Meaning Holism and Indeterminacy of Reference in Ontologies

Adrien BARTON[a, b,1], Paul FABRY[b] and Jean-François ETHIER[b,1]
[a] *Institut de recherche en informatique de Toulouse (IRIT), CNRS, Université de Toulouse, France*
[b] *Groupe de recherche interdisciplinaire en informatique de la santé (GRIIS), Université de Sherbrooke, Canada (QC)*
ORCiD ID: Adrien BARTON https://orcid.org/0000-0001-5500-6539
Paul FABRY https://orcid.org/ 0000-0002-3336-2476
Jean-François ETHIER https://orcid.org/0000-0001-9408-0109

Abstract. According to meaning holism, the meanings of all of the words in a language are interdependent. If this was true, then the very practice of building largely interconnected set of ontologies would be threatened. We examine here the extent of the severity of meaning holism for ontology engineering, based on several definitions of the meaning of a class term in an ontology, with regard to the classical analytic/synthetic distinction. We show that meaning holism is not as pervasive in ontologies as traditionally assumed in philosophy of language, and that a conception of meaning of a class term as a collection of statements expressing necessary conditions on this term limits meaning holism further. Still, meaning holism presents substantial challenges for ontology engineering and requires mitigation strategies. We also investigate the related phenomenon of indeterminacy of reference and show how anchoring formal ontologies in natural language can mitigate this problem, even if not fully control it.

Keywords. Meaning holism, Indeterminacy of reference, Analytic-synthetic distinction

1. Introduction

Ontologies aim to represent the general categories and relations within a given domain. They serve as means to facilitate semantic interoperability, enabling agents to share the meanings of the terms they use – consider e.g. Guarino et al. [1] characterization of ontologies as accounting "for the intended meaning of the vocabulary used by a logical language". However, a significant challenge concerning meanings has been largely discussed in philosophy since Quine [2] under the umbrella term of "meaning holism", which is defined in the Stanford Encyclopedia of Philosophy as follows [3]:

> (H) "The determinants of the meanings of our terms are interconnected in a way that leads a change in the meaning of any single term to produce a change in the meanings of each of the rest."

If meaning holism were as severe as formulated above, then adding any new term or altering the meaning of any existing term within an ontology would change the meanings of all terms within the ontology. However, ontologies are continually evolving to encompass a broader scope of reality or to keep pace with scientific advancements.

[1]Corresponding Authors: Adrien Barton, adrien.barton@irit.fr; Jean-François Ethier, jf.ethier@usherbrooke.ca

Consequently, this phenomenon would pose a very significant obstacle to the practical and sound use of ontologies.

This issue should be particularly prominent in large sets of ontologies that are interconnected by the use of upper-level or mid-level ontologies, such as the OBO Foundry [4]. For example, consider OGMS, a general ontology for medical science, which is based on the upper ontology BFO and serves as the foundation for several ontologies addressing specific families of diseases (e.g., IDO for infectious diseases, CVDO for cardiovascular diseases, etc.). If meaning holism were as extreme as described earlier, it would imply that creating a new class of specific disease in IDO or CVDO would alter the meanings of OGMS:Disease and BFO:Continuant, potentially leading to ripple effects across all other connected ontologies. This issue, related to the so-called problem of "unstability" [3], would then pose a significant challenge to the practice of ontologies. This phenomenon has arguably not received the attention it deserves in the literature so far. Vindicating the common practice of developing evolving, interconnected ontologies requires to analyze the real extent of the issue of meaning holism.

This raises two critical questions that this paper will address: 1) How can we define meaning in applied ontologies? 2) What is the extent of meaning holism in ontologies, compared to its characterization in HOL? Answering those questions is important in order to prepare the groundwork for mitigating this problem through the implementation of appropriate methodological rules in ontology engineering. The next section will introduce relevant considerations from philosophy of language and ontology pertaining to the issue.

2. Philosophy of Language, Meaning and Reference

Correctly addressing the issue of meaning holism necessitates delineating the distinction between meaning and reference, which in turn will facilitate addressing the related problem of reference indeterminacy. As we will explore, the analytic/synthetic distinction lies at the core of these phenomena.

2.1. Meaning and Reference

In the aftermath of Frege's distinction between "sense" and "reference" [5], philosophical theories have traditionally differentiated the reference of a term (the portion of reality to which it points) from its sense (the manner in which this term points to that portion of reality). For instance, "evening star" and "morning star" share the same reference (the planet Venus) but have different senses.

According to classical "internalist" conceptions of meaning, the meaning of a term aligns closely with the notion of sense and is typically located within the language where the term is used or within the cognitive structure of its users. Conversely, opposing theories of meaning, known as "externalist" perspectives, such as Putnam's [6], posit that the meaning of a term is partly determined by the external world. In this paper, we will adopt a classical, internalist conception of meaning as belonging to the language, as ontologies are often presented as a tool providing the meaning of their terms in order to foster semantic interoperability [7]. Furthermore, we will not take into account the pragmatics of language relying on contextual factors that can be relevant to meaning in a broader sense [8] – as the problem of meaning holism would appear even in the absence of any pragmatic consideration.

2.2. Indeterminacy of Reference

A phenomenon linked to meaning holism, and also challenging for the field of ontologies, is known as the "indeterminacy of reference" (or "inscrutability of reference"), famously articulated by Quine [9]. Consider a scenario where I find myself isolated in a remote linguistic community whose language I am entirely unfamiliar with. A native speaker of this language points to a rabbit and utters the word "Gavagai." In this situation, I am unable to ascertain whether "Gavagai" refers to the rabbit, a part of the rabbit, the history of the rabbit, or perhaps even a suggestion for dinner. In a nutshell, various theories about the reference of words are compatible with the empirical evidence about speakers of a language.

We can analyze the relationship between the indeterminacy of reference and meaning holism as follows: as exemplified by the Gavagai scenario, determining the reference of a single term solely from the behavior of speakers of a language is impossible. Consequently, we must rely on the (internal) meanings of terms to establish their reference. However, according to classical meaning holism (H), the meaning of a term is intricately linked to the meanings of other terms. Therefore, without comprehensive knowledge of the meanings of all terms in a language, it becomes challenging if not impossible to determine the meaning of a single term and consequently its reference. According to this analysis, there is a same underlying cause to both meaning holism and the indeterminacy of reference: the interconnectedness of meanings within language. However, as we will show later, the issue of indeterminacy of reference can persist even when the meanings of all terms are known. To comprehend this further, we must first introduce the fundamental distinction between analytic and synthetic statements.

2.3. The Analytic-Synthetic Distinction

As we discussed earlier, meanings can be encapsulated within the language and thus within statements. Throughout Western philosophy, there has been considerable interest in the distinction between analytic and synthetic statements[2], dating back to Kant [10]. This distinction can be elucidated as follows:

- An analytic statement is one whose truth is determined purely by the meanings of its terms. For example, "Bachelors are unmarried men."
- A synthetic statement, on the other hand, is one whose truth is determined not solely by the meanings of its terms. A classic example is "Bachelors are happy."

The former statement is analytic because it is true by definition of the term "bachelor". However, the latter is synthetic because its truth is not determined by the definitions of the terms; it is contingent upon empirical observation.

It is important to note that the analytic/synthetic distinction is distinct from the *a priori*/*a posteriori* distinction. Kant defined *a priori* beliefs as those justifiable independently of experience, like "2+2=4", by contrast to *a posteriori* beliefs (say, "the

[2] There are various accounts of analyticity that can concern statements, beliefs or judgments. Since we are here interested in analyticity in ontologies, we will focus on the analytic character of statements.

speed of light is 299,792,458 m/s"). However, Kant also argued for the existence of synthetic *a priori* statements – a claim that has been debated by later philosophers, and on which we will not take a stance here: the synthetic statements we will consider in this paper will be *a posteriori*.

Convention has traditionally been viewed as the cornerstone of analyticity, with Carnap [11] positing that meaning postulates are essentially conventional stipulations. However, Quine [9] challenged this notion by arguing that the selection of such meaning postulates is arbitrary. For instance, when defining Newtonian forces, one might interpret "F = ma" either as a defining feature or as an empirical observation that holds true due to the laws governing the physical world. In his renowned metaphor of the "web of belief" [12], Quine acknowledged that language is influenced by both conventions and empirical facts but rejected the idea that some statements are purely conventional while others are purely factual. The notion of analyticity was one of the two dogmas of logical empiricism that he famously criticized [2].

The discourse on analyticity has predominantly revolved around the status of logic, with original attempts to portray logic (and sometimes mathematics) as entirely analytic. Another domain where the concept of analyticity has been extensively debated is natural language. According to Quine, neither synonymy nor linguistic meaning, upon which Frege's conception of analyticity was based, could be clearly defined [2]. Quine argued that there is no fact of the matter regarding whether two expressions possess the same meaning or not. Partially supporting Quine's assertion, Putnam [13] argued that even seemingly analytic statements, like "Cats are animals", could potentially be revised if, for instance, we were to discover that cats are actually robots deployed by Martians to manipulate humanity[3].

However, there are certain statements that appear to be strong candidates for being analytic, such as "A fortnight is a period of fourteen days" [14]. Putnam [13] attempted to salvage at least some analytic truths by introducing the notion of "one-criterion" concepts, such as *pediatrician, bachelor* or *widow*, for which there is only "one way" to tell whether they apply[4].

Although such issues have been little discussed in applied ontology to our knowledge, a notable exception is Hastings & Neuhaus [16] claiming that "[t]he representation of empirical knowledge is a secondary task of an ontology, its primary task is to establish a vocabulary for a formal language, which may be used to express empirical knowledge." This viewpoint suggests that analyticity lies at the core of ontology development, while the expression of *a posteriori* statements is considered a secondary concern. This perspective is to be contrasted to claims such as: "Ontology is concerned with representing the results of science at the level of general theory (the generalizations and laws of science)" [7], which advocate for ontology to reflect our best scientific knowledge of the world.

[3] In ontologies, this conclusion could be eschewed, since terms are usually not identified with their labels. If the term C_001 with label "Cat" is defined (say) as referring to the subkind of animal that meows, then one might conclude, after realizing that what we named "cats" are robots, that there are no instances of C_001 on Earth; one could then introduce a term C_002 defined as the subkind of robot that meows, and transfer the label "cat" from C_001 to C_002. Natural language statements of instantiations such as "Willow Biden is a cat" would be kept, while the analytic formal definition of C_001 as animal that meows would not have been altered; and particulars formerly considered as instances of C_001 would now be considered as instances of C_002.

[4] This is related to the distinction between *stipulative* definitions, which introduce novel meanings for terms, and *descriptive* definitions, which seek to represent existing usages of a term, on which we will not elaborate here [15].

2.4. How Ontologies Can Capture Meaning

Guarino et al. [1] compare extensional and intensional accounts of meaning and identify Carnapian "meaning postulates" with the axioms of an ontology. If one believes in the analytic/synthetic distinction, though, not every axiom should be seen as an expression of meaning.

In the OBO Foundry, there are at least two annotation properties aimed at capturing meaning. The first property is "Definition" (IAO_0000115), which is described as "explaining the meaning of a class or property." Arp et al. [7] state that "the definition of a term captures what we can think of as the essential features of the entities that are instances of the designated type." The second property is "Elucidation" (IAO_0000600), defined as "illustrative examples, statements of recommended usages, and axioms." These are particularly used for high-level terms for which necessary and sufficient conditions are difficult to determine.

In OWL, we could envision generalizing such annotation properties in order to tag statements as analytic or synthetic. To borrow a famous example by Quine [2], suppose that we 1) define stipulatively the term "VH" in an OWL ontology as a "Vertebrate with a heart", and 2) state that it was found that, as a matter of fact, VHs are exactly the vertebrates with a kidney. Then we might introduce the two following statements, the first one being tagged analytic and the second one as synthetic:

(AX_H) VH EquivalentTo (Vertebrate and has_part some Heart)
(AX_K) VH EquivalentTo (Vertebrate and has_part some Kidney)

On the opposite, if one would want to define VH as a "Vertebrate with a kidney" and express that as a matter of fact, the VHs are exactly the vertebrates with a heart, one would tag AX_H as synthetic and AX_K as analytic.

In the following, we will start by characterizing meanings as provided by definitions in section 3, which will enable to characterize the phenomenon of indeterminacy of reference in section 4. Then, in section 5, we will assess the real extent of meaning holism based on more general conceptions of meaning, according to which it is provided by some analytic statements – as well as according to views refusing the analytic/synthetic distinction.

3. Meaning and Definitions

3.1. Natural and Formal Statements

An ontology introduces terms and statements in both natural and formal languages (see e.g. [17] for a justification of the importance of both). In this paper, we will concentrate on terms that refer to a class, that we will call "class terms" (by contrast to terms that refer to a relation or to a particular). For instance, in OWL 2, the logical language used is the description logic SROIQ(D) [18]. OWL class terms are IRIs, to which natural language labels can be associated. Additionally, they can be linked with natural language statements through annotation properties, and appear within axioms in description logic.

In line with the internalist perspective described earlier, we will consider that the meaning of each class term is constituted by some of those statements[5]. We will distinguish between the "formal" meaning of a term, expressed by formal statements in a formal language, and the "natural language" meaning of a term, conveyed through textual statements in natural language – where the latter is usually more expressive than formal languages. Since the articulation between formal and natural languages is a matter of debate and natural language constructs are harder to define than for formal language, we will concentrate in this paper on the former; note however that natural language might be an integral, indispensable part of an ontology, as argued by Neuhaus and Smith [19]: an ontology in which the correspondence between formal term and natural language would be totally cut would likely be impossible to understand by anyone. Formal definitions emerge as the primary candidate statements for expressing the formal meaning of a term.

3.2. Meaning and Definitions

In a first conception named "MEAN_0", the meaning_0 of a term is identified by its definition, which is an analytic formal or natural language statement expressing a necessary and sufficient condition (NSC, the *definiens*) for the term (the *definiendum*) that does not mention the *definiendum*. In particular, we assert:

> (MEAN_0) The formal meaning_0 of an ontological term is the formal definition of this term.

In practice, both natural language definitions and formal definitions can be found within an ontology. In OWL, definitions take the form of an axiom 'A EquivalentTo Expr' (where Expr is an anonymous class that does not mention A) – but not all such axioms are definitions, as illustrated by the example of synthetic axioms such as AX_K above: they might express a coincidence between two classes that is due to natural regularities.

Definitions can minimize the risk for the reader to import their own idiosyncrasies when interpreting terms, as ontological terms are intended to facilitate communication among agents that may start with different understandings. Seppälä et al. [20] argue that definitions of terms have (at least) two roles: ensuring the inferential function and the referential function of those terms. As explained above in the second paragraph of section 2.2, the inferential dimension provided by their meanings can help to determine their reference. While one might aspire to solely use formal language and discount any use of natural language, this approach is not realistic, as we are going to see in Section 4.

3.3. Primitive Terms and Circularity in Formal Definitions

When constructing formal definitions within ontologies, terms are used to define other terms, and these defining terms may themselves be formally defined using additional terms. At some level, this process must lead to one of the following scenarios: circularity,

[5] In this paper, we identify first here the meaning of a term with its definition, and later in Section 5 with a collection of statements. Note that this is a first approximation designed to expose the phenomena of indeterminacy of reference and meaning holism and characterize their extent and limits. In a more sophisticated account, meanings might rather be identified with language-independent propositions, as suggested by Neuhaus [17]. We will also not consider the details of how natural language meanings and formal meanings can interact to lead to a unified conception of meaning.

(that is: the formal definition of the term t_0 uses a term whose definition uses a term … whose definition uses this term t_0); primitiveness, when there is no NSC formal statement associated with some terms in the ontology; or a combination of both.

Term	Label	Natural Language Definition	Formal Definition
F_002	Chair	"A chair is an entity in which inheres a chair function."	F_002 EquivalentTo (R_005^{-1} some F_003)
F_003	Chair function	"A chair function is a function that inheres in a chair."	F_003 EquivalentTo [F_004 and (R_005 some F_002)]
F_004	Function	*Primitive*	*Primitive*
R_005	inheres_in	*Primitive*	*Primitive*

Figure 1. Terms, labels and definitions in the OWL ontology O_1

To illustrate, consider the OWL ontology O_1 above involving both primitiveness and circularity. First, F_004 and R_005 are primitive. Second, F_002 is defined in terms of F_003, and F_003 is defined in terms of F_002; thus, their definitions are circular.

As we shall see, both cases of primitivity and circularity present challenges regarding the "indeterminacy of reference" within ontologies.

4. Indeterminacy of Reference

To elucidate the issue of indeterminacy of reference in ontologies, it is essential to explain first how ontologies handle reference.

4.1. Reference in Ontologies

We will here introduce informal extensional notions of interpretation, model and intended reference in an ontology, in the spirit of Guarino et al.'s [1] first extensional account of an ontology[6], as this account is tractable enough to characterize and discuss the notions of indeterminacy of reference and meaning holism.

In the view developed here, ontological terms refer to the world thanks to *interpretations* mapping each term with a portion of reality of our world (as defined in [21]). For example, realists typically interpret class terms as referring to universals, while nominalists interpret them as referring to collections of particulars [22] – we will use in this paper the neutral word "class" to remain agnostic between those two approaches.

Let's introduce further the notion of theory of an ontology:

> **Formal theory of an ontology**:$=_{def}$ The collection[7] of formal statements explicitly formulated[8] within this ontology.

[6] Guarino et al. [1] later propose in their paper an intensional conception of an ontology, according to which interpretation map towards a function that associates to each possible world a portion of reality; such an intension might be called an "intended meaning", to be contrasted with the "expressed meaning" that is to be found in the ontology. In the remainder of the paper, "meaning" will refer to such an expressed meaning – namely collection of statements in an ontology. See also Neuhaus [17] for a different view of ontologies according to which interpretations map towards language-independent propositions. Note that an interpretation requires an interpretant, as expressed by the triplet sign/interpretant/portion of the world that is at the core of semiotics. We will not discuss this notion here.

[7] We use the term "collection" as a general primitive term that does not commit to the notions of set or mereological sum. For some possible analysis of collections and related notions, see Masolo et al. [23]

[8] Note that theorems in an ontology do not belong to its formal theory, but to the deductive closure of its formal theory – see section 5 for more details.

Here, we understand by "statement" an assertion intended to be interpreted as being about the referents of its terms (including both logical axioms and what Neuhaus [17] called "assertive annotations"; by contrast e.g. to metadata assertions such as author attribution or date of creation specification). For example, the formal theory of an OWL ontology written in SROIQ(D) is constituted by the T-Box, R-Box and A-Box axioms to be found in the ontology; its natural language theory would typically be its set of annotations characterizing the classes, relations and individuals referred to by its terms. The kind of interpretations we are primarily interested in ontological theories are *models*:

> **Formal model of an ontological theory**:$=_{def}$ An interpretation of the terms of an ontology making true every statement of the ontology's theory.

For instance, consider an OWL ontology O_1 endorsing terms "A" and "B", with the theory only stating 'A SubClassOf B'. An interpretation mapping A to the class *Cat* and B to the class *Animal* is a model of O_1, but an interpretation mapping A to *Animal* and B to *Cat* is not. Axioms constrain valid interpretations by limiting the number of models.

Note that different agents may endorse different interpretations of a term in a language. For example, this nightstand might belong to John's interpretation of the term "table" (he would say: "This nightstand is a table") whereas it might not according to Mary's interpretation (she would say: "This nightstand is not a table"). The portion of reality pointed by a term according to the interpretation of the term's creator is called here the "intended reference"[9]. Since ontology creators cannot point directly to the world, they build a theory conveying the meaning of those terms in order to help readers exclude models with non-intended references.

Note that analytic and synthetic statements do not play the same role when using ontologies to make judgment of instantiations. As a matter of fact, analytic statements effectively constrain the reference of a term, whereas synthetic statement express a regularity that is contingent upon how the world is. In the example of section 2.4, if AX_H is tagged as analytic and AX_K as synthetic, then the reference of "VH" is the class of vertebrate with a heart. One might take the risk to classify a particular organism with a kidney as an instance of VH on the basis of AX_K, but one might be wrong, since synthetic axioms express empirical and thus falsifiable knowledge. Thus, synthetic statements can merely act as heuristic devices when making judgments of instantiation. This role difference between analytic and synthetic statements can provide a rationale for labeling statements as either analytic or synthetic in ontological engineering, a practice that seems to be currently uncommon or even entirely absent to our knowledge.

As we are now going to see, several interpretations are possible even for a well-defined term: this is the problem of indeterminacy of reference.

4.2. Indeterminacy of Reference in Ontologies

When considering solely the formal statements within a theory, numerous interpretations of primitive terms within an ontology can arise. For instance, in the ontology O_1, even if a singular interpretation of the primitive terms R_005 and F_004 is assumed, there can be several models. For instance, within a BFO-inspired ontology, F_002 and F_003 could be interpreted as the classes *Table* and *Table function*, *Chair* and *Chair function*, *Door* and *Door function*, and so forth.

[9] See [17] for further considerations on the notion of "intended interpretation".

The issue of indeterminacy of reference is obvious in the current, straightforward cases of O_1. The more intricate the theory, the greater the challenge in aligning it with reality through various mappings. However, this phenomenon remains theoretically possible in all cases. And as a matter of fact, those acquainted with ontological practice would recognize that disagreements can arise from specific terms, even within a comprehensive ontology.

Thus, one cannot ascertain whether two ontology users refer to the same portion of reality, even when they employ the same language and endorse the same theory. In other words, two individuals might accept identical statements and consequently delineate reality in isomorphic ways, yet there may still be discrepancies in their references: certain terms may denote distinct portions of reality based on their interpretations, which could vary slightly or significantly.

Garbacz [24] illustrates a special case of the indeterminacy of reference, wherein perfectly symmetrical treatments of terms make it impossible to differentiate their intended reference. In such cases, if a model of the theory interprets the term t_1 as the class A and the term t_2 as the class B, there exists another model of the theory that interprets t_1 as B and t_2 as A.

The indeterminacy can be partially alleviated by linking our ontological language to natural language through natural language definitions (see also [16] for considerations on the importance of natural language in ontologies). However, if one follows Quine (cf. section 2), natural languages themselves are susceptible to the problem of indeterminacy of reference. Consequently, the indeterminacy of reference in natural language will contaminate the ontological language.

The communicative challenges of ontology have been emphasized by Guarino et al. [1] ("Since however meaning postulates cannot fully characterize the ontological commitment of primitive terms, one may recognize that sharing of conceptualizations is at best partial." and "it is important […] that the basic primitives [ontologies] are built on are sufficiently well-chosen and axiomatized to be generally understood") and by Neuhaus [17] ("The authors of an ontology version specify an annotated theory in such a way that the users are able to grasp its intended interpretation."). The difficulty is, however, how to determine whether primitives are indeed "generally understood", and how we can certify that users grasp the "intended interpretation": since we typically use language to improve such an understanding or communicate such an interpretation, the indeterminacy of reference of both formal and natural languages precludes any certitude that we do indeed share a common understanding or interpretation.

Overall, the indeterminacy of reference permeates all languages, including ontological ones. This uncertainty persists even when agents employ perfectly identical ontological statements, making it unclear whether they are referring to the same reality using the same terms. The best we can do is mitigating this phenomenon by providing well-chosen additional statements, without certainty that we can fully control it.

5. Meaning Holism

Not all class terms in an ontology have definitions: in some ontologies, some terms are characterized by a set of necessary conditions [NC] without any necessary and sufficient formal or natural language definitions. For example, instead of AX_H, an ontology may accept the following statement tagged as analytic: 'VH SubClassOf (Vertebrate and has_part some Heart)', leaving it open whether all vertebrates with a heart instantiate VH.

Of course, this incomplete characterization would leave the reference of VH indeterminate – but as we saw earlier, the problem of reference indeterminacy also appears for well-defined terms. This phenomenon motivates a deeper exploration of how to define meaning in ontologies and the extent of meaning holism.

5.1. Meaning and Inferential Role

Assume first that the analytic/synthetic distinction is valid, *pace* Quine. We define the analytic formal theory of the ontology as the collection of statements tagged as analytic in the formal theory of the ontology (it is thus a *fiat* decision of the ontology creator which statements are analytic, the same way it is up to them which statements are definitional). We can then consider the deductive closure of this theory, namely, the collection of statements that can be deductively inferred from them using the underlying logic. While formal statements in an ontology adhere to well-defined languages like FOL, DL, CLIF or OWL, natural language statements rely on inferential rules that are less strictly defined (although generally more expressive), introducing additional ambiguity to what belongs to the deductive closure of those statements and thus to the natural language meanings of the terms[10]. For this reason, we will concentrate here on the formal meanings of ontological terms. Given the definition of analyticity, any statement in the deductive closure of the analytic theory is also analytic.

Then, we define the formal $meaning_1$ of an ontological term as follows:

> ($MEAN_1$) The formal $meaning_1$ of a class term in an ontology O is the collection of formal statements expressing NCs (including NSCs) or sufficient conditions [SCs] on this term in the deductive closure of the analytic theory.

In particular, $MEAN_1$ can easily be operationalized in OWL as follows:

> $(MEAN_1)^{OWL}$ The formal meaning of a class term A in an OWL ontology is the collection of all statements of the form 'Expr SubClassOf A', 'A SubClassOf Expr' and 'A EquivalentTo Expr' in the deductive closure of the formal analytic theory.

5.2. Meaning Holism in Formal Ontologies and a First Limit

Let's now illustrate meaning holism with an example. Consider an initial theory containing only the analytic statement AX_H. Suppose now that we modify the meaning of "Vertebrate" by adding one of the following analytic axioms: 'Vertebrate SubClassOf Animal' or 'Mammal SubClassOf Vertebrate'. As a result, the deductive closure of the new analytic theory will now include the theorem 'VH SubClassOf (Animal and has_part some Heart)' or '(Mammal and has_part some Heart) SubClassOf VH' and thus, in either case, the $meaning_1$ and inferential $role_1$ of "VH" are changed. Hence, adopting $MEAN_1$ or IR_1 implies that the meanings and inferential roles of certain terms in an ontology are interconnected, as explained by meaning holism. As we will see, however, meaning holism is not as pervasive as claimed by HOL.

[10] Facing the same difficulty, Neuhaus [17] considers that the propositions asserted by an ontology encompass the formulas entailed by its logical theory and its annotations, but does not consider any notion of natural language entailment.

To show this, let's consider an example supposed to illustrate meaning holism [3]:

> "For instance, a word like "squirrel" might be inferentially connected to, say, "animal" which is in turn connected to "Koala" […] and through similar chains, every word will be related inferentially to (and thus semantically entangled with) every other term in the language (especially when one considers connections like that between, say, "is a squirrel" and "is not a building" or any other thing we take squirrels not to be)."

Let's examine a possible OWL ontology formalizing this example, where the following axioms are all analytic:

(AX_1) Squirrel SubClassOf Animal
(AX_2) Koala SubClassOf Animal
(AX_3) Disjoint_with(Squirrel,Building)

Interestingly, this example shows that meaning holism does not operate as systematically as stated by H when we endorse $(MEAN_1)^{OWL}$ or $(IR_1)^{OWL}$. Let's start with a theory limited to AX_1. When adding AX_2 to the theory (that is, in the theory (AX_1, AX_2)), no CN or CS on "Squirrel" is added to the deductive closure of the theory[11]. Thus, the $meaning_1$ of "Squirrel" and its inferential $role_1$ remain unchanged (however, adding either AX_3 or the axiom 'Black_squirrel SubClassOf Squirrel' would alter the $meaning_1$ and inferential $role_1$ of "Squirrel").

Therefore, meaning holism is not as systematic as claimed in H when one adopts $MEAN_1$ within formal ontologies: the $meanings_1$ of some terms can be changed without altering the $meaning_1$ of some other term. The same applies to inferential $roles_1$. This observation is noteworthy, considering that using inferential $roles_1$ as proxy for meanings would amount to a very broad theory of meaning, taking into account the deductive closure of the entire theory, rather than just its analytic component. However, the fact that adding a taxonomical axiom would change the $meaning_1$ of the parent term (e.g., 'Black_squirrel SubClassOf Squirrel' changing the meaning of "Squirrel") still poses a significant challenge for interconnected ontologies, such as those structured into foundational, mid-level, and domain ontologies like OBO Foundry ontologies. But, as we are going to see, not only IR_1 but also $MEAN_1$ are arguably too broad conceptions, and we can restrict them by adopting instead a "top-down" approach to meaning.

5.3. Second Limit on Meaning Holism: A Top-down Conception of Meaning and Inferential Role

The new conception of meaning is based on the intuition that the meaning of a class term is determined by the general analytic claims concerning it – claims that apply to any instances of that class, i.e., by *necessary* analytic conditions on that term. Indeed, those "involved in knowledge representation […] are mainly interested in what is general" [25]. Additionally, as highlighted by Arp et al. above [7], definitions encapsulate what is essential in a term – and what is essential is general. Seppälä et al. [20] emphasizes the

[11] This example also illustrates why $MEAN_1$ restricts the meaning of a term to *statements expressing necessary and/or sufficient conditions on this term*, rather than encompassing *all statements using this term* in the deductive closure of the analytic theory. Adding AX_2 to AX_1 entails adding a theorem mentioning "Squirrel", namely '(Squirrel or Koala) SubClassOf Animal'. However, intuitively, this theorem does not augment the meaning of Squirrel compared to AX_1. Additionally, if B is a term in the ontology, then '(Squirrel and B) SubClassOf Animal' would be a theorem of the theory, but arguably it does not convey anything further about the meaning of "Squirrel" beyond what AX_1 already conveys.

primacy of necessary conditions over sufficient ones, though ideally, introduced conditions should be both necessary and sufficient.

In this "top-down" conception of meaning, the meaning of a term is constituted by its necessary conditions (NC), including its necessary and sufficient conditions (NSC), but not by its sufficient conditions (SC).

Another restriction must be made. Adding tautologies such as 'A SubClassOf Thing', 'A SubClassOf (B or not-B)' or 'A SubClassOf (A or B)' should not change the meaning of "A". Also, adding the axiom 'C EquivalentTo (A or B)' would imply the axiom 'A SubClassOf C', which intuitively does not change the meaning of "A". Since such an axiom is equivalent to the tautology 'A SubClassOf (A or B)' in the presence of the former axiom, we should also exclude axioms that are equivalent to a tautology in the deductive closure of the ontology (which we call "non-tautological axioms").

This can be formulated as follows:

> ($MEAN_2$) The formal meaning$_2$ of a class term in an ontology O is the collection of non-tautological axioms expressing NC (including NSC) on this term within the deductive closure of O's analytic theory.

In particular, this can be easily operationalized in OWL:

> $(MEAN_2)^{OWL}$ The formal meaning of a class term A in an OWL ontology O is the collection of non-tautological axioms of the form 'A SubClassOf Expr' and 'A EquivalentTo Expr' (where Expr is a named or anonymous class) within the deductive closure of O's analytic theory.

In cases where a non-tautological axiom of the form 'A SubClassOf Expr' or 'A EquivalentTo Expr' belongs to the deductive closure of the analytic theory and "B" appears in "Expr", we will say that the meaning$_2$ of "A" *immediately depends on* the meaning$_2$ of "B", and we can introduce the relation "*depends on*" as the transitive closure of the relation "*immediately depends on*". When the meaning$_2$ of "A" depends on the meaning$_2$ of "B", the interpretation of the reference of "B" will constrain the interpretation of the reference of "A". If "B" undergoes a change of meaning$_2$, such as adding an axiom introducing a named or anonymous superclass, or changing such a superclass, then the meaning$_2$ of "A" might also be altered (although it might remain unaltered in some cases), indicating an upward meaning dependency – and the interpretation of the reference of "A" might thus be further constrained too.

In $MEAN_2$, the meaning of a term is determined by its necessary conditions, while changes or additions of sufficient conditions generally do not alter the meanings of other terms (e.g. adding 'A SubClassOf B' to the ontology's theory generally does not change the meaning of B). However, there are some clarifications and caveats to consider.

First, it's important to clarify that $MEAN_2$ does not preclude cases where the meaning of a class depends on one of its subclasses. For example, in an ontology with the following axioms:

(AX_4) A SubClassOf B
(AX_5) B SubClassOf (R some A)

the meaning of "B" depends on the meaning of "A" due to AX_5 (but not in virtue of AX_4).

Secondly, there is a caveat to consider: necessary conditions on a class can impose necessary conditions on another class that is not a subclass of it. For instance, if the analytic axiom ‘A SubClassOf not-B’ (indicating that A and B are disjoint) is added to the theory, then the equivalent statement ‘B SubClassOf not-A’ appears in the deductive closure of the theory, and thus the meaning of “B” according to $MEAN_2$ has been changed. Additionally, if class A is covered by the class “C or D” (i.e., ‘A SubClassOf (C or D)’ is in the theory) and the axioms ‘C SubClassOf Expr’ and ‘D SubClassOf Expr’ are added, then the axiom ‘A SubClassOf Expr’ is added in the deductive closure of the theory, altering the meaning of “A” according to $MEAN_2$.

6. Discussion and Conclusion

The problem of indeterminacy of reference pervades any language, including ontological ones, marked by circular definitions or primitive terms, leading to the possibility of unintended interpretations. It can be alleviated, though arguably not fully controlled, by incorporating carefully selected additional statements, formal or natural.

Analytic statements effectively constrain the reference of terms, whereas synthetic statements can be used at most as heuristic tools for judgments of instantiation: this motivates the introduction of the analytic/synthetic distinction into ontological engineering, a practice largely overlooked today.

Meaning holism as classically formulated in H would make the practice of ontological engineering nearly impossible. We saw that this phenomenon is not as strong and general as formulated in H when we operationalize meaning with $MEAN_1$ or IR_1, but would still make the practice of interconnected ontologies as in e.g. the OBO Foundry methodology basically impossible. Fortunately, one can devise a reasonable, more restricted theory of meaning, namely the top-down conception $MEAN_2$, which fits well with the consideration of ontologies as characterizing what is general in the world, and limits meaning holism. Moreover, it can be easily operationalized in OWL.

An open question is whether $MEAN_2$ is still a too large conception of meaning and should be further restricted. Consider the OWL axiom ‘A SubClassOf (R only B)’, which would be, according to $MEAN_2$, part of the meaning of A. Intuitively, this axiom does not say something about any instance of A, but only about the instances of A that are in relation R with something - namely, it states that this something must be a B. Thus, it is logically equivalent to: ‘(A and (R some Thing)) SubClassOf (R only B)’. Such axioms might have to be excluded from the meaning of A, and instead assigned to the meaning of any named class equivalent to (A and (R some Thing)).

Note that if we refuse the analytic/synthetic distinction, we can introduce the notion of “inferential role” as an operational proxy for meaning as the collection of formal statements expressing NCs on a term *within the deductive closure of O’s theory* (rather than within its deductive closure). This might be a good proxy if statements in an ontology are mostly analytic in nature, as suggested by Neuhaus & Hastings [16] or as seems implicitly assumed by Guarino et al. [1].

While $MEAN_2$ limits meaning holism, still any change of meaning of a term B might change the meaning of the terms whose meaning depend on B. This phenomenon is not enough controlled in the current practice of applied ontology using interconnected ontologies. One must be cautious not to alter the $meaning_2$ of class terms by introducing or changing analytic necessary conditions—neither directly nor indirectly through disjunction axioms or axioms subclassing mutually covering classes, as explained in section 5.3. This holds for class terms that one did not author (as one might not have the same reference as the authors of the term [26]), or on class terms that one did author but have already been made public and thus might have been reused by someone else. In such a case, alternative strategies should be used, such as the introduction of a new term or suggesting changes to the author of the term. Future work will focus on developing a rigorously structured versioning system for terms to address this aspect of holism.

$MEAN_2$ specifically applies to class terms. In OWL, one might consider adapting it to object property terms by considering the axioms in the R-Box, but the only axioms in the R-Box in SROIQ(D) are purely taxonomic axioms (using SubPropertyOf), domain/range axioms and axioms describing properties such as symmetry, irreflexivity, inverse property, etc. – which only very partially characterize relations. Therefore, further research is needed to delve deeper into the question of the meaning of object property terms in OWL and more generally of relation terms.

This analysis should also investigate the import of natural language statements in constraining ontologies. The analysis presented here should be operationalized in ontologies written in other languages than OWL, such as FOL or CLIF. Future work should also analyze further the analytic/synthetic distinction in ontology engineering, the adoption or rejection of which would have consequences on whether a conception of meaning like $MEAN_2$ should be used, or instead an operational substitute like IR_2. In particular, the distinction between analytic and synthetic true statements might be characterized logically by analyzing the former as statements true in every possible world, and the latter as statements true in merely every epistemically possible world. The status of OntoClean [27] metaproperties in determining the meaning of classes terms should be analyzed. The phenomena of indeterminacy of reference and meaning holism could also be analyzed in more formal frameworks of the nature of ontologies (e.g. considering that classes terms are associated to intensions [1], namely functions that associate to each possible world a portion of reality in this world; or identifying meanings with collections of propositions as in [17], rather than as collections of statements). Future work should also control other aspects that complicate the connection between meaning and reference, such as the possibility of making errors when expressing the meaning of a term in regard of its intended reference (as analyzed in [26]). Finally, the import of those issues for the Semantic Web should be analyzed: is such an endeavor possible at all given meaning holism?[12]

References

[1] N. Guarino, D. Oberle, S. Staab, What is an ontology?, *Handbook on Ontologies* (2009), pp. 1–17.

[2] W.V.O. Quine, Two Dogmas of Empiricism, in: *From a Logical Point of View: Nine Logico-Philosophical Essays*, Harvard University Press, 1980, pp. 20–46.

[3] H. Jackman, Meaning Holism, *The Stanford Encyclopedia of Philosophy* (Winter 2020 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2020/entries/meaning-holism/>.

[4] R. Jackson, N. Matentzoglu, J.A. Overton, R. Vita, J.P. Balhoff, P.L. Buttigieg, et al., OBO Foundry in 2021: operationalizing open data principles to evaluate ontologies, *Database* **2021** (2021), baab069.

[5] G. Frege, Über sinn und bedeutung, *Zeitschrift Für Philosophie Und Philosophische Kritik* **100** (1892), 25–50.

[6] H. Putnam, The meaning of "meaning," in: *Reprinted in Mind Language and Reality*, Cambridge University Press, New York, 1975, pp. 215–271.

[12] We thank Fabian Neuhaus for helpful feedback and discussions on an earlier of this paper, as well as two anonymous reviewers and Laure Vieu for their valuable suggestions.

[7] R. Arp, B. Smith, A.D. Spear, *Building Ontologies with Basic Formal Ontology*, The MIT Press, 2015.
[8] K. Korta, J. Perry, Pragmatics (Spring 2020 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2020/entries/pragmatics/>.
[9] W.V.O. Quine, *Word and Object*, MIT press, 2013.
[10] I. Kant, *Kritik der Reinen Vernunft*, Akademie Verlag GmbH, 1998.
[11] R. Carnap, Meaning postulates, in: *Meaning and Necessity*, University of Chicago Press, Chicago, 1956, pp. 222–229.
[12] W.V.O. Quine, Carnap and Logical Truth. In: *The Ways of Paradox and Other Essays*, Harvard University Press, Cambridge, Mass., 1976, pp. 107–132.
[13] H. Putnam, It ain't necessarily so, *The Journal of Philosophy* **59** (1962), 658–671.
[14] G. Fey, The Analytic/Synthetic Distinction, *The Stanford Encyclopedia of Philosophy* (Spring 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/spr2023/entries/analytic-synthetic/>.
[15] A. Gupta, S. Mackereth, Definitions, *The Stanford Encyclopedia of Philosophy* (Fall 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/fall2023/entries/definitions/>.
[16] F. Neuhaus, J. Hastings, Ontology development is consensus creation, not (merely) representation, *Applied Ontology* **17**:4 (2022), 495–513.
[17] F. Neuhaus, What is an Ontology?, (2018), arXiv preprint arXiv:1810.09171.
[18] I. Horrocks, O. Kutz, U. Sattler, The Even More Irresistible SROIQ., *Kr* **6** (2006), 57–67.
[19] F. Neuhaus, B. Smith, Modelling Principles and Methodologies–Relations in Anatomical Ontologies, in: *Anatomy Ontologies for Bioinformatics: Principles and Practice*, Springer, 2008, pp. 289–306.
[20] S. Seppälä, A. Ruttenberg, B. Smith, The functions of definitions in ontologies, in: R. Ferrario, W. Kuhn (Eds.), *Formal Ontology in Information Systems. Proceedings of the 10th International Conference (FOIS 2016)*, IOS Press, 2016, pp. 37–51.
[21] B. Smith, W. Ceusters, Aboutness: Towards foundations for the information artifact ontology, in: *Proceedings of the 6th International Conference on Biomedical Ontology*, CEUR-WS.org, Lisbon, Portugal, 2015, pp. 1–5.
[22] D.M. Armstrong, *A Theory of Universals: Volume 1: Nominalism and Realism*, Cambridge University Press, 1980.
[23] C. Masolo, L. Vieu, S. Borgo, D. Porello, Pluralities, collectives, and composites, in: B. Brodaric and F. Neuhaus (Eds.), *Formal Ontology in Information Systems: Proceedings of the 11th International Conference (FOIS 2020)*, IOS Press, 2020, pp. 186–200.
[24] P. Garbacz, Does your ontology make a (sense) difference?, in: Maureen Donnelly & Giancarlo Guizzardi (Eds.), *Formal Ontology in Information Systems: Proceedings of the 7th International Conference (FOIS 2012)*, IOS Press, pp. 177–190.
[25] K. Munn, B. Smith, *Applied Ontology: An Introduction*, Ontos, 2008.
[26] P. Fabry, A. Barton, J.-F. Ethier, Version control for interdependent ontologies: challenges and first propositions, in: *Proceedings of the International Conference on Biomedical Ontologies 2023 (ICBO 2023)*, in press.
[27] N. Guarino, C.A. Welty, An overview of OntoClean, *Handbook on Ontologies* (2009), 201–220.