

A framework for a modular multi-concept lexicographic closure semantics (an abridged report) ^{*}

Laura Giordano and Daniele Theseider Dupré

DISIT - Università del Piemonte Orientale, Italy
laura.giordano@uniupo.it, dtd@uniupo.it

Abstract. We define a modular multi-concept extension of the lexicographic closure semantics for defeasible description logics with typicality. The idea is that of distributing the defeasible properties of concepts into different modules, according to their subject, and of defining a notion of preference for each module based on the lexicographic closure semantics. The preferential semantics of the knowledge base can then be defined as a combination of the preferences of the single modules. The range of possibilities, from fine grained to coarse grained modules, provides a spectrum of alternative semantics.

1 Introduction

Kraus, Lehmann and Magidor’s preferential logics for non-monotonic reasoning [43, 44], have been extended to description logics, to deal with inheritance with exceptions in ontologies, allowing for non-strict forms of inclusions, called *typicality or defeasible inclusions*, with different preferential and ranked semantics [30, 15] as well as different closure constructions such as the rational closure [19, 18, 33], the lexicographic closure [20], the relevant closure [17], and MP-closure [29].

In this paper we define a modular multi-concept extension of the lexicographic closure for reasoning about exceptions in ontologies. The idea is very simple: different modules can be defined starting from a defeasible knowledge base, containing a set \mathcal{D} of typicality inclusions (or defeasible inclusions) describing the prototypical properties of classes in the knowledge base. We will represent such defeasible inclusions as $\mathbf{T}(C) \sqsubseteq D$ [30], meaning that “typical C ’s are D ’s” or “normally C ’s are D ’s”, corresponding to conditionals $C \sim D$ in KLM framework.

A set of modules m_1, \dots, m_n is introduced, each one concerning a subject, and defeasible inclusions belong to a module if they are related with its subject. By subject, here, we mean any concept of the knowledge base. Module m_i with subject C_i does not need to contain just typicality inclusions of the form $\mathbf{T}(C_i) \sqsubseteq D$, but all defeasible inclusions in \mathcal{D} which are concerned with subject C_i are admitted in m_i . We call a collection of such modules a *modular multi-concept knowledge base*.

This modularization of the defeasible part of the knowledge base does not define a partition of the set \mathcal{D} of defeasible inclusions, as an inclusion may belong to more

^{*} Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

than one module. For instance, the typical properties of employed students are relevant both for the module with subject *Student* and for the module with subject *Employee*. The granularity of modularization has to be chosen by the knowledge engineer who can fix how large or narrow is the scope of a module, and how many modules are to be included in the knowledge base (for instance, whether the properties of employees and students are to be defined in the same module with subject *Person* or in two different modules). At one extreme, all the defeasible inclusions in \mathcal{D} can be put together in a module associated with subject \top (Thing). At the other extreme, which has been studied in [35], a module m_i is a defeasible TBox containing *only* the defeasible inclusions of the form $\mathbf{T}(C_j) \sqsubseteq D$ for some concept C_i . In this paper we remove this restriction considering general modules, containing arbitrary sets of defeasible inclusions, intuitively pertaining some subject.

In [35], following Gerard Brewka’s framework of Basic Preference Descriptions for ranked knowledge bases [12], we have assumed that a specification of the relative importance of typicality inclusions for a concept C_i is given by assigning ranks to typicality inclusions. However, for a large module, a specification by hand of the ranking of the defeasible inclusions in the module would be awkward. In particular, a module may include all properties of a class as well as properties of its exceptional subclasses (for instance, the typical properties of penguins, ostriches, etc. might all be included in a module with subject *Bird*). A natural choice is then to consider, for each module, a lexicographic semantics which builds on the rational closure ranking to define a preference ordering on domain elements. This preference relation corresponds, in the propositional case, to the lexicographic order on worlds in Lehmann’s model theoretic semantics of the lexicographic closure [45]. This semantics already accounts for the specificity relations among concepts inside the module, as the lexicographic closure deals with specificity, based on ranking of concepts computed by the rational closure.

Based on the ranked semantics of the single modules, a compositional (preferential) semantics of the knowledge base is defined by combining the multiple preference relations into a single global preference relation $<$. This gives rise to a modular multi-concept extension of Lehmann’s preference semantics for the lexicographic closure. When there is a single module, containing all the typicality inclusions in the knowledge base, the semantics collapses to a natural extension to DLs of Lehmann’s semantics, which corresponds to Lehmann’s semantics for the fragment of \mathcal{ALC} without universal and existential restrictions.

We introduce a notion of entailment for modular multi-concept knowledge bases, based on the proposed semantics, which satisfies the KLM properties of a preferential consequence relation. This notion of entailment has good properties inherited from lexicographic closure: it deals properly with irrelevance and specificity, and it is not subject to the “blockage of property inheritance” problem, i.e., the problem that property inheritance from classes to subclasses is not guaranteed, which affects the rational closure [47]. In addition, separating defeasible inclusions in different modules provides a simple solution to another problem of the rational closure and its refinements (including the lexicographic closure), that was recognized by Geffner and Pearl [27], namely, that “conflicts among defaults that should remain unresolved, are resolved anomalously”, giving rise to too strong conclusions. The preferential (not necessarily ranked) nature

of the global preference relation $<$ provides a simple way out to this problem, when defeasible inclusions are suitably separated in different modules.

2 The description logics \mathcal{ALC} and its extension with typicality

In this section we recall the syntax and semantics of the description logic \mathcal{ALC} [2] and an extension of \mathcal{ALC} with typicality [33].

Let N_C be a set of concept names, N_R a set of role names and N_I a set of individual names. The set of \mathcal{ALC} concepts (or, simply, concepts) can be defined inductively:

- $A \in N_C$, \top and \perp are concepts;
- if C and D are concepts, and $r \in N_R$, then $C \sqcap D$, $C \sqcup D$, $\neg C$, $\forall r.C$, $\exists r.C$ are concepts.

A knowledge base (KB) K is a pair $(\mathcal{T}, \mathcal{A})$, where \mathcal{T} is a TBox and \mathcal{A} is an ABox. The TBox \mathcal{T} is a set of concept inclusions (or subsumptions) $C \sqsubseteq D$, where C, D are concepts. The ABox \mathcal{A} is a set of assertions of the form $C(a)$ and $r(a, b)$ where C is a concept, a and b are individual names in N_I and r a role name in N_R .

An \mathcal{ALC} interpretation is defined as a pair $I = \langle \Delta, \cdot^I \rangle$ where: Δ is a domain—a set whose elements are denoted by x, y, z, \dots —and \cdot^I is an extension function that maps each concept name $C \in N_C$ to a set $C^I \subseteq \Delta$, each role name $r \in N_R$ to a binary relation $r^I \subseteq \Delta \times \Delta$, and each individual name $a \in N_I$ to an element $a^I \in \Delta$. It is extended to complex concepts as follows:

$$\begin{aligned} \top^I &= \Delta, & \perp^I &= \emptyset, & (\neg C)^I &= \Delta \setminus C^I, \\ (\exists r.C)^I &= \{x \in \Delta \mid \exists y.(x, y) \in r^I \text{ and } y \in C^I\}, & (C \sqcap D)^I &= C^I \cap D^I, \\ (\forall r.C)^I &= \{x \in \Delta \mid \forall y.(x, y) \in r^I \Rightarrow y \in C^I\}, & (C \sqcup D)^I &= C^I \cup D^I. \end{aligned}$$

The notion of satisfiability of a KB in an interpretation and the notion of entailment are defined as follows:

Definition 1 (Satisfiability and entailment). Given an \mathcal{ALC} interpretation $I = \langle \Delta, \cdot^I \rangle$:

- I satisfies an inclusion $C \sqsubseteq D$ if $C^I \subseteq D^I$;
- I satisfies an assertion $C(a)$ (resp., $r(a, b)$) if $a^I \in C^I$ (resp., $(a^I, b^I) \in r^I$).

Given a KB $K = (\mathcal{T}, \mathcal{A})$, an interpretation I satisfies \mathcal{T} (resp. \mathcal{A}) if I satisfies all inclusions in \mathcal{T} (resp. all assertions in \mathcal{A}); I is a model of K if I satisfies \mathcal{T} and \mathcal{A} .

A subsumption $F = C \sqsubseteq D$ (resp., an assertion $C(a)$, $r(a, b)$), is entailed by K , written $K \models F$, if for all models $I = \langle \Delta, \cdot^I \rangle$ of K , I satisfies F .

Given a knowledge base K , the *subsumption problem* is the problem of deciding whether an inclusion $C \sqsubseteq D$ is entailed by K .

In the following we will refer to an extension of \mathcal{ALC} with typicality inclusions, that we will call $\mathcal{ALC} + \mathbf{T}$ as in [30], and to the *rational closure* of $\mathcal{ALC} + \mathbf{T}$ knowledge bases $(\mathcal{T}, \mathcal{A})$ [33]. In addition to standard \mathcal{ALC} inclusions $C \sqsubseteq D$ (called *strict inclusions* in the following), in $\mathcal{ALC} + \mathbf{T}$ the TBox \mathcal{T} also contains typicality inclusions of the form $\mathbf{T}(C) \sqsubseteq D$, where C and D are \mathcal{ALC} concepts. Let us recall the notions of preferential, ranked and canonical model of a defeasible knowledge base $(\mathcal{T}, \mathcal{A})$.

Definition 2 (Interpretations for $\mathcal{ALC} + \mathbf{T}$). A preferential interpretation \mathcal{N} is any structure $\langle \Delta, <, \cdot^{\mathcal{N}} \rangle$ where: Δ is a domain; $<$ is an irreflexive, transitive and well-founded relation over Δ ; $\cdot^{\mathcal{N}}$ is a function that maps all concept names, role names and

individual names as defined above for \mathcal{ALC} interpretations, and provides an interpretation to all \mathcal{ALC} concepts as above, and to typicality concepts as follows: $(\mathbf{T}(C))^I = \min_{<}(C^I)$, where $\min_{<}(S) = \{u : u \in S \text{ and } \nexists z \in S \text{ s.t. } z < u\}$.

When relation $<$ is required to be also modular (i.e., for all $x, y, z \in \Delta$, if $x < y$ then $x < z$ or $z < y$), \mathcal{N} is called a ranked interpretation.

Preferential interpretations for description logics were first studied in [30], while ranked interpretations (i.e., modular preferential interpretations) were first introduced for \mathcal{ALC} in [15]. A preferential (ranked) model of an $\mathcal{ALC} + \mathbf{T}$ knowledge base K is a preferential (ranked) $\mathcal{ALC} + \mathbf{T}$ interpretation $\mathcal{N} = \langle \Delta, <, \cdot^I \rangle$ that satisfies all inclusions in K , where: a strict inclusion or an assertion is satisfied in \mathcal{N} if it is satisfied in the \mathcal{ALC} model $\langle \Delta, \cdot^I \rangle$, and a typicality inclusion $\mathbf{T}(C) \sqsubseteq D$ is satisfied in \mathcal{N} if $(\mathbf{T}(C))^I \sqsubseteq D^I$. Preferential entailment in $\mathcal{ALC} + \mathbf{T}$ is defined in the usual way: for a knowledge base K and a query F (a strict or defeasible inclusion or an assertion), F is preferentially entailed by K ($K \models_{\mathcal{ALC}+\mathbf{T}} F$) if F is satisfied in all preferential models of K .

A canonical model for K is a preferential (ranked) model containing, roughly speaking, as many domain elements as consistent with the knowledge base specification K . Given an $\mathcal{ALC} + \mathbf{T}$ knowledge base $K = (\mathcal{T}, \mathcal{A})$ and a query F , let us define \mathcal{S}_K as the set of all \mathcal{ALC} concepts (and subconcepts) occurring in K or in F , together with their complements. We consider all the sets of concepts $\{C_1, C_2, \dots, C_n\} \subseteq \mathcal{S}_K$ consistent with K , i.e., s.t. $K \not\models_{\mathcal{ALC}+\mathbf{T}} C_1 \sqcap C_2 \sqcap \dots \sqcap C_n \sqsubseteq \perp$.

Definition 3 (Canonical model). A preferential model $\mathcal{M} = \langle \Delta, <, I \rangle$ of K is canonical with respect to \mathcal{S}_K if it contains at least a domain element $x \in \Delta$ s.t. $x \in (C_1 \sqcap C_2 \sqcap \dots \sqcap C_n)^I$, for each set $\{C_1, C_2, \dots, C_n\} \subseteq \mathcal{S}_K$ consistent with K .

For finite, consistent $\mathcal{ALC} + \mathbf{T}$ knowledge bases, existence of finite (ranked) canonical models has been proved in [33] (Theorem 1). In the following, we consider finite $\mathcal{ALC} + \mathbf{T}$ knowledge bases, and we can restrict our consideration to finite preferential models.

3 Modular multi-concept knowledge bases

In this section we introduce a notion of a multi-concept knowledge base, starting from a set of strict inclusions \mathcal{T} , a set of assertions \mathcal{A} , and a set of typicality inclusions \mathcal{D} , each one of the form $\mathbf{T}(C) \sqsubseteq D$, where C and D are \mathcal{ALC} concepts.

Definition 4. A modular multi-concept knowledge base K is a tuple $\langle \mathcal{T}, \mathcal{D}, m_1, \dots, m_k, \mathcal{A}, s \rangle$, where \mathcal{T} is an \mathcal{ALC} TBox, \mathcal{D} is a set of typicality inclusions, such that $m_1 \cup \dots \cup m_k = \mathcal{D}$, \mathcal{A} is an ABox, and s is a function associating each module m_i with a concept, $s(m_i) = C_i$, the subject of m_i .

The idea is that each m_i is a module defining the typical properties of the instances of some concept C_i . The defeasible inclusions belonging to a module m_i with subject C_i are the inclusions that intuitively pertain to C_i . We expect that all the typicality inclusions $\mathbf{T}(C) \sqsubseteq D$, such that C is a subclass of C_i , belong to m_i , but not only. For instance, for a module m_i with subject $C_i = \text{Bird}$, the typicality inclusion $\mathbf{T}(\text{Bird} \sqcap \text{Live_at_SouthPole}) \sqsubseteq \text{Penguin}$, meaning that the birds living at the south

pole are normally penguins, is clearly to be included in m_i . As penguins are birds, also inclusion $\mathbf{T}(Penguin) \sqsubseteq Black$ is to be included in m_i , and, if $\mathbf{T}(Bird) \sqsubseteq FlyingAnimal$ and $\mathbf{T}(FlyingAnimal) \sqsubseteq BigWings$ are defeasible inclusions in the knowledge base, they both may be relevant properties of birds to be included in m_i . For this reason we will not put restrictions on the typicality inclusions that can belong to a module. We will see that the semantic construction for a module m_i will be able to ignore the typicality inclusions which are not relevant for subject C_i .

The modularization m_1, \dots, m_k of the defeasible part \mathcal{D} of the knowledge base does not define a partition of \mathcal{D} , as the same inclusion may belong to more than one module m_i . For instance, the typical properties of employed students are relevant for both concept *Student* and concept *Employee* and should belong to their related modules (if any). Also, a granularity of modularization has to be chosen and, as we will see, this choice may have an impact on the global semantics of the knowledge base. At one extreme, all the defeasible inclusions in \mathcal{D} are put together in the same module, e.g., the module associated with concept \top . At the other extreme, which has been studied in [35], a module m_i contains *only* the defeasible inclusions of the form $\mathbf{T}(C_i) \sqsubseteq D$, where C_i is the subject of m_i (and in this case, the inclusions $\mathbf{T}(C) \sqsubseteq D$ with C subsumed by C_i are not admitted in m_i). In this regard, the framework proposed in this paper could be seen as an extension of the proposal in [35] to allow coarser grained modules, while here we do not allow for user-defined preferences among defaults.

Let us consider an example of multi-concept knowledge base.

Example 1. Let K be the knowledge base $\langle \mathcal{T}, \mathcal{D}, m_1, m_2, m_3, \mathcal{A}, s \rangle$, where $\mathcal{A} = \emptyset$, \mathcal{T} contains the strict inclusions:

$$\begin{array}{lll} Adult \sqsubseteq \exists has_SSN.\top & Employee \sqsubseteq Adult & PhdStudent \sqsubseteq Student \\ PhDStudent \sqsubseteq Adult & Has_no_Scholarship \equiv \neg \exists hasScholarship.\top & \\ PrimarySchoolStudent \sqsubseteq Children & Driver \sqsubseteq \exists has_DrivingLicence.\top & \\ PrimarySchoolStudent \sqsubseteq HasNoClasses & Driver \sqsubseteq Adult & \end{array}$$

and the defeasible inclusions in \mathcal{D} are distributed in the modules m_1, m_2, m_3 as follows.

Module m_1 has subject *Employee*, and contains the defeasible inclusions:

$$\begin{array}{ll} (d_1) \mathbf{T}(Employee) \sqsubseteq \neg Young & (d_2) \mathbf{T}(Employee) \sqsubseteq \exists has_boss.Employee \\ (d_3) \mathbf{T}(ForeignerEmployee) \sqsubseteq \exists has_Visa.\top & \\ (d_4) \mathbf{T}(Employee \sqcap Student) \sqsubseteq Busy & (d_5) \mathbf{T}(Employee \sqcap Student) \sqsubseteq \neg Young \end{array}$$

Module m_2 has subject *Student*, and contains the defeasible inclusions:

$$\begin{array}{ll} (d_6) \mathbf{T}(Student) \sqsubseteq \exists has_classes.\top & (d_7) \mathbf{T}(Student) \sqsubseteq Young \\ (d_8) \mathbf{T}(Student) \sqsubseteq Has_no_Scholarship & (d_9) \mathbf{T}(HighSchoolStudent) \sqsubseteq Teenager \\ (d_{10}) \mathbf{T}(PhDStudent) \sqsubseteq \exists hasScholarship.Amount & \\ (d_{11}) \mathbf{T}(PhDStudent) \sqsubseteq Bright & \end{array}$$

together with (d_4) and (d_5) . Module m_3 has subject *Vehicle*, and we omit its content.

Observe that, in previous example, (d_4) and (d_5) belong to both modules m_1 and m_2 .

4 A lexicographic semantics of modular multi-concept KBs

In this section, we define a semantics of modular multi-concept knowledge bases, based on Lehmann's lexicographic closure semantics [45]. The idea is that, for each module

m_i , a semantics can be defined using lexicographic closure semantics, with some minor modification.

Given a modular multi-concept knowledge base $K = \langle \mathcal{T}, \mathcal{D}, m_1, \dots, m_k, \mathcal{A}, s \rangle$, we let $\text{rank}(C)$ be the rank of concept C in the rational closure ranking of the knowledge base $(\mathcal{T} \cup \mathcal{D}, \mathcal{A})$, according to the *rational closure* construction in [33]. In the rational closure ranking, concepts with higher ranks are more specific than concepts with lower ranks. While we will not recall the rational closure construction, let us consider again Example 1. In Example 1, the rational closure ranking assigns to concepts *Adult*, *Employee*, *ForeignEmployee*, *Driver*, *Student*, *HighSchoolStudent*, *PrimarySchoolStudent* the rank 0, while to concepts *PhDStudent* and *Employee* \sqcap *Student* the rank 1. In fact, *PhDStudent* are exceptional students, as they have a scholarship, while employed students are exceptional students, as they are not young: they are exceptional subclasses of class *Student*.

Based on the concept ranking, the rational closure assigns a rank to typicality inclusions: the rank of $\mathbf{T}(C) \sqsubseteq D$ is equal to the rank of concept C . For each module m_i of a knowledge base $K = \langle \mathcal{T}, \mathcal{D}, m_1, \dots, m_k, \mathcal{A}, s \rangle$, we aim to define a canonical model, using the lexicographic order based on the rank of typicality inclusions in m_i . In the following we will assume that the knowledge base $(\mathcal{T} \cup \mathcal{D}, \mathcal{A})$ is consistent in the logic $\mathcal{ALC} + \mathbf{T}$, that is, it has a preferential model. This also guarantees the existence of (finite) canonical models [33].

Let us define the *projection of the knowledge base K on module m_i* as the knowledge base $K_i = \langle \mathcal{T} \cup m_i, \mathcal{A} \rangle$. K_i is an $\mathcal{ALC} + \mathbf{T}$ knowledge base. Hence a preferential model $\mathcal{N}_i = \langle \Delta, <_i, \cdot^I \rangle$ of K_i is defined as in Section 2 (but now we use $<_i$, instead of $<$, for the preference relation in \mathcal{N}_i , for $i = 1, \dots, k$).

In his seminal work on the lexicographic closure, Lehmann [45] defines a model theoretic semantics of the lexicographic closure construction by introducing an order relation among propositional models, considering which defaults are violated in each model, and introducing a seriousness ordering \prec among sets of violated defaults. For two propositional models w and w' , $w \prec w'$ (w is preferred to w') is defined in [45] as:

$$w \prec w' \text{ iff } V(w) \prec V(w') \quad (1)$$

w is preferred to w' when the defaults $V(w)$ violated by w are less serious than the defaults $V(w')$ violated by w' . As we will recall below, the seriousness ordering also depends on the number of defaults violated by w and by w' for each rank.

In a similar way, in the following, we introduce a ranked relation $<_i$ on the domain Δ of a model of K_i . Let us first define, for a preferential model $\mathcal{N}_i = \langle \Delta, <_i, \cdot^I \rangle$ of K_i , what it means that an element $x \in \Delta$ *violates* a typicality inclusion $\mathbf{T}(C) \sqsubseteq D$ in m_i .

Definition 5. *Given a module m_i of K , with $s(m_i) = C_i$, and a preferential model $\mathcal{N}_i = \langle \Delta, <_i, \cdot^I \rangle$ of K_i , an element $x \in \Delta$ violates a typicality inclusion $\mathbf{T}(C) \sqsubseteq D$ in m_i if $x \in C_i^I$, $x \in D^I$ and $x \notin D^I$.*

Notice that, the set of typicality inclusions violated by a domain element x in a model only depends on the interpretation \cdot^I of \mathcal{ALC} concepts, and on the defeasible inclusions in m_i . Furthermore, differently from the usual notion of violation in Lehmann's semantics, for a module m_i with subject C_i , we do not consider the violations of domain

elements $x \notin C_i^I$ (i.e., the domain elements x which are not C_i -instances are assumed not to violate any default in m_i). Let $V_i(x)$ be the set of the defeasible inclusions of m_i violated by domain element x , and let $V_i^h(x)$ be the set of all defeasible inclusions in m_i with rank h which are violated by domain element x .

In order to compare alternative sets of defaults, in [45] the seriousness ordering \prec among sets of defaults is defined by associating with each set of defaults $D \subseteq K$ a tuple of numbers $\langle n_0, n_1, \dots, n_r \rangle$, where r is the *order* of K , i.e. the least finite i such that there is no default with the finite rank r or rank higher than r (but there is at least one default with rank $r - 1$). The tuple is constructed considering the ranks of defaults in the rational closure. n_0 is the number of defaults in D with rank ∞ and, for $1 \leq i \leq k$, n_i is the number of defaults in D with rank $r - i$ (in particular, n_r is the number of defaults in D with rank 0). Lehmann defines the strict modular order \prec among sets of defaults from the natural lexicographic order over the tuples $\langle n_0, n_1, \dots, n_r \rangle$. This order gives preference to those sets of defaults containing a larger number of more specific defaults. As we have seen from equation (1), \prec is used by Lehmann to compare sets of violated defaults and to prefer the propositional models whose violations are less serious.

We use the same criterion for comparing domain elements, introducing a seriousness ordering \prec_i for each module m_i . Considering that the defaults with infinite rank must be satisfied by all domain elements, we will not need to consider their violation in our definition (that is, we will not consider n_0 in the following).

The set $V_i(x)$ of defaults from module m_i which are violated by x , can be associated with a tuple of numbers $t_{i,x} = \langle |V_i^{r-1}(x)|, \dots, |V_i^0(x)| \rangle$, where $V_i^l(x)$ is the number of defaults in D with rank l which are violated by x . Following Lehmann, we let $V_i(x) \prec_i V_i(y)$ iff $t_{i,x}$ comes before $t_{i,y}$ in the natural lexicographic order on tuples (restricted to the violations of defaults in m_i), that is:

$$V_i(x) \prec_i V_i(y) \quad \text{iff} \quad \exists l \text{ such that } |V_i^l(x)| < |V_i^l(y)| \\ \text{and, } \forall h > l, |V_i^h(x)| = |V_i^h(y)|$$

Definition 6. A preferential model $\mathcal{N}_i = \langle \Delta, \prec_i, \cdot^I \rangle$ of $K_i = \langle \mathcal{T} \cup m_i, \mathcal{A} \rangle$, is a lexicographic model of K_i if $\langle \Delta, \cdot^I \rangle$ is an ALCC model of $\langle \mathcal{T}, \mathcal{A} \rangle$ and \prec_i satisfies the following condition: $x \prec_i y$ iff $V_i(x) \prec_i V_i(y)$.

Informally, \prec_{C_j} gives higher preference to domain elements violating less typicality inclusions of m_i with higher rank. In particular, all $x, y \notin C_i^I$, $x \sim_{C_i} y$, i.e., all $\neg C_i$ -elements are assigned the same preference wrt \prec_i , the least one, as they trivially satisfy all the typicality properties in m_i . As in Lehmann's semantics, in a lexicographic model $\mathcal{N}_i = \langle \Delta, \prec_i, \cdot^I \rangle$ of K_i , the preference relation \prec_i is a strict *modular* partial order, i.e. an irreflexive, transitive and modular relation. As well-foundedness trivially holds for finite interpretations, a lexicographic model \mathcal{N}_i of K_i is a ranked model of K_i .

A multi-concept model for K can be defined as a multi-preference interpretation with a preference relation \prec_i for each module m_i .

Definition 7 (Multi-concept interpretation). Let $K = \langle \mathcal{T}, \mathcal{D}, m_1, \dots, m_k, \mathcal{A}, s \rangle$ be a multi-concept knowledge base. A multi-concept interpretation \mathcal{M} for K is a tuple $\langle \Delta, \prec_1, \dots, \prec_k, \cdot^I \rangle$ such that, for all $i = 1, \dots, k$, $\langle \Delta, \prec_i, \cdot^I \rangle$ is a ranked ALCC + T interpretation, as defined in Section 2.

Definition 8 (Multi-concept lexicographic model). Let $K = \langle \mathcal{T}, \mathcal{D}, m_1, \dots, m_k, \mathcal{A}, s \rangle$ be a multi-concept knowledge base. A multi-concept lexicographic model $\mathcal{M} = \langle \Delta, <_1, \dots, <_k, \cdot^I \rangle$ of K is a multi-concept interpretation for K , such that, for all $i = 1, \dots, k$, $\mathcal{N}_i = \langle \Delta, <_i, \cdot^I \rangle$ is a lexicographic model of $K_i = \langle \mathcal{T} \cup m_i, \mathcal{A} \rangle$.

A canonical multi-concept lexicographic model of K is multi-concept lexicographic model of K such that Δ and \cdot^I are the domain and interpretation function of some canonical preferential model of $\langle \mathcal{T} \cup \mathcal{D}, \mathcal{A} \rangle$, according to Definition 3 (see [36]).

Observe that, restricting to the propositional fragment of the language (which does not allow universal and existential restrictions nor assertions), for a knowledge base K without strict inclusions and with a single module m_1 , with subject \top , containing all the typicality inclusions in K , the preference relation $<_1$ corresponds to Lehmann's lexicographic closure semantics, as its definition is based on the set of all defeasible inclusions in the knowledge base.

5 The combined lexicographic model of a KB

For multiple modules, each $<_i$ determines a ranked preference relation which can be used to answer *queries over module m_i* (i.e. queries whose subject is C_i). If we want to evaluate the query $\mathbf{T}(C) \sqsubseteq D$ (are all typical C elements also D elements?) in module m_i (assuming that C concerns subject C_i), we can answer the query using the $<_i$ relation, by checking whether $\min_{<_i}(C^I) \subseteq D^I$. For instance, in Example 1, the query “are all typical Phd students young?” can be evaluated in module m_2 . The answer would be positive, as the property of students of being normally young is inherited by PhD Student. The evaluation of a query in a specific module is something considered in context-based formalisms, such as in the CKR framework [8], where a language construct $eval(X, c)$ allows for evaluating a concept (or role) X in context c .

The lexicographic orders $<_i$ and $<_j$ (for $i \neq j$) do not need to agree. For instance, in Example 1, for two domain elements x and y , we might have that $x <_1 y$ and $y <_2 x$, as x is more typical than y as an employee, but less typical than x as a student. To answer a query $\mathbf{T}(C) \sqsubseteq D$, where C is a concept which is concerned with more than one subject in the knowledge base (e.g., are typical employed students young?), we need to *combine the relations* $<_i$.

A simple way of combining the modular partial order relations $<_i$ is to use Pareto combination. Let \leq_i be defined as follows: $x \leq_i y$ iff $y \not<_i x$. As $<_i$ is a modular partial order, \leq_i is a total preorder. Given a canonical multi-concept lexicographic model $\mathcal{M} = \langle \Delta, <_1, \dots, <_k, \cdot^I \rangle$ of K , we define a global preference relation $<$ on Δ as follows:

$$x < y \text{ iff } \begin{array}{l} (i) \text{ for some } i = 1, \dots, k, x <_i y \text{ and} \\ (ii) \text{ for all } j = 1, \dots, k, x \leq_j y, \end{array} \quad (*)$$

The resulting relation $<$ is a partial order, while modularity may not hold for $<$.

Definition 9. Given a canonical multi-concept lexicographic model $\mathcal{M} = \langle \Delta, <_1, \dots, <_k, \cdot^I \rangle$ of K , the combined lexicographic interpretation of \mathcal{M} , is a triple $\mathcal{M}^P = \langle \Delta, <, \cdot^I \rangle$, where $<$ is the global preference relation defined by (*).

We call \mathcal{M}^P a *combined lexicographic model of K* (shortly, an m_i^c -model of K).

Proposition 1. *A combined lexicographic model \mathcal{M}^P of K is a preferential interpretation satisfying all the strict inclusions and assertions in K .*

A combined lexicographic model \mathcal{M}^P of K is a preferential interpretation as those defined for $\mathcal{ALC} + \mathbf{T}$ in Definition 2 (and, in general, it is not a ranked interpretation). However, preference relation $<$ in \mathcal{M}^P is not an arbitrary irreflexive, transitive and well-founded relation. It is obtained by first computing the lexicographic preference relations $<_i$ for modules, and then by combining them into $<$. As \mathcal{M}^P satisfies all strict inclusions and assertions in K but is not required to satisfy all typicality inclusions $\mathbf{T}(C) \sqsubseteq D$ in K , \mathcal{M}^P is *not* a preferential $\mathcal{ALC} + \mathbf{T}$ model of K defined in Section 2.

Consider a situation in which there are two concepts, *Student* and *YoungPerson*, that are very related in that students are normally young persons and young persons are normally students (i.e., $\mathbf{T}(\textit{Student}) \sqsubseteq \textit{YoungPerson}$ and $\mathbf{T}(\textit{YoungPerson}) \sqsubseteq \textit{Student}$) and suppose there are two modules m_1 and m_2 such that $s(m_1) = \textit{Student}$ and $s(m_2) = \textit{YoungPerson}$. The two concepts may have different (and even contradictory) prototypical properties, for instance, normally students are quiet ($\mathbf{T}(\textit{Student}) \sqsubseteq \textit{Quiet}$), but normally young persons are not quiet ($\mathbf{T}(\textit{YoungPerson}) \sqsubseteq \neg \textit{Quiet}$). Considering the preference relations $<_1$ and $<_2$, associated with m_1 and m_2 in a canonical multi-concept lexicographic model, there may be two young persons and student Bob and John, such that $\textit{bob} <_1 \textit{john}$ and $\textit{john} <_2 \textit{bob}$, as Bob is quiet and John is not. Then, John and Bob are incomparable in the global relation $<$. Both of them, depending on the other prototypical properties of students and young persons, might be minimal, among students, wrt the global preference relation $<$.

In general, for a knowledge base K and a module m_i , with $s(m_i) = C_i$, the inclusion $\min_{<}(C_i^I) \subseteq \min_{<_i}(C_i^I)$ may not hold and, for this reason, a combined lexicographic interpretation may fail to satisfy all typicality inclusions. In this respect, canonical multi-concept lexicographic models are more liberal than KLM-style preferential models for typicality logics [31], where $\min_{<}(\textit{Student}^I) \sqsubseteq \textit{Quiet}^I$ must hold for the typicality inclusion to be satisfied. As a consequence, the knowledge base above has no preferential model according to the semantics in Section 2.

In [36] the notion of m_i^c -model of K has also been strengthened, by considering \mathbf{T} -compliant m_i^c -models (or $m_i^c\mathbf{T}$ -models) of K , which further satisfy the condition that, for all the typicality inclusions $\mathbf{T}(C) \sqsubseteq D$ in K , $(\mathbf{T}(C))^I = \min_{<}(C^I) \subseteq D^I$.

A notion of *multi-concept lexicographic entailment* (m_i^c -entailment) can be defined as usual: a query F is m_i^c -entailed by K ($K \models_{m_i^c} F$) if, for all m_i^c -models $\mathcal{M}^P = \langle \Delta, <, \cdot^I \rangle$ of K , F is satisfied in \mathcal{M}^P . Notice that a query $\mathbf{T}(C) \sqsubseteq D$ is satisfied in \mathcal{M}^P when $\min_{<}(C^I) \subseteq D^I$. Similarly, a notion of $m_i^c\mathbf{T}$ -entailment can be defined: $K \models_{m_i^c\mathbf{T}} F$ if, for all $m_i^c\mathbf{T}$ -models $\mathcal{M}^P = \langle \Delta, <, \cdot^I \rangle$ of K , F is satisfied in \mathcal{M}^P .

As, for any multi-concept knowledge base K , the set of $m_i^c\mathbf{T}$ -models of K is a subset of the set of m_i^c -models of K , and there is some K for which the inclusion is proper, $m_i^c\mathbf{T}$ -entailment is stronger than m_i^c -entailment. It can be proven that both notions of entailment satisfy the KLM postulates of preferential consequence relations, which can be reformulated for a typicality logic, considering that typicality inclusions $\mathbf{T}(C) \sqsubseteq D$ [30] stand for conditionals $C \sim D$ in KLM preferential logics [43, 44]. See also [7] for the formulation of KLM postulates in the Propositional Typicality Logic.

The notions of m_i^c -entailment and $m_i^c\mathbf{T}$ -entailment are not stronger than Lehmann's lexicographic closure in the propositional case. Let us consider again Example 1.

Example 2. Let us add another module m_4 with subject *Citizen* to the knowledge base K , plus the following additional axioms in \mathcal{T} :

$$\textit{Italian} \sqsubseteq \textit{Citizen} \quad \textit{French} \sqsubseteq \textit{Citizen} \quad \textit{Canadian} \sqsubseteq \textit{Citizen}$$

Module m_4 has subject *Citizen*, and contains the defeasible inclusions:

$$(d_{17}) \mathbf{T}(\textit{Italian}) \sqsubseteq \textit{DriveFast} \quad (d_{18}) \mathbf{T}(\textit{Italian}) \sqsubseteq \textit{HomeOwner}$$

Suppose the following typicality inclusion is also added to module m_2 :

$$(d_{19}) \mathbf{T}(\textit{PhDStudent}) \sqsubseteq \neg \textit{HomeOwner}$$

What can we conclude about typical Italian PhD students? We can see that neither the inclusion $\mathbf{T}(\textit{PhDStudent} \sqcap \textit{Italian}) \sqsubseteq \textit{HomeOwner}$ nor the inclusion $\mathbf{T}(\textit{PhDStudent} \sqcap \textit{Italian}) \sqsubseteq \neg \textit{HomeOwner}$ is m_i^c -entailed by K . In fact, as $<_2$ -minimal and $<_4$ -minimal *PhDStudent* \sqcap *Italian*-elements are incomparable with respect to $<$, the $<$ -minimal Italian PhD students will include them all. Thus, $\min_{<}((\textit{PhDStudent} \sqcap \textit{Italian})^I) \not\sqsubseteq \textit{HomeOwner}^I$ and $\min_{<}((\textit{PhDStudent} \sqcap \textit{Italian})^I) \not\sqsubseteq (\neg \textit{HomeOwner})^I$.

The home owner example is a reformulation of the example used by Geffner and Pearl to show that the rational closure of conditional knowledge bases sometimes gives too strong conclusions, as “conflicts among defaults that should remain unresolved, are resolved anomalously” [27]. Informally, if defaults (d_{18}) and (d_{19}) are conflicting for Italian Phd students before adding any default which makes PhD students exceptional wrt Students (in our formalization, default (d_{10})), they should remain conflicting after this addition. On the contrary, in the propositional case, both the rational closure [44] and Lehmann's lexicographic closure [45] would entail that normally Italian Phd students are not home owners. This conclusion is unwanted, and is based on the fact that (d_{18}) has rank 0, while (d_{19}) has rank 1 in the rational closure ranking. On the other hand, $\mathbf{T}(\textit{PhDStudent} \sqcap \textit{Italian}) \sqsubseteq \neg \textit{HomeOwner}$ is neither m_i^c -entailed from K , nor $m_i^c\mathbf{T}$ -entailed from K . Both notions of entailment, when restricted to the propositional case, cannot be stronger than Lehmann's lexicographic closure.

Geffner and Pearl's Conditional Entailment [27] does not suffer from the above mentioned problem as it is based on (non-ranked) preferential models. The same problem, which is related to the representation of preferences as levels of reliability, has also been recognized by Brewka [11] in his logical framework for default reasoning, leading to a generalization of the approach to allow a partial ordering between premises. The example above shows that our approach using ranked preferences for the single modules, but a non-ranked global preference relation $<$ for their combination, does not suffer from this problem, provided a suitable modularization is chosen.

6 Further issues: Reasoning with a hierarchy of modules and user-defined preferences

The approach considered in Section 4 does not allow to reason with a hierarchy of modules, but it considers a flat collection of modules m_1, \dots, m_k , each module concerning some subject C_i . As we have seen, a module m_i may contain defeasible inclusions referring to subclasses of C_i , such as *PhDStudent* in the case of module m_2 with subject

Student. When defining the preference relation $<_i$ the lexicographic closure semantics already takes into account the specificity relation among concepts within the module (e.g., the fact that *PhDStudent* is more specific than *Student*).

However, nothing prevents us from defining two modules m_i (with subject C_i) and m_j (with subject C_j), such that concept C_j is more specific than concept C_i . For instance, as a variant of Example 1, we might have introduced two different modules m_2 with subject *Student* and m_5 with subject *PhDStudent*. As concept *PhDStudent* is more specific than concept *Student* (in particular, $PhDStudent \sqsubseteq Student$ is entailed from the strict part of knowledge base \mathcal{T} in \mathcal{ALC}), the specificity information should be taken into account when combining the preference relations. More precisely, preference $<_5$ should override preference $<_2$ when comparing *PhDStudent*-instances.

This is the principle followed by Giordano and Theseider Dupré [35] to define a global preference relation, in the case when each module with subject C_i only contains typicality inclusions of the form $\mathbf{T}(C_i) \sqsubseteq D$. A more sophisticated way to combine the preference relations $<_i$ into a global relation $<$ is used to deal with this case with respect to Pareto combination, by exploiting the specificity relation among concepts. While we refer therein for a detailed description of this more sophisticated notion of preference combination, let us observe that this solution could be as well applied to the modular multi-concept knowledge bases considered in this paper, provided an irreflexive and transitive notion of specificity among modules is defined.

Another aspect that has been considered in the previously mentioned paper is the possibility of assigning ranks to the defeasible inclusions associated with a given concept. While assigning a rank to all typicality inclusions in the knowledge base may be awkward, often people have a clear idea about the relative importance of the properties for some specific concept. For instance, we may know that the defeasible property that students are normally young is more important than the property that student normally do not have a scholarship. For small modules, which only contain typicality inclusions $\mathbf{T}(C_i) \sqsubseteq D$ for a concept C_i , the specification of user-defined ranks of the C_i 's typical properties is a feasible option and a ranked modular preference relation can be defined from it, by using Brewka's $\#$ strategy from his framework of Basic Preference Descriptions for ranked knowledge bases [12]. This alternative may coexist with the use of the lexicographic closure semantics built from the rational closure ranking for large modules.

According to the choice of fine grained or coarse grained modules, to the choice of the preferential semantics for each module (e.g., based on user-specified ranking or on Lehmann's lexicographic closure, or on the rational closure, etc.), and to the presence of a specificity relation among modules, alternative preferential semantics for modularized multi-concept knowledge bases can emerge.

7 Conclusions and related work

In this paper, we have proposed a modular multi-concept extension of the lexicographic closure semantics, based on the idea that defeasible properties in the knowledge base can be distributed in different modules, for which alternative preference relations can be computed. Combining multiple preferences into a single global preference allows

a new preferential semantics and a notion of multi-concept lexicographic entailment (m_i^e -entailment) which, in the propositional case, is not stronger than the lexicographic closure. This work has been first presented in [36].

m_i^e -entailment satisfies the KLM postulates of a preferential consequence relation. It retains some good properties of the lexicographic closure, being able to deal with irrelevance, with specificity within the single modules, and not being subject to the “blockage of property inheritance” problem. The combination of different preference relations provides a simple solution to a problem, recognized by Geffner and Pearl, that the rational closure of conditional knowledge bases sometimes gives too strong conclusions, as “conflicts among defaults that should remain unresolved, are resolved anomalously” [27]. This problem also affects the lexicographic closure, which is stronger than the rational closure. Our approach using ranked preferences for the single modules, but a non-ranked preference $<$ for their combination, does not suffer from this problem, provided a suitable modularization is chosen. As Geffner and Pearl’s Conditional Entailment [27], also some non-monotonic DLs, such as $\mathcal{ALC} + \mathbf{T}_{min}$, a typicality DL with a minimal model preferential semantics [32], and the non-monotonic description logic \mathcal{DL}^N [5], which supports normality concepts based on a notion of overriding, do not suffer from the problem above.

Reasoning about exceptions in ontologies has led to the development of many non-monotonic extensions of Description Logics (DLs), incorporating non-monotonic features from most of NMR formalisms in the literature. In addition to those already mentioned in the introduction, let us recall the work by Straccia on inheritance reasoning in hybrid KL-One style logics [48] the work on defaults in DLs [3], on description logics of minimal knowledge and negation as failure [24], on circumscriptive DLs [6], the generalization of rational closure to all description logics [4], as well as the combination of description logics and rule-based languages [26, 25, 46, 42, 40, 9].

The lexicographic closure for DLs has been first investigated by Casini and Straccia [19, 21]. Our multipreference semantics is related with the multipreference semantics for \mathcal{ALC} developed by Gliozzi [39], which is based on the idea of refining the rational closure construction considering the preference relations $<_{A_i}$ associated with different aspects. We follow a different route concerning the definition of the preference relations associated with modules, and the way of combining them in a single preference relation.

Starting from Brewka’s framework of basic preference descriptions [12], multiple preferences have also been used under different approaches: in system ARS, a refinement of System Z developed by Kern-Isberner and Ritterskamp [41], through preference fusion; by Gil [28] to define a multipreference formulation of the typicality DL $\mathcal{ALC} + \mathbf{T}_{min}$, mentioned above; by Britz and Varzinczak [16, 14], by associating multiple preferences to roles; in the first-order logic setting, by Delgrande and Rantsaudis [23]; in ranked \mathcal{EL}^\perp knowledge bases, by Giordano and Theseider Dupré [35].

Bozzato et al. present extensions of the CKR (Contextualized Knowledge Repositories) framework by Bozzato et al. [8, 9] in which defeasible axioms are allowed in the global context and exceptions can be handled by overriding and have to be justified in terms of semantic consequence, considering sets of clashing assumptions for each defeasible axiom. An extension of this approach to deal with general contextual hierarchies has been studied by the same authors [10], by introducing a coverage relation among

contexts, and defining a notion of preference among clashing assumptions, which is used to define a preference relation among justified CAS models, based on which CKR models are selected. An ASP based reasoning procedure, that is complete for instance checking, is developed for *SRQIQ-RL*.

For the lightweight description logic \mathcal{EL}_{\perp}^+ , an Answer Set Programming (ASP) approach has been proposed [35] for defeasible inference in a multipreference extension of \mathcal{EL}_{\perp}^+ , in the specific case in which each module only contains the defeasible inclusions $\mathbf{T}(C_i) \sqsubseteq D$ for a single concept C_i , where the ranking of defeasible inclusions is specified in the knowledge base, following the approach by Gerhard Brewka in his framework of Basic Preference Descriptions for ranked knowledge bases [12]. A specificity relation among concepts is also considered. The ASP encoding exploits *asprin* [13], by formulating multipreference entailment as a problem of computing preferred answer sets, which is proved to be Π_2^P -complete. A similar encoding has been developed for defeasible reasoning with weighted conditional \mathcal{EL}_{\perp}^+ knowledge bases (in the two-valued case) [37], a formalism that has been introduced for capturing the logical semantics of Multilayer Perceptrons [38].

For \mathcal{EL}_{\perp}^+ knowledge bases, we aim at extending this ASP encoding to deal with the modular multi-concept lexicographic closure semantics proposed in this paper, as well as with a more general framework, allowing for different choices of preferential semantics for the single modules and for different specificity relations for combining them. For lightweight description logics of the \mathcal{EL} family [1], the ranking of concepts determined by the rational closure construction can be computed in polynomial time in the size of the knowledge base [34, 22]. This suggests that we may expect a Π_2^P upper-bound on the complexity of multi-concept lexicographic entailment.

Acknowledgement: This research is partially supported by INDAM-GNCS Projects 2019 and 2020.

References

1. F. Baader, S. Brandt, and C. Lutz. Pushing the \mathcal{EL} envelope. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI 2005)*, pages 364–369, Edinburgh, Scotland, UK, August 2005.
2. F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, and P.F. Patel-Schneider. *The Description Logic Handbook - Theory, Implementation, and Applications*. Cambridge, 2007.
3. F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. *J. Autom. Reasoning*, 14(1):149–180, 1995.
4. P. A. Bonatti. Rational closure for all description logics. *Artif. Intell.*, 274:197–223, 2019.
5. P. A. Bonatti, M. Faella, I. Petrova, and L. Sauro. A new semantics for overriding in description logics. *Artif. Intell.*, 222:1–48, 2015.
6. P. A. Bonatti, M. Faella, and L. Sauro. Defeasible inclusions in low-complexity dls. *J. Artif. Intell. Res. (JAIR)*, 42:719–764, 2011.
7. Richard Booth, Giovanni Casini, Thomas Meyer, and Ivan Varzinczak. On rational entailment for propositional typicality logic. *Artif. Intell.*, 277, 2019.
8. L. Bozzato, T. Eiter, and L. Serafini. Contextualized knowledge repositories with justifiable exceptions. In *DL 2014*, volume 1193 of *CEUR Workshop Proceedings*, pages 112–123.
9. L. Bozzato, T. Eiter, and L. Serafini. Enhancing context knowledge repositories with justifiable exceptions. *Artif. Intell.*, 257:72–126, 2018.

10. L. Bozzato, T. Eiter, and L. Serafini. Justifiable exceptions in general contextual hierarchies. In *CONTEXT 2019, Trento, Italy, November 20-22, 2019, Proceedings*, volume 11939 of *Lecture Notes in Computer Science*, pages 26–39. Springer, 2019.
11. G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence. Detroit, MI, USA, August 1989*, pages 1043–1048, 1989.
12. G. Brewka. A rank based description language for qualitative preferences. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI'2004, Valencia, Spain, August 22-27, 2004*, pages 303–307, 2004.
13. G. Brewka, J. P. Delgrande, J. Romero, and T. Schaub. asprin: Customizing answer set preferences without a headache. In *Proc. AAAI 2015*, pages 1467–1474, 2015.
14. A. Britz and I. Varzinczak. Contextual rational closure for defeasible ALC (extended abstract). In *Proc. 32nd International Workshop on Description Logics, Oslo, Norway, June 18-21, 2019*, 2019.
15. K. Britz, J. Heidema, and T. Meyer. Semantic preferential subsumption. In G. Brewka and J. Lang, editors, *KR 2008*, pages 476–484, Sidney, Australia, September 2008. AAAI Press.
16. K. Britz and I. J. Varzinczak. Rationality and context in defeasible subsumption. In *Proc. 10th Int. Symp. on Found. of Information and Knowledge Systems, FoIKS 2018, Budapest, May 14-18, 2018*, pages 114–132, 2018.
17. G. Casini, T. Meyer, K. Moodley, and R. Nortje. Relevant closure: A new form of defeasible reasoning for description logics. In *JELIA 2014*, LNCS 8761, pages 92–106. Springer, 2014.
18. G. Casini, T. Meyer, I. J. Varzinczak, and K. Moodley. Nonmonotonic Reasoning in Description Logics: Rational Closure for the ABox. In *DL 2013*, volume 1014 of *CEUR Workshop Proceedings*, pages 600–615, 2013.
19. G. Casini and U. Straccia. Rational Closure for Defeasible Description Logics. In T. Janhunen and I. Niemelä, editors, *JELIA 2010*, volume 6341 of *LNCS*, pages 77–90, Helsinki, Sept. 2010. Springer.
20. G. Casini and U. Straccia. Lexicographic Closure for Defeasible Description Logics. In *Proc. of Australasian Ontology Workshop, vol.969*, pages 28–39, 2012.
21. G. Casini and U. Straccia. Defeasible inheritance-based description logics. *Journal of Artificial Intelligence Research (JAIR)*, 48:415–473, 2013.
22. G. Casini, U. Straccia, and T. Meyer. A polynomial time subsumption algorithm for nominal safe ELO_{\perp} under rational closure. *Inf. Sci.*, 501:588–620, 2019.
23. J. Delgrande and C. Rantsoudis. A preference-based approach for representing defaults in first-order logic. In *18th Int. Workshop on Non-Monotonic Reasoning, NMR2020, Sept. 12th - 14th, 2020*.
24. F. M. Donini, D. Nardi, and R. Rosati. Description logics of minimal knowledge and negation as failure. *ACM Transactions on Computational Logic (ToCL)*, 3(2):177–225, 2002.
25. T. Eiter, G. Ianni, T. Lukasiewicz, and R. Schindlauer. Well-founded semantics for description logic programs in the semantic web. *ACM Trans. Comput. Log.*, 12(2):11, 2011.
26. T. Eiter, G. Ianni, T. Lukasiewicz, R. Schindlauer, and H. Tompits. Combining answer set programming with description logics for the semantic web. *Artif. Intell.*, 172(12-13):1495–1539, 2008.
27. Hector Geffner and Judea Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artif. Intell.*, 53(2-3):209–244, 1992.
28. Oliver Fernandez Gil. On the Non-Monotonic Description Logic $ALC+T_{\min}$. *CoRR*, abs/1404.6566, 2014.
29. L. Giordano and V. Gliozzi. Reasoning about exceptions in ontologies: An approximation of the multipreference semantics. In *ECSQARU 2019, Belgrade, Serbia, September 18-20, 2019, Proceedings*, volume 11726 of *LNCS*, pages 212–225. Springer, 2019.

30. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. Preferential Description Logics. In *LPAR 2007*, volume 4790 of *LNAI*, pages 257–272, Yerevan, Armenia, October 2007.
31. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. ALC+T: a preferential extension of Description Logics. *Fundamenta Informaticae*, 96:1–32, 2009.
32. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. A NonMonotonic Description Logic for Reasoning About Typicality. *Artif. Intell.*, 195:165–202, 2013.
33. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. Semantic characterization of rational closure: From propositional logic to description logics. *Artif. Intell.*, 226:1–33, 2015.
34. L. Giordano and D. Theseider Dupré. Defeasible Reasoning in SROEL from Rational Entailment to Rational Closure. *Fundam. Inform.*, 161(1-2):135–161, 2018.
35. L. Giordano and D. Theseider Dupré. An ASP approach for reasoning in a concept-aware multipreferential lightweight DL. *Theory and Practice of Logic programming, TPLP*, 10(5):751–766, 2020.
36. L. Giordano and D. Theseider Dupré. A framework for a modular multi-concept lexicographic closure semantics. In *18th Int. Workshop on Non-Monotonic Reasoning, NMR2020, September 12th - 14th, 2020*. <https://nmr2020.dc.uba.ar/WorkshopNotes.pdf>.
37. L. Giordano and D. Theseider Dupré. Weighted conditional \mathcal{EL}^+ knowledge bases with integer weights: an ASP approach. In *ICLP 2021*, 2021. To appear.
38. L. Giordano and D. Theseider Dupré. Weighted defeasible knowledge bases and a multi-preference semantics for a deep neural network model. In *JELIA 2021, May 17-20*, volume 12678 of *LNCS*, pages 225–242. Springer, 2021.
39. V. Gliozzi. Reasoning about multiple aspects in rational closure for DLs. In *Proc. AI*IA 2016, Genova, Italy, Nov 29 - Dec 1, 2016*, pages 392–405, 2016.
40. G. Gottlob, A. Hernich, C. Kupke, and T. Lukasiewicz. Stable model semantics for guarded existential rules and description logics. In *Proc. KR 2014*, 2014.
41. G. Kern-Isberner and M. Ritterskamp. Preference fusion for default reasoning beyond system Z. *J. Autom. Reasoning*, 45(1):3–19, 2010.
42. M. Knorr, P. Hitzler, and F. Maier. Reconciling owl and non-monotonic rules for the semantic web. In *ECAI 2012*, page 474479, 2012.
43. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):167–207, 1990.
44. D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, 1992.
45. D. J. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.
46. B. Motik and R. Rosati. Reconciling Description Logics and rules. *Journal of the ACM*, 57(5), 2010.
47. J. Pearl. System Z: A Natural Ordering of Defaults with Tractable Applications to Non-monotonic Reasoning. In *TARK (3rd Conference on Theoretical Aspects of Reasoning about Knowledge)*, pages 121–135, Pacific Grove, CA, USA, 1990. Morgan Kaufmann.
48. U. Straccia. Default inheritance reasoning in hybrid kl-one-style logics. In R. Bajcsy, editor, *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI 1993)*, pages 676–681, Chambéry, France, August 1993. Morgan Kaufmann.