uncertain completion rule base $\{(b \rightarrow \neg y, min(\alpha_1, \alpha_2)\}$, hence, in the context $FC = \{(b, 1)\}$, the completion rule is consistent with FC and D, so it allows us to deduce f with certainty $min(\alpha_1, \alpha_2)$ and l with certainty α_3. However using method 1 or 2 would have permitted to get a better lower bound of the necessity measure of f, namely α_1. This is the price paid for the computational simplicity of the method. In the context $FC = \{(b \wedge y, 1)$ we cannot add the completion rule since it is inconsistent with FC so we can conclude $\neg f$ with certainty α_2 and l with certainty α_3.

Note that the possibilistic setting also allows to process uncertain factual context, namely formulas in FC may have certainty levels less than 1.

Application to persistence modeling

The ability of handling uncertain default rules is useful for representing dynamical systems. Indeed, default reasoning can help solving the "frame" and "qualification" problems. The "frame problem" expresses the impossibility to enumerate every fluent which is not changed by an action. The "qualification problem" refers to the difficulty to exactly define all the preconditions of an action. An idea common to many proposals for solving the frame problem is to use default comportment descriptions for expressing persistence. Stating default transitions may be also useful for coping with the qualification problem. Besides, the available knowledge about the way a real system under study can evolve may be incomplete. This is why uncertainty should also be represented, at least in a qualitative way.

In this section, the variables set $\mathscr{V}$, on which the representation language $\mathscr{L}$ is built, may contain occurrences of action. More formally, let $\mathscr{A}$ be the set of action symbols. We consider that the variables set $\mathscr{V}$ contains in addition to the symbol representing facts all the symbols $do(a)$ where $a \in \mathscr{A}$, representing action occurrences. When there is ambiguity, variables may be indexed by a number representing the time point in which it is considered. We denote by f_t the formula f in which all variables are indexed by time point t. The evolution of the world is described by uncertain default rules of the form $(a_t \rightsquigarrow b_{t+1}, \alpha)$ meaning that if a is true at time t then b is generally true at time $t+1$ with a certainty level of α.

In order to handle the frame problem, we choose to define a frame axiom. Among all kinds of fluents we can distinguish persistent fluents (for which a change of value is surprising) from non persistent ones (which are also called dynamic (Sandewall 1995)). Alternative fluents represent another type of fluents (which should change their value at

[1] $\Pi(b|a)$ is defined as the largest solution of the equation $\Pi(a \wedge b) = min(\Pi(b|a), \Pi(a))$ applying the minimal specificity principle.

It yields: $\Pi(b|a) = \begin{cases} 1 & \text{if } \Pi(a \wedge b) > \Pi(a \wedge \neg b) \\ \Pi(a \wedge b) & \text{otherwise} \end{cases}$.

$$\text{Then } N(b|a) = 1 - \Pi(\neg a|b) = \begin{cases} 0 & \text{if } N(a \rightarrow \neg b) > N(a \rightarrow b) \\ N(a \rightarrow b) & \text{otherwise} \end{cases}$$

each time point); alternative fluents are non persistent fluents, their behavior can easily be described by transition rules of the form $f_t \rightsquigarrow \neg f_{t+1}$ and $\neg f_t \rightsquigarrow f_{t+1}$. Here, we consider that a set of non persistent literals NP is defined. Note that occurrences of actions are clearly non persistent fluents: $\{do(a)|a \in \mathscr{A}\} \subseteq NP$.

Definition 7 (frame axiom) *$\forall f \in \mathscr{V}$, if $f \notin NP$ then $(f_t \rightsquigarrow f_{t+1}, p(f))$ and if $\neg f \notin NP$ then $(\neg f_t \rightsquigarrow \neg f_{t+1}, p(\neg f))$ where $p(f)$ is the persistence degree of the fluent f.*

The persistence degree depends on the nature of the fluent, for instance, the fluent $asleep$ is persistent but it is less persistent than $deaf$. We will see in the following section that our formalism can also encode decreasing persistence.

Given the description of an evolving system composed of a set of uncertain default transition rules Δ describing its behavior (Δ contains pure dynamic laws and default persistence rules (coming from the frame axiom)) and a possibilistic knowledge base FC_t that describes the initial state of the world, we can study the problem of predicting the next state FC_{t+1} of the world. The following example inspired from (Mailhé & Prade 2004) shows how to describe a coffee machine behavior with uncertain default transition rules.

Example 6 *Let us consider a coffee machine that may be working (w), have enough money in it (m), have a goblet under the tap (g). Its normal behavior is roughly described by:*

$$\begin{array}{llll} \varphi_1 : m_t & \rightsquigarrow & g_{t+1} \wedge \neg m_{t+1} & 0.9 \\ \varphi_2 : m_t \wedge \neg w_t & \rightsquigarrow & \neg g_{t+1} & 0.9 \end{array}$$

where φ_1 means that if the machine has money in it then in the next step a goblet is under the tap and the money is spent. This first rule describes the intended coffee machine behavior supposing that it is working correctly. But it admits an exception described by φ_2. The agent is able to perform only one action on this machine: "give money" (gm). This action has an uncertain effect since giving money may fail if the coin is faked money (f).

$$\begin{array}{llll} \varphi_3 : do(gm)_t & \rightsquigarrow & m_{t+1} & 0.8 \\ \varphi_4 : do(gm)_t \wedge f_t & \rightsquigarrow & \neg m_{t+1} & 0.7 \end{array}$$

We consider m as the only non persistent fluent (as soon as it is true, it becomes false because of the rule φ_1): $NP = \{m\}$. Hence, persistence is encoded as follows (persistence degrees are put at a high value, but strictly less than 1):

$$\begin{array}{ll|ll} \varphi_5 : g_t \rightsquigarrow g_{t+1} & 0.9 & \varphi_9 : \neg g_t \rightsquigarrow \neg g_{t+1} & 0.9 \\ \varphi_6 : w_t \rightsquigarrow w_{t+1} & 0.9 & \varphi_{10} : \neg w_t \rightsquigarrow \neg w_{t+1} & 0.9 \\ \varphi_7 : f_t \rightsquigarrow f_{t+1} & 0.9 & \varphi_{11} : \neg f_t \rightsquigarrow \neg f_{t+1} & 0.9 \\ \varphi_8 : \neg m_t \rightsquigarrow \neg m_{t+1} & 0.9 & & \end{array}$$

In the initial state the agent is not absolutely sure that the coffee machine is working but he puts money in it; he thinks that is coin is not faked. $FC_t = \{(do(gm)_0, 1), (\neg m_0, 1), (\neg g_0, 1), (\neg f_0, 0.9)\}$. there is no money in the machine and no goblet.

¿From a set of uncertain default transition rules of the form $(a_t \rightsquigarrow b_{t+1}, \alpha)$, we can apply the methods presented in the previous section in order to obtain a set D of uncertain transition rules of the form $(a_t \rightarrow b_{t+1}, \alpha)$. ¿From D and a knowledge base FC_t describing the initial state, the next state can be computed syntactically as follows:

$$\begin{array}{ll} FC_{t+1} = & \{(b_{t+1}, \alpha) | \exists (a_t, \gamma) \text{ s.t. } (a_t \rightarrow b_{t+1}, \beta) \in D \\ & \text{and } FC_t \vdash_\pi (a_t, \gamma) \text{ and } \alpha = \min(\beta, \gamma)\} \end{array}$$

More generally, the resulting state can be computed by considering the extended set of rules D' corresponding to all the possible states of knowledge about the initial state of the system (Mailhé & Prade 2004):

$$D' = \{(\vee_I(\wedge_J a_i) \rightarrow \vee_I(\wedge_J b_i), \min_{i \in I \cup J} \alpha_i) | \forall (a_i \rightarrow b_i, \alpha_i) \in D\}$$

where I and J are any independent sets of indices of rules in D.

Example 7 *System Z gives three strata for example 6: $\Delta_0 = \{\varphi_1, \varphi_5, \ldots, \varphi_{11}\}$, $\Delta_1 = \{\varphi_2, \varphi_3\}$ and $\Delta_2 = \{\varphi_4\}$. Applying the first method leads to compute $Inc(FC_t \cup D_\pi)$ where D_π is the possibilistic knowledge base associated with Δ. Then delete all the rules of D_π that have a smaller specificity level. Only three rules are kept:*

$$\begin{array}{llll} \varphi_2 : m_t \wedge \neg w_t & \rightarrow & \neg g_{t+1} & 0.9 \\ \varphi_3 : do(gm)_t & \rightarrow & m_{t+1} & 0.8 \\ \varphi_4 : do(gm)_t \wedge f_t & \rightarrow & \neg m_{t+1} & 0.7 \end{array}$$

Hence, we can deduce $(m_{t+1}, 0.8)$ meaning that the machine has money in it in the next state.

The above example shows a drawback of the first method: all the persistence rules are drowned. Hence we are not able to determine the value of the fluents that are not concerned by transitions. The third method has not this drawback and preserves the following larger rule base where the modified parts of rules are in bold:

Example 8

$$\begin{array}{ll} \varphi_4 : do(gm)_t \wedge f_t \rightarrow \neg m_{t+1} & 0.7 \\ \varphi_2 : m_t \wedge \neg w_t \rightarrow \neg g_{t+1} & 0.9 \\ \varphi_3 : do(gm)_t \wedge \neg \mathbf{f_t} \rightarrow m_{t+1} & 0.8 \\ \varphi_1 : m_t \wedge \mathbf{w_t} \wedge \neg(\mathbf{do(gm)_t} \wedge \neg \mathbf{f_t}) \rightarrow g_{t+1} \wedge \neg m_{t+1} & 0.9 \\ \varphi_5 : g_t \wedge \neg(\mathbf{m_t} \wedge \neg \mathbf{w_t}) \rightarrow g_{t+1} & 0.9 \\ \varphi_6 : w_t \rightarrow w_{t+1} & 0.9 \\ \varphi_7 : f_t \rightarrow f_{t+1} & 0.9 \\ \varphi_8 : \neg m_t \wedge \neg(\mathbf{do(gm)_t} \wedge \neg \mathbf{f_t}) \rightarrow \neg m_{t+1} & 0.9 \\ \varphi_9 : \neg g_t \rightarrow \neg g_{t+1} & 0.9 \\ \varphi_{10} : \neg w_t \rightarrow \neg w_{t+1} & 0.9 \\ \varphi_{11} : \neg f_t \rightarrow \neg f_{t+1} & 0.9 \end{array}$$

Note that exceptions to persistence laws correspond to occurrences of actions, as expected. If the initial knowledge base FC_t is $\{(do(gm)_t, 1), (\neg m_t, 1), (\neg g_t, 1)\}$, completion rules are: $\{(do(gm)_t \rightarrow \neg f_t, min(0.8, 0.7) = 0.7)$, $(m_t \rightarrow w_t, 0.9)$, $(m_t \rightarrow \neg(do(gm)_t \wedge \neg f_t), 0.8)$, $(g_t \rightarrow \neg(m_t \wedge \neg w_t), 0.9)\}$. So at time point $t+1$, FC_{t+1} contains $(m_{t+1}, 0.7)$, $(\neg g_{t+1}, 0.9)$, $(\neg f_{t+1}, 0.9)$, meaning that there is money (with a certainty degree of 0.7) in the machine, no goblet and the coin is not faked (with a certainty degree of 0.9).

An interest of the third method is that the deduction can be iterated without recompilation of the default base (whereas it would be necessary with the second method).

Fuzzy default rules

Let us outline another application of the handling of uncertain default rules. In contrast with the previous sections, the certainty levels associated with rules are no longer fixed once for all but will be variable weights.

This setting enables us to handle fuzzy default rules of the form 'the more b, the more it is certain that a implies c is true" (Benferhat, Dubois, & Prade 1999). For instance, "the younger a bird is, the more certain it cannot fly". This kind of rule can be encoded by $(b \leadsto \neg f, \mu_y)$ where μ_y is a certainty level depending on how young is the bird. For instance, if $tweety$ is a bird of known age then the plausible consequence $(\neg fly(tweety), \mu_y(age(tweety)))$ can be obtained.

Again, note that this certainty level μ_y should not be confused with the specificity level of the rule. In example 6, we can imagine a rule of this kind: "the more strongly you hit the coffee machine, the more it is certain that it will not work in the next state" encoded by $(do(hm)_t \wedge w_t \leadsto \neg w_{t+1}, \mu_{\text{strongly}})$. Indeed, such a rule describes an exceptional situation with respect to the more general rule saying that by default if the coffee machine is working in t it still works in $t+1$ $(w_t \leadsto w_{t+1}, \alpha)$.

However, since the levels may be variable weights, this indicates that in some sense they then depend on some contextual information (e.g., the age of the bird in the above example). As referring to a context, this should play a role not only in the evaluation of a certainty level, but also in the preliminary step of our approach that analyses the specificity level of the default rules. This point is illustrated by the following example:

Example 9 *Consider the two defaults:*
$\varphi_1 \quad (b \leadsto \neg f, \mu_y)$
$\varphi_2 \quad (b \leadsto f, \alpha)$
stating that the younger a bird, the more certain it cannot fly, and that birds generally fly.

The above default knowledge base cannot be stratified. Indeed, the youngness feature of the context cannot be identified in the condition part of the default φ_1, although it should be recognized as more specific.

This problem can be solved by the following trick. The general idea is to introduce an uninstanciated literal to capture this context feature. Namely, the default φ_1 could be rewritten into:

$$(b \wedge (age = x) \leadsto \neg f, \mu_y(x))$$

which is a first order universally quantified uncertain default rule. System Z stratification method extends to this generalized type of default, as well as possibilistic logic inference to the associated type of uncertain logic formula.

The possibility to affect variable levels to a rule may be also useful in order to express decreasing persistence (Dupin de Saint-Cyr & Lang 1997): the more the time has gone by the less it is certain that a fluent has kept its value. A decreasing persistence rule is generally of the form $(m_t \leadsto m_{t+n}, f(m, n))$ where the level attached to the rule depends on the fluent quality (highly persistent or dynamic) and of the length of the time interval.

Related Works

There has been very few works handling both defeasibility and uncertainty, up to the noticeable exception of system Z^+(Goldszmidt & Pearl 1996) that handles default rules having strengths modeled in the setting of Spohn ordinal condition functions, and their exploitation by maximum entropy principle, taking advantage of the probabilistic interpretation of Spohn functions (Bourne & Parsons 2003). In system Z^+, a default rule $(a \leadsto b)$ is extended with a parameter representing the degree of strength or firmness of the rule and denoted by $(a \rightarrow^{\delta} b)$. This is interpreted as a constraint of the form $\kappa(a \wedge b) < \kappa(a \wedge \neg b) + \delta$ where κ is a Spohn kappa function associating any set of interpretations with an integer value that expresses impossibility (thus 0 means full possibility and ∞ means full impossibility). Translated in possibilistic terms, it amounts to deal with constraints of the form $\Pi(a \wedge b) > k.\Pi(a \wedge \neg b)$ with $k \geq 1$, using the standard transformation between kappa functions and possibility measures (Dubois & Prade 1991). Thus, the k's are like uncertainty odds. In Z^+, the ranking of defaults is obtained by comparing sums of strength degrees, somewhat mixing the ideas of specificity and strength. Separate scales for specificity and certainty are not used is this approach, so certainty levels are introduced in the computation of the levels reflecting specificity ordering. This leads to an interaction between the two notions. For instance, encoding our example 1 in a Z^+ formalism, we get:
$r_1: \quad b \rightarrow^{\delta_1} f$
$r_2: \quad b \wedge y \rightarrow^{\delta_2} \neg f$
, where δ_1 and δ_2 are non negative integers. System Z^+ generates the following ranking on the two interpretations $\{b, y, f\}$ and $\{b, y, \neg f\}$: $\kappa(\{b, y, f\}) = Z^+(r_2) = \delta_1 + \delta_2 + 1$ and $\kappa(\{b, y, \neg f\}) = Z^+(r_1) = \delta_1$. Thus in Z^+, the strengths of the defaults are combined for determining their respective specificity level, and paradoxically, not really for computing the certainty levels of the conclusions. The approach presented here distinguishes more carefully between specificity and certainty.

As shown on the following example, the way system Z^+ handles defeasibility and certainty in a mixed way may not yield the expected conclusion always.

Example 10 *Consider the following default base stating that birds generally fly, birds generally have not webbed feet, young birds generally do not fly, and that duck birds generally have webbed feet.*
$\varphi_1: \quad b \rightarrow^{\delta_1} f$
$\varphi_2: \quad b \rightarrow^{\delta_2} \neg wf$
$\varphi_3: \quad b \wedge y \rightarrow^{\delta_3} \neg f$
$\varphi_4: \quad b \wedge d \rightarrow^{\delta_4} wf$

System Z^+ associates to these defaults the following respective ranks $\delta_1, \delta_2, \delta_1 + \delta_3 + 1, \delta_2 + \delta_4 + 1$. Assume that the values of the δ_i's are such that $\delta_1 < \delta_1 + \delta_3 + 1 < \delta_2 < \delta_2 + \delta_4 + 1$ (which does not correspond to a refinement of the Z ordering!). Then, from a young duck bird, System Z^+ concludes that it has webbed feet but cannot conclude that it cannot fly as System Z will do.

Another interesting approach handling both defeasibility and uncertainty has been proposed in (Lukasiewicz 2005) in a setting where probabilistic logic is combined with default reasoning. Lukasiewicz proposes a framework which can handle simultaneously strict propositional rules, probabilistic formulas and default formulas. A basic difference with our proposal is that default formulas are classical default rules, meanwhile in this paper a new kind of default rules that are also pervaded with uncertainty is considered.

Nicolas, Garcia and Stéphan (Nicolas, Garcia, & Stéphan 2005) also present an approach that deals with defeasibility and uncertainty in a possibilistic framework. But, they combine possibilistic logic with Answer Set Programming rather than using the same setting for default and uncertainty handling. Certainty levels are used in order to help to restore consistency of a logic program by removing rules that are below a level of inconsistency. As our first method, this approach does not avoid the drowning problem, while our two other methods do. Besides, our third method where defaults are rewritten by mentioning explicit exceptions is close to techniques used in circumscription.

Using an uncertain framework in order to describe an evolving system has been done by many authors, for instance in a probabilistic setting. But reasoning in this setting implies to dispose of many a priori probabilities, this is why using defeasibility may help to reduce the size of information for representing the system. Besides, it is a common idea to define a frame axiom in terms of default rules (see (Lang, Lin, & Marquis 2003) for an overview). But, as far as we know, frame rules are either considered as default rules (see (Giunchiglia *et al.* 2004; Baral & Lobo 1997) for instance), or are associated with low priority levels (see (Kakas, Miller, & Toni 2001)), but do not involve both default and uncertainty feature.

Conclusion

We propose a representation language which allows us to handle rules which are both uncertain and by default. This tool is useful in the context of dynamic systems since it helps solving the "frame" and "qualification" problems, thanks to default transition rules. The suggested use of the approach for handling fuzzy default rules may also find applications for handling default inheritance in fuzzy description logic in a possibilistic logic setting (Dubois, Mengin, & Prade 2006). The proposed approach, specially the one based on the rewriting algorithm, could be cast in a logic programming setting to solve the drowning problem pointed out in (Nicolas, Garcia, & Stéphan 2005).

Annex A: Background on possibility theory

Possibility theory (Dubois & Prade 1988) associates to a formula f two measures, namely its possibility $\Pi(f)$ which measures how unsurprising the formula f is ($\Pi(f) = 0$ means that f is bound to be false) and its dual necessity $N(f) = 1 - \Pi(\neg f)$ ($N(f) = 1$ means that f is bound to be true). Necessity obeys to the characteristic axiom $N(f \wedge g) = \min(N(f), N(g))$. A possibilistic knowledge base is a set $K = \{(\varphi_i, \alpha_i), i = 1 \ldots n\}$, where φ_i is a propositional formula of $\mathbf{L}$ and its *certainty* level (or weight) α_i is such that $N(\varphi_i) \geq \alpha_i$, N being a necessity measure. The resolution rule (Dubois, Lang, & Prade 1994) is valid in possibilistic logic: $(a \vee b, \alpha); (\neg a \vee c, \beta) \vdash (b \vee c, \min(\alpha, \beta))$, where $\vdash$ denotes the syntactic inference of possibilistic logic. Classical resolution is retrieved when all weights are equal to 1. The resolution rule allows us to compute the maximal certainty level that can be attached to a formula according to the constraints expressed by the base K. This can be done by adding to K the clauses obtained by refuting the proposition to evaluate, with a necessity level equal to 1. Then it can be shown that any lower bound obtained on $\perp$, by resolution, is a lower bound of the necessity of the proposition to evaluate. Let $Inc(K) = \max\{\alpha \mid K_\alpha \vdash \perp\}$ with $K_\alpha = \{f | (f, \beta) \in K \text{ and } \beta \geq \alpha\}$, with the convention $max(\varnothing) = 0$. In case of partial inconsistency of K ($Inc(K) > 0$), a refutation carried out in a situation where $Inc(K \cup \{(\neg f, 1)\}) = \alpha > Inc(K)$ yields the nontrivial conclusion (f, α), only using formulas whose certainty levels are strictly greater than the level of inconsistency of the base. This is the syntactic possibilistic entailment, noted $\vdash_\pi$.

Annex B: Background on default rules

A default rule is an expression $a \rightsquigarrow b$ where a and b are propositional formulas of $\mathscr{L}$ and $\rightsquigarrow$ is a new symbol. $a \rightsquigarrow b$ translates, in the possibility theory framework, into the constraint $\Pi(a \wedge b) > \Pi(a \wedge \neg b)$ which expresses that having b true is strictly more possible than having it false when a is true (Benferhat, Dubois, & Prade 1992).

The use of default rules has two main interests. First, it simplifies the writing: it allows us to express a rule without mentioning every exceptions to it. Second, it allows us to reason with incomplete descriptions of the world: if nothing is known about the exceptional character of the situation, it is assumed to be normal, and reasoning can be completed.

(Kraus, Lehmann, & Magidor 1990; Gärdenfors & Makinson 1994) have developed an approach for handling reasoning with default rules based on postulates stating the characteristic properties of a non-monotonic consequence relations. In this setting, two inferences are defined: a cautious one named "preferential" and a more adventurous one named "rational closure inference".

Pearl (Pearl 1990) provides an algorithm which gives a stratification of a set of default rules in a way that reflects the specificity of the rules. Roughly speaking, the first stratum contains the most specific rules i.e., which do not admit exceptions (at least, expressed in the considered default base), the second stratum has exceptions only in the first stratum and so on.

Definition 8 (System Z stratification) *A default rule $a \rightsquigarrow b$ is tolerated by a set of default rules Δ if it exists an interpretation ω such that $\omega \models a \wedge b$ and $\forall a_i \rightsquigarrow b_i \in \Delta$, $\omega \models \neg a_i \vee b_i$. This definition allows us to stratify Δ into $(\Delta_0, \Delta_1, \ldots, \Delta_n)$ such that Δ_0 contains the set of rules of Δ tolerated by Δ, Δ_1 contains the set of rules of $\Delta \setminus \Delta_0$ tolerated by $\Delta \setminus \Delta_0$ and so on. The number $Z(r)$ corresponds to the rank of the stratum in which the rule r is.*

It has been shown (Benferhat, Dubois, & Prade 1992) that each default rule $r = a \rightsquigarrow b$ of a default base Δ, can be associated with a possibilistic formula $(a \rightarrow b, \sigma)$, where σ

represents its specificity level $\sigma = \frac{Z(r)+1}{n+2}$, n being the index of the last stratum in the system Z stratification of Δ. Applying possibilistic inference to the possibilistic base associated with a default base in this sense is equivalent to compute the rational closure inference (Kraus, Lehmann, & Magidor 1990; Gärdenfors & Makinson 1994) of the original default base (Benferhat, Dubois, & Prade 1992).

References

Baral, C., and Lobo, J. 1997. Defeasible specifications in action theories. In *IJCAI-97*, 1441–1446.

Benferhat, S., and Dupin de Saint-Cyr, F. 1996. Contextual handling of conditional knowledge. In *Proc. of the 6th Conf. on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, volume 3, 1369–1374.

Benferhat, S.; Dubois, D.; and Prade, H. 1992. Representing default rules in possibilistic logic. In *KR-92*, 673–684.

Benferhat, S.; Dubois, D.; and Prade, H. 1997. Nonmonotonic reasoning, conditional objects and possibility theory. In *Artificial Intelligence*, volume 92. 259–276.

Benferhat, S.; Dubois, D.; and Prade, H. 1998. Practical handling of exception-tainted rules and independence information in possibilistic logic. *Applied intelligence* 9:101–127.

Benferhat, S.; Dubois, D.; and Prade, H. 1999. Towards fuzzy default reasoning . In *Proceedings of the 18th International Conference of the North American Fuzzy Information Processing Society , New York, USA*, 23–27. IEEE Service Center.

Bourne, A., and Parsons, S. 2003. Extending the maximum entropy approach to variable strength defaults. *Annals of Maths and AI* 39:123–146.

Dubois, D., and Prade, H. 1988. *Possibility Theory: An Approach to Computerized Processing of Uncertainty*. New York: Plenum Press.

Dubois, D., and Prade, H. 1991. Epistemic entrenchment and possibilistic logic. *Artificial Intelligence* 50(2):223–239.

Dubois, D., and Prade, H. 1998. Possibility theory:qualitative and quantitative aspects . In Gabbay, D., and Smets, P., eds., *Quantified Representation of Uncertainty and Imprecision* , volume 1 of *Handbook of Defeasible Reasoning and Uncertainty Management Systems*. The Netherlands: Kluwer Academic Publishers. 169–226.

Dubois, D.; Lang, J.; and Prade, H. 1994. Possibilistic logic. In Gabbay, D.; Hogger, C.; Robinson, J.; and Nute, D., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3* , volume 3. Oxford University Press. 439–513.

Dubois, D.; Mengin, J.; and Prade, H. 2006. *Fuzzy logic and the semantic Web*. Elsevier. chapter Possibilistic Uncertainty and Fuzzy Features in Description Logic. A preliminary Discussion, 101–113.

Dupin de Saint-Cyr, F., and Lang, J. 1997. Reasoning about unpredicted change and explicit time. In *LNAI-1244*, 223–236.

Gärdenfors, P., and Makinson, D. 1994. Nonmonotonic inference based on expectations. *Artificial Intelligence* 65:197–245.

Giunchiglia, E.; Lee, J.; Lifschitz, V.; McCain, N.; and Turner, H. 2004. Nonmonotonic causal theories. *Artif. Intell.* 153:49–104.

Goldszmidt, M., and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision and causal modeling. *Artificial Intelligence* 84:57–112.

Kakas, A.; Miller, R.; and Toni, F. 2001. E-res: reasoning about actions, events and observations. In *LPNMR-01*, 254–266.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207.

Lang, J.; Lin, F.; and Marquis, P. 2003. Causal theories of action: A computational core. In *IJCAI-03*, 1073–1078.

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.

Lukasiewicz, T. 2005. Weak nonmonotonic probabilistic logics. *Artificial Intelligence* 168(1-2):119–161.

Mailhé, B., and Prade, H. 2004. Updating of a possibilistic knowledge base by crisp or fuzzy transition rules. *Proc. of the 9th KR* 338–346.

Nicolas, P.; Garcia, L.; and Stéphan, I. 2005. A possibilistic inconsistency handling in anwer set programming. In *Proc. of ESCQARU'05*, 402–414. Springer.

Pearl, J. 1990. System Z : a natural ordering of defaults with tractable applications for default reasoning. In *Proc. of theoretical aspects of reasoning about knowledge*, 121–135.

Sandewall, E. 1995. *Features and Fluents*. Oxford University Press.

Touretzky, D. 1984. Implicit ordering of defaults in inheritance systems. In *Proc. of AAAI'84*.