

Polymorphism as a Model for Ambiguity: the Case of Nominal Modification

Walid S. Saba

We suggest modeling concepts as types in a strongly-typed ontology that reflects our commonsense view of the world and the way we talk about it in ordinary language. In such a framework, certain types of ambiguities in natural language are explained by the notion of polymorphism. In this paper we suggest such a typed compositional semantics for nominal compounds of the form (Adj Noun) where adjectives are modeled as higher-order polymorphic functions. In addition to (Adj Noun) compounds our proposal seems also to suggest a plausible explanation for well known adjective ordering restrictions.

1 Introduction

Over two decades ago a "quite revolution", as Charniak (1995) once called it, overwhelmingly replaced knowledge-based approaches in natural language processing (NLP) by quantitative (e.g., statistical, corpus-based, machine learning) methods. In recent years, however, the terms *ontology*, *semantic web* and *semantic computing* have been in vogue, and regardless of how these terms are being used (or misused) we believe that this 'semantic counter revolution' is a positive trend since corpus-based approaches to NLP, while useful in some language processing tasks - see (Ng and Zelle, 1997) for a good review - cannot account for compositionality and productivity in natural language, not to mention the complex inferential patterns that occur in ordinary language use. The inferences we have in mind here can be illustrated by the following example:

- (1) *Pass that car will you.*
 (a) *He is really annoying me.*
 (b) *They are really annoying me.*

Clearly, speakers of ordinary language can easily infer that 'he' in (1a) refers to the person driving [that] car, while 'they' in (1b) is a reference to the people riding [that] car. Such inferences, we believe, cannot theoretically be learned (how many such examples will be needed, and what exactly would constitute a negative example in this context?), and are thus beyond the capabilities of any quantitative approach. On the other hand, and although it is our firm belief that purely quantitative approaches cannot be the only paradigm for NLP, dissatisfaction with purely engineering approaches to the construction of large knowledge bases for NLP (e.g., Lenat and Ghua, 1990) are somewhat justified. While language 'understanding' is, for the most part, a commonsense 'reasoning' process at the pragmatic level, as example (1) illustrates, the knowledge structures that an NLP system must utilize should have sound linguistic and ontological underpinnings and must be formalized if we ever hope to build scalable systems (or, as John McCarthy once said, if we ever hope to build systems that we can actually understand!). As

we have argued elsewhere (Saba, 2007), therefore, we believe that both trends are partly misguided and that the time has come to enrich logical semantics with an ontological structure that reflects our commonsense view of the world and the way we talk about it in ordinary language. Specifically, we argue that very little progress within logical semantics have been made in the past several years due to the fact that these systems are, for the most part, mere symbol manipulation systems that are devoid of any content. What we suggest instead is a semantics that is grounded in a strongly-typed ontology - an ontology that reflects our commonsense view of reality and the way we talk about it in ordinary language.

In this paper we suggest exactly such a semantics and we subsequently demonstrate the utility of this approach by tackling one particular challenge in the semantics of natural language. Specifically, in this paper we will first introduce the notions of intersective vs. non-intersective adjectives, as well as the notion of adjective-ordering restrictions. In section 2 we will introduce a strongly typed system that reflects our commonsense view of the world and the way we talk about it in ordinary spoken language. In the rest of the paper we will suggest how such a strongly-typed compositional system can possibly utilize such information to explain the adjective-ordering restriction phenomenon as well as the type of ambiguity that occurs in nominal modification.

2 Ambiguity in Nominal Modification

The ambiguity in nominal modification we are concerned with here can be illustrated by the sentence in (2), which could be uttered by someone who believes that: (i) *Olga* is a dancer and a beautiful person; or (ii) *Olga* is beautiful as a dancer (i.e., *Olga* is a dancer and she dances beautifully).

- (2) *Olga is a beautiful dancer.*

As suggested by Larson (1998), there are two possible routes to explain this ambiguity: one could assume that a noun such as *dancer* is a simple one place predicate of type $\langle e, t \rangle$ and 'blame' this ambiguity on the adjective; alterna-

tively, one could assume that the adjective is a simple one place predicate and blame the ambiguity on some sort of complexity in the structure of the head noun (Larson calls these alternatives *A-analysis* and *N-analysis*, respectively).

In an *A-analysis*, an approach predominantly advocated by Siegel (1976), adjectives are assumed to belong to two classes, termed predicative and attributive, where predicative adjectives (e.g. *red*, *small*, etc.) are taken to be simple functions from entities to truth-values and are extensional, and thus intersective: $\|Adj\ Noun\| = \|Adj\| \cap \|Noun\|$. Attributive adjectives (e.g., *former*, *rightful*, etc.), on the other hand, are functions from common noun denotations to common noun denotations - i.e., they are predicate modifiers of type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$, and are thus intensional and non-intersective (but are subsective: $\|Adj\ Noun\| \subseteq \|Noun\|$). On this view, the ambiguity in (2) is explained by posting two distinct lexemes (*beautiful1* and *beautiful2*) for the adjective *beautiful*, one of which is an attributive while the other is a predicative adjective. In keeping with Montague's (1970) edict that similar syntactic categories must have the same semantic type, for this proposal to work, all adjectives are initially assigned the type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$ where intersective adjectives are considered to be subtypes obtained by triggering an appropriate meaning postulate. For example, assuming the lexeme *beautiful1* is marked as +INTERSECTIVE, the meaning postulate $\exists P \forall Q \forall x [Beautiful(Q)(x) \leftrightarrow P(x) \wedge Q(x)]$ would then yield an intersective meaning when *P* is *beautiful1*; and where a phrase such as 'a beautiful dancer' is interpreted as follows¹:

$$\|a\ beautiful1\ dancer\| \\ \Rightarrow \lambda P[(\exists x)(Dancer(x) \wedge Beautiful(x) \wedge P(x))]$$

$$\|a\ beautiful2\ dancer\| \\ \Rightarrow \lambda P[(\exists x)(Beautiful(Dancer(x)) \wedge P(x))]$$

While it does explain the ambiguity in (2), several reservations have been raised regarding this proposal. As Larson (1995; 1998) notes, however, this approach entails considerable duplication in the lexicon as this means that there are 'doublets' for all adjectives that can be ambiguous between an intersective and a non-intersective meaning. Another objection, raised by McNally and Boleda (2004), is that in an *A-analysis* there are no obvious ways of determining the context in which a certain adjective can be considered intersective. For example, they suggest that the most natural reading of

Look at Olga dance. She is beautiful

is the one where *beautiful* is describing Olga's dancing, although it does not modify any noun and is thus wrongly considered intersective by modifying Olga. While valid in other contexts, in our opinion this observation does not necessarily hold in this specific example since the resolution of 'she' must ultimately consider all entities in the discourse, including, presumably, the dancing activity that would be introduced by a Davidsonian representation of 'Look at Olga dance' (this issue is discussed further below). A more

¹ Note that as an alternative to meaning postulates that specialize intersective adjectives to $\langle e, t \rangle$, one can perform a type-lifting operation from $\langle e, t \rangle$ to $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$ (see Partee, 2007).

promising alternative to the *A-analysis* of the ambiguity in (2) has been proposed by Larson (1995, 1998), who suggests that 'beautiful' in (2) is a simple intersective adjective of type $\langle e, t \rangle$ and that the source of the ambiguity is due to a complexity in the structure of the head noun. Specifically, Larson suggests that a deverbal noun such as *dancer* should have a Davidsonian representation such as

$$(\forall x)(Dancer(x) \equiv (\exists e)(Dancing(e) \wedge Agent(e, x)))$$

That is, any *x* is a dancer iff *x* is the agent of some dancing activity (Larson's notation is slightly different). In this analysis, the ambiguity in (2) is attributed to an ambiguity in what 'beautiful' is modifying, in that it could be said of Olga or her dancing activity. That is, (2) is to be interpreted as follows:

$$\|Olga\ is\ a\ beautiful\ dancer\| \\ \Rightarrow (\exists e)(Dancing(e) \wedge Agent(e, Olga) \\ \wedge (Beautiful(e) \vee Beautiful(Olga)))$$

In our opinion, Larson's proposal is plausible on several grounds. First, in Larson's *N-analysis* there is no need for impromptu introduction of a considerable amount of lexical ambiguity. Second, and for reasons that are beyond the ambiguity of *beautiful* in (1), there is ample evidence that the structure of a deverbal noun such as 'dancer' must admit a reference to an abstract object, namely a dancing activity; as, for example, in the resolution of 'that' in (3).

(3) *Olga is an old dancer.*
She has been doing that for thirty years.

Furthermore, and in addition to a plausible explanation of the ambiguity in (2), Larson's proposal seems to provide a plausible explanation for why 'old' in (4a) seems to be ambiguous while the same is not true of 'elderly' in (4b): 'old' could be said of Olga or her teaching; while 'elderly' is not an adjective that is ordinarily said of objects that are of type activity

(4) a. *Olga is an old teacher.*
b. *Olga is an elderly teacher.*

With all its apparent appeal, however, Larson's proposal is still lacking. For one thing, and while it presupposes that some sort of type matching is what ultimately results in rejecting the subsective meaning of 'elderly' in (4b), the details of such processes are more involved than Larson's proposal suggests. For example, while it explains the ambiguity of 'beautiful' in (2), it is not quite clear how an *N-Analysis* can explain why 'beautiful' does not seem to admit a subsective meaning in (5).

(5) *Olga is a beautiful street dancer.*

In fact, 'beautiful' in (5) seems to be modifying Olga for the same reason the sentence in (6a) seems to be more natural than that in (6b).

(6) a. *Maria is a clever young girl.*
b. *Maria is a young clever girl.*

The sentences in (6) exemplify what is known in the literature as adjective ordering restrictions (AORs). However, despite numerous studies of AORs (e.g., see Wulff, 2003; Teodorescu, 2006), the slightly differing AORs that have been suggested in the literature have never been formally justified. What we hope to demonstrate below however is that the apparent ambiguity of some adjectives and adjective-ordering restrictions are both related to the nature of the ontological categories that these adjectives apply to in ordinary spoken language. While the general assumptions in Larson's (1995; 1998) N-Analysis seem to be valid, it will be demonstrated here that nominal modification seems to be more involved than has been suggested thus far. In particular, it seems that a proper semantics for nominal modification requires a much richer type system than currently employed in formal semantics. Before we proceed, therefore, in the next section we will briefly introduce a type system akin to that suggested by Sommers (1963); a system that forms the foundation of a semantics that is grounded in an ontology that reflects our commonsense view of reality.

3 Ontological Concepts as Types

We assume a Platonic universe that includes everything that can be spoken about in ordinary language, in a manner akin to that suggested by Hobbs (1985). However, in our formalism concepts belong to two distinct categories: (i) ontological concepts, such as *Animal*, *Substance*, *Entity*, *Artifact*, *Event*, *State*, etc., which are assumed to exist in a subsumption hierarchy, and where the fact that an object of type *Human* is ultimately an object of type *Entity* is expressed as $\text{Human} \prec \text{Entity}$; and (ii) logical concepts, which are the properties (that can be said) of and relations (that may hold) between ontological concepts. Since adjectives are our immediate concern, consider the following illustrating the difference between ontological and logical concepts:

- (7) a. *Dedicated*($x :: \text{Human}$)
 b. *Clever*($x :: \text{Animal}$)
 c. *Imminent*($x :: \text{Event}$)
 d. *Old*($x :: \text{Entity}$)
 e. *Beautiful*($x :: \text{Entity}$)

These predicates are supposed to reflect the fact that, in ordinary spoken language, *Dedicated* is a property that is ordinarily said of objects that must be of type *Human* (7a); that *Clever* could be said of objects of type *Animal* (7b); *Imminent* is a property that is said of objects that must be of type *Event* (7c), etc. In addition to logical and ontological concepts, there are also proper nouns, which are the names of objects; objects that could be of any type. A proper noun such as *sheba* is interpreted as $\|sheba\| \Rightarrow \lambda P[(\exists^1 x)(\text{Noo}(x :: \text{Thing}, sheba) \wedge P(x))]$ where $\text{Noo}(x :: \text{Thing}, s)$ holds between some x (which could be any thing), and some s if (the label) s is the name of x , and t is presumably the type of objects that P applies to (to simplify notation we often write $\|sheba\| \Rightarrow \lambda P[(\exists^1 sheba :: \text{Thing})(P(sheba))]$).

Consider now the following, where we have assumed that $\text{Thief}(x :: \text{Human})$, i.e., that *Thief* is a property that is

ordinarily said of objects that must be of type *Human*, and where $\text{Be}(x, y)$ is true when x and y are the same objects:

$$(8) \|sheba \text{ is a thief}\| \\ \Rightarrow (\exists^1 sheba :: \text{Thing})(\exists x)(\text{Thief}(x :: \text{Human}) \\ \wedge \text{Be}(x, sheba))$$

That is, there is a unique object named *sheba* (which could be any *Thing*) and some x such that x is a *Thief* (and must therefore be of type *Human*) and such that *sheba* is that x . Note now that *sheba* is associated with more than one type in a single scope, and this necessitates a type unification, where a type unification ($s \bullet t$) between two types s and t , and where $Q \in \{\exists, \forall\}$ is defined (for now) as follows:

$$(9) (s \prec t) \supset (Qx :: (s \bullet t))(P(x)) \equiv (Qx :: s)(P(x)) \\ (t \prec s) \supset (Qx :: (s \bullet t))(P(x)) \equiv (Qx :: t)(P(x)) \\ \neg(s \prec t) \wedge \neg(t \prec s) \\ \supset (Qx :: (s \bullet t))(P(x)) \equiv (Qx :: \perp)(P(x))$$

where $P(x :: \perp) = \perp$ and $(t \bullet \perp) = (\perp \bullet t) = \perp$. Since $\text{Human} \prec \text{Thing}$ the type unification required in (8) now proceeds as follows:

$$\|sheba \text{ is a thief}\| \\ \Rightarrow (\exists^1 sheba :: (\text{Human} \bullet \text{Thing}))(\exists x)(\text{Thief}(x) \wedge \text{Be}(sheba, x)) \\ \Rightarrow (\exists^1 sheba :: \text{Human})(\exists x)(\text{Thief}(x) \wedge \text{Be}(sheba, x))$$

Finally, and since $\text{Be}(sheba, x)$ we could replace x by *sheba* obtaining the following:

$$\|sheba \text{ is a thief}\| \\ \Rightarrow (\exists^1 sheba :: (\text{Human} \bullet \text{Thing}))(\exists x)(\text{Thief}(x) \wedge \text{Be}(sheba, x)) \\ \Rightarrow (\exists^1 sheba :: \text{Human}) \\ (\text{Thief}(sheba) \wedge \text{Be}(sheba, sheba)) \\ \Rightarrow (\exists^1 sheba :: \text{Human})(\text{Thief}(sheba) \wedge \text{True}) \\ \Rightarrow (\exists^1 sheba :: \text{Human})(\text{Thief}(sheba))$$

In the final analysis, therefore, 'Sheba is a thief' is interpreted as follows: there is a unique object named *sheba*, an object that must be of type *Human*, and such that *sheba* is a thief². Finally, note the clear distinction between ontological concepts (such as *Human*), which Cocchiarella (2001) calls first-intension concepts, and logical (or second-intension) concepts, such as $\text{Thief}(x)$. That is, what ontologically exist are objects of type *Human*, not thieves, and *Thief* is a mere property that we have come to use to talk of objects of type *Human*. Moreover, logical concepts such as *Thief* are assumed to be defined by virtue of some logical expression, such as $(\forall x :: \text{Human})(\text{Thief}(x) \equiv \phi)$ where the exact nature of ϕ might very well be susceptible to temporal, cultural, and other contextual factors, depending on what, at a certain point in time, a certain community considers an *Thief* to be.

² The removal of $\text{Be}(sheba, x)$ essentially means that the copular 'is' was in this case interpreted as the 'is' of identity. This was due to the fact that in this case a subsumption relation exists between the types of the relevant objects. In other contexts, such as 'Liz is aging', 'Sheba is angry', etc., where it seems that we are dealing with the 'is' of predication, removing $\text{Be}(sheba, x)$ involves introducing some implicit relation between the different types that do not unify (*Human/Process*, *Human/State*), essentially resulting in interpretations such as 'Liz is-going-through-the-process-of aging', and 'Sheba is-in-a-state-of anger'

What is of particular interest to us here is that logical concepts such as *Thief* (or *Dancer*, *Writer*, etc.), are defined by logical expressions that admit abstract objects such as activities, processes, states, etc., each of which could be the object of modification.

4 Types and Nominal Modification

In this section we use the type system described above and the notion of type unification to properly formalize the intuitions behind Larson's proposal for nominal modification. Subsequently we show that our formalism explains the relationship between intersective/non-intersective adjectives and adjective-ordering restrictions.

4.1 Formalizing Larson's Proposal

First we begin by showing that the apparent ambiguity of 'beautiful' in (2) is due to the fact that beautiful applies to a generic type that subsumes many others. Consider the following, where we assume $Beautiful(x :: Entity)$; that is, $Beautiful$ is a property that can be said of any $Entity$:

$$\begin{aligned} & \|Olga \text{ is a beautiful dancer}\| \\ \Rightarrow & (\exists e :: Activity)(\exists olga :: Human) \\ & (Dancing(e) \wedge Agent(e, olga) \\ & \wedge (Beautiful(e :: Entity) \vee Beautiful(olga :: Entity))) \end{aligned}$$

Note now that, in a single scope, e is considered to be an object of type $Activity$ as well as an object of type $Entity$, while $Olga$ is considered to be a $Human$ and an $Entity$. This, as discussed above, requires a pair of type unifications:

$$\begin{aligned} & \|Olga \text{ is a beautiful dancer}\| \\ \Rightarrow & (\exists e :: Activity)(\exists olga :: Human) \\ & (Dancing(e) \wedge Agent(e, olga :: Human) \\ & \wedge (Beautiful(e :: (Activity \bullet Entity)) \\ & \vee Beautiful(olga :: (Human \bullet Entity)))) \end{aligned}$$

Since $(Activity \prec Entity)$ and $(Human \prec Entity)$, all the type unifications succeed, resulting in the following:

$$\begin{aligned} & \|Olga \text{ is a beautiful dancer}\| \\ \Rightarrow & (\exists e :: Activity)(\exists olga :: Human) \\ & (Dancing(e) \wedge Agent(e, olga) \\ & \wedge (Beautiful(e) \vee Beautiful(olga))) \end{aligned}$$

In the final analysis 'Olga is a beautiful dancer' is interpreted as follows: $Olga$ is the agent of some dancing $Activity$, and either $Olga$ or her dancing is $Beautiful$ (or both, of course). However, consider now the following:

$$\begin{aligned} (10) \quad & \|Olga \text{ is an elderly teacher}\| \\ \Rightarrow & (\exists e :: Activity)(\exists olga :: Human) \\ & (Teaching(e) \wedge Agent(e, olga) \\ & \wedge (Elderly(e :: (Activity \bullet Human)) \\ & \vee Elderly(olga :: (Human \bullet Human)))) \end{aligned}$$

Note now that e is considered to be an object of type $Activity$ as well as an object of type $Human$. Since $(Activity \bullet Human) = \perp$ the type unification in this

case fails resulting in the following:

$$\begin{aligned} & \|Olga \text{ is an elderly teacher}\| \\ \Rightarrow & (\exists e :: Activity)(\exists olga :: Human) \\ & (Teaching(e) \wedge Agent(e, olga) \\ & \wedge (\perp \vee Elderly(olga :: Human))) \\ \Rightarrow & (\exists e :: Activity)(\exists olga :: Human) \\ & (Teaching(e) \wedge Agent(e, olga) \wedge Elderly(olga)) \end{aligned}$$

The exercise we have just gone through clearly shows that a classification of adjectives as intersective (extensional) and non-intersective (intensional) is not necessary. Instead, it seems that the embedding of an ontological type structure - a structure that reflects our commonsense view of reality and the way we talk about it in ordinary language - along with type unification can systematically provide a proper explanation for nominal modification without unnecessarily complicating our logical formalisms.

4.2 Adjective Ordering Restrictions

Consider again the logical concepts given in (7). Note that $Beautiful$ can be said of objects of type $Entity$, and thus it can be said of a Cat , a $Person$, a $City$, a $Movie$, a $Dance$, an $Island$, etc. Therefore, $Beautiful$ can be thought of as a polymorphic function that applies to objects at several levels and where the semantics of this function depend on the type of the object, as illustrated in figure 1 below. Thus, and although $Beautiful$ applies to objects of type $Entity$, in saying 'a beautiful car', for example, the meaning of $Beautiful$ that is accessed is that defined in the type $Physical$ (which could in principal be inherited from a supertype). Moreover, and as is well known in the theory of programming languages, one can always perform type casting upwards, but not downwards (e.g., one can always view a Car as just an $Entity$, but the converse is not true)³. Thus, assuming $Red(x :: Physical)$ and $Beautiful(x :: Entity)$; that is, assuming that 'red' can be said of $Physical$ objects and 'beautiful' can be said of any $Entity$, and using the notation $A1(A2(x :: t2) :: t1)$ to represent adjectives $A1$ and $A2$ that are assumed to apply to objects of type $t1$ and $t2$, respectively, then, for example, the type casting that will be required in (11a) is valid, while that in (11b) is not

$$\begin{aligned} (11) \quad & \text{a. } Beautiful(Red(x :: Physical) :: Entity) \\ & \text{b. } Red(Beautiful(x :: Entity) :: Physical) \end{aligned}$$

This, in fact, is precisely why 'Jon owns a beautiful red car' is more natural than 'Jon owns a red beautiful car'. In general, a sequence $A1(A2(x :: t2) :: t1)$ is a valid sequence iff $(t2 \prec t1)$. Note that this notion of type casting is independent from that of type unification, in that the unification does indeed succeed here in both cases in (11). However, before we perform type unification, the direction of the type casting must be valid. The importance of this interaction will become apparent below.

4.3 How an Ambiguous Adjective gets one Meaning

Let us explain the example in (5), where we argued that Larson's proposal cannot explain why 'beautiful', which is

³ Technically, the reason we can always cast up is that we can always ignore additional information. Casting down, which entails adding information, is however undecidable.

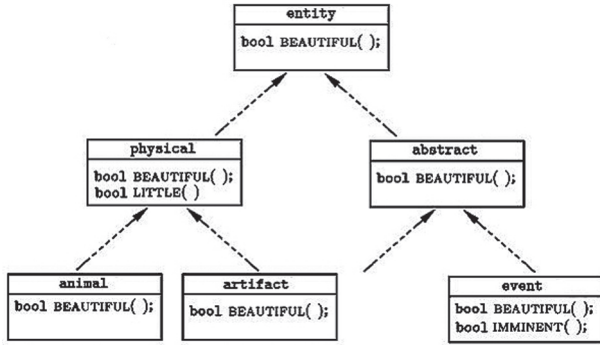


Figure 1: Adjectives as higher-order functions

considered to be ambiguous in (1), does not admit a subjective meaning in (5).

Consider the following:

$$\begin{aligned}
 (12) \quad & \|Olga \text{ is a beautiful young dancer}\| \\
 & \Rightarrow (\exists e :: \text{Activity})(\exists olga :: \text{Human}) \\
 & \quad (\text{Dancing}(e) \wedge \text{Agent}(e, olga :: \text{Human}) \\
 & \quad \wedge (\text{Beautiful}(\text{Young}(e :: \text{Activity}) :: \\
 & \quad \quad \text{Physical}) :: \text{Entity}) \\
 & \quad \vee (\text{Beautiful}(\text{Young}(olga :: \text{Human}) :: \\
 & \quad \quad \text{Physical}) :: \text{Entity}))
 \end{aligned}$$

Note now that the casting required is valid in both cases. In other words, the order of adjectives is valid. This means that we can now perform the required type unifications. We first note however that since $(\text{Activity} \bullet \text{Physical}) = \perp$, the term involving this type unification in (12) is reduced to \perp , and the term $(\beta \vee \perp)$ to β , hence:

$$\begin{aligned}
 (13) \quad & \|Olga \text{ is a beautiful young dancer}\| \\
 & \Rightarrow (\exists e :: \text{Activity})(\exists olga :: \text{Human}) \\
 & \quad (\text{Dancing}(e) \wedge \text{Agent}(e, olga :: \text{Human}) \\
 & \quad \wedge (\text{Beautiful}(\text{Young}(olga))))
 \end{aligned}$$

Note here that since beautiful was preceded by young, it could have not been applicable to an abstract object of type Activity, but was instead reduced to that defined at the level of Physical, and subsequently to that defined at the type Human. A valid question that comes to mind here is how then do we express the thought 'Olga is a young dancer and she dances beautifully'. The answer is that we usually make a statement such as this:

$$(14) \quad Olga \text{ is a young and beautiful dancer.}$$

In this case we are essentially overriding the sequential processing of the adjectives, and thus the adjective-ordering restrictions (or, the type-casting rules!) are no more applicable. That is, (14) is essentially equivalent to two sentences that are processed in parallel:

$$\|Olga \text{ is a young dancer}\| \wedge \|Olga \text{ is a beautiful dancer}\|$$

Note now that 'beautiful' would again have an intersective and a subjective meaning, although 'young' will only apply to Olga due to type constraints.

5 Concluding Remarks

If the main business of semantics is to explain how linguistic constructs relate to the world, then semantic analysis of natural language text is, indirectly, an attempt at uncovering the semiotic ontology of commonsense knowledge, and particularly the background knowledge that seems to be implicit in all that we say in our everyday discourse. While this intimate relationship between language and the world is generally accepted, semantics (in all its paradigms) has traditionally proceeded in one direction: by first stipulating an assumed set of ontological commitments followed by some machinery that is supposed to, somehow, model meanings in terms of that stipulated structure of reality. With the gross mismatch between the trivial ontological commitments of our semantic formalisms and the reality of the world these formalisms purport to represent, it is not surprising therefore that challenges in the semantics of natural language are rampant. However, as correctly observed by Hobbs (1985), semantics could become nearly trivial if it was grounded in an ontological structure that is "isomorphic to the way we talk about the world". The obvious question however is 'how does one arrive at this ontological structure that implicitly underlies all that we say in everyday discourse?' One plausible answer is the (seemingly circular) suggestion that the semantic analysis of natural language should itself be used to uncover this structure. In this regard we strongly agree with Dummett (1991):

We must not try to resolve the metaphysical questions first, and then construct a meaning-theory in light of the answers. We should investigate how our language actually functions, and how we can construct a workable systematic description of how it functions; the answers to those questions will then determine the answers to the metaphysical ones.

What this suggests, and correctly so, in our opinion, is that in our effort to understand the complex and intimate relationship between ordinary language and everyday commonsense knowledge, one could, as also suggested in (Bateman, 1995), "use language as a tool for uncovering the semiotic ontology of commonsense" since ordinary language is the best known theory we have of everyday knowledge. To avoid this seeming circularity (in wanting this ontological structure that would trivialize semantics; while at the same time suggesting that semantic analysis should itself be used as a guide to uncovering this ontological structure), we suggested here performing semantic analysis from the ground up, assuming a minimal (almost a trivial and basic) ontology, in the hope of building up the ontology as we go guided by the results of the semantic analysis. The advantages of this approach are: (i) the ontology thus constructed as a result of this process would not be invented, as is the case in most approaches to ontology (e.g., Lenat, and Guha (1990) and Sowa (1995), but would instead be discovered from what is in fact implicitly assumed in our use of language in everyday discourse; (ii) the semantics of several natural language phenomena should as a result become trivial, since the semantic analysis was itself the source of the underlying knowledge structures (in a sense, the semantics would have been done before we even started!) In this paper we have shown that nominal modification can be adequately treated in a semantics embedded in such a strongly-typed ontology; an ontol-

ogy that reflects our commonsense view of the world and the way we talk about it in ordinary language. While our concern in this paper was the semantics of [*Adj Noun*] nominals, our proposal seems to also provide an explanation for some well-known adjective-ordering restrictions.

6 Acknowledgements

I am grateful to the valuable feedback of the reviewers and participants of KI-2008, as well as that of two anonymous reviewers of this journal and those of Romeo Issa of PRAGMATECH.

References

- [1] Charniak, E. (1995), Natural Language Learning, *ACM Computing Surveys*, 27(3), pp. 317-319.
- [2] Bateman, J. A. (1995), On the Relationship between Ontology Construction and Natural Language: A Socio-Semiotic View, *International Journal of Human-Computer Studies*, 43, pp. 929-944.
- [3] Cocchiarella, N. B. (2001), Logic and Ontology, *Axiomathes*, 12, pp. 117-150.
- [4] Dummett, M. (1991), *Logical Basis of Metaphysics*, Duckworth, London.
- [5] Hobbs, J. (1985), Ontological Promiscuity, In *Proc. of the 23rd Annual Meeting of the Association for Computational Linguistics*, pp. 61-69, Chicago, Illinois, 1985.
- [6] Larson, R. (1995), Olga is a Beautiful Dancer, Presented at the Winter Meetings of the Linguistic Society of America, New Orleans.
- [7] Larson, R. (1998), Events and Modification in Nominals, In D. Strolovitch and A. Lawson (Eds.), *Proceedings from Semantics and Linguistic Theory (SALT) VIII*, pp. 145-168.
- [8] Lenat, D. B. and Guha, R.V. (1990), *Building Large Knowledge-Based Systems: Representation and Inference in the CYC Project*. Addison-Wesley.
- [9] McNally, L. and Boleda, G. (2004), Relational Adjectives as Properties of Kinds, In O. Bonami and P. Cabredo Hpfherr (Eds.), *Empirical Issues in Formal Syntax and Semantics*.
- [10] Montague, R. (1970), English as a Formal Language, In R. Thomason (Ed.), *Formal Philosophy - Selected Papers of Richard Montague*, New Haven, Yale University Press.
- [11] Ng, Tou H. and Zelle, J. (1997), Corpus-based Approaches to Semantic Interpretation in Natural language, *AI Magazine*, Winter 1997, pp. 45-64.
- [12] Partee, B. (2007), Compositionality and Coercion in Semantics - the Dynamics of Adjective Meanings, In G. Bouma et al (Eds.), *Cognitive Foundations of Interpretation*, Amsterdam: Royal Netherlands Academy of Arts and Sciences, pp. 145-161.
- [13] Saba, W. (2007), Language, logic and ontology: Uncovering the structure of commonsense knowledge, *International Journal of Human-Computer Studies*, 65(7), pp. 610-623.
- [14] Seigel, E. (1976), *Capturing the Adjective*, Ph.D. thesis, University of Massachusetts.
- [15] Sowa, J.F., 1995. *Knowledge Representation: Logical Philosophical and Computational Foundations*. PWS Publishing Company, Boston.
- [15] Sommers, F. (163), Types and Ontology, *Philosophical Review*, lxxii, pp. 327-363
- [16] Teodorescu, A. (2006), Adjective ordering restrictions revisited, In D. Baumer et. al. (Eds.), *Proc. of the 25th West Coast Conference on Formal Linguistics*, pp. 399-407.
- [17] Wulff, S. (2003), A Multifactorial Analysis of Adjective Order in English, *International Journal of Corpus Linguistics*, pp. 245-282.

Contact

Dr. Walid S. Saba
 PRAGMATECH/United Development Corp
 Doha, Qatar, P.O.Box 7256
 Phone: +974 44 63 444
 walid.saba@udcqtatar.com



Walid Saba currently the Chief Information Officer of Pragmatech, a technology company that specializes in human language technologies and business intelligence solutions, received his PhD in Computer Science from Carleton University in 1999. He spent over 16 years in industry where he was a principal software engineer at the American Institutes for Research, a knowledge engineer at Bell Labs, and a consulting software engineer at Metlife, Nortel Networks and Cognos, Inc. He has also taught Computer Science at the University of Windsor, the New Jersey Institute of Technology (NJIT), the American University of Beirut, and the American University of Technology. His research interests are in computational semantics, ontology and intelligent agents.