Modelling object typicality in Description Logics

Katarina Britz[1,3], Johannes Heidema[2], and Thomas Meyer[1,3]

[1] KSG, Meraka Institute, CSIR, South Africa
[2] Dept of Mathematical Sciences, University of South Africa
[3] School of Computing, University of South Africa
arina.britz@meraka.org.za
johannes.heidema@gmail.com
tommie.meyer@meraka.org.za

Abstract. We present a semantic model of typicality of concept members in description logics (DLs) that accords well with a binary, globalist cognitive model of class membership and typicality. We define a general preferential semantic framework for reasoning with object typicality in DLs. We propose the use of feature vectors to rank concept members according to their defining and characteristic features, which provides a modelling mechanism to specify typicality in composite concepts.

1 Introduction

The study of natural language concepts in cognitive psychology has led to a range of hypotheses and theories regarding cognitive constructions such as concept inclusion, composition, and typicality. Description logics (DLs) have been very successful in modelling some of these cognitive constructions, for example IS-A and PART-OF. In this paper, we focus on the semantic modelling of typicality of concept members in such a way that it accords well with empirically well-founded cognitive theories of how people construct and reason about concepts involving typicality. We do not attempt to survey all models of concept typicality, but briefly outline some aspects of the debate:

According to the *unitary model* of concept typicality and class membership, variations in both graded class membership and typicality of class members reflect differences in similarity to a concept prototype. Class membership and typicality are determined by placing some criterion on the similarity of objects to the concept prototype [10, 11]. According to the *binary model* of concept typicality and class inclusion, typicality and concept membership reflect essentially different cognitive processes. Concepts have *defining features* providing necessary and sufficient conditions for class membership, as well as *characteristic features* indicating typicality within that class [4, 17, 18]. According to the *localist* view of concepts, the meaning of a compound concept is a function of the meanings of its semantic constituents. According to the *globalist* view, the meanings of concepts are entrenched in our world knowledge, which is context-dependent and cannot be decomposed into, or composed from, our understanding of basic building blocks [12, 16]. Concept typicality can therefore not be determined

[1] A version of this paper will appear in the LNCS series. The original publication will be available at www.springerlink.com.

from concept definition alone, but requires a world view to provide context relative to which typicality may be determined.

Description logics cannot resolve any of these debates, but we can use DLs to model some aspects of them. In particular, we can model typicality of concept members based on their characteristic features. We can also model compositional aspects of typicality. Other aspects, such as the graded class membership that underpins the unitary model, and non-compositionality of compound class membership in the globalist view, cannot be modelled using DLs, or at least not in an intuitively natural way. In [21] a model of graded concept membership was proposed, but this presented a marked departure from classical DL reasoning. We therefore restrict our attention to the binary model, with a compositional model of class membership, where being a member of a class is an all-or-nothing affair, and membership of compound concepts are determined by membership of their atomic constituents or defining features, while characteristic features contribute to induce degrees of typicality within a class.

DLs have gained wide acceptance as underlying formalism in intelligent knowledge systems over complex structured domains, providing an unambiguous semantics to ontologies, and balancing expressive power with efficient reasoning mechanisms [1]. The nature of DL reasoning has traditionally been *deductive*, but there have been a fair number of proposals to extend DLs to incorporate some form of *defeasible* reasoning, mostly centered around the incorporation of some form of default rules, e.g. [5].

In a previous paper [3], we presented a general preferential semantic framework for defeasible subsumption in DLs, analogous to the KLM preferential semantics for propositional entailment [2, 13]. We gave a formal semantics of defeasible subsumption, as well as a translation of defeasible subsumption to classical subsumption within a suitably rich DL language. This was done by defining a preference order on objects in a knowledge base, which allowed for defeasible terminological statements of the form "All the most preferred objects in C are also in D".

In practice, an ontology may call for different preference orders on objects, and correspondingly, multiple defeasible subsumption relations within a single knowledge base. An object may be typical (or preferable) with respect to one property, but not another. For example, a guppy may be considered a typical pet fish, even though it is neither a typical fish, nor a typical pet [17]. So we may want a pet typicality order on pets, a fish typicality order on fish, and some way of combining these orders, or other relevant characteristics, into a pet fish typicality order. That is, we want to order objects in a given class according to their typicality with respect to the chosen features of that class. The subjective world view adopted in the fish shop may be different from that adopted in an aquarium, or a pet shop, hence the features deemed relevant may differ in each case, and this has to be reflected in the respective typicality orders.

Relative to a particular interpretation of a DL, any concept C partitions all objects in the domain according to their class membership into those belonging to C, and those not belonging to C. This yields a two-level preference order, with all objects in C preferred to all those not in C. This order may be refined further to distinguish amongst objects in C, but even the basic two-level order suffices to define an important class of preferential subsumption relations, namely those characterising the stereotypical reasoning of [14].

A preference order on objects may be employed to obtain a notion of defeasible subsumption that relaxes the deductive nature of classical subsumption. To this end, we introduce a parameterised defeasible subsumption relation $\sqsubsetsim_j$ to express terminological statements of the form $C \sqsubsetsim_j D$, where C and D are arbitrary concepts, and $\sqsubsetsim_j$ is induced by a preference order $\leq_j$. If $\leq_j$ prefers objects in A to objects outside of A, we say that C is preferentially subsumed by D relative to A iff all objects in C that are typical in A (i.e. preferred by the typicality order corresponding to A), are also in D. When translated into DL terminology, the proposal of [14] reads as follows: Given concepts C, D and S such that S represents a best stereotype of C, C is preferentially subsumed by D relative to S if all stereotypical objects in C also belong to D.

The rest of the paper is structured as follows: We first fix some standard semantic terminology on DLs that will be useful later on. After giving some background on rational preference orders, we introduce the notion of an ordered interpretation, and present a formal semantics of parameterised defeasible subsumption. This is a natural extension of the work presented in [3], and provides a way of reasoning defeasibly with the IS-A relationship between concepts relative to a given concept. We then put forward two approaches to the definition of a derived typicality order on concepts, namely *atomic composition* and *feature composition*. We argue that feature composition is the more general approach, and is not as vulnerable to arguments against compositionality as is the case with atomic composition. We show how feature vectors may be used to determine typicality compositionally, taking into account semantic context.

2 Preliminaries

2.1 DL terminology

In the standard set-theoretic semantics of concept descriptions, concepts are interpreted as subsets of a domain of interest, and roles as binary relations over this domain. An interpretation I consists of a non-empty set Δ^I (the *domain* of I) and a function $\cdot^I$ (the *interpretation function* of I) which maps each atomic concept A to a subset A^I of Δ^I, and each atomic role R to a subset R^I of $\Delta^I \times \Delta^I$. The interpretation function is extended to arbitrary concept descriptions (and role descriptions, if complex role descriptions are allowed in the language) in the usual way.

A DL knowledge base consists of a *Tbox* which contains *terminological axioms*, and an *Abox* which contains *assertions*, i.e. facts about specific named objects and relationships between objects in the domain. Depending on the expressive power of the DL, a knowledge base may also have an *Rbox* which contains *role axioms*. Tbox statements are *concept inclusions* of the form $C \sqsubseteq D$, where C and D are (possibly complex) concept descriptions. $C \sqsubseteq D$ is also called a *subsumption statement*, read "C is subsumed by D". An interpretation I *satisfies* $C \sqsubseteq D$, written $I \Vdash C \sqsubseteq D$, iff $C^I \subseteq D^I$. $C \sqsubseteq D$ is *valid*, written $\models C \sqsubseteq D$, iff it is satisfied by all interpretations. Rbox statements include *role inclusions* of the form $R \sqsubseteq S$, and assertions used to define *role properties* such as asymmetry. Objects named in the Abox are referred to by a finite number of *individual names*. These names may be used in two types of assertional statements – *concept assertions* of the form $C(a)$ and *role assertions* of the form $R(a, b)$, where C is a concept description, R is a role description, and a and b are individual names.

To provide a semantics for Abox statements it is necessary to add to every interpretation a *denotation function* which satisfies the unique names assumption, mapping each individual name a to a different element a^I of the domain of interpretation Δ^I. An interpretation I satisfies the assertion $C(a)$ iff $a^I \in C^I$; it satisfies $R(a,b)$ iff $(a^I, b^I) \in R^I$. An interpretation I satisfies a DL knowledge base $\mathcal{K}$ iff it satisfies every statement in $\mathcal{K}$. A DL knowledge base $\mathcal{K}$ *entails* a DL statement ϕ, written as $\mathcal{K} \models \phi$, iff every interpretation that satisfies $\mathcal{K}$ also satisfies ϕ.

2.2 Preferential semantics

In a preferential semantics for a propositional language, one assumes some order relation on propositional truth valuations (or on interpretations or worlds or, more generally, on states) to be given. The intuitive idea captured by the order relation is that interpretations higher up (greater) in the order are more typical in the context under consideration, than those lower down. For any given class C, we assume that all objects in the application domain that are in (the interpretation of) C are more typical of C than those not in C. This is a technical construction which allows us to order the entire domain, instead of only the members of C. This leads us to take as starting point a finite set of preference orders $\{\leq_j : j \in \mathcal{J}\}$ on objects in the application domain, with index set $\mathcal{J}$. If $\leq_j$ prefers any object in C to any object outside of C, we call $\leq_j$ a C*-order*.

To ensure that the subsumption relations generated are *rational*, i.e. satisfy a weak form of strengthening on the left, the *rational monotonicity* postulate (see [6, 15], we assume the preference orders to be *modular partial orders*, i.e. reflexive, transitive, antisymmetric relations such that, for all a, b, c in Δ^I, if a and b are incomparable and a is strictly below c, then b is also strictly below c.

Modular partial orders have the effect of stratifying the domain into layers, with any two elements in the same layer being unrelated to each other, and any two elements in different layers being related to each other. (We could also have taken the preference order to be a total preorder, i.e. a reflexive, transitive relation such that, for all a, b in Δ^I, a and b are comparable. Since there is a bijection between modular partial orders and total preorders on Δ^I, it makes no difference here which formalism we choose.)

We further assume that the order relations have no infinite chains (and hence, in Shoham's terminology [20, p.75], are bounded, which is the dual of well-founded, which in turn implies, in the terminology of [13], that the order relations are smooth). In the presence of transitivity, this implies that, for any $j \in \mathcal{J}$, nonempty $X \subseteq \Delta^I$ and $a \in X$, there is an element $b \in X$, $\leq_j$-maximal in X, with $a \leq_j b$.

3 Preferential subsumption

We now develop a formal semantics for preferential subsumption in DLs. We assume a DL language with a finite set of preference orders $\{\preceq_j : j \in \mathcal{J}\}$ in its signature. We make the preference orders on the domain of interpretation explicit through the notion of an *ordered interpretation*: $(I, \{\leq_j : j \in \mathcal{J}\})$ is the interpretation I with preference orders $\{\leq_j : j \in \mathcal{J}\}$ on the domain Δ^I. The preference orders on domain elements may be constrained by means of role assertions of the form $a \preceq_j b$ for $j \in \mathcal{J}$, where the interpretation of $\preceq_j$ is $\leq_j$, that is, $\preceq_j^I = \leq_j$:

Definition 1. *An ordered interpretation* $(I, \{\leq_j : j \in \mathcal{J}\})$ *consists of an interpretation* I *and finite, indexed set of modular partial orders* $\{\leq_j : j \in \mathcal{J}\}$ *without infinite chains over their domain* Δ^I.

Definition 2. *An ordered interpretation* $(I, \{\leq_j : j \in \mathcal{J}\})$ *satisfies an assertion* $a \preceq_j b$ *iff* $a^I \leq_j b^I$.

We do not make any further assumptions about the DL language, but assume that concept and role assertions and constructors, and classical subsumption are interpreted in the standard way, ignoring the preference orders of ordered interpretations.

We first introduce the notion of satisfaction by an ordered interpretation, thereafter we relax the semantics of concept inclusion to arrive at a definition of satisfaction of a parameterised preferential subsumption relation $\sqsubsim_j$ by an ordered interpretation. Finally, we define what it means for a preferential subsumption statement to be entailed by a knowledge base.

3.1 Satisfaction of preferential subsumption statements

Definition 3. *An ordered interpretation* $(I, \{\leq_j : j \in \mathcal{J}\})$ *satisfies* $C \sqsubseteq D$, *written* $(I, \{\leq_j : j \in \mathcal{J}\}) \Vdash C \sqsubseteq D$, *iff* I *satisfies* $C \sqsubseteq D$.

The preferential semantics of $\sqsubsim_j$ is then defined as follows:

Definition 4. *An ordered interpretation* $(I, \{\leq_j : j \in \mathcal{J}\})$ *satisfies the preferential subsumption* $C \sqsubsim_j D$, *written* $(I, \{\leq_j : j \in \mathcal{J}\}) \Vdash C \sqsubsim_j D$, *iff* $C_j^I \subseteq D^I$, *where*

$$C_j^I = \{x \in C^I : \text{there is no } y \in C^I \text{ such that } x \leq_j y \text{ and } x \neq y\}.$$

For brevity, we shall at times write $\leq_{\mathcal{J}}$ instead of $\{\leq_j : j \in \mathcal{J}\}$. Preferential subsumption satisfies the following three properties:

- Supraclassicality: If $(I, \leq_{\mathcal{J}}) \Vdash C \sqsubseteq D$ then $(I, \leq_{\mathcal{J}}) \Vdash C \sqsubsim_j D$ for all $j \in \mathcal{J}$.
- Nonmonotonicity: $(I, \leq_{\mathcal{J}}) \Vdash C \sqsubsim_j D$ does not necessarily imply $(I, \leq_{\mathcal{J}}) \Vdash C \sqcap C' \sqsubsim_j D$ for any $j \in \mathcal{J}$.
- Defeasibility: $(I, \leq_{\mathcal{J}}) \Vdash C \sqsubsim_j D$ does not necessarily imply $(I, \leq_{\mathcal{J}}) \Vdash C \sqsubseteq D$ for any $j \in \mathcal{J}$.

It also satisfies the familiar properties of rational preferential entailment [13, 15] (when expressible in the DL under consideration): Reflexivity, And, Or, Left Logical Equivalence, Left Defeasible Equivalence, Right Weakening, Cautious Monotonicity, Rational Monotonicity, and Cut.

3.2 Entailment of preferential subsumption statements

Satisfaction for defeasible subsumption is defined relative to a fixed, ordered interpretation. We now take this a step further, and develop a general semantic theory of entailment relative to a knowledge base using ordered interpretations. Note that, although the knowledge base may contain preferential subsumption statements, entailment from the knowledge base is classical and monotonic.

Definition 5. *The preferential subsumption statement* $C \precsim_j D$ *is valid, written* $\models C \precsim_j D$*, iff it is satisfied by all ordered interpretations* $(I, \{\leq_j : j \in \mathcal{J}\})$*.*

Definition 6. *A DL knowledge base* $\mathcal{K}$ *entails the preferential subsumption statement* $C \precsim_j D$*, written* $\mathcal{K} \models C \precsim_j D$*, iff every ordered interpretation that satisfies* $\mathcal{K}$ *also satisfies* $C \precsim_j D$*.*

The following properties of $\precsim_j$ are direct consequences of its corresponding properties relative to a fixed, ordered interpretation:

$\precsim_j$ is supraclassical: If $\mathcal{K} \models C \sqsubseteq D$ then also $\mathcal{K} \models C \precsim_j D$.
$\precsim_j$ is nonmonotonic: $\mathcal{K} \models C \precsim_j D$ does not necessarily imply that $\mathcal{K} \models C \sqcap C' \precsim_j D$.
$\precsim_j$ is defeasible: $\mathcal{K} \models C \precsim_j D$ does not necessarily imply that $\mathcal{K} \models C \sqsubseteq D$.

The other properties of $\precsim_j$ mentioned earlier relative to a fixed, ordered interpretation extend analogously in the context of entailment relative to a knowledge base. For example, reflexivity of $\precsim_j$ relative to $\mathcal{K}$ reads $\mathcal{K} \models C \precsim_j C$.

4 Derived typicality of concept membership

In the previous section we presented a semantic framework to model typicality of concept membership: $\leq_j$ is a C-order if it ranks any object in C higher than any object outside of C. In a DL with value restrictions, we can write this as: $C \sqsubseteq \forall \preceq_j .C$. We now address the question of derived typicality C-orders. We distinguish between two possible approaches to resolve this problem:

1. *Atomic composition*: Here we use the atomic constituents or defining features of the compound concept C as building blocks. We combine their respective typicality orders recursively, depending on the operators used in the syntactic construction of C. Say $C \equiv A \sqcap B$, and typicality orders $\leq_j$ and $\leq_k$ are defined such that $\leq_j$ is an A-order and $\leq_k$ is a B-order respectively. We may then form a new typicality order for C by composing $\leq_j$ and $\leq_k$ according to some composition rule for $\sqcap$.
2. *Feature composition*: Here we identify the relevant features of the concept C. For each object a belonging to C, we form a feature vector characterising a. These feature vectors are then used to determine the typicality of a in C.

Irrespective of the composition rules applied, atomic composition is vulnerable to the same criticisms that have been levied against localist, compositional cognitive models of typicality of concept membership [16].

Feature composition is also compositional, but, in contrast with atomic composition, it is not localist. That is, the typicality of a member of a concept may be influenced by characteristic features that do not constitute part of the definition of the concept. For example, the diet of penguins may be a relevant characteristic feature in determining their typicality, but atomic composition cannot take this into account when determining typicality unless this feature forms part of the definition of a penguin.

Atomic composition may be viewed as a restricted version of feature composition, since any defining feature may be considered a relevant feature. Hence, we will only consider feature composition further. We consider the definition of feature vectors, their normalisation, and their composition.

4.1 Feature vectors

The features of a concept come in two guises: They are either *characteristic features*, co-determining typicality of objects in the concept, or they are *defining features* of the concept. In a DL extended with suitable preferential subsumption relations, characteristic features may be introduced on the right-hand side of preferential subsumption statements. For example, in the axioms given below, if $\precsim_1$ is derived from the $Penguin$-order $\leq_1$, then $\forall eats.Fish$ is a characteristic feature of $Penguin$. Defining features are introduced on the right hand-side of classical subsumption statements. For example, in the following axioms, $Seabird$ is a defining feature of $Penguin$, so are $Bird$ and $\exists eats.Fish$. Similarly, $Bird$ and $\exists eats.Fish$ are both defining features of $Seabird$: $Seabird \equiv Bird \sqcap \exists eats.Fish; Penguin \sqsubseteq Seabird; Penguin \precsim_1 \forall eats.Fish.$

The question arises whether relevant features should be determined algorithmically through some closure operator, or whether their identification is a modelling decision. While defining features can easily be derived from the knowledge base, this is not obvious in the case of characteristic features. We therefore view the choice of relevant features as a modelling decision, in accordance with a globalist view of concepts as context sensitive. The choice of features relevant for a particular concept, and their respective preference orders, are therefore determined by a subjective world view and have to be re-evaluated in each new context. The following development assumes a fixed ordered interpretation, even when some order is defined in terms of others.

Definition 7. *A feature vector is an n-tuple of concepts $\langle C_1^I, \ldots, C_n^I \rangle$ with corresponding preference vector $\langle \leq_1, \ldots, \leq_n \rangle$ such that $\leq_j$ is a C_j-order, for $1 \leq j \leq n$, and weight vector $\langle w_1, \ldots, w_n \rangle$ such that $w_j \in \mathbb{Z}$, for $1 \leq j \leq n$.*

We do not place any formal relevance restriction on the choice of elements of a feature vector, as this is a modelling decision. We may even, for example, have two feature vectors for $Fish$, one for use in the fish shop, and one for the pet shop. We may also define different preference orders for the same concept, for use in different contexts. For example, miniature, colourful fish may be typical in a pet shop, but not even relevant in a fish shop.

Next, we consider the normalisation of preference orders, which paves the way for their composition.

Definition 8. *Let $\langle C_1^I, \ldots, C_n^I \rangle$ be a feature vector with corresponding preference vector $\langle \leq_1, \ldots, \leq_n \rangle$. The level of an object $x \in \Delta^I$ relative to preference order $\leq_j$, written $level_j(x)$, is defined recursively as follows:*

$$level_j(x) := \begin{cases} 1 \text{ if } x \text{ is } \leq_j \text{-minimal in } C_j^I; \\ 0 \text{ if } x \text{ is } \leq_j \text{-maximal in } \Delta^I \setminus C_j^I; \\ max\{level_j(y) : y <_j x\} + 1 \text{ for non-minimal objects in } C_j^I; \\ min\{level_j(y) : x <_j y\}) - 1 \text{ for non-maximal objects in } \Delta^I \setminus C_j^I. \end{cases}$$

Definition 8 maps objects in the domain to integers. We note that the absence of infinite $\leq_j$-chains ensures that $level_j$ is defined on the whole of Δ^I. Given any feature C_j^I in the feature vector, Definition 8 assigns a positive level to all objects in C_j, and a non-positive level to all objects not in C_j^I. In the case where $\leq_j$ is a two-level order, $level_j(x) = 1$ for $x \in C_j^I$, and $level_j(x) = 0$ for $x \notin C_j^I$.

It is not difficult to see (given the modularity of the preference orders) that this mapping preserves the relative order of elements in the corresponding preference order:

Proposition 1. *For any $x, y \in \Delta^I$, $x \leq_j y$ iff $level_j(x) \leq level_j(y)$.*

We now have the required apparatus to compose the chosen preference orders of a feature vector. We define the typicality of objects relative to a given concept, based on its relevant features. The weight vector may be used in two ways – to normalise the preference orders so that they have the same range, or to adjust the relative importance of each feature. Normalisation can be done without intervention from the modeller, and resonates better with the qualitative approach to typicality followed so far in the paper.

The intuition of Definition 9 is that it ranks those objects that conform better to the features of C in terms of typicality on a higher level. The function f first maps each object in the domain to a non-negative integer. This induces a modular C-order, say $\leq_k$, on objects in the domain.

Definition 9. *Given concept C with feature vector $\langle C_1^I, \ldots, C_n^I \rangle$, preference vector $\langle \leq_1, \ldots, \leq_n \rangle$ and weight vector $\langle w_1, \ldots, w_n \rangle$, let $f : \Delta^I \rightarrow \mathbb{Z}_0^+$, such that$f(a) := Max\{1, \sum_{j=1}^{n}(level_j(a) \times w_j)\}$ if $a \in C^I$ and 0 otherwise, for any object $a \in \Delta^I$. The associated preference relation $\leq_k$ on Δ^I given by: $a \leq_k b$ iff $f(a) \leq f(b)$, for some $k \in \mathcal{J}$, is the typicality C-order induced by the features, preferences and weights.*

Our choice for f is not arbitrary, but there are alternatives, such as taking the maximum of the input preferences instead of their sum. By choosing different functions for different connectives, atomic composition can be simulated using feature vectors.

4.2 Example

We conclude this section with an illustrative example. Suppose we have the following terminological statements:

$$Penguin \sqsubseteq Bird \sqcap Flightless \sqcap Aquatic \tag{1}$$
$$Penguin \sqsubseteq\!\!\!\sim_1 \forall habitat.Southern \tag{2}$$
$$Southern \sqsubseteq \neg Equatorial \tag{3}$$
$$GalapagosPenguin \sqsubseteq Penguin \tag{4}$$
$$Penguin \sqsubseteq \forall \preceq_1 .Penguin \tag{5}$$
$$\exists habitat.Equatorial \sqsubseteq \forall \preceq_2 .\exists habitat.Equatorial \tag{6}$$

Line (2) of the TBox states that the habitat of typical penguins is restricted to the southern regions. Note that we cannot derive from (2) and (4) that the habitat of typical Galapagos penguins is restricted to the southern regions. Lines (5-6) ensure that $\preceq_1$ and $\preceq_2$ are indeed, respectively, a $Penguin$-order and an $\exists habitat.Equatorial$-order. In the ordered interpretation I satisfying this Tbox, and where $\preceq_1^I$ partitions objects into typical penguins, atypical penguins, and non-penguins, we have that:

$$level_1(a) := \begin{cases} 2 \text{ if } a \text{ is typical in } Penguin^I; \\ 1 \text{ if } a \text{ is atypical in } Penguin^I; \\ 0 \text{ otherwise.} \end{cases}$$

Suppose further that $\preceq_2^I$ is the modular default $\exists habitat.Equatorial$-order that partitions this concept into two classes. Then $level_2(a) := 1$ if $a \in \exists habitat.Equatorial^I$, 0 otherwise.

We now construct a feature vector for $GalapagosPenguin$. We choose $Penguin$ as relevant defining feature, and $\exists habitat.Equatorial$ as relevant characteristic feature. That is, a Galapagos penguin is a penguin whose distinctive characteristic is that it occurs in the equatorial region. The feature vector for $GalapagosPenguin$ is therefore $\langle Penguin^I, \exists habitat.Equatorial^I \rangle$. Its preference vector is $\langle \leq_1, \leq_2 \rangle$, and as weight vector we choose $\langle 1, 2 \rangle$ in order to normalise the ranges of $\leq_1$ and $\leq_2$. The resulting derived $GalapagosPenguin$-order is $\leq_3$, obtained from:

$$f_3(a) := \begin{cases} 4 \text{ if } a \text{ is typical in } Penguin^I \text{ and } a \in \exists habitat.Equatorial^I; \\ 3 \text{ if } a \text{ is atypical in } Penguin^I \text{ and } a \in \exists habitat.Equatorial^I; \\ 2 \text{ if } a \text{ is typical in } Penguin^I \text{ and } a \in \forall habitat.(\neg Equatorial)^I; \\ 1 \text{ if } a \text{ is atypical in } Penguin^I \text{ and } a \in \forall habitat.(\neg Equatorial)^I; \\ 0 \text{ otherwise.} \end{cases}$$

Note that the first case, i.e. where $f_3(a) = 4$, does not hold for any object a, as it contradicts terminological axiom (2) in the knowledge base. The following preferential subsumption statement holds in I: $GalapagosPenguin \sqsubsetsim_3 \exists habitat.Equatorial$.

So, typically, Galapagos penguins are found in the equatorial region, not exclusively in the southern regions. Of course, in this example we could simply have stated this, but the point is that defining and characteristic features may be used to derive compositionally the typicality of objects in a class based on chosen relevant features. Our example gives a simple illustration of this claim.

5 Related work

Notions of typicality have been studied in a wide variety of contexts, most of them beyond the scope of this paper. In the context of ontologies, Yeung and Leung [21] proposed a model of graded membership, but their representation is not directly in terms of DLs. Giordano et al. [7, 8] define a nonmonotonic extension of the description logic $\mathcal{ALC}$ to reason about typicality, while Grossi et al. [9] use contexts, modelled as sets of DL models, to describe a version of typicality. In order to be able to determine similarity between objects, Sheremet et al. [19] extend a DL with the constructors of the similarity logic $\mathcal{SL}$.

6 Conclusion

We presented a semantic framework for modelling object typicality in description logics. In [3] we showed how reasoning with a single typicality order on the domain of interpretation (and the induced defeasible subsumption relation) can be reduced to reasoning in a sufficiently expressive DL. This translation is also applicable when reasoning with typicality of individual concept members, as presented in this paper.

We also presented a proposal for deriving new typicality orders from existing ones using feature vectors. Our proposal is compositional, and rooted in a globalist cognitive stance on the semantics of typicality. The determination of compositional rules is therefore a modelling decision, unlike compound class membership, the meaning of which can be completely determined from the meanings of its atomic constituents. Implementation of feature vectors in a DL setting is a topic for further research.

References

1. F. Baader, I. Horrocks, and U. Sattler. Description logics. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 135–180. Elsevier, Amsterdam, 2008.
2. K. Britz, J. Heidema, and W. Labuschagne. Semantics for dual preferential entailment. *Journal of Philosophical Logic*, 38:433–446, 2009.
3. K. Britz, J. Heidema, and T. Meyer. Semantic preferential subsumption. In *Proceedings of KR2008*, pages 476–484. AAAI Press, 2008.
4. N. Chater, K. Lyon, and T. Meyers. Why are conjunctive categories overextended? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(3):497–508, 1990.
5. F.M. Donini, D. Nardi, and R. Rosati. Description logics of minimal knowledge and negation as failure. *ACM Transactions on Computational Logic*, 3(2):177–225, 2002.
6. M. Freund, D. Lehmann, and P. Morris. Rationality, transitivity and contraposition. *Artificial Intelligence*, 52(2):191–203, 1991.
7. L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. Preferential description logics. In N. Dershowitz and A. Voronkov, editors, *Proceedings of LPAR 2007*, volume 4790 of *LNAI*, pages 257–272. Springer-Verlag, 2007.
8. L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. Reasoning about typicality in preferential description logics. In S. Hölldobler, C. Lutz, and H. Wansing, editors, *Proceedings of JELIA 2008*, volume 5293 of *LNAI*, pages 192–205. Springer-Verlag, 2008.
9. D. Grossi, F. Dignum, and J-J.C. Meyer. Context in categorization. In L. Serafini and P. Bouquet, editors, *Proceedings of CRR'05*, volume 136 of *CEUR Workshop Proceedings*. CEUR-WS, 2005.
10. J.A. Hampton. Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(1):12–32, 1988.
11. J.A. Hampton. Concepts as prototypes. *The Psychology of Learning and Motivation*, 46:79–113, 2006.
12. H. Kamp and B. Partee. Prototype theory and compositionality. *Cognition*, 57:129–191, 1995.
13. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
14. D. Lehmann. Stereotypical reasoning: Logical properties. *Logic Journal of the IGPL*, 6(1):49–58, 1998.
15. D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.

16. K. Lyon and N. Chater. Localist and globalist approaches to concepts. In K.J. Gilhooly, M.T.G. Keane, R.H.Logie, and G. Erdos, editors, *Lines of Thinking*. John Wiley & Sons Ltd, 1990.
17. D.N. Osherson and E.E. Smith. Gradedness and conceptual combination. *Cognition*, 12:299–318, 1982.
18. D.N. Osherson and E.E. Smith. On typicality and vagueness. *Cognition*, 64:189–206, 1997.
19. M. Sheremet, D Tishkovsky, F. Wolter, and M. Zakhharyaschev. A logic for concepts and similarity. *Journal of Logic and Computation*, 17(3):415–452, 2007.
20. Y. Shoham. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. The MIT Press, Cambridge, MA, 1988.
21. C. Yeung and H. Leung. Ontology with likeliness and typicality of objects in concepts. In D.W. Embley, A. Olivé, and S. Ram, editors, *Proceedings of ER 2006*, volume 4215 of *LNCS*, pages 98–111. Springer-Verlag, Berlin Heidelberg, 2006.