Free Description Logic for Ontologists

Fabian NEUHAUS [a,b,1] and Oliver KUTZ [a] and Guendalina RIGHETTI [a]

[a] *Conceptual and Cognitive Modelling Research Group (CORE), KRDB Research Centre for Knowledge and Data, Faculty of Computer Science, Free University of Bozen-Bolzano, Italy*

[b] *Otto-von-Guericke University of Magdeburg, Magdeburg, Germany*

Abstract. Free Logic deviates from classical first-order logic in that singular terms need not to refer, or not refer to existing entities, i.e. those in the scope of the usual first-order quantifiers. We thus face the basic questions whether sentences referring to such non-existing objects can be true, or rather denied a truth value. Moreover, the ontologist needs to decide whether she wants to endorse the existence of non-existing, or 'fictional', objects in the meta-theory, or rather deny reference. We here explore the various possible answers to these questions in the paradigm of dual-domain semantics and analyse the choices in the context of basic Description Logic languages. We finally sketch a treatment of definite descriptions under the different choices.

Keywords. Free Logic, Existential Commitment, Description Logic, Fictional Objects, Definite Descriptions

1. Introduction

According to conceptualists, ontologies are about concepts, including possibly fictional concepts. Rather obviously, fictions often use names that fail to refer to anything existing: this is the case, for instance, of the name 'The Hulk' within a statement such as 'The Hulk is green'. From a modelling point of view, this may not cause problems, or inconsistencies, if we assume the point of view of the story. But the situation may become more difficult if we need to distinguish between what is true in the story and what is true in the real world, e.g. compare 'The Hulk is a person and is green' with 'The Hulk is stronger than Muhammad Ali'. If we take conceptualism seriously, ontologies and knowledge bases may then include singular terms that refer to something that exists (e.g., 'Pope Francis') singular terms that do not refer to anything in the domain of quantification (e.g., 'The Hulk' or 'Santa Claus') and singular terms of which we are uncertain whether they refer to something (e.g., 'Homer').

The representation of scientific theories and empirical knowledge also requires the consideration of non-existing or possibly non-existing entities. As [1] pointed out, this is particularly evident within scientific domains, where hypothesis and theories often involve entities that are unknown to exist, or not observable, or even idealisations that hardly exist, but that are still fundamental in the development and success of the theory.

As argued by the authors, one example is the *Higgs boson*, which was predicted by the Standard Model of particle physics, even if, up to 2012, no instances of the predicted entity had been found. A representation of the Standard Model in 2011 would have been incomplete without the Higgs boson, in spite of the fact that its existence was hypothetical at this time. In contrast, the non-existence of the *ideal gas* is not open for debate, since the ideal gas is a gas that consists of particles that lack spatial extension (point particles) and that do not attract or repel each other. While the ideal gas does not exist, it is subject to the ideal gas law $PV = nRT$, which describes the relationship between pressure, volume, and temperature of a given amount of an ideal gas. For many purposes real gases (like hydrogen and oxygen) may be treated as the ideal gas, and, thus, the ideal gas law is frequently used for calculations in Chemistry. For this reason, a representation of chemical knowledge would be incomplete without the representation of the ideal gas.

Fictions, scientific hypotheses, and idealisations are not the only origin of non-referring terms. A more mundane source are errors during the modelling process. For example, if one adds an IRI as name for an individual to an OWL ontology, and later realises that this was a mistake, there is no elegant way to deal with the situation. If one just deletes the IRI from the ontology, then references to the ontology that use the IRI may break. However, if one leaves the IRI in the ontology (possibly annotated as "obsolete"), then from a logical point of view a corresponding entity does exists in the universe of discourse and may play a role for automatic reasoning (see e.g. [2]).

These three examples illustrate that there is need for allowing non-referring singular terms in a knowledge representation language. Classical first-order logic presupposes that all individual constants refer to something in the universe of discourse. Description Logics (DLs) are one of the most important formalism for knowledge representation and provide the logical foundation of the most widely used ontology language, OWL. Most modern DLs are designed as fragments of classical first-order logic. From this FOL heritage DLs inherit the presupposition that names always refer to something in the universe of discourse, and that this universe is always non-empty.

In this paper, we develop the first steps towards a general framework to handle non-referring singular terms in the area of Description Logic. Free Logic is the branch of logic that studies logical systems free of *existential presuppositions*, in particular the presupposition that singular terms denote something in the domain of quantification. Although an important subject in its own right [3, 4], it is of particular importance for instance in the context of counterpart theory or more generally the combination of modality and quantification [5, 6] where existence of an object in a possible world is a central concept. For instance, objects may go in and out of existence along the flow of time.[2]

In particular, we discuss here three different philosophical intuitions about the truth-values of sentences that include non-referring singular terms. *Positive free logics* allow atomic sentences with non-referring singular terms to be true. This is a natural position if one intends to represent information about fictional entities or idealisations; e.g., in atomic sentence like *Santa has a beard* or *The ideal gas is a gas*. In *negative free logics*, on the other hand, the truth of an atomic sentence requires that all arguments refer to something in the domain of quantification; hence, atomic sentences about Santa and the ideal gas are always false. A negative free logic is the natural choice if one considers non-referring singular terms as errors. Finally, according to *gapped free logics* which we also

[2]Specifically, [5] used a free quantificational logic in order to provide a generic completeness proof for different theories of cross-world identity.

call *Fregean* (and which sometimes are called ‘neutral’ [4]), a semantic precondition for the evaluation of the truth-value of a sentence is that all singular terms refer to something. Thus, sentences that fail that condition lack a truth-value; there is a *truth-value gap*. This option is closest in spirit to the classical Fregean position, as also noted by [7, 8], according to which a sentence only has a truth-value if all singular terms that occur in the sentence have a referent [9].

We present three different semantics for $\mathcal{ALCO}$ without existential presuppositions, which reflect the three different philosophical choices (positive, negative, Fregean) and discuss their implications. The different choices are implemented based on a dual domain semantics, where we distinguish between an inner domain, which represents the standard domain of discourse, and an outer domain, where individual names are interpreted that do not denote anything that exists. The dual domain semantics may, arguably, be seen as a less natural fit for negative or Fregean free logic than some of its alternatives (see Section 2). However, it enables us to compare the implications of the different philosophical stances within one framework.

The rest of the paper is organised as follows: In the next section we provide an introduction to Free Logic as extensions of classical FOL. In Section 3 we introduce three alternative semantics for a free description logic representing the different philosophical stances on the truth-value of sentences with non-referring terms. Section 4 sketches how these semantics may be applied to an extension of the language with definite descriptions.

2. Free Logic: The basic landscape

We summarise some of the basic semantic distinctions that can be made. More detail can be found in [4, 3, 10].

Classical Logic requires that each singular term (i.e., a constant or functional expression) refers to an object in the domain of interpretation. Consequently, $\exists x(x = t)$ is a logical truth in classical logic. Following Quine’s famous maxim “To be is to be the value of a variable” [11], the existential quantifier is usually read as *there exists*. Hence, classical FOL seems to entail that t exists, regardless of the choice of “t”, including “Santa Claus” and “$\frac{1}{0}$”. Thus, the existential presupposition of classical FOL concerning singular terms leads to unintended ontological commitments. Another existential presupposition of classical FOL is the assumption that something exists and, thus, the universe of discourse is not empty.

Free Logic rejects the first and, typically, both of these existential presuppositions.[3] In particular, Free Logic allows to handle formulas which contain singular terms which do not refer to anything in the domain of quantification; in these cases we speak of *empty terms*, following [4], or of *non-referring terms*. In order to do this, existence is treated as a predicate of objects/individuals: an existence predicate $E!$ is introduced in the language in order to explicitly declare whether a term (or the object it denotes) *exists* in the domain or not. It is usually defined as follows: $E!t =_{def} \exists x(x = t)$. Quantification in free logic must then range over all and only those objects that satisfy the $E!$ predicate.

[3] Sometimes ‘free logic’ is just associated with the rejection of the existential presupposition of singular terms, while logical systems that allow an empty domain of quantification are called ‘inclusive’. In this paper we will consider only free logics that are also inclusive in this sense.

Free Logic is an extension of classical FOL in the sense that it preserves the semantics of classical FOL for sentences that contain no non-referring singular terms. One way to categorise the different approaches to the semantics of free logic is according to their philosophical position about the truth-value of sentences with non-referring singular terms. This leads to the distinction between free logic (NFL), positive free logic (PFL), and Fregean free logic (FFL).

Let us imagine we want to formalise statements about Pegasus. Since Pegasus does not exist, there is no Pegasus in the domain of quantification of our ontology. In other words, the singular term "Pegasus" is empty (i.e., does not refer to anything).

PFL claims that some atomic sentences involving non-referring terms are true. Candidates are sentences about fictions like "Pegasus is a horse" and "Pegasus is the offspring of Poseidon" or about idealisations "The ideal gas is a gas". In PFL these sentences may be true without requiring the existence of Pegasus or the ideal gas. The equation *pegasus* = *pegasus* is typically considered to be logically true in PFL.

In contrast, according to NFL, every atomic formula involving "Pegasus" (or any other non-referring term) is false, including the equation *pegasus* = *pegasus*. Therefore, for any atomic sentence $P^n(t_1,\dots,t_n)$ and any of its arguments t_i, the sentence $P^n(t_1,\dots,t_n) \rightarrow E!t_i$ is a logical truth in NFL. Since the negation in NFL is classical, it follows that $\neg P^n(t_1,\dots,t_n)$ is true, if some t_i is empty.

Lastly, FFL endorses the Fregean position that the reference of all singular terms is a precondition for assigning truth-values to atomic sentences. Hence in FFL atomic formulas involving non-referring terms are truth-valueless, or, as it is sometime said, have a truth-value gap. Thus, any atomic formula involving Pegasus lacks a truth-value.

The three philosophical stances regarding truth-values (or lack thereof) of sentences involving non-referring terms may be realised by different semantic approaches. The choice of a suitable semantics for a given philosophical position is, arguably, one of the most interesting and important questions in this context [12].

A *dual domain semantics* distinguishes between two domains, the inner domain and the outer domain. The inner domain is the domain of existing entities and, thus, the domain of quantification. The outer domain is used to interpret empty terms like "Santa Claus". In the following, we will distinguish between the *reference* of a singular term and its *denotation*. In a dual domain semantics all singular terms denote some entity, but only some of them refer to something; namely the entities that denote something in the inner domain. This distinction is particularly useful for positive free logics, since it enables, for example, to explain why "Santa Claus is Father Christmas" is true, while "Santa Claus is the Easter Bunny" is false. While all these terms are without reference, "Santa Claus" and "Father Christmas" denote the same fictional entity, while "Easter Bunny" denotes a different one.

Single domain semantics avoid the addition of additional entities. Instead, their interpretation function is partial and, thus, the denotation of non-referring terms is undefined. For this reason, the interpretation function on its own does not determine the truth-value of all atomic sentences. One possible way to ensure bivalence for all sentences is to adopt a *convention* that completes the interpretation function. In NFL the choice is to assign all atomic sentences with a non-referring term the truth-value **False**. However, it is important to note that any function that assigns truth-values to atoms with non-referring terms serves the same purpose: restoring a total assignment of truth values. An alternative approach is to relax the principle of bivalence and leave the truth-value of sentences with non-referring terms undefined. Given this strategy, the partial interpretation function results in a partial valuation function of sentences, because there are truth-value gaps. Hence, this choice aligns naturally with FFL.

A lack of truth-value for some atomic formulas yields the question how to handle complex formulae and compositionality. For instance, given two formulas ϕ and ψ, where ϕ is truth-valueless while ψ is true, what should the truth value of $\phi \vee \psi$ be? [8] proposes an approach where truth-value gaps are 'infectious' in the sense that a complex sentence lacks a truth-value as soon as it contains some non-referring singular term. A different kind of approach uses supervaluations to 'close' some of the truth-value gaps [13]. For this purpose, one extends a given interpretation I by assigning arbitrary truth-values to the atomic formulas with non-referring terms. A formula is supertrue (superfalse) in I iff it is true (false) under all such extensions of I. E.g., if t does not exists in the interpretation I, then $P(t) \vee \neg P(t)$ would lack a truth-value in I according to the Fregean semantics, but it would evaluate to true under the supervaluational semantics.

2.1. Definite descriptions

One of the main applications of Free Logic is the treatment of definite descriptions. Definite descriptions are expressions which consist of a combination of the definite article 'the' and a noun phrase, such as 'The main application of Free Logic' or 'The man wearing a blue t-shirt'.

In analytical philosophy, the treatment of definite descriptions is a long-established subject. Consider the following statement: 'The present king of France is bald'. According to Russell [14], the logical form of this sentence consists of two assertions: (i) there is one and only one entity who is a king of France and (ii) this entity is bald. Following this intuition, according to Russell definite descriptions should not be represented in FOL as kind of singular terms, but as part of a formula that represents the whole sentence: $\exists x(\text{FKing}(x) \wedge \forall(y)(\text{FKing}(y) \rightarrow x = y) \wedge \text{Bald}(x))$. It follows that in cases where a definitive description fails to apply to a unique individual (e.g., because there is no present King of France), the sentence containing that definite description is false. Note that according to Russell's analysis, definite descriptions are not treated as singular terms, but instead they are evaluated contextually as part of a complex first-order formula.

Given the goal of representing definite description in DLs, Russell's approach is then not feasible, because DLs lack the quantificational apparatus.

In contrast to Russell, Frege treats definite descriptions as a type of singular term. A sentence which contains a definite description (or, more generally, a singular term) with no reference (as in the case of 'The present king of France') would in turn have no reference, and would then lack a truth value, being neither true nor false [9].

Since Free Logic studies the logical behaviour of non-referring singular terms, the study of definite descriptions in the context of free logic suggests itself. In free logic systems definite descriptions are represented as singular terms. Given some formula ϕ containing the free variable x, the singular term $\iota x\phi$ represents the definite description *The x such that* ϕ. E.g., our running example would be represented as $\text{Bald}(\iota x\, \text{FKing}(x))$. Typically, free logics with definite descriptions obey Lambert's law [4]:

$$\forall y(y = \iota x\phi \leftrightarrow \forall x(\phi \leftrightarrow x = y)), \quad x \text{ free in } \phi$$

Depending on their philosophical stance on non-referring singular terms, free logics differ with respect to their treatment of non-referring definite descriptions. For example, in NPL atomic formulas involving non-referring definite description are false, leading to a position similar to the one of Russell.

3. Free Description Logics

In the following we will present free-logical variants of description logic, based semantically on the dual-domain approach.

3.1. The Description Logic $\mathcal{ALCO}$

We use description logics (DLs) as well-known examples of languages for describing concepts. We briefly introduce the well-known DL $\mathcal{ALCO}$, i.e. $\mathcal{ALC}$ enriched with the nominal construct; for a full introduction to DL, see [15]. The syntax of $\mathcal{ALCO}$ concepts is based on three pairwise disjoint sets, the set of *concept names* N_C, the set of *role names* N_R, and the set N_I for individual names. The set of $\mathcal{ALCO}$ *singular terms* is identical with N_I.[4] The set of $\mathcal{ALCO}$ *concepts* is generated by the grammar

$$C ::= \top \mid \bot \mid A \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \forall R.C \mid \exists R.C \mid \{t\}$$

where $A \in N_C$, $R \in N_R$, t is a singular term.

A *TBox* is a finite set of concept inclusions (GCIs) of the form $C \sqsubseteq D$ where C and D are concepts. It is used to store terminological knowledge regarding the relationships between concepts. An *ABox* is a finite set of assertions, i.e., formulas of the following form:[5]

$$t_1 = t_2 \qquad t_1 \neq t_2 \qquad C(t_1) \qquad {\sim}C(t_1) \qquad t_1 R t_2 \quad {\sim}(t_1 R t_2)$$

where t_1, t_2 are $\mathcal{ALCO}$ singular terms, $R \in N_R$, and C is a $\mathcal{ALCO}$ concept. A *knowledge base* $\mathcal{K}$ is a set of inclusions and assertions. Note that the negated assertions use a different negation symbol in order to distinguish the negation of sentences from the negation of concepts. E.g., the concept $\neg A$ is the negation of the concept A, which may be used in the positive assertion $\neg A(b)$ (in the sense that we affirm the predication of $\neg A$) or in the negative assertion ${\sim}\neg A(b)$ (the statement $\neg A(b)$ is false). In classical $\mathcal{ALCO}$, ${\sim}\neg A(b)$ is equivalent to $A(b)$, but as we will see below, this is not the case in free logic. These two types of negation are well known in philosophical logic [16].

The semantics of $\mathcal{ALCO}$ is defined through *interpretations* $I = (\Delta^I, \cdot^I)$, where Δ^I is a non-empty *domain*, and $\cdot^I$ is a function mapping every individual name to an element of Δ^I, each concept name to a subset of the domain, and each role name to a binary relation on the domain. The semantics of complex concepts is as follows: $\top^I = \Delta^I$, $\bot^I = \emptyset$, and

[4] We will extend the grammar of $\mathcal{ALCO}$ by adding definite descriptions as singular terms below.

[5] Note that these assertions are syntactic sugar in the context of $\mathcal{ALCO}$, but will be important in the free version of the logic. in the case of $\mathcal{ALC}$ on the other hand, these assertions are in fact not expressible within the language. To see the reduction in $\mathcal{ALCO}$, consider for instance the subsumption $\{a\} \sqsubseteq \neg\exists R.\{b\}$, which holds in a model iff $\neg(aRb)$, etc.

$\{t\}^I = \{t^I\}$, for any singular term t. the Boolean operations on concepts are given by the usual set-theoretic intersection (conjunction $\sqcap$), union (disjunction $\sqcup$), and complement (negation $\neg$) of the extension of concepts. The extensions of universal and existential restrictions are defined as

$$(\forall R.C)^I = \{d \in \Delta^I \mid \text{for all } e \in \Delta^I : (d,e) \in R^I \text{ implies } e \in C^I\}, \text{ and}$$
$$(\exists R.C)^I = \{d \in \Delta^I \mid \text{exists } e \in \Delta^I : (d,e) \in R^I \text{ and } e \in C^I\}$$

The satisfaction relationship $\models$ between interpretation and formulas is defined as follows:

$$I \models C \sqsubseteq D \text{ iff } C^I \subseteq D^I \qquad I \models C(t_1) \text{ iff } t_1^I \in C^I$$
$$I \models t_1 R t_2 \text{ iff } (t_1^I, t_2^I) \in R^I \qquad I \models t_1 = t_2 \text{ iff } t_1^I = t_2^I$$
$$I \models {\sim}\phi \text{ iff } I \not\models \phi$$

The interpretation $\mathcal{I}$ is a *model* of the TBox $\mathcal{T}$ (of the knowledge base $\mathcal{K}$) if it satisfies all the elements of $\mathcal{T}$ ($\mathcal{K}$, respectively). Given two concepts C and D, we say that C is *subsumed* by D w.r.t. the TBox $\mathcal{T}$ ($C \sqsubseteq_{\mathcal{T}} D$) if $C^I \subseteq D^I$ for every model I of $\mathcal{T}$. We write $C \equiv_{\mathcal{T}} D$ when $C \sqsubseteq_{\mathcal{T}} D$ and $D \sqsubseteq_{\mathcal{T}} C$. We say that $\mathcal{T}$ (resp. $\mathcal{K}$) is consistent if there exists an interpretation that is a model of $\mathcal{T}$ (resp. $\mathcal{K}$). A concept C is satisfiable wrt. $\mathcal{T}$ ($\mathcal{K}$) if there exists an interpretation I that is a model of $\mathcal{T}$ ($\mathcal{K}$) for which C^I is not empty. A knowledge base $\mathcal{K}$ *entails* an inclusion or an assertion ϕ, $\mathcal{K} \models \phi$, if ϕ is satisfied in every model of $\mathcal{K}$.

3.2. Dual Domain Semantics

An advantage of the dual-domain semantics that we will introduce below is that it allows to handle not only the case of non-existent objects having non-trivial properties (i.e. positive free logic), but equally well the frameworks of negative, inclusive, or neutral, i.e. Fregean semantics. It can therefore serve as a semantic foundation for all these paradigms. Informally, we distinguish between an inner domain and an outer domain. The inner domain represents the universe of discourse, i.e., the set of entities that quantifiers range over. The outer domain, extending the inner one, is utilised to analyse the semantics of terms that do not refer to anything that exists, see Fig. 1.

In the following we define the basic dual-domain semantics for the language of $\mathcal{ALCO}^{\mathsf{f}}$. The basic intuition is that concepts, in general, have extensions within the global, outer domain of reference, which includes fictional objects. However, explicit quantification in the language is restricted to the inner domain of existing objects.

The free logics we will consider only disagree in how they evaluate the truth of statements containing non-referring terms. The notions of dual-domain interpretation and of concept extensions are independent of this.

Definition 1 (Dual-domain interpretations) *A dual-domain interpretation, or* DDI*, is a triple* $I = (\Delta^I, \Omega^I, \cdot^I)$*, where* $\Delta^I \subseteq \Omega^I$ *is a (possibly empty)* inner domain*,* Ω^I *a non-empty* outer domain*, and* $\cdot^I$ *is a function mapping every individual name* a *to an element of* Ω^I*, each concept name to a subset of* Ω^I*, and each role name to a binary relation on* $\Omega^I \times \Omega^I$.

We clarify the terminology regarding terms:

Definition 2 (Denotation–Reference-Emptiness) *Let I be a DDI and let t be an individual name. We call an element $d \in \Omega$ the denotation of t if $I(t) = d$. If $d \in \Delta$, we call the individual name t referring, and otherwise (if $d \in \Omega\backslash\Delta$) non-referring (or also empty).*

The extension of complex concepts is defined independently of the assignment of truth-values to statements, see Definition 3. This definition is inspired by the standard dual-domain semantics of free first-order logic: concepts denote subsets of the outer domain, $\top$ plays the role of the existence predicate $E!$, and the quantifiers are restricted to quantify over the inner domain. However, the Boolean operations are defined as usual in DL.

Definition 3 (Extension of concepts) *Given a DDI I, concepts A and B, and a name a the extension of complex concepts is defined inductively as follows:*

$$(\bot)^I = \varnothing \qquad (\top)^I = \Delta^I \qquad (\{a\})^I = \{a^I\}$$
$$(\neg A)^I = \Omega\backslash A^I \qquad (A \sqcap B)^I = A^I \cap B^I \qquad (A \sqcup B)^I = A^I \cup B^I$$
$$(\forall R.C)^I = \{d \in \Omega^I \mid \textit{for all } e \in \Delta^I : (d,e) \in R^I \textit{ implies } e \in C^I\}$$
$$(\exists R.C)^I = \{d \in \Omega^I \mid \textit{exists } e \in \Delta^I : (d,e) \in R^I \textit{ and } e \in C^I\}$$

We next define three different semantics to interpret ABox and TBox statements in the language of $\mathcal{ALCO}$. To this end, we define three *(partial) valuation functions*, V_I^+, V_I^-, V_I^g that assign truth values **True** or **False** to $\mathcal{ALCO}$ statements, but which are sometimes *undefined*. Note that we are therefore still operating within a bivalent semantics which however allows for truth-value gaps. Moreover, we write V_I^* for any of V_I^+, V_I^-, V_I^g, in case a definition is uniform across all three semantics, as in the case of defining the subsumption.

Definition 4 (TBox Statements / GCI)

$$V_I^*(A \sqsubseteq B) = \textbf{True} \iff \textit{for all } e \in \Delta^I : \textit{ if } e \in A^I \textit{ then } e \in B^I;$$
$$\textit{otherwise } V_I^*(A \sqsubseteq B) = \textbf{False}$$

Since GCIs represent quantified statements of the form *All X are Y*, the semantics of GCIs is defined as quantification over the inner domain, i.e., the existing entities. Note that V_I^* is a total function on the set of TBox statements. That is, under all three semantics there are no truth value gaps on TBox axioms.
The positive semantics for ABox assertions is just like the classical one. Moreover, it should be clear that the semantics collapses to the standard semantics under the assumption $(\neg\top)^I = \varnothing = \bot^I$, which, however, is not expressible with the above definition of GCI since it only quantifies over the existing objects.

Definition 5 (Positive semantics of $\mathcal{ALCO}^{\mathbf{f}^+}$)

$$V_I^+(C(t)) = \textbf{True} \textit{ iff } t_1^I \in C^I \qquad \textit{otherwise } V_I^+(C(t)) = \textbf{False}$$

$$V_I^+(R(t_1,t_2)) = \textbf{True} \textit{ iff } (t_1^I,t_2^I) \in R^I \qquad \textit{otherwise } V_I^+(R(t_1,t_2)) = \textbf{False}$$

$$V_I^+(t_1 = t_2) = \textbf{True} \textit{ iff } t_1^I = t_2^I \qquad \textit{otherwise } V_I^+(t_1 = t_2) = \textbf{False}$$

$$V_I^*(\sim\phi) = \textbf{True} \textit{ iff } V_I^*(\phi) = \textbf{False} \qquad V_I^*(\sim\phi) = \textbf{False} \textit{ iff } V_I^*(\phi) = \textbf{True}$$

In $\mathcal{ALCO}^{\mathsf{f}+}$ the valuation function is therefore a total and bivalent function to $\{\textbf{True},\textbf{False}\}$ on the above ABox statements. The semantics for the negative assertions is classical and is the same for $\mathcal{ALCO}^{\mathsf{f}+}$, $\mathcal{ALCO}^{\mathsf{f}-}$, and $\mathcal{ALCO}^{\mathsf{f}g}$. Next, in the negative semantics, atomic ABox assertions can only be true of existing things, and are false otherwise:

Definition 6 (Negative Semantics of $\mathcal{ALCO}^{\mathsf{f}-}$)

$$V_I^-(C(t)) = \textbf{True} \textit{ iff } t^I \in C^I \textit{ and } t^I \in \Delta \qquad \textit{otherwise } V_I^-(C(t)) = \textbf{False}$$

$$V_I^-(R(t_1,t_2)) = \textbf{True} \textit{ iff } (t_1^I,t_2^I) \in R^I \textit{ and } t_1^I,t_2^I \in \Delta \qquad \textit{otherwise } V_I^-(R(t_1,t_2)) = \textbf{False}$$

$$V_I^-(t_1 = t_2) = \textbf{True} \textit{ iff } t_1^I = t_2^I \textit{ and } t_1^I,t_2^I \in \Delta \qquad \textit{otherwise } V_I^-(t_1 = t_2) = \textbf{False}$$

In $\mathcal{ALCO}^{\mathsf{f}-}$, the valuation function V_I^- is also total. The gapped semantics defined next corresponds to the Fregean semantics discussed earlier and does include truth value gaps, i.e. V_I^g is partial.

Definition 7 (Gapped Semantics of $\mathcal{ALCO}^{\mathsf{f}g}$)

$$V_I^g(C(t)) = \textbf{True} \textit{ iff } t^I \in C^I \textit{ and } t^I \in \Delta$$

$$V_I^g(C(t)) = \textbf{False} \textit{ iff } t^I \notin C^I \textit{ and } t^I \in \Delta \qquad \textit{otherwise } V_I^g(C(t)) \text{ is undefined}$$

$$V_I^g(R(t_1,t_2)) = \textbf{True} \textit{ iff } (t_1^I,t_2^I) \in R^I \textit{ and } t_1^I,t_2^I \in \Delta$$

$$V_I^g(R(t_1,t_2)) = \textbf{False} \textit{ iff } (t_1^I,t_2^I) \notin R^I \textit{ and } t_1^I,t_2^I \in \Delta \qquad \textit{otherwise } V_I^g(R(t_1,t_2)) \text{ is undefined}$$

$$V_I^n(t_1 = t_2) = \textbf{True} \textit{ iff } t_1^I = t_2^I \textit{ and } t_1^I,t_2^I \in \Delta$$

$$V_I^g(t_1 = t_2) = \textbf{False} \textit{ iff } t_1^I \neq t_2^I \textit{ and } t_1^I,t_2^I \in \Delta \qquad \textit{otherwise } V_I^g(t_1 = t_2) \text{ is undefined}$$

For all three semantics the notions of model, satisfiability, entailment etc. are defined in the standard way (see Section 3.1 on classical $\mathcal{ALCO}$).

3.3. Some basic properties of free $\mathcal{ALCO}$ without definite descriptions

In free DL based on dual domains, $\top$ plays the role of the existence predicate. Hence, $\top(b)$ may be read as 'b exists'. What is a trivial truth in the classical version of $\mathcal{ALCO}$ ($\top(b)$ is always true and its negation always false) turns into a statement that illustrates the basic semantic distinctions between positive, negative, and Fregean paradigms, as shown in Table 1.

Table 1. The behaviour of some basic $\mathcal{ALCO}$ assertions under the different semantics. '—' represents that the situation cannot happen. ↓ is used for an undefined truth-value.

Condition on I		$\mathcal{ALCO}$	$\mathcal{ALCO}^{\mathsf{f}+}$	$\mathcal{ALCO}^{\mathsf{f}-}$	$\mathcal{ALCO}^{\mathsf{fg}}$
$t^I \in \Delta \backslash t^I \notin \Delta$	$I \models \top(t)$	**True** \—	**True** **False**	**True** **False**	**True** \↓
$t^I \in \Delta \backslash t^I \notin \Delta$	$I \models \neg\top(t)$	**False** \—	**False** **True**	**False** **False**	**False** \↓
$t^I \in \Delta \backslash t^I \notin \Delta$	$I \models \sim\top(t)$	**False** \—	**False** **True**	**False** \ **True**	**False** \ ↓
for all I	$(\neg\top)^I = (\bot)^I$	yes	no	no	no
	$\models \neg\top \sqsubseteq \bot$	yes	yes	yes	yes
	$\models \exists r.C \equiv \neg\forall r.\neg C$	yes	yes	yes	yes
	$C(a), C \sqsubseteq D \models D(a)$	yes	no	yes	yes
	$C(a) \models \top(a)$	yes	no	yes	yes
	$\neg C(a) \models \top(a)$	yes	no	yes	yes
	$\sim C(a) \models \top(a)$	yes	no	no	yes
	$\neg\top(a) \models \exists R.\{a\} \sqsubseteq \bot$	yes	yes	yes	yes
	$R(a,b) \models (\exists R.\{b\})(a)$	yes	no	yes	yes
	$\models a = a$	yes	yes	no	no
	$\top(a) \models a = a$	yes	yes	yes	yes

Since the existing entities (the inner domain) are a subset of the entities that are considered in an interpretation (the outer domain), note that if $\Delta^I \subset \Omega^I$ (a proper subset), then $\bot^I \neq (\neg\top)^I$. Indeed, one can consider $(\neg\top)^I$ as the set of fictitious (= non-existing) entities in I. Thus, in positive free $\mathcal{ALCO}^{\mathsf{f}+}$ we may represent 'Santa has a Beard and does not exists' as $\neg\top \sqcap \text{Beard}(\text{santa})$. In spite of the fact that $\bot^I \neq (\neg\top)^I$ (for some interpretations I), $\neg\top \sqsubseteq \bot$ is still logically true in $\mathcal{ALCO}^{\mathsf{f}+}$, because the semantics of $\sqsubseteq$ quantifies only over existing entities; i.e., the entities in the inner domain Δ. (We will discuss this point in more detail below.)

Since the Boolean operators $\neg, \sqcap, \sqcup$ are all defined with respect to the outer domain Ω^I, they behave classically (e.g., $(\neg\neg C)^I = C^I, (\neg(C \sqcap D))^I = (\neg C \sqcup \neg D)^I$, $(C \sqcup \neg C)^I = \Omega^I$).

A characteristic feature of classical logic as well as the classical version of description logics such as $\mathcal{ALCO}$ is the duality and interdefinability of the existential and universal quantifiers. It is therefore instructive to notice that the dualities remain valid also in our four basic scenarios. In order to better understand the behaviour of a dual domain semantics, we will discuss below a few examples and use them to illustrate the different semantics.

$\{A(a), B(a), A \sqsubseteq \neg B\}$ is satisfiable in $\mathcal{ALCO}^{\mathsf{f}+}$. Because subsumption is defined only over the elements of the inner domain, but it specifies nothing about the element of $\Omega^I \backslash \Delta^I$. Conversely, the negation is defined over the outer domain, and behaves classically. So, for instance, if all the elements of the set $(\neg B)^I$ lie in Δ, and the elements of the set $(A)^I$ lie both in Δ^I and in $\Omega^I \backslash \Delta^I$, $A \sqsubseteq \neg B$ could be true while A shares some elements with B in $\Omega^I \backslash \Delta^I$. To better visualise the situations, see the Fig. 1.

Analogously, we have that $C(a), C \sqsubseteq D \not\models D(a)$ in $\mathcal{ALCO}^{\mathsf{f}+}$. The subsumption requires that the element of Δ^I that are in $(C)^I$ must be elements of the set $(D)^I$. But it could be the case that $a^I \in \Omega^I \backslash \Delta^I$. In this case $\neg D(a)$ would be true in $\mathcal{ALCO}^{\mathsf{f}+}$. On the other hand, if we assume the existence of a we have a different result: $C(a), \top(a), C \sqsubseteq D \models D(a)$ in $\mathcal{ALCO}^{\mathsf{f}+}$. Here, since we assume $\top(a)$, a^I must belong to Δ, and, according

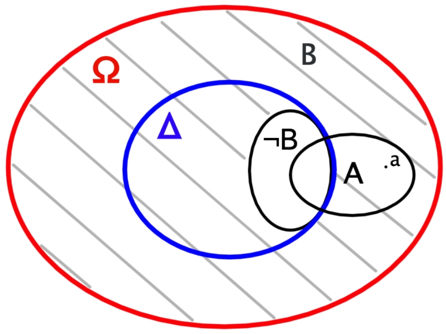

Figure 1. $\{A(a), B(a), A \sqsubseteq \neg B\}$ is satisfiable in $\mathcal{ALCO}^{\mathsf{f}^+}$.

to the definition of the subsumption, we get the result $D(a)$. Please notice that here we are using the $\top$ exactly as the predicate $E!$ is normally used in Free Logic.

In contrast to $\mathcal{ALCO}^{\mathsf{f}^+}$, in $\mathcal{ALCO}^{\mathsf{f}^-}$ and $\mathcal{ALCO}^{\mathsf{fg}}$ the truth of $C(a)$ requires a to be in Δ. Thus in these logics, $C(a) \models \top(a)$. Consequently, $C(a), C \sqsubseteq D \models D(a)$ in $\mathcal{ALCO}^{\mathsf{f}^-}$ and $\mathcal{ALCO}^{\mathsf{fg}}$.

In $\mathcal{ALCO}^{\mathsf{f}^-}$ the two negations differ in their behaviour. In $\neg C(a)$ the complex concept $\neg C$ is asserted of the individual a. Hence, in $\mathcal{ALCO}^{\mathsf{f}^-}$ $\neg C(a)$ may only be true, if a exists. In contrast, $\sim C(a)$ is the negation of $C(a)$. Thus, if $C(a)$ is false (possibly because a does not exist), $\sim C(a)$ is true.[6] Hence, in $\mathcal{ALCO}^{\mathsf{f}^-}$ $\neg C(a)$ entails $\top(a)$, but $\sim C(a)$ does not. Note that $\neg C(a) \models \sim C(a)$. E.g., *Alan is unhappy* entails *It is not the case that Alan is happy*.

Since the quantifiers are quantifying over the inner domain Δ, it follows that $(\exists R.\{a\})^I = \emptyset$, if $a^I \notin \Delta$. Thus, $\neg\top(a) \models \exists R.\{a\} \equiv \bot$. For the same reason $R(a,b) \not\models (\exists R.\{b\})(a)$ in $\mathcal{ALCO}^{\mathsf{f}^+}$.

Since a may not exist, $a = a$ is not a logical truth in $\mathcal{ALCO}^{\mathsf{f}^-}$ and $\mathcal{ALCO}^{\mathsf{fg}}$, unless $\top(a)$ is assumed.

4. Notes on Definite Descriptions in DLs

This section sketches some of the difficulties encountered when adding the idea of definite description to description logics. $\mathcal{ALCO}^{\iota}$ extends $\mathcal{ALCO}$ syntactically by adding the basic machinery for definite descriptions as follows. For any given (complex) concept C of $\mathcal{ALCO}$, we introduce ιC as a $\mathcal{ALCO}^{\iota}$ *definite description*. Intuitively, we want ιC to pick out, in accordance with Lambert's Law, the unique object that is a C, if it exists.

The set of $\mathcal{ALCO}^{\iota}$ *singular terms* is then the union of N_I with the set of $\mathcal{ALCO}^{\iota}$ *definite descriptions*.

If ϕ is a $\mathcal{ALCO}^{\iota}$ statement, let $\mathfrak{D}(\phi)$ be the set of definite descriptions ιC that occur in ϕ. An interpretation I meets the *definite descriptions presupposition* (ddp) of a

[6] A similar observation was made by Lambert in [17].

Table 2. Some examples for definite descriptions in free $\mathcal{ALCO}$. Here, C, D are (possibly distinct) arbitrary concepts, and the ddp is assumed for all interpretations.

	$\mathcal{ALCO}^{\iota,\mathsf{f}^+}$	$\mathcal{ALCO}^{\iota,\mathsf{f}^-}$	$\mathcal{ALCO}^{\iota,\mathsf{f}^g}$
$D(\iota C), C(a) \models a = \iota C$	yes	yes	yes
$D(\iota C) \models \top(\iota C)$	no	yes	yes
$\sim D(\iota C) \models \top(\iota C)$	no	no	yes
$\top(\iota C) \models C(\iota C)$	yes	yes	yes
$\neg\bot(\iota C) \models C(\iota C)$	yes	no	no

statement ϕ iff for all $\iota C \in \mathfrak{D}(\phi)$ there exists some $e \in \Omega$ such that $C^I = \{e\}$. If I meets the definite descriptions presupposition of ϕ, then we also say I is *suitable* for ϕ.

A free semantics for $\mathcal{ALCO}^\iota$ should respect the basic law regulating definite descriptions, i.e. Lambert's Law discussed above. This leads to the following twofold extension of the various semantics for free $\mathcal{ALCO}$:

- $V_I^*(\phi)$ is only defined if I is suitable for ϕ;
- $(\iota C)^I = e$ iff $C^I = \{e\}$, if I is suitable for ϕ

Because of the suitability requirements, the valuation functions V_I^+, V_I^-, V_I^g on $\mathcal{ALCO}^\iota$ are partial functions. Further, because of the uniqueness of ιC, $\{C \sqsubseteq \bot, C(\iota C)\}$ is unsatisfiable and $C(\iota C), C(a) \models a = \iota C$. Since definite descriptions are singular terms the differences between $\mathcal{ALCO}^{\mathsf{f}^+}$, $\mathcal{ALCO}^{\mathsf{f}^-}$ and $\mathcal{ALCO}^{\mathsf{f}^-}$ with respect to existential presuppositions apply to them (see lines 2 and 3 of Table 2). $\top(\iota C) \models C(\iota C)$ is a logical truth in all systems and $\neg\bot(\iota C) \models C(\iota C)$ in $\mathcal{ALCO}^{\mathsf{f}^+}$.

[18] extend $\mathcal{ALCO}$ with definite descriptions similarly, but use the approach of partial interpretations on names. The system is called $\mathcal{ALCO}^\iota$. In this scenario, a nominal $\{t\}$ is stipulated to denote the empty concept $\varnothing$ whenever the term t does not denote (is not defined), both if t is a name and a definite description. This follows the negative free logic paradigm and has the advantage to avoid truth-value gaps. However, we obtain simultaneously $\{t\} \sqsubseteq C$ and $\{t\} \sqsubseteq \neg C$, for any C and undefined term t.

$\mathcal{ALCO}^\iota$ is reduced to $\mathcal{ALCO}_U$, i.e. standard $\mathcal{ALCO}$ over total interpretations with the universal modality, and therefore has an EXPTIME-complete satisfiability problem [19].

5. Discussion and Outlook

We could here only sketch the basic semantical distinctions that arise in the context of applying free logic paradigms and definite descriptions to description logics. An intriguing landscape emerges when analysing the interactions between domain of quantification, negation, and definite descriptions and nominals.

Of course, standard free-logical semantics could be simply applied to the standard translation of DL into FOL, obtaining some free-logical DL. However, the more interesting free-logical phenomena emerge when taking the status of DL as a *concept language* seriously. This becomes perhaps most obvious when noting the non-equivalence in negative free logic of sentence negation vs. concept membership negation. Similarly, DLs require a new approach to definite descriptions since Russell's approach to unfold them into more complex FOL expressions is simply not available.

One of the somewhat hidden use cases of free logic is the validation of ontologies. Many ontologies do not contain explicit existential commitments. E.g., they may contain a hierarchy of classes of animals, but nothing in the ontology requires the existence of these animals. Thus, an ontology may be consistent in spite of the fact that in all interpretations that satisfy the ontology only a small minority of concept names is non-empty. However, the interpretations that are *intended* by ontology developers satisfy the vast majority (or even all) concept names of an ontology. In the consistency proof for Dolce we captured this intention by distinguishing between those concept names that may be empty and those that have to be non-empty. To use this information we represented Dolce in a multi-sorted logic. Essentially we introduced a free and inclusive semantics for the 'possibly empty' categories whilst applying classical semantics to the remaining ones [20].

Studying the expressive power and reducibility between systems, as also begun by [18], is important future work. In addition to the sentence and predicate negation as sketched above, another important language extension concerns the distinction between general concept inclusions that quantify on the meta-level only over the existing entities, as presented in this paper, and those that quantify over the entire domain of existing and non-existing objects. The latter allows inclusions that are valid universally, both for real and fictional/non-existing entities, preventing at the same time fictional counterexamples (for instance, all triangles will have three edges both in the inner and in the outer domain). In contrast, the former subsumption applies only to the inner domain, allowing for fictional counterexamples, e.g. Santa Claus is a man, all men are mortal, but Santa Claus is not mortal.

References

[1] Dumontier M, Hoehndorf R. Realism for scientific ontologies. In: Formal Ontology in Information Systems Proceedings of the Sixth International Conference (FOIS 2010); 2010. p. 387–399.

[2] Borgida A, Toman D, Weddell G. Concerning Referring Expressions in Query Answers. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17; 2017. p. 4791–4795.

[3] Lambert K. Free Logics: Their Foundations, Character, and Some Applications Thereof. Academia Verlag, Sankt Augustin; 1997.

[4] Nolt J. Free Logic. In: Zalta EN, editor. The Stanford Encyclopedia of Philosophy. fall 2018 ed. Metaphysics Research Lab, Stanford University; 2018. .

[5] Kracht M, Kutz O. The Semantics of Modal Predicate Logic I. Counterpart-Frames. In: Wolter F, Wansing H, de Rijke M, Zakharyaschev M, editors. Advances in Modal Logic. Volume 3. New Jersey, London, Singapore, Hong Kong: World Scientific Publishing Company; 2002. p. 299–320.

[6] Kracht M, Kutz O. Logically Possible Worlds and Counterpart Semantics for Modal Logic. In: Jacquette D, editor. Philosophy of Logic, Handbook of the Philosophy of Science. vol. 5. Amsterdam: Elsevier; 2007. p. 943–996.

[7] Skyrms B. Supervaluations: Identity, Existence, and Individual Concepts. The Journal of Philosophy. 1968;65(16):477–482.

[8] Lehmann S. Strict Fregean free logic. Journal of Philosophical Logic. 1994;23(3):307–336.

[9] Frege G. Über Sinn und Bedeutung. Zeitschrift für Philosophie und philosophische Kritik. 1892;100:25–50.

[10] Morscher E, Simons P. Free logic: a fifty-year past and an open future. In: New Essays in Free Logic. Springer; 2001. p. 1–34.

[11] Quine WV. On what there is. The review of metaphysics. 1948:21–38.

[12] Bencivenga E. Free Logics. In: Gabbay DM, Guenthner F, editors. Handbook of Philosophical Logic. Dordrecht: Springer Netherlands; 2002. p. 147–196.

[13] Van Fraassen BC. Singular terms, truth-value gaps, and free logic. The Journal of Philosophy. 1966;63(17):481–495.

[14] Russell B. On Denoting. Mind. 1905;14(56):479–493.

[15] Baader F, Horrocks I, Lutz C, Sattler U. An Introduction to Description Logic. Cambridge University Press; 2017.

[16] Haack S. Philosophy of Logics. Cambridge University Press; 1978.

[17] Lambert K. Free logic and definite descriptions. In: New Essays in Free Logic. Springer; 2001. p. 37–47.

[18] Artale A, Mazzullo A, Ozaki A, Wolter F. On Free Description Logics with Definite Descriptions. In: Proc. of the 33rd International Workshop on Description Logics. CEUR-WS; 2020. .

[19] Passy S, Tinchev T. An essay in combinatory dynamic logic. Information and Computation. 1991;93:263–332.

[20] Kutz O, Mossakowski T. A Modular Consistency Proof for Dolce. In: Twenty-Fifth Conference on Artificial Intelligence (AAAI-11). San Francisco, California: AAAI Press; 2011. .