

# Towards Foundational Semantics

## *Ontological Semantics Revisited*

Philipp CIMIANO <sup>a,1</sup> and Uwe REYLE <sup>b</sup>

<sup>a</sup> *Institute AIFB, University of Karlsruhe*

<sup>b</sup> *IMS, University of Stuttgart*

**Abstract.** In line with Nirenburg and Raskin's paradigm of ontological semantics, we adhere to the basic tenet that natural language semantics needs to be captured with respect to an explicitly formalized ontology. Many researchers in computational semantics, however, have neglected the ontological aspects of meaning representation, and even more have neglected aspects of meaning representation related to domain-independent ontologies, i.e. foundational or upper-level ontologies. In this paper we argue for a stronger integration of foundational ontologies in computational semantics. We show that relying on foundational ontologies can, on the one hand, lead to a clean separation between domain-specific and domain-independent components of natural language processing systems. On the other hand, we show how the interplay between foundational, domain ontologies and lexical semantics resources can elegantly account for disambiguation as well as allow to draw non-trivial inferences. Further, a temporal theory compliant with the foundational ontology is absolutely necessary for supporting temporal reasoning in natural language understanding.

**Keywords.** NL semantics, ontologies, foundational ontologies, lexical resources

## 1. Introduction

In the computational linguistics community, on the one hand, huge manual efforts have been and are still being devoted to developing large lexical semantic resources such as WordNet<sup>2</sup>, FrameNet<sup>3</sup> or PropBank<sup>4</sup>. WordNet is in essence a lexical database linking words to their meanings, FrameNet basically provides case frames and their roles for situations and events occurring in the world, and the aim of PropBank is to provide argument structures for verbs, nouns etc. In the Semantic Web and Knowledge Engineering communities, on the other hand, a lot of effort has been spent on developing foundational [16] or general ontologies [12,19], domain ontologies<sup>5</sup> and ontology languages [21]. While the above mentioned lexical resources are widely used within natural language processing, neither ontologies nor their interplay with the above mentioned lexical resources have received much attention. Within computational semantics, for example,

---

<sup>1</sup>The first author acknowledges financial support from the BMBF project SmartWeb, funded by the German Ministry of Research, as well as the projects SEKT and X-Media funded by the European Union.

<sup>2</sup><http://wordnet.princeton.edu/>

<sup>3</sup><http://framenet.icsi.berkeley.edu/>

<sup>4</sup>[http://www.cis.upenn.edu/~mpalmer/project\\_pages/ACE.htm](http://www.cis.upenn.edu/~mpalmer/project_pages/ACE.htm)

<sup>5</sup>See for example the DAML ontology library at <http://www.daml.org/ontologies/>

a large body of work has addressed the construction of logical form (LF) from natural language input. However, aspects of meaning related to domain theories or ontologies have been neglected to a large extent. For the interpretation of the logical form, a logical theory or ontology axiomatizing the meaning of the symbols used is nevertheless crucial.

In line with Nirenburg and Raskin's *Ontological Semantics* framework, we thus adhere to the basic tenet that natural language semantics needs to be captured with respect to an explicitly formalized ontology. Further, we argue for a novel direction in computational semantics, i.e. what we will call *foundational semantics*. Foundational semantics differs from ontological semantics in that it is concerned with identifying that abstract meaning layer which remains constant across domains and applications. In this respect our approach differs crucially from the ontological semantics framework of Nirenburg and Raskin, who are not concerned with domain-independent aspects of meaning.

From a theoretical point of view, foundational semantics aims at identifying the core components of the domain-independent meaning layer as well as to clarify their interplay, thus contributing to the understanding of the principles of semantic construction. From a practical point of view, the commitment to the principles of foundational semantics is expected to have a clear impact on the engineering of natural language processing systems, allowing to modularize their design and foster their adaption to new domains by clearly separating domain-specific from domain-independent components. When using a foundational ontology, the meaning of (question) pronouns, prepositions, adverbs and other closed-class words can in fact be captured in a domain-independent manner, thus fostering the reuse of such a domain-independent lexicon across domains and applications. Talking about foundational semantics is thus in our view tantamount to talking about domain-independent meaning representation. The core ingredients of foundational semantics are thus, on the one hand, a foundational ontology allowing to express elementary things about the world, but also linguistic components such as a lexical ontology, linking language to the world (e.g. WordNet) as well as lexical semantic resources such as FrameNet or PropBank, providing case frames with their corresponding roles as well as subcategorization structures for verbs, adjectives, nouns etc.

In this paper we provide a first step towards clarifying how the different components of foundational semantics interact with each other, but also with domain-specific ontologies to construct a logical form which is interpretable with respect to the logical theories or ontologies in question. We focus in this paper in particular on the role that foundational ontologies can play in meaning construction and we show how the different resources interplay together for the purposes of lexical disambiguation and reasoning. The novelty of our paper lies exactly in its exploratory nature as it is, to our knowledge, the first paper devoted to exploring the relation between foundational ontologies and natural language semantics.

As we will need to get concrete, we need to use one specific foundational ontology. For pragmatic reasons, we will commit to the DOLCE foundational ontology. However, this choice does not reflect any ontological commitment from our side. With respect to what will be said in this paper, any foundational ontology can be reused as long as it is reasonably axiomatized. In what follows we give a brief overview of DOLCE, which will be necessary for the understanding of the remainder of this paper. Further, in Section 3 we discuss how the meaning of closed-class words can be specified with respect to the foundational ontology, and in Section 4 we show how the different resources interplay for the purposes of disambiguation and reasoning. Finally, in Section 5 we discuss the

importance of temporal reasoning for natural language understanding before discussing some related work and concluding.

## 2. Foundational Ontologies - DOLCE

Recently, there has been considerable research on foundational ontologies, especially in the context of the Semantic Web (compare [16]). One of the envisioned scenarios in the Semantic Web is that computer agents are able to understand content as well as to negotiate with other agents autonomously. A successful negotiation, however, presupposes that both parties agree on the meaning of the issues under consideration and to which they legally commit. Therefore, it is an absolute must that meaning is formalized in a reasonably unambiguous way. To address these needs, foundational ontologies have become interesting in the context of the Semantic Web initiative as their aim is to provide such a (reasonably) unambiguous axiomatization of meaning independently of a certain domain. Foundational ontologies are typically also called *general* or *upper level* ontologies. The crucial characteristics of a foundational ontology are (compare [16]): (i) strong axiomatization, (ii) explicit ontological commitment, and (iii) minimality. The first point, strong axiomatization, directly relates to the need for the unambiguous specification of meaning necessary for allowing a sound negotiation between agents. Strong axiomatization contrasts with many so called *light-weight* ontologies developed nowadays, mainly consisting of a taxonomy, thus leaving a lot of margin for interpretation of the concepts. Explicit ontological commitment means that the foundational ontology should make its basic design choices explicit. Such design choices typically reflect basic logico-philosophical choices related to the representation of time, space, modality etc. Finally, minimality means that a foundational ontology should commit to as few ontological choices as possible to allow for a wide use and applicability.

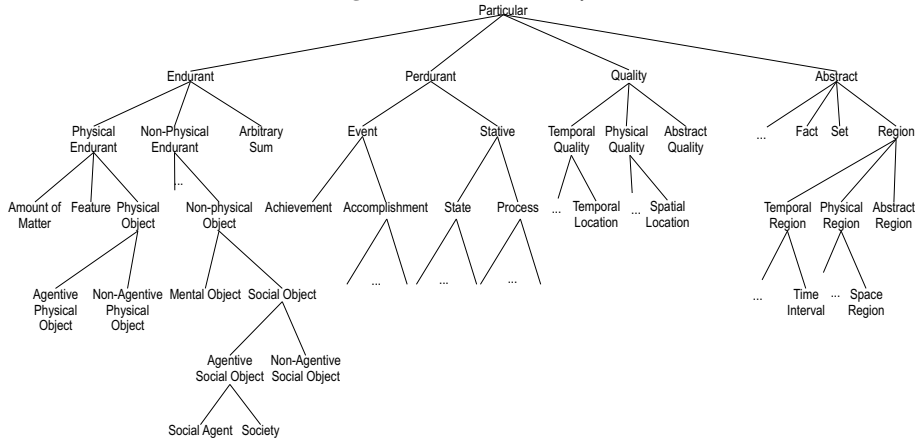
A more or less recent example for a foundational ontology is the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE), which has been developed by the Laboratory of Applied Ontology in the context of the WonderWeb project (compare [16]). As the name suggests, DOLCE has a strong cognitive bias in the sense that it does not aim at representing the world as it is with respect to logico-philosophical considerations, but as it is perceived by humans. It is certainly out of the scope of this paper to discuss the basic ontological commitments of DOLCE with respect to time and space. The interested reader should consult [16]. The basic class hierarchy adopted in DOLCE is depicted in Figure 1. A crucial distinction in DOLCE is the one between *perdurants* and *endurants*. Endurants are entities which exist in time (bound to a certain interval) and undergo change in shape, parts etc. Examples are persons, cars, theories, etc. Perdurants are entities which happen in time, e.g. events such as a party, a concert. All entities have *qualities* such as color, shape, size, etc. as well as concrete *quale*, i.e. values of these qualities at a certain time point. Qualities are related to their quales through the predicate  $ql(q, t)$ . In particular, DOLCE also distinguishes between spatial and temporal quale, i.e.  $ql_{SL}$  and  $ql_{TL}$ .

Further, DOLCE also provides fundamental relations between perdurants, such as temporal overlap, which is defined as follows:

$$x \circ_T y \leftrightarrow \exists t, t' (ql_T(t, x) \wedge ql_T(t', y) \wedge O(t, t'))$$

Overlap is defined in terms of the atomic predicate parthood, i.e.

**Figure 1. DOLCE Taxonomy**



$$O(x, y) \leftrightarrow \exists z(P(z, x) \wedge P(z, y))$$

Further, a very useful property is homogeneity, it is defined as follows in DOLCE:

$$HOM(\Phi) \leftrightarrow \Phi \sqsubseteq PD \wedge \Box \forall x, y(\Phi(x) \wedge P_T(y, x) \rightarrow \Phi(y))$$

i.e. a homogeneous property holds for all its temporal parts. Hereby,  $\Phi \sqsubseteq PD$  essentially states that  $\Phi$  is subsumed by  $PD$ <sup>6</sup>, i.e.  $\Phi$  is a kind of perdurant (see [16] for details). DOLCE also provides predicates for expressing temporal inclusion between entities which have temporal quaes (denoted by  $\subseteq_T$ ) and spatio-temporal inclusion between entities which have spatial quaes, i.e.,  $\subseteq_{ST}$ .

Finally, for the purposes of this paper we will assume a temporal order between temporal regions. However, we will not make assumptions about whether this order should be a partial or a total one. There exist different possibilities to axiomatize a temporal order  $\leq_T$  between temporal regions. However, we will not discuss any further the different possibilities for defining such a temporal order. A standard choice would for example be an interval-based temporal logic such as presented in [1] or [13].

### 3. Domain Independence

As argued in the introduction, a clean separation between domain-specific and domain independent meaning is very desirable to foster the reuse of a system across domains. In this section we discuss how the meanings of certain closed-class words with constant meaning across domains can be specified with respect to a foundational ontology. We discuss this using a question answering system as an example.

In question answering systems, wh-pronouns such as *which*, *what*, *where*, *who* or *when* have a constant meaning across domains. The same holds for temporal and locative

<sup>6</sup>In this paper we use the description logic notation  $\sqsubseteq$  to denote subsumption. Further we use the signs  $\sqcup$  and  $\sqcap$  to denote concept union and intersection, respectively. We assume this notation as an abbreviation of the corresponding first-order formulas given in [3].

prepositions such as *in*, *at*, *after*, etc. It would be thus desirable to capture the meaning of these words with respect to a foundational ontology such as DOLCE. To illustrate our proposal, let us consider the following example questions to a natural language interface:

- Who killed John F. Kennedy?
- Where was John F. Kennedy murdered?
- When was John F. Kennedy murdered?
- Who was murdered on November 22, 1963?
- Who was murdered in Dallas?
- Which american president was murdered after Kennedy was killed?

Here, the wh-pronouns and prepositions would have the meaning specified in Figure 2. There, APO stands for *Agentive Physical Object*, S for *Space Region* and TR for *Temporal Region*.

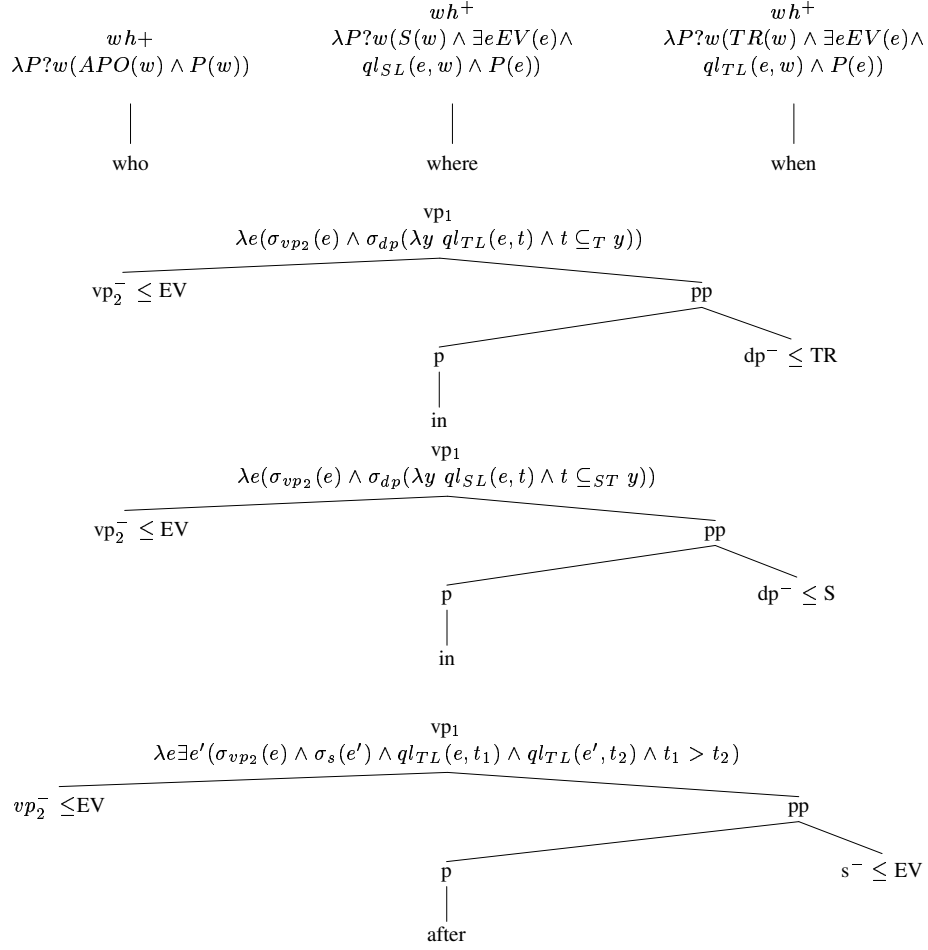
As syntactical backbone we build on Logical Description Grammars (LDG) [18], a lexicalized formalism inspired by Lexicalized Tree Adjoining Grammars (LTAG) [14], in which the basic syntactic units are so called elementary trees representing extended lexical projections of words which encapsulate logical arguments. Nodes in these elementary trees are marked positively or negatively and parsing boils down to identifying positively and negatively nodes of compatible syntactic categories, respecting precedence and dominance in the tree. Negatively marked nodes hereby typically correspond to argument positions which need to be filled with lexical content provided by positively marked nodes. The root node of the elementary trees for the wh-pronouns in Figure 2 are thus marked positively. The semantics is specified using the lambda calculus and constructed en par with the identification of nodes. The lambda expressions are thus composed with each other by means of functional application as specified by the elementary trees yielding an overall interpretation of a sentence (or question) as a result (see [5] for a detailed description of the use of the lambda calculus for semantic construction). In our notation, the lambda expression constituting the semantics of a node is given under it and refer to the semantics of other nodes below in the tree.

Semantically, wh-pronouns behave like a determiner in the sense that they typically combine with a property to yield a complete formula. Prepositions behave differently in the sense that they combine with a determiner phrase (*dp*) and a verb phrase (*vp<sub>2</sub>*) to yield a further verb phrase (*vp<sub>1</sub>*) the semantics of which is – in essence – the result of attaching the temporal or spatial condition imposed by the preposition to the event variable of *vp<sub>2</sub>*. Note here that the temporal and spatial conditions imposed by the preposition are specified with respect to DOLCE predicates. It is also important to mention that the different meanings of *in* (spatial vs. temporal) pose different constraints on the *dp*, i.e. they require a temporal region (TR) or spatial region (S), respectively. The extension of LDG allowing to pose type constraints on the nodes as well as the corresponding notation (specifying the exact types after the node with a colon ‘:’ or the subsuming type with ‘≤’) were already introduced in [7].

Here ‘?x’ is a question operator which specifies which variables are bound within the logical query. Our example questions would thus be interpreted as the following formal queries to a knowledge base:

1. ?x (APO(x) ∧ ∃e kill(e, x, JFK))
2. ?x (S(x) ∧ ∃e (ql<sub>SL</sub>(e, x) ∧ murdered(e, JFK) ∧ EV(e)))
3. ?x (TR(x) ∧ ∃e (ql<sub>TL</sub>(e, x) ∧ murdered(e, JFK) ∧ EV(e)))


**Figure 2.** Semantics of wh-pronouns and prepositions specified w.r.t. DOLCE



4.  $?x (APO(x) \wedge \exists e (murdered(e, x) \wedge ql_{TL}(e, t) \wedge t \subseteq_T t' \wedge month(t', 11) \wedge day(t', 22) \wedge year(t', 1963)))$
5.  $?x (APO(x) \wedge \exists e (murdered(e, x) \wedge ql_{SL}(e, t) \wedge t \subseteq_{ST} Dallas))$
6.  $?x (president(x, USA) \wedge \exists e, e' (murdered(e, x) \wedge ql_{TL}(e, t_1) \wedge killed(e', JFK) \wedge ql_{TL}(e', t_2) \wedge t_1 >_T t_2))$

We have thus shown how the meaning of wh-pronouns as well as spatial and temporal prepositions can be captured with respect to a foundational ontology and thus reused across domains.

Figure 3. Elementary trees for disambiguation of examination



#### 4. Interplay of Resources

Having shown how the meaning of closed-class words can be captured with respect to a foundational ontology, we turn to the issue of describing the interplay of different resources, i.e. foundational ontologies, domain ontologies, case frames and selectional restrictions specified with respect to the ontologies for the purpose of disambiguation. Further, we also show how this interplay can also yield non-trivial inferences as a result, provided that domain-specific knowledge is considered.

##### 4.1. Lexical Disambiguation

Selectional restrictions of verbs pose type constraints on their potential arguments and have thus a natural application in the disambiguation of the meaning of verbs as well as of their arguments. They are naturally expressed in terms of concept hierarchies, where the realm of relevant concepts ranges from domain-specific ones to those found in upper-level ontologies. In case the different meanings of a word correspond to different foundational categories, it even suffices to directly represent selectional restrictions with respect to categories such as provided by DOLCE.

Take, for example, a nominalization like *examination*, which is ambiguous between an event reading and a physical object reading. Combinations with verbal phrases or adjectives may disambiguate the noun depending on the concept the verbal phrase or adjective selects. An adjective like *brief* will identify the event reading and a verbal phrase like *being on the table* the physical object reading. The lexicon entry for the adjective *brief* and the representation of the verbal phrase *is on the table* would look as in Figure 3, where the ontological selectional restriction w.r.t. DOLCE on a node is given after the colon.

Using these entries, we could thus clearly distinguish between the event and object reading of *examination* in these contexts. In other cases, these distinctions are more subtle as is the case of the verb *to force* which has, for example, a *compel-* and a *break\_open-* reading. The latter one requires the object to be of type physical object (PO), and not ANIMATE.

Following Dowty ([8]) we assume that the participants of an event are given by thematic roles. Thematic roles are functions from perdurants to entities that are implicated in these perdurants. The thematic roles that we will consider in this paper are AGENT,<sup>7</sup> CAUSER, THEME, and INSTR. Their values are constrained by the following set of axioms, in which  $\theta$  ranges over thematic roles and the DOLCE participation relation  $PC$  states that  $x$  is involved in the occurrence of  $y$ .

<sup>7</sup>We use AGENT to include agents that are not necessarily capable of intentions.

$$\theta(e, x) \rightarrow PO(x) \wedge PD(e) \wedge PC(x, e)$$

Thus, thematic roles are specializations of the participation relation  $PC$  of DOLCE. We assume that the thematic roles are mutually exclusive (without stating the corresponding axioms explicitly here).

A further and more interesting set of axioms that involves thematic roles deals with ontological constraints that are determined for each thematic role by the type of perdurant the verb denotes, i.e. with the verb's selectional restrictions. Suppose that  $e$  is a perdurant denoted by a verb like *force* with its two meanings, *compel* and *break\_open*, then that meaning is selected in context for which the corresponding implication is fulfilled:

$$\begin{aligned} & \text{compel}(e) \wedge AGENT(e, x) \wedge THEME(e, y) \rightarrow \\ & \text{Human}(x) \wedge \text{Human}(y) \\ & \text{break\_open}(e) \wedge AGENT(e, x) \wedge THEME(e, y) \rightarrow \\ & \text{Human}(x) \wedge (PO \sqcap \neg ANIMATE)(y) \end{aligned}$$

The formulation of selectional restrictions on thematic roles thus leads to disambiguation of verbal meanings.

Prepositions are typically also ambiguous. *With*-pps are for example ambiguous between an instrumental reading as in (1.a), a co-agentive reading as in (1.b) and a noun-modifying reading. The corresponding elementary trees for the different readings of *with* we consider are shown in Figure 4. Now let us consider the following examples:

- (1) a. The doctor cured Peter with penicillin.
- b. The doctor cured Peter with the internist.

Let us assume that an instrumental reading in a cure event poses the constraint that the instrument in question is either an Amount of Matter (M), Light, Heat or some Process (PRO), i.e.

$$\text{cure}(e) \wedge INSTR(e, x) \wedge \neg \text{Human}(x) \rightarrow (M \sqcup \text{Light} \sqcup \text{Heat} \sqcup \text{PRO})(x)$$

This allows us to interpret the penicillin as the instrument of curing, but not the internist, which requires a co-agentive interpretation. This shows how the correct meaning of prepositions can be selected as a byproduct of fulfilling the logical conditions imposed by thematic roles.

#### 4.2. Inferencing

To see how verb meanings, logical conditions on their thematic roles and ontological knowledge interact with each other to yield non-trivial inferences, let us consider the sentence:

- (2) The doctor cured Peter with Belladonna.

World knowledge about Belladonna says that it contains the toxic substance Atropine. Further, Atropine leads to poisoning if ingested in a quantity of more than 3 mg for adults and 1 mg for children, i.e.



- (3)  $\forall x (Belladonna(x) \rightarrow M(x) \wedge \exists s \text{ contains}(x, s) \wedge Atropine(s))$   
 $\forall e, a, b (\text{ingest}(e, a, b) \wedge \exists y (\text{age}(a, y) \wedge y \geq 18) \wedge$   
 $\exists s, m (\text{contains}(b, s) \wedge Atropine(s) \wedge \text{mass}(s, m) \wedge m \geq 3mg)$
- (4)  $\rightarrow \text{poisoned}(Res(e), a)$   
 $\forall e, a, b (\text{ingested}(e, a, b) \wedge \exists y \text{ age}(a, y) \wedge y < 18 \wedge$   
 $\exists m, s (\text{contains}(b, s) \wedge Atropine(s) \wedge \text{mass}(s, m) \wedge m \geq 1mg)$
- (5)  $\rightarrow \text{poisoned}(Res(e), a)$

where  $Res(e)$  denotes the resultative state of an event  $e$ . The core of the meaning of *cure* is a change of state,  $e$ , of the organism or organ,  $y$ , from being affected by some disease,  $z$ , to the state  $s$  of not being affected. This event may be represented by means of the BEC(ome) operator as  $BEC(e, \lambda s \neg SUFFER\_FROM(s, y, z))$ , i.e.

- $$\text{cure}(e) \wedge AGENT(e, x) \wedge THEME(e, y) \wedge$$
- $$Human(x) \wedge (BodyPart \sqcup Organism)(x) \rightarrow$$
- $$\exists ec, e, e', z (CAUSE(ec, e', e) \wedge AGENT(e', x) \wedge$$
- (6)  $BEC(e, \lambda s \neg SUFFER\_FROM(s, y, z))$

The meaning of the  $BEC(e, \lambda s P(s))$  operator essentially is that the event  $e$  brings about a state  $s$  in which the condition  $P$  holds. For a more detailed description of the BEC(ome) operator, the interested reader is referred to Dowty [9].

If a substance is the instrument or agent of a curing event, then it is either ingested, inhaled, injected or applied to the skin, i.e.

$$\forall e, a, p, m \text{ cure}(e, a, p) \wedge INSTR(e, m) \wedge M(m) \rightarrow$$

$$\exists e' \subseteq_T e \wedge (\text{ingest}(e', p, m) \vee \text{inhale}(e', p, m) \vee$$

$$\text{injectedInto}(e', m, p) \vee \text{appliedToSkinOf}(e', m, p))$$

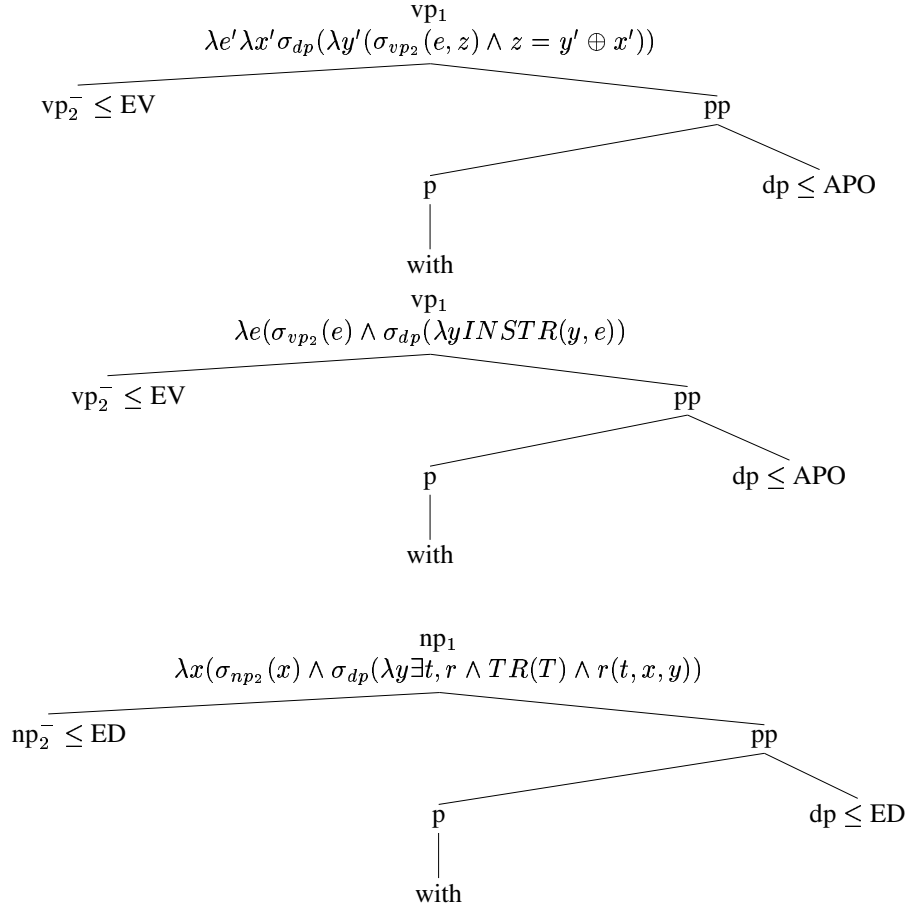
Further, assuming that being poisoned yields a contradiction with the resultative state of cure as well as assuming that Peter is an adult and ingested the Belladonna, we could derive that he was treated with a dose of Belladonna below 3mg, as otherwise he would have been poisoned. This shows how additional conditions on thematic roles and world knowledge can lead to non-trivial inferences.

## 5. Temporal Reasoning

To demonstrate how important temporal reasoning is for natural language understanding, let us discuss the following contrastive examples already discussed in [2], [11] and recently in [6]:

- a. John arrived at the oasis. The camels are standing under the palms.
- b. John arrived at the oasis. The camels were standing under the palms.

**Figure 4.** Three readings of with



While in the first case it is possible to interpret the camels as the means of transport by which John arrived, in the second discourse this interpretation is not possible due to the use of the imperfect *were standing* (compare [11,6]). The reason which is typically assumed is that the standing  $e'$  temporally overlaps with the arrival  $e$ , i.e.  $e \circ_T e'$  in terms of DOLCE, thus leading to a contradiction. It is definitely not the issue of this paper to explain what linguistic conditions lead exactly to pose that  $e$  and  $e'$  overlap (see [15] for a discussion of the temporal implications of imperfect vs. the simple past). Interesting for our purposes is the event structure as well as spatio-temporal consequences of *arrive* and *stand*. First of all, we will assume the following logical representation of sentence b):

$$\begin{aligned} & \exists e, e', j, o, c, p \text{ arrive\_at}(e) \wedge AGENT(e, j) \wedge john(j) \\ & THEME(e, o) \wedge event(e) \wedge oasis(o) \wedge camels(c) \wedge \\ & under(e', c, p) \wedge state(e') \wedge palms(p) \wedge e \circ_T e' \end{aligned}$$

Further, arriving implies a preparatory traveling phase which is part of the nucleus of arrive, and traveling implies a means of transport spatio-temporally correlated with the traveler:

$$\begin{aligned} &\forall e, a, o (arrive\_at(e) \wedge AGENT(e, a) \wedge THEME(e, o) \rightarrow \exists e' \\ &\quad (travel\_to(e') \wedge AGENT(e', a) \wedge THEME(e', o) \wedge e' \in prep(e))) \\ &\forall e, a (travel\_to(e) \wedge AGENT(e, a) \rightarrow \\ &\quad \exists m (MeansOfTransport(m) \wedge loc(e, m) = loc(e, a))) \end{aligned}$$

And *loc* is functional, i.e.  $\forall e, a, b, c (loc(e, a) = b \wedge loc(e, a) = c \rightarrow b = c)$ .

With respect to our example, it is thus the case that the following holds for the preparatory traveling phase  $e'$ :  $loc(e', j) \neq o$  and thus  $loc(e', c) \neq o$  (assuming that the camels are the mode of transport). Assuming that  $HOM(loc)$  this should be the case for any temporal part of the preparatory phase.

Overlapping an event also implies overlapping its preparatory phase:

$$\forall e, e', e'' (e \circ_T e' \wedge e'' \in prep(e') \rightarrow e \circ_T e'')$$

This means that the standing in which  $loc(e, c) = o$  holds overlaps with the preparatory phase in which  $loc(e', c) \neq o$  holds, yielding a logical inconsistency due to homogeneity and the functional definition of *loc*.

## 6. Related Work and Conclusion

We discuss in this section the work of Nirenburg and Raskin [17], Bateman [4], as well as Fillmore et al. [10].

The ontological semantics framework of Nirenburg and Raskin shares many aspects with our proposal of foundational semantics. First, both approaches share the commitment to an explicitly represented ontology. Second, Nirenburg and Raskin are also concerned with the specification of selectional restrictions for disambiguation purposes. However, they are not concerned with separating domain-specific from domain-independent meaning representation.

Bateman has also considered upper-level ontologies for natural language processing, in particular in the context of generation tasks. The Penmann Upper Model is in fact an upper-level ontology built on the basis of linguistic concepts. Defining concepts on a linguistic basis in fact eases the generation of natural language to express these concepts. It remains unclear, however, if the Penmann Upper Model is also suitable for natural language understanding purposes.

In the context of the FrameNet project, the aim of Fillmore et al. is to provide case frame semantics for verbs, specifying their core and non-core roles for application within text understanding [10]. However, FrameNet does not specify additional logical conditions which a frame element or slot needs to fulfill as in our approach. We have shown that specifying such conditions is necessary to rule out inconsistent readings as well as to support inferencing. Recently, Scheffczyk et al. have also discussed how to link FrameNet to existing general ontologies [20].

Summarizing, we have argued in this paper for the benefits of using a foundational ontology such as DOLCE for the purpose of capturing natural language semantics. We have in particular shown how foundational ontologies can (i) foster reusability of a system across domains, (ii) play an important role in disambiguation, (iii) provide a basis to draw non-trivial inferences as well as (iv) support temporal reasoning for NLP applications. From the perspective of our foundational semantics proposal, we have provided a first step towards clarifying its ingredients and examining their interplay.

## References

- [1] J. Allen and G. Ferguson. Actions and events in temporal logic. *Journal of Logic Computation*, 4(5):531–579, 1994.
- [2] N. Asher and A. Lascarides. Bridging. *Journal of Semantics*, 15, 1999.
- [3] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [4] John A. Bateman. Upper modeling: organizing knowledge for natural language processing. In *5th. International Workshop on Natural Language Generation, 3-6 June 1990*, 1990.
- [5] P. Blackburn and J. Bos. *Representation and Inference for Natural Language – A First Course in Computational Semantics*. CSLI Publications, 2005.
- [6] P. Cimiano. Ingredients of a first-order account of bridging. In *Proceedings of the 5th International Workshop on Inference in Computational Semantics (ICOS-5)*, 2006.
- [7] P. Cimiano and U. Reyle. Ontology-based semantic construction, underspecification and disambiguation. In *Proceedings of the Prospects and Advances in the Syntax-Semantic Interface Workshop*, pages 33–38, 2003.
- [8] D. Dowty. On the semantic content of the notion of "thematic role". In G. Chierchia, B. Partee, and R. Turner, editors, *Properties, Types, and Meanings*, volume 2, pages 69–129. Kluwer Academic Publishers, 1989.
- [9] D.R. Dowty. *Word Meaning and Montague Grammar*. Dordrecht, 1979.
- [10] C.J. Fillmore and C.F. Baker. Frame semantics for text understanding. In *Proceedings of the NAACL Workshop on WordNet and Other Lexical Resources*, 2001.
- [11] C. Gardent and K. Konrad. Interpreting definites using model generation. *Journal of Language and Computation*, 1(2):193–209, 2000.
- [12] R.V. Guha and D.B. Lenat. CYC: A midterm report. *AI Magazine*, 11(3):32–59, 1990.
- [13] J.R. Hobbs and F. Pan. An ontology of time for the semantic web. *ACM Transactions on Asian Language Information Processing (TALIP)*, 3(1):66–85, 2004.
- [14] A.K. Joshi and Y. Schabes. Tree-adjointing grammars. In *Handbook of Formal Languages*, volume 3, pages 69–124. Springer, 1997.
- [15] H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer, 1993.
- [16] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. Ontology library (final). WonderWeb deliverable D18.
- [17] M. McShane, S. Nirenburg, and S. Beale. An implemented, integrative approach to ontology-based NLP and interlingua. Technical report, Institute for Language and Information Technologies, University of Maryland, Baltimore County, March 2005.
- [18] Reinhard Muskens. Talking about trees and truth-conditions. *Journal of Logic, Language and Information*, 10(4):417–455, 2001.
- [19] I. Niles and A. Pease. Towards a standard upper ontology. In *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS)*, pages 17–19, 2001.
- [20] J. Scheffczyk, C.F. Baker, and S. Narayanan. Ontology-based reasoning about lexical resources. In *Proceedings of the OntoLex Workshop at the 5th International Conference on Lexical Resources and Evaluation (LREC)*, 2006.
- [21] S. Staab and R. Studer, editors. *Handbook on Ontologies*. International Handbooks on Information Systems. Springer, 2004.