

# Semantics

David Beaver and Joey Frazee

## 1 Introduction

Semantics is concerned with meaning: what meanings are, how meanings are assigned to words, phrases and sentences of natural and formal languages, and how meanings can be combined and used for inference and reasoning. The goal of this chapter is to introduce computational linguists and computer scientists to the tools, methods, and concepts required to work on natural language semantics.

Semantics, while often paired with pragmatics, is nominally distinct. On a traditional view, semantics concerns itself with the compositional buildup of meaning from the lexicon to the sentence level whereas pragmatics concerns the way in which contextual factors and speaker intentions affect meaning and inference (see, e.g., Potts *pear* in this volume). Although the semantics-pragmatics distinction is historically important, and continues to be widely adopted, in practice it is not clearcut. Work in semantics inevitably involves pragmatics and *vice versa*. Furthermore, it is not a distinction which is of much relevance for applications in computational linguistics.

This chapter is organized as follows. In sections 2 and 2.2 we introduce foundational concepts and discuss ways of representing the meaning of sentences, and of combining the meaning of smaller expressions to produce sentential meanings. In section 3 we discuss the representation of meaning for larger units, especially with respect to anaphora, and introduce two formal theories that go beyond sentence meaning: Discourse Representation Theory and Dynamic Semantics. Then, in section 4 we discuss temporality, introducing event semantics and describing standard approaches to the semantics of time. Section 5 concerns the tension between the surface-oriented statistical methods characteristic of mainstream computational linguistics and the more abstract methods typical of formal semantics and includes discussion of phenomena for which it seems particularly important to utilize insights from formal semantics. Throughout the chapter we assume familiarity with basic notions from logic (propositional logic and first-order predicate logic), computer science and, to a lesser extent, computational linguistics (e.g., algorithms, parsing, syntactic categories and tree representations).

## 2 Sentential semantics

### 2.1 Representation & logical form

Aristotle and the medieval post-Aristotelian tradition apart, work on formal semantic representation only began in earnest with Boole's (1854) semantics for propositional logic and Frege's (1879) development of first-order predicate logic (FOPL). Frege's work provided a precise and intuitive way of characterizing the meaning of sentences.<sup>1</sup> The influence of this development has been vast,

as seen in the fact that for many years, introductory courses on logic and semantics have commonly included exercises in translating sentences of natural language into statements of propositional or first-order logic. Thus, for example, (1a) might be represented as (1b):

- (1) a. Fischer played a Russian.
- b.  $\exists x (\text{russian}(x) \wedge \text{played}(\text{Fischer}, x))$

Subtle typographical distinctions we have made between (1a) and (1b) relate to a crucial theme in the development of semantics, made completely explicit in the work of Tarski (1944). This is the distinction between object language and metalanguage. For example, the sentence in (1a) and all the words that make it up are expressions of the object language, the language from which we are translating. On the other hand, the translations are given in a semantic metalanguage. Thus far, the semantic metalanguage has been FOPL. We use a *sans serif* font to mark expressions of the metalanguage, so that, for exMple, **played** is our way of representing in the semantic metalanguage the meaning of the English word *played*.<sup>2</sup>

Representations like that in (1b), so-called *Logical Forms* (LFs), immediately provide a way of approaching a range of computational tasks. For example, consider a database application, say a database of information about chess tournaments. While it is far from obvious how to query a database with an arbitrary sentence of English, the problem of checking whether a sentence of FOPL is satisfied by a record in that database is straightforward: (1) translate the English sentence into a sentence of FOPL, and (2) verify the sentence's translation. Thus we can check whether (1a) holds in the database by breaking it down into these sub-problems.

We can also consider whether one sentence of natural language follows from another using a similar procedure. The computational process for deciding whether (2a) follows from (1a) could involve translating both sentences into FOPL, (1b,2b), and then verifying the inference with an automated theorem prover.

- (2) a. Fischer played someone.
- b.  $\exists x (\text{played}(\text{Fischer}, x))$

Here let us note immediately that while FOPL offers us a strategy for dealing with such problems, it is far from a complete solution. In the case of the above examples, we have reduced one problem (natural language inference) to two problems (translation into FOPL, and inference), neither of which itself is trivial. The simplicity of the sentences in examples (1) and (2) gives off the appearance that translation into FOPL is easy, or uninteresting, but sentences often have multiple LFs as well as LFs that are not necessarily intuitive. This makes the particulars of the choice of an LF representation critical, as noted by Russell (1905) a century ago.

As regards the unintuitiveness of LFs, Russell argued that a definite description like *The American* in (3a) does not simply denote or refer to an individual (e.g., Bobby Fischer); rather, it is a quantificational element. So for Russell (3a) ought to be represented as in (3b).

- (3) a. The American won.
- b.  $\exists x (\text{american}(x) \wedge \forall y (\text{american}(y) \rightarrow x = y) \wedge \text{won}(x))$

Instead of *The American* being *Fischer* (in the right context), it's a quantifier that imposes the existence of an individual that (1) is, uniquely, American and (2) has the property denoted by whatever predicate it combines with (i.e., **won** in the above example). The question of whether this LF is a

good representation for definite descriptions is a subject of debate, with many (from Strawson 1950 on) arguing that the full meaning of sentences like (3a) cannot be captured in classical FOPL at all.<sup>3</sup>

Whatever position is taken about the meaning of definite descriptions, it is hard to escape Russell’s conclusion that surface forms bear a complex relationship to their LFs. This is especially evident when we look at sentences with multiple LFs. For example, a sentence like (4a) may exhibit a *scope ambiguity* whereby it may either have an LF like (4b) where the universal quantifier takes *wide* scope or an LF like (4c) where it takes *narrow* scope.

- (4) a. Every Russian played an American.  
b.  $\forall x (\text{russian}(x) \rightarrow \exists y (\text{american}(y) \wedge \text{play}(x,y)))$   
c.  $\exists y (\text{american}(y) \wedge \forall x (\text{russian}(x) \rightarrow \text{play}(x,y)))$

The wide scope LF corresponds to the reading where every Russian played at least one American, but possibly different ones. The narrow scope LF corresponds to the reading where every Russian played the same American. The LFs show that this distinction can be represented in FOPL, but how do we get from the sentence (4a) to its LFs (4b) and (4c)?

## 2.2 Compositional semantics

Above we have sketched a common way of representing the meanings of sentences, via translation to FOPL, but we have provided no general method for deriving these translations. This raises two questions: (1) how should we represent the meanings of smaller expressions (e.g., verbs like *played* and noun-phrases like *Every Russian*), and (2) how should the meanings of smaller expressions combine to yield sentence representations?

In the examples so far, some parts of the original sentences have direct counterparts in LF, others do not. For example, while *Russian* in (4a) has `russian` as its translation, the expression *Every Russian* has no counterpart in either of the sentence’s LFs (4b,4c). It’s tempting to say that the translation of *Every Russian* is “ $\forall x (\text{russian}(x) \rightarrow$ ”, but this is not a well-formed expression of FOPL and so it has no meaning in the usual semantics of FOPL. More troubling, though, is that FOPL on its own does not provide a method for deriving the meaning of an expression like “ $\forall x (\text{russian}(x) \rightarrow$ ” from the meaning of its parts, *Every* and *Russian*.

A method for assigning meanings to all syntactic units that make up a sentence and deriving the meanings of that sentence, algorithmically, from those parts, however, is available in Richard Montague’s (1973; 1970a; 1970b) seminal papers. Montague contended that there is no substantive difference between natural languages and formal languages (e.g., languages of philosophical logic such as FOPL, and programming languages) and that both can be analyzed in the same way (Montague 1970b,b; see also Halvorsen and Ladusaw 1979). As unlikely as this seems (indeed, Frege, Russell and others were drawn to FOPL thinking that natural language was too messy to be analyzable in precise terms), Montague (1970b,b) showed that significant fragments of English are *directly* interpretable in a precise and systematic fashion. Translating object an language into a formal language for interpretation (see, e.g., Montague 1973) is usually more straightforward, however, and we will follow this strategy throughout the paper.

Indirect interpretation is primarily achieved using the *Lambda Calculus* (Church, 1932).<sup>4</sup> Object language terms are mapped to typed lambda terms which are subsequently used to assign meanings. For example, *Russian* is paired with the function `russian`, which maps Boris Spassky to 1 or true and Bobby Fischer and Deep Blue to 0 or false. The type of this function is  $\langle e, t \rangle$ , the type of properties,

which maps entities  $e$  to truth-values  $t$ . Quantifiers like *every* are likewise analyzed as functions, but the corresponding function **every** is a mapping not from entities but from a pair of properties like **russian** and **won** to truth-values. Its type is  $\langle\langle e,t\rangle, \langle\langle e,t\rangle, t\rangle\rangle$ , a function from properties  $\langle e,t\rangle$  to other properties  $\langle e,t\rangle$  to truth-values  $t$ .

The interpretation algorithm can then proceed compositionally from two rules: (1) functional application, the combination of functions with their arguments, and (2) Beta-reduction, lambda evaluation via substitution (see Figure 1). So the meaning of an expression consisting of subexpressions  $\beta$  and  $\gamma$  with commensurate types will be calculated as the meaning of  $\beta$  applied to the meaning of  $\gamma$ , which we write  $\llbracket\textit{beta}\rrbracket(\llbracket\textit{gamma}\rrbracket)$ . We get the meaning of *Every Russian* by applying **every** to **russian** and we get the meaning of *Every Russian won* by applying **every(russian)** to **won**. See Figures 2 and 3.

---

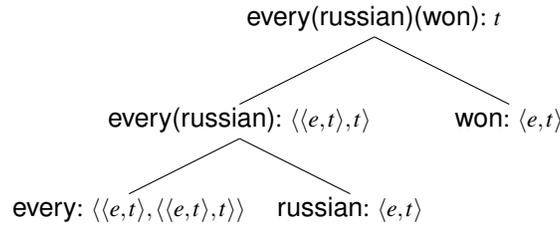
**Beta-reduction:** applying a function  $\lambda x.\alpha$  to a term  $y$  — $\lambda x.\alpha(y)$ —when  $y$  does not contain any free occurrences of and is of the same type as  $x$  —results in or “reduces” to the term  $\alpha$  except that every instance of  $x$  in  $\alpha$  has been replaced by  $y$ .

**Examples:**

- $\lambda x.P(x)(a) \rightsquigarrow P(a)$
- $\lambda R.\exists x (R(x) \wedge Q(x))(P) \rightsquigarrow \exists x (P(x) \wedge Q(x))$
- $\lambda P\lambda Q.\forall x (P(x) \rightarrow Q(x))(R)(W) \rightsquigarrow \forall x (R(x) \rightarrow W(x))$

**Figure 1:** Beta-reduction rule of the Lambda Calculus and examples. Beta-reduction drives semantic composition by functional application in the Montagovian model.

---



**Figure 2:** Parse tree and semantic types for *Every Russian won*. Each node of the tree is labeled with an expression of the meta-language and its corresponding semantic type.

---

---


$$\begin{aligned} \llbracket\textit{Every}\rrbracket &= \lambda P\lambda Q.\forall x (P(x) \rightarrow Q(x)) \\ \llbracket\textit{Russian}\rrbracket &= \lambda x.\textit{russian}(x) \\ \llbracket\textit{won}\rrbracket &= \lambda x.\textit{won}(x) \\ \llbracket\textit{Every Russian}\rrbracket &= \llbracket\textit{Every}\rrbracket(\llbracket\textit{Russian}\rrbracket) = \lambda Q.\forall x (\textit{russian}(x) \rightarrow Q(x)) \\ \llbracket\textit{Every Russian won}\rrbracket &= \llbracket\textit{Every Russian}\rrbracket(\llbracket\textit{won}\rrbracket) = \forall x (\textit{russian}(x) \rightarrow \textit{won}(x)) \end{aligned}$$

**Figure 3:** Derivation of *Every Russian won*.

---

Figure 2 illustrates the algorithm. Expressions combine with other expressions of compatible types,

and as they combine they produce larger expressions which combine with other type-compatible expressions. The algorithm continues, recursively, until it produces an expression of type  $t$ ; that is, a sentence. For example, the quantifier **every** with type  $\langle\langle e,t\rangle,\langle\langle e,t\rangle,t\rangle\rangle$  combines with an expression of type  $\langle e,t\rangle$  (i.e., the type-system allows *Every* to combine with properties like *Russian* but not names or descriptions like *Fischer* or *The American*), and on doing so, the resulting expression again combines with a type  $\langle e,t\rangle$  predicate (e.g., the verb *won*) yielding a sentence. Each step is an instance of functional application and Beta-reduction. Figure 3 shows the derivation. The meaning of *Every*,  $\lambda P\lambda Q.\forall x (P(x) \rightarrow Q(x))$ , applies to the meaning of *Russian*, *russian*, giving  $\lambda Q.\forall x (\text{russian}(x) \rightarrow Q(x))$ . This subsequently applies to the meaning of *won*, *won*, giving the expected LF.

This method is significant linguistically and computationally because it shows that there is an algorithmic way of relating expressions of natural language with meanings, showing that language is not too messy to allow principled semantic analysis using tools like FOPL.

### 3 Discourse Semantics

The relationship between sentence form and LF becomes especially problematic when considering discourse-level expressions, *anaphoric expressions* like pronouns that connect back to earlier sentence elements.

#### 3.1 Anaphoric expressions

A standard view of anaphoric expressions is that they are like bound variables in logic. For example, translating *Every cat chased its tail* as  $\forall x \text{cat}(x) \rightarrow \text{chased}(x, \text{tail-of}(x))$ , *its* is translated as the variable  $x$ . A naive application of that translation strategy, however, goes awry.

The best known problematic examples are Geach’s (1962) termed ‘donkey sentences’. In a classic example, given in (5), the pronouns *he* and *it* refer back to the earlier introduced farmer and donkey:

(5) If a farmer owns a donkey then he beats it.

But, the naive translation, (6a), is clearly not the desired meaning. The problem is that  $x$  and  $y$  are unbound in their final occurrence. The binding failure in (6a) occurs because the implication operator outscopes the existential operators. So it might seem that the problem could be solved by giving the existentials wider scope. However, this results in another incorrect translation, (6b). The problem with (6b) is that it is true if there is a farmer and a donkey that he doesn’t own, regardless of whether any farmers beat their donkeys. In order to represent the meaning of (5) an LF like (6c) is needed. But in (5), the indefinite NPs use universal quantifiers. This is unsatisfying because, without a general principle telling us when to translate indefinites as existentials, and when to translate them as universals, we do not have a deterministic translation procedure.

- (6) a.  $(\exists x, y \text{farmer}(x) \wedge \text{donkey}(y) \wedge \text{owns}(x, y)) \rightarrow \text{beats}(x, y)$   
 b.  $\exists x, y \text{farmer}(x) \wedge \text{donkey}(y) \wedge (\text{owns}(x, y) \rightarrow \text{beats}(x, y))$   
 c.  $\forall x, y (\text{farmer}(x) \wedge \text{donkey}(y) \wedge \text{owns}(x, y) \rightarrow \text{beats}(x, y))$

### 3.2 Discourse Representation Theory

Kamp (1981) and Heim (1982) saw the problem of donkey pronouns as part of a broader issue with discourse anaphora. Once again, the problem can be cast in terms of scope. In (7a), *she* refers back to *A Hungarian*, but if we translate sentence sequencing using logical conjunction, then a direct translation of (7a) as (7b) leaves  $x$  unbound.

- (7) a. A Hungarian won. She was a prodigy.  
 b.  $\exists x (\text{hungarian}(x) \wedge \text{won}(x)) \wedge \text{prodigy}(x)$

This led Kamp and Heim (building on earlier insights of Fauconnier (1975); Karttunen (1969); Stalnaker (1972) and others) to initiate a radical shift in the focus of work in semantics. On the new view, attention was refocused from sentence level to discourse level. The goal of semantic theory was no longer to merely provide a static representation of linguistic expressions but also to account for the dynamic effect of language, the information conveyed, and the resulting mental representations of interlocutors. The idea that led to the resolution of the problematic cases introduced above is that indefinite NPs are not existential quantifiers but rather expressions that create or introduce *discourse referents*, mental representations of the entities under discussion. For example, in the mind of the hearer, *A Hungarian* prompts the creation of a discourse reference with the property of being Hungarian. The meaning of an indefinite, then, is understood as being fundamentally dynamic, and pronouns are seen as naming those references. So for example, in (7a), *she* is interpreted as the name of the discourse referent introduced by *A Hungarian*.

Kamp and Heim’s proposals are similar so we will focus here on Kamp’s presentation, *Discourse Representation Theory* (DRT), which is standardized in Kamp and Reyle 1993 and implemented in the wide-coverage semantic parser Boxer Bos 2008; Blackburn and Bos 2005, 2000.<sup>5</sup> Kamp’s DRT departs from Montague Grammar as regards both the meaning representation language the way that meaning representations are constructed. A DRT meaning representation is a *Discourse Representation Structure* (DRS), which consists of a set *discourse referents* and a set of conditions on discourse referents. DRSs are commonly presented using a two-dimensional “box notation” with the discourse referents on top, and the conditions below, such that a DRS for the first sentence of (7a) is as in (8), with one discourse referent on top, and two conditions below. The same structure may be given a more compact linear representation as  $[x \mid \text{hungarian}(x), \text{won}(x)]$ .

(8)

x
hungarian(x) won(x)

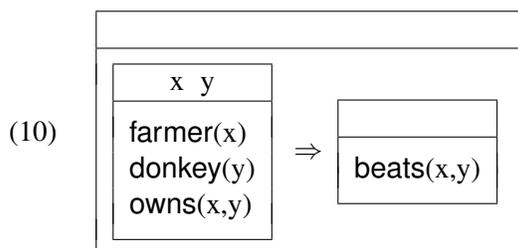
As discourse unfolds, information from successive sentences is added to the DRS, so that (9) is the full representation of (7a).

(9)

x
hungarian(x) won(x) prodigy(x)

The semantics of simple DRSs, whether in two-dimensional or linear form, is straightforward, essentially the same as the corresponding FOPL with all discourse referents existentially quantified,

and conjunctions between conditions; for example  $\exists x (\text{hungarian}(x) \wedge \text{won}(x) \wedge \text{prodigy}(x))$  in (9). However, the semantics of the language departs from FOPL for conditionals and quantifiers, and it's this that solves the problem of binding in donkey sentences. In (5) the conditional introduces a duplex condition, that involves two new boxes, each with what we might term an “attic” space for extra discourse referents:



It should be noted that in (5), the discourse referents associated with the NPs *a farmer* and *a donkey* (i.e.,  $x$  and  $y$ ) are introduced in the attic space of the sub-DRS where the DRS conditions for the NPs (i.e.,  $\text{farmer}(x)$  and  $\text{donkey}(y)$ ) are found.

Crucially, the semantics of implication in DRT is *dynamic*. To understand this, we need first the idea that the meaning of an expression is defined relative to a context, and second the idea that different subparts of a complex expression may be evaluated in different contexts, where changes in context are understood as resulting from the effects of evaluating other expressions. Thus an indefinite NP like *A farmer* has a dynamic effect, since a later pronoun like *he* will be evaluated in a context which provides a discourse referent for a farmer. In the semantics of the DRS language, context change is mediated by assignment functions and connectives are given a dynamic semantics in the sense that their arguments are evaluated with respect to different assignment functions. Implicational conditions are also given such a dynamic semantics.

For example if we evaluate (10) relative to a model  $M$  and assignment function  $f$ , we first find all the assignment functions which potentially differ from  $f$  by mapping the referent  $x$  onto a farmer in the model, and  $y$  onto a donkey which is owned by the farmer in the model. Then we check that for *any* such assignment, the righthand condition also holds. Instead of evaluating the righthand box (in linear form:  $[ \mid \text{beats}(x,y) ]$ ) relative to  $M$  and  $f$ , we evaluate it relative to  $M$  and  $g$ , for all the different assignments  $g$  that satisfy the lefthand box.<sup>6</sup>

By defining the semantics of conditionals so the righthand box is evaluated relative to all contexts which satisfy the lefthand box, we get the same effect as universal quantification. The DRS (10) is satisfied if in every farmer-donkey ownership pair, the farmer beats the donkey; (10) is truth-conditionally equivalent to the FOPL formula in (6c). This reveals a notable characteristic of Kamp's DRT (and equally of Heim's account): whereas indefinites are traditionally translated as existentials, indefinites in DRT lack quantificational force. Instead, indefinites introduces discourse referents, and quantificational force is determined by the position the referents occupy in the DRS.

### 3.3 Dynamic Semantics

Kamp's DRT exhibits two sorts of dynamic behavior: first, representations are built and extended dynamically, and second, those representations are interpreted dynamically in the sense that sub-DRSs in implications and other operators are evaluated relative to dynamically updated assignment functions. In the 1980s it was unclear to many semanticists whether the representation-building

aspect of the dynamics was necessary, especially since the system proposed by Heim 1982 is much less dependent on the specific representations she used than is Kamp’s DRT. One important difference concerns conjunction. Let us start with what the two proposals have in common, namely the following intuition: a context is updated with a conjunction of two formulae by first updating with the first, and then updating with the second. This intuition can be seen as the defining generalization of what came to be known as *Dynamic Semantics*.

To understand Heim’s proposal, we can utilize the concept of a possible world, which is analogous to a first-order model, a total description of the way the things might be. It is helpful to make a short detour to an earlier proposal given by Stalnaker (1972, 1979, 2002) as an account of assertion. Stalnaker’s idea was that as conversational participants talk to each other, they narrow in on the way (they jointly think) the world is. Thus conversational *common ground* can be modeled as the set of worlds compatible with what the participants have agreed on. So, context is a set of worlds, and assertion of a proposition is a process that involves taking the intersection of that context set with the set of worlds in which the proposition is true. The conjunction of propositions, discourse, is the same as asserting the first, then the second, and so on.

A Heimian context is more complex than a Stalnakerian context. Where Stalnaker uses a set of worlds, Heim uses a set of pairs of worlds and (partial) assignment functions. Context is still a set of alternatives, but each records not only what the world is like, but also what the discourse referents are. For Stalnaker, successive update was a fact about the way conversation works, but for Heim, and likewise Kamp, it is a fact about the meaning of conjunction as an instruction to update context with the left and then the right conjunct. Having seen the crucial difference between Heim and Stalnaker, we can now also see how Heim differs from Kamp: for Kamp, the context a conjunction operates on is a DRS, while Heim offers the idea of a conjunction operating on a set of world-assignment pairs.

The system proposed by Heim 1982, *File Change Semantics*, has much in common with later dynamic semantic proposals, such as *Dynamic Predicate Logic* (DPL; Groenendijk and Stokhof 1991a), Update Semantics (Veltman, 1996) and Compositional DRT (Muskens, 1996), all of which can be seen as extensions of foundational work in the logic of programs (Hoare, 1969; Pratt, 1976).

Groenendijk and Stokhof’s (1991a) reinterpretation of the language of FOPL makes the point of dynamic semantics particularly clear. DPL provides a semantics which has exactly the right logical properties such that the earlier examples of discourse anaphora like (11a) (repeated from 7a) and donkey anaphora like (12a) (repeated from 5) can be translated in the most obvious way, as in (11b) and (12b).

- (11) a. A Hungarian won. She was a prodigy.  
 b.  $\exists x (\text{hungarian}(x) \wedge \text{won}(x)) \wedge \text{prodigy}(x)$
- (12) a. If a farmer owns a donkey then he beats it.  
 b.  $(\exists x, y \text{ farmer}(x) \wedge \text{donkey}(y) \wedge \text{owns}(x, y)) \rightarrow \text{beats}(x, y)$

The logical properties needed are *Discourse Equivalence* and *Donkey equivalence* (13). These properties allow existential quantifiers to take non-standard scope, binding variables across both

conjunctions and implications.

- (13) a. *Discourse Equivalence*:  
 $(\exists x\varphi) \wedge \psi \equiv \exists x(\varphi \wedge \psi)$   
 b. *Donkey equivalence*:  
 $(\exists x\varphi) \rightarrow \psi \equiv \forall x(\varphi \rightarrow \psi)$

Given the semantics for implication in DRT, it is easy to see how these non-standard scope effects are achieved. In DPL the work of the existential quantifier is disentangled from implication and done instead in the semantics of conjunction.

A dynamic semantics for the language of FOPL is given in (14), to be interpreted along with standard equivalences for implication and universal quantification,  $\varphi \rightarrow \psi \equiv_{\text{def}} \neg(\varphi \wedge \neg\psi)$ , and  $\forall x \varphi \equiv_{\text{def}} \neg\exists x \neg\varphi$ . In the semantics in (14) we have simplified relative to Heim's system by using sets of assignments instead of sets of assignment-world pairs. The context is represented as  $\sigma$ , and the update of  $\sigma$  with a formula  $\varphi$  is written  $\sigma[\varphi]$ . We can think of  $\sigma$  as an input context, and  $\sigma[\varphi]$ , the result of applying  $\varphi$  to  $\sigma$  as an output. The first clause then says that a simple predication applied to  $n$  variables,  $P(x_1, \dots, x_n)$ , is interpreted as a filter, outputting only those assignments from the input which classically satisfy the predication.<sup>7</sup> The second clause says that conjunction is interpreted as sequential update by each of the conjuncts. The third clause says to update with an existential  $\exists x \varphi$  in two stages. In the first stage, the input set of assignments is replaced with a new set of assignments just like those in the input, except that  $x$  is allowed to take any value, regardless of what the input assignments mapped it to. This new set of assignments is now updated with  $\varphi$ . The fourth clause says that a set of assignments can be updated with the negation of a formula  $\varphi$  just in case none of the individual assignments could be successfully updated with  $\varphi$ .

(14) Basic clauses for dynamic semantics:

$$\begin{aligned} \sigma[P(x_1, \dots, x_n)] &= \{f \in \sigma \mid P(x_1, \dots, x_n) \text{ is classically satisfied by } M, f\} \\ \sigma[\varphi \wedge \psi] &= (\sigma[\varphi])[\psi] \\ \sigma[\exists x \varphi] &= \left\{ g \left| \begin{array}{l} \text{some assignment in } \sigma \text{ agrees with } g \text{ on all vari-} \\ \text{ables except possibly } x \end{array} \right. \right\} [\varphi] \\ \sigma[\neg\varphi] &= \{g \in \sigma \mid \{g\}[\varphi] = \emptyset\} \end{aligned}$$

The formulation of dynamic semantics for FOPL in (14) is one of many possibilities.<sup>8</sup> There are not only many formulations of dynamic semantics, but also alternative systems that maintain an essentially classical logic, and achieve the dynamic effects through other mechanisms. For example, Dekker's (1994) Predicate Logic with Anaphora keeps a static logic, and introduces an extra mechanism for the interpretation of anaphora. And more recently, a number of scholars, building on the the semantics of continuations in programming languages, have shown that the effects of dynamic logic can be reproduced using type shifting mechanisms (Shan and Barker, 2006; de Groote, 2006; Barker and Shan, 2008; Asher and Pogodalla, 2011).

### 3.4 Semantic Constraints on Anaphora Resolution

In the computational linguistic community, there has been more interest in the resolution of anaphora than in its interpretation. Resolution problems are often framed in terms of *coreference*, ignoring the

possibility of non-referential anaphoric expressions. In donkey sentences like (5), the pronouns *he* and *it* don't refer to particular farmers and donkeys, and similarly in the earlier *Every cat chased its tail*, *its* doesn't refer to a particular cat. A completely general account of anaphora resolution should incorporate cases like these where the semantic relationship between anaphor and antecedent is not coreference, but another kind of quantificational dependency.<sup>9</sup> Semantic accounts of anaphora such as given by DRT contribute to the general problem in two ways: (1) DRT provides a way of interpreting anaphora, and (2) it provides constraints on possible resolutions of anaphoric expressions in terms of *accessibility*.

The fact that there are semantic constraints on anaphora resolution is shown by (15a-15c): the *He* of *He was good* can be interpreted as Fischer for any of the three examples, but can only be resolved to *the Russian* in the first:

- (15) a. Fischer played a Russian. He was good.  
 b. If Fischer was in Iceland, he played a Russian. He was good.  
 c. Fischer did not play a Russian. He was good.

In standard DRT, anaphora resolution takes place on partially formed DRSs. So for (15a), an initial DRS like (16) is created, partially formed in the sense of having a “?” where a discourse referent should be, and then anaphora resolution replaces “?” with a discourse referent. In (16) *x* and *y* are both *accessible* to “?”, and so it can be replaced with *x* or *y*, allowing either reading of (15a).

(16)

<i>x y</i>
named( <i>x</i> ,“Fischer”)
russian( <i>y</i> )
played( <i>x</i> , <i>y</i> )
good(?)

For (15b), the initial DRS is as in (17). While the indefinite NP *a Russian*, here, creates a discourse referent which is embedded within the implication, the proper name *Fischer* creates a discourse referent at the top level. In DRT, the discourse referents for proper names are *promoted* to the top box.<sup>10</sup> Two question-marks (indexed for convenience) must now be resolved, and in order to resolve them, a general constraint on accessibility is defined:

ACCESSIBILITY CONSTRAINT: from a site within a box, a discourse referent is accessible *iff* that referent is introduced in that box, or along a path that heads either leftwards across connectives, or outwards.

For “?<sub>1</sub>”, that means that both *x* and *y* are accessible. A further completely standard syntactic constraint prevents ordinary non-reflexive pronoun arguments of a verb from being resolved to other arguments —this is *Principle B* of Government and Binding Theory. Thus “?<sub>1</sub>” has to be resolved to *x*. For “?<sub>2</sub>”, the accessibility constraint applies. Since *x* is introduced in the same box as “?<sub>2</sub>”, but *y* is introduced in a box which is inside that one (as opposed to outside, as the constraint requires), only *x* is accessible to “?<sub>2</sub>”. Thus the correct prediction is made that (15b) is unambiguous, and that



In addition to being the denotation of such sentences, events also play the role of discourse referents (Asher 1993, Lambrecht 1996, Humphreys, Gaizauskas & Azzam 1997). For example, the first of each of the sentences in (20) describes an event that is referred to anaphorically in the subsequent sentence.

- (20) a. Fischer moved quickly. Spassky didn't.  
 b. Deep Blue beat Garry Kasparov. It happened in 1997.  
 c. Abe Turner was stabbed in the back with a knife. A co-worker at Chess Review did it.

Examples like these show that events are more than an ancillary part of meaning. In translating sentences like (19) and (20), we are not free to only consider predicates ranging over nominal entities (as we have so far); if we do, two kinds of problems can arise.

First, representing (19) and (20) with LFs like in (21) fails to capture the intuition that these expressions are about events. The LFs in (21) specify what properties are true of individuals but say nothing about the properties of the events.

- (21) a.  $\text{moved-quickly}(\text{Fischer}) \wedge \neg \text{did}(\text{Spassky})$   
 b.  $\text{beat}(\text{Deep Blue}, \text{Garry Kasparov}) \wedge \text{happened}(1997)$   
 c.  $\exists x \text{ stabbed}(x, \text{Abe Turner}, \text{the back}, \text{a knife}) \wedge \text{did}(\text{a co-worker}, \text{Chess Review})$

There is also no guarantee that the predicates in the lefthand and righthand conjuncts of (21) refer to the same action. The functions *did* and *happened* are not required to refer to their logical antecedents; if they do, it is by luck or stipulation.

The second issue that arises is that for event-denoting sentences, a non-event-based semantics fails to predict the right entailments (Davidson, 1967). The actions in (20) happen in particular ways: *quickly*, *in 1997*, and *in the offices of Chess Review*. But, manner, time and place modifiers like these are additional information about the event, and so if the event is true with such a modifier it's true without it. The sentences on the lefthand side of the 'turnstile' ( $\models$ ) in (22) entails the sentences on the righthand side (i.e., whenever the first is true the second is true).

- (22) a. Fischer moved quickly.  $\models$  Fischer moved.  
 b. Deep Blue beat Garry Kasparov in 1997.  $\models$  Deep Blue beat Garry Kasparov.  
 c. Abe Turner was stabbed in the back with a knife.  $\models$   
     Abe Turner was stabbed in the back.  
     Abe Turner was stabbed with a knife.  
     Abe Turner was stabbed.

However, with LFs like (23a-23c), which lack a way to grab onto the event, we cannot explain the entailments.

- (23) a.  $\text{moved-quickly}(\text{Fischer}) \not\models \text{moved}(\text{Fischer})$   
 b.  $\text{beat}(\text{Deep Blue}, \text{Garry Kasparov}, 1997) \not\models \text{beat}(\text{Deep Blue}, \text{Garry Kasparov})$   
 c.  $\exists x \text{ stabbed}(x, \text{Abe Turner}, \text{the back}, \text{a knife}) \not\models$   
      $\exists x \text{ stabbed}(x, \text{Abe Turner}, \text{the back})$   
      $\exists x \text{ stabbed}(x, \text{Abe Turner}, \text{a knife})$   
      $\exists x \text{ stabbed}(x, \text{Abe Turner})$

In DRT, introducing event terms into the semantic representation means that there are event-type discourse referents and DRS conditions that specify what properties the event referents have, typically conditions on how the event took place (manner, place, time, etc.), and who was involved in it

and how (thematic role conditions such as **agent** and **theme** that link events to other discourse referents).<sup>11</sup> Each of these conditions is given separately (that they are separate is crucial to explaining the entailment issues discussed above). For example, (24b), the DRS for (24a):

(24) a. Fischer moved quickly.

	$x e$
b.	<p>named(<math>x</math>, "Fischer")          move(<math>e</math>)          agent(<math>e, x</math>)          quick(<math>e</math>)</p>

(24b) says that (24a) introduces a discourse referent for a *moving* event  $e$ , asserting that  $e$  is performed by Fischer and that  $e$  occurs quickly. In this way, verbs in event-denoting sentences are more than predicates of their arguments. The contribution of an event-denoting verb to the meaning of a sentence is both an event argument and the conditions on what the event is like. This makes it easy to represent event anaphora and talk about entailments between event-denoting sentences.

The sentences and DRSs in (25) and (26) illustrate event anaphora. If the DRS for sentence (25a), (25b), serves as the context for the interpretation of (26a), then in the resulting DRS, (26b), the first event is accessible to and resolvable to the pronoun *it* in the second sentence, and any conditions applying to either apply to both. So if  $it$  is assigned to  $e'$  and resolved to  $e$ , the condition  $e' \subseteq 1997$  also applies to  $e$ ; as desired, (26a) says that (25a) occurred in 1997.

(25) a. Deep Blue beat Garry Kasparov.

	$x y e$
b.	<p>named(<math>x</math>, "Deep Blue")          named(<math>y</math>, "Kasparov")          beat(<math>e</math>)          agent(<math>e, x</math>)          theme(<math>e, y</math>)</p>

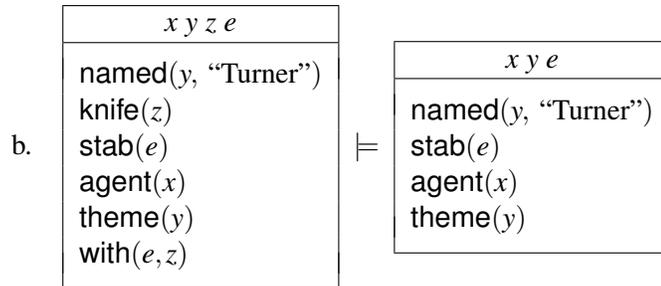
(26) a. It happened in 1997.

	$x y e e'$
b.	<p>named(<math>x</math>, "Deep Blue")          named(<math>y</math>, "Kasparov")          beat(<math>e</math>)          agent(<math>e, x</math>)          theme(<math>e, y</math>)</p> <p><math>e' \subseteq 1997, e' = ?</math></p>

Entailment relations are also easy to account for in the DRT analysis of event-denoting sen-

tences; for example:

(27) a. Abe Turner was stabbed with a knife.  $\models$  Abe Turner was stabbed.



One sentence entails another if whenever the first is true, the second is true, so to verify (27a) we check that whenever the lefthand side is true, the righthand side is too, which we can confirm with the DRSs in (27b).<sup>12</sup> The DRS on the lefthand side of (27b) is true if there are some real-world entities that match up with discourse referents  $x$ ,  $y$  and  $z$  so that  $y$  is Abe Turner,  $z$  is a knife, and  $e$  is an event of *stabbing* of  $y$  by some person  $x$  with  $z$ . Clearly, whenever these can be satisfied, the DRS on the righthand side will be true. This is because the DRS conditions of the latter are all among the DRS conditions of the former, and what satisfies the first set of conditions can be used to satisfy the second set.

While the need for events in semantic representation was originally just a philosophical concern, motivated by the aforementioned considerations, it is of practical significance too. Event semantic representations are needed for resolving both coreference resolution in the general case, and inference patterns in textual entailment tasks are often issues of what events have happened and how they relate to each other (see Mitkov; Padó and Dagan, this volume, as well as Hirschman and Chinchor 1997; NIST 2007; Dagan et al. 2006; Androutsopoulos and Malakasiotis 2010).

## 4.2 Tense

At a high level, language can convey two kinds of information about time: (1) tense and (2) aspect. Tense, which is discussed here, is the (grammatical) relationship between when a sentence is said and the time of the event described by it. Aspect (Vendler 1957, 1967; Comrie 1976; Smith 1991; see Allen 1984; Moens and Steedman 1988; Dorr and Olsen 1997 for computationally-oriented discussion), on the other hand, concerns the dynamics of events; e.g., whether they are ongoing, whether they have starting and end points, how they develop, and how they are viewed. While it is tempting to analyze tense just as past, present and future, the comparison below between simple tense logic and an analysis in terms speech time, event time & reference time makes it clear that there is more going on.

### Tense logic

In **tense logic** (Prior, 1957, 1967), tense is represented by applying temporal operators to basic, present tense sentences (Montague, 1973; Dowty et al., 1981). The meaning of a tensed sentence on this approach is its present tense translation with the relevant temporal operator applied:

- (28) a. Deep Blue is a chess-playing computer.  
 b. **P**(Deep Blue is a chess-playing computer)  
 c. **F**(Deep Blue is a chess-playing computer)

In (28), taking the present tense (28a) as basic, the past tense *Deep Blue was a chess-playing computer* is got by applying the predicate **P** (*past*). Similarly, the future tense comes from (28a) by applying **F** (*future*). The truth conditions of past and future tense sentences are given in terms of the truth conditions of the present. If the event time of the simple present sentence stands in the correct relation to the speech time, then the tensed sentence is true:

- (29) a.  $\varphi$  is true iff  $[[\varphi]]$  is true at  $t = \text{now}$   
 b. **P**( $\varphi$ ) is true iff  $[[\varphi]]$  is true at time  $t < \text{now}$   
 c. **F**( $\varphi$ ) is true iff  $[[\varphi]]$  is true at time  $t > \text{now}$

For example:

- (30) a. chess-playing-computer(*Deep Blue*) is true iff it's true at time  $t = \text{now}$   
 b. **P**(chess-playing-computer(*Deep Blue*)) is true iff  
 chess-playing-computer(*Deep Blue*) is true at time  $t < \text{now}$   
 c. **F**(chess-playing-computer(*Deep Blue*)) is true iff  
 chess-playing-computer(*Deep Blue*) is true at time  $t > \text{now}$

### Speech time, event time & reference time

While tense logic has the advantage of great simplicity, it was earlier noted by Reichenbach (1947) that tensed sentences depend on more than the relationship between speech time and event time. E.g., in (31a) and (31b), Fischer's move has to occur relative to Spassky's move and not just in the past or future.

- (31) a. Fischer moved after Spassky (moved). R < E < S  
 b. Fischer will move before Spassky (moves). S < E < R

There are three relevant time points: (1) event time (E), (2) speech time (S) which is when a sentence is said, and (3) reference time (R), which is how E and S are ordered. In (31a) and (31b) the event time is the time at which Fischer moves and the reference time is the time at which Spassky moves. When Fischer moves after Spassky,  $R < E$ . When Fischer moves before Spassky,  $E < R$ . In past tense, event time and reference time precede the speech time,  $E, R < S$ , and in future tense, they follow speech time,  $S < E, R$ . In these terms, simple present has the orderings  $E = R = S$ .

Tense logic can be and has been generalized to account for relationships among arbitrary time points such as event time, reference time, and speech time (Kamp, 1968; Allen, 1984), but accounting for these tense relationships is more perspicuous in the DRT formalism. Just as event terms are used for event-denoting sentences, terms corresponding to times are used in the analysis of tense

(Kamp and Reyle, 1993). For example, (31a) and (31b):

(32) a. Fischer moved after Spassky (moved).

	$x$	$y$	$e$	$e'$	$t$	$t'$	now
a.			named( $x$ , “Fischer”)				
			move( $e$ )				
			agent( $e$ , $x$ )				
			$e \subseteq t$ , $t < \text{now}$				
b.			named( $y$ , “Spassky”)				
			move( $e'$ )				
			agent( $e'$ , $y$ )				
			$e' \subseteq t'$ , $t' < \text{now}$				
			$t > t'$				

(33) a. Fischer will move before Spassky (moves).

	$x$	$y$	$e$	$e'$	$t$	$t'$	now
a.			named( $x$ , “Fischer”)				
			move( $e$ )				
			agent( $e$ , $x$ )				
			$e \subseteq t$ , $t > \text{now}$				
b.			named( $y$ , “Spassky”)				
			move( $e'$ )				
			agent( $e'$ , $y$ )				
			$e' \subseteq t'$ , $t' > \text{now}$				
			$t < t'$				

In (32), Fischer’s move  $e$  occurs within the time interval  $t$  which precedes the speech time (i.e., it’s in the past), but it also occurs after the time  $t'$  of the reference event  $e'$ , Spassky’s move. In contrast, in (33), the time  $t$  of  $e$ , follows the speech time (i.e., it’s in the future) but precedes the time  $t'$  of the reference event  $e'$ . Note that in the semantics of tensed DRSs, intervals of time ( $t, t'$  above), rather than events, are ordered; it’s only through their association with time intervals that events can then be said to occur before, after, or at the same time as each other. It’s not uncommon, however, for relations between events to be encoded directly in temporally annotated corpora (see, e.g., Verhagen et al., 2009).

### 4.3 Temporal anaphora

With the addition of temporal discourse referents, DRT predicts that there should be temporal analogues to nominal anaphora, something which does arise as first noted by Partee (1973) and later

discussed in Partee 1984 and Hinrichs 1986. Moreover, there are temporal correlates to indefinite reference and donkey anaphors, discussed below.

We saw in the discussion on tense that the time of one sentence can serve as an antecedent time for a subsequent sentence; in particular as the reference time. For example, when the second sentence in (34a) is interpreted in the context of the first, there is a reference event time discourse referent  $r$  which must be resolved to an antecedent event time. This is akin to indefinite reference (Partee, 1984):

- (34) a. Fischer moved. Spassky resigned.

$x \ e \ t \ \text{now} \ y \ e' \ t' \ r$
named( $x$ , "Fischer") move( $e$ ) agent( $e, x$ ) $e \subseteq t, t < \text{now}$
b.
named( $y$ , "Spassky") resign( $e'$ ) agent( $e', y$ ) $e' \subseteq t', t' < \text{now}, t' > r, r = ?$

In (34b), the time  $t$  of Fischer's moving event  $e$  is accessible to  $r$ , allowing the resolution  $r = t$ . The reference event time, acting like a pronoun in nominal reference, is introduced and resolved to the accessible antecedent event time.

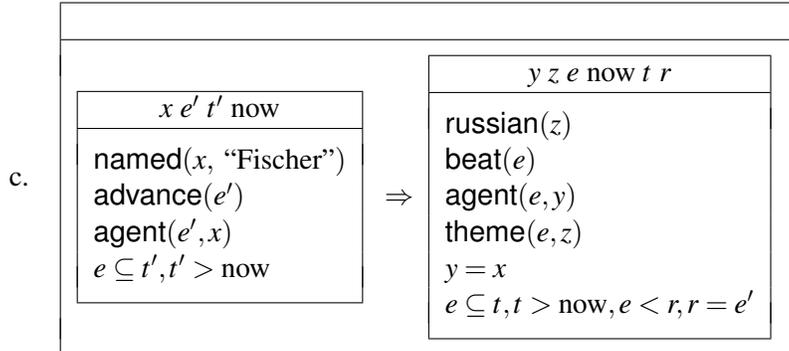
There are also temporal variants of donkey sentences. In the FOPL representations of (35a) and (36a), (35b) and (36b), the event time  $t'$  of  $e'$ , which originates in the antecedent, can't be bound in the consequent. This is just like the problem of nominal donkey anaphora except that it's in the domain of temporal discourse referents. The solution to temporal donkeys in DRT is straightforward:

- (35) a. If Fischer advanced he had (already) beat a Russian.

- b.  $(\exists e', t' (\text{advance}(e') \wedge \text{agent}(e', \text{Fischer}) \wedge e' \subseteq t' \wedge t' < \text{now})) \rightarrow$   
 $(\exists z, e, t (\text{beat}(e) \wedge \text{agent}(e, \text{Fischer}) \wedge \text{theme}(e, z) \wedge \text{russian}(z) \wedge e \subseteq t \wedge t < \text{now} \wedge t < t'))$

c.	<table border="1" style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding: 5px;"> <math>x \ e' \ t' \ \text{now}</math> </td> <td style="padding: 5px;"> <math>y \ z \ e \ \text{now} \ t \ r</math> </td> </tr> <tr> <td style="padding: 5px;">           named(<math>x</math>, "Fischer")            advance(<math>e'</math>)            agent(<math>e', x</math>)  <math>e' \subseteq t', t' &lt; \text{now}</math> </td> <td style="padding: 5px;"> <math>\Rightarrow</math> </td> <td style="padding: 5px;">           russian(<math>z</math>)            beat(<math>e</math>)            agent(<math>e, y</math>)            theme(<math>e, z</math>)  <math>y = x</math>  <math>e \subseteq t, t &lt; \text{now}, t &lt; r, r = ?</math> </td> </tr> </table>	$x \ e' \ t' \ \text{now}$	$y \ z \ e \ \text{now} \ t \ r$	named( $x$ , "Fischer") advance( $e'$ ) agent( $e', x$ ) $e' \subseteq t', t' < \text{now}$	$\Rightarrow$	russian( $z$ ) beat( $e$ ) agent( $e, y$ ) theme( $e, z$ ) $y = x$ $e \subseteq t, t < \text{now}, t < r, r = ?$
$x \ e' \ t' \ \text{now}$	$y \ z \ e \ \text{now} \ t \ r$					
named( $x$ , "Fischer") advance( $e'$ ) agent( $e', x$ ) $e' \subseteq t', t' < \text{now}$	$\Rightarrow$	russian( $z$ ) beat( $e$ ) agent( $e, y$ ) theme( $e, z$ ) $y = x$ $e \subseteq t, t < \text{now}, t < r, r = ?$				

- (36) a. If Fischer advances he will have (already) beat a Russian.  
 b.  $(\exists e', t' (\text{advance}(e') \wedge \text{agent}(e', \text{Fischer}) \wedge e' \subseteq t' \wedge t' > \text{now})) \rightarrow$   
 $(\exists z, e, t (\text{beat}(e) \wedge \text{agent}(e, \text{Fischer}) \wedge \text{theme}(e, z) \wedge \text{russian}(z) \wedge e \subseteq t \wedge t > \text{now} \wedge e <$   
 $e'))$



Just as the quantificational force of a nominal indefinite is determined by its accessibility from an anaphor, whether reference event time can be resolved to an indefinite antecedent event time is determined by accessibility constraints. The analysis of nominal and temporal anaphora is entirely uniform.

## 5 Deep and shallow semantic methods

It seems paradoxical that while many NLP tasks (e.g., textual inference, question answering, sentiment classification, and natural language generation) are inherently semantic, the best systems often make little to no use of methods used by formal semanticists. A standard dichotomy is that between *deep* and *shallow* methods. Deep methods are those which represent the meaning of texts using a linguistically motivated analysis and techniques such as those discussed in this chapter, parsing each sentence, identifying its logical or semantic structure and interpreting it compositionally. At the other extreme are shallow methods that represent the meaning of texts in terms of surface-level features of text, such as n-grams, the presence or absence of keywords, sentence length, punctuation and capitalization, and rough document structure. One approach of this sort is found in vector space models such as Latent Semantic Analysis (LSA; Landauer et al. 1998).

In vector space models like LSA, word meaning is captured by word-word frequencies, document-word frequencies, and vector similarities. For example, the meaning of *chess* is represented by the vector of weights for words that co-occur with it. These weights would be large for words like *win* and *bishop* but small for words like *dog* and *climb*. Sentences and documents, then, are represented by sets of vectors for and similar to the word vectors in the text. So the meaning of a news report about Bobby Fischer is the set of vectors for and similar to the words in the story; for example, the vectors for *chess*, *American*, *strategy*, etc.

Shallow methods like LSA are semantic in that they characterize the *aboutness* of sentences and documents; however, there are ways in which they do not measure up as general theories of meaning. First, they do not account for the semantic significance of high-frequency, closed class words like quantifiers and pronouns, the cornerstone of formal semantics. Second, they ignore the role of syntax, and so scope and binding, in semantic composition. This has a substantial effect on

sentential meaning. Last, they do not and cannot characterize propositional content because they do not take on quantification, anaphora, scope and binding.

The importance of shallow methods despite the impoverished representation of meaning is largely a product of the ease with which such methods are implemented, but it also a result of the evaluation criteria for many NLP tasks, which put more emphasis on robustness and breadth of coverage than on handling the edge cases studied by formal semanticists. A question then arises: when do deep semantic methods pay off in NLP? We will not attempt to answer that question in generality, but discuss a class of problematic cases for shallow methods that involve embedding of expressions within the scope of predicates and operators. For example, in considering Fischer's famous remark (37), it is important that a sentiment classification or textual inference application take into account the fact that *like chess*, an expression of positive sentiment towards the game, is embedded in the negative template *don't . . . anymore*.

(37) I don't like chess any more.

NLP practitioners have long been aware of the problems created by negation, and sometimes circumvent them by building systems that are sensitive to just the most obvious cases (e.g., *not*, *don't*, *nothing*, etc.). However, negation is just one example of an embedding operator that drastically affects meaning. In the remainder of this section we discuss other embedding operators that illustrate some of the many difficulties facing shallow approaches.

## 5.1 Intensionality and Non-veridicality

### Intensionality

The *extension* of an expression is the object(s) that it picks out, so, for example, the extension of *wild horses* is the set of all wild horses, and the extension of *unicorn* is, presumably, the empty set. A counterpart to this is the intension<sup>13</sup> (Carnap 1947) of an expression, its concept, which abstracts over its possible extensions. An *intensional* context is one in which it is necessary to go beyond extensions in order to capture the meaning. For example, in the sentence *Fred wishes there were unicorns*, we can only understand his wishes by considering what the world would be like *if* there were unicorns.

Intensional contexts can be identified using a substitutivity test introduced by Frege (1892): suppose that two terms have the same extension in the actual world, but pick out different concepts; for example, *the eleventh World Chess Champion* vs. *Bobby Fischer*; then a sentence creates an intensional context if swapping the terms changes the truth of the sentence. According to this, the verb *meet* does not create an intensional context. If Fischer is the eleventh World Chess Champion, (38a) is true iff (38b) is true. However, *hope* does create an intensional context, since holding fixed the course of chess history, (39a) can be true and (39b) false.

(38) a. Nixon met the eleventh World Chess Champion.  
b. Nixon met Fischer.

(39) a. Brezhnev hoped the eleventh World Chess Champion would be Russian.  
b. Brezhnev hoped Fischer would be Russian.

By this test, intensional contexts are created by a wide range of constructions and expressions: belief and desire (e.g., *discover*, *regret*), speech reports (e.g., *say that*), probability, possibility, and

necessity (e.g., *perhaps*, *should*), causal and inferential relationships (e.g., *caused*, *so* and *therefore*), etc. There is a correspondingly huge literature on intensionality in both linguistic semantics (Montague, 1970a; Kratzer, 1981) and philosophy of language, of which modal logic (Blackburn et al., 2001) is just one strand.

### **(Non-)veridicality**

Intensionality is closely related to the notion of (non-)veridicality. A *veridical* operator is one that preserves truth: if  $O$  is veridical, then  $O(\phi)$  entails  $\phi$ . For example, the phrase *it's undoubtedly the case that* is veridical, so *it's undoubtedly the case that you'll enjoy this game* entails *you'll enjoy this game*. On the other hand, *it is doubtful whether* is non-veridical, so *it is doubtful whether you'll enjoy this game* does not entail *you'll enjoy this game*. Negation is also non-veridical, and so are expressions like *maybe*, *I dreamed that*, *it is rumored that*, *it is unlikely that*, and many more. While many intensional operators are non-veridical, *it's undoubtedly the case that* is both intensional and veridical, as are many so-called *factive* verbs like *regret* and *know*. So while intensionality involves considerations of ways the world might have been, it is independent of non-veridicality. In many languages, non-veridicality is also signaled by special grammatical forms, such as the subjunctive mood. A subjunctive clause, such as *he were channelling Fischer* in *it's as though he were channelling Fischer* typically is not believed by the speaker.

Given the frequency and range of constructions that signal intensionality and non-veridicality across the world's languages<sup>14</sup>, NLP tasks like textual entailment are not isolated from such problems. For this reason, research facing this issue head-on is active (see, e.g., Condoravdi et al. 2003; Bobrow et al. 2005; MacCartney and Manning 2007; MacCartney 2009; Nairn et al. 2006; Schubert et al. 2010).

## **5.2 Monotonicity**

### **Inferential (non-)monotonicity**

The term (non-)monotonicity has several uses, one of which relates to (non-)veridicality. The general idea will be familiar: loosely speaking, a monotonic function is one that always goes in the same direction. Its sense in semantics is a direct extension of this, where the relevant function is entailment. However, monotonicity is applied to entailment in two ways. In AI and philosophical logic, entailment is monotonic if adding premises doesn't invalidate conclusions. Non-monotonicity in this sense is 'inferential non-monotonicity', exemplified by situations like (40a) and (40b) in which additional information invalidates or changes the meaning of an argument.

- (40) a. Fischer beat Spassky. Therefore Fischer won a game.  
b. Fischer beat Spassky. Spassky was his donkey. Therefore Fischer won a game.

In (40), the inferential non-monotonicity effects relate to the lexical ambiguity of *beat*. However, inferential non-monotonicity can also result from default expectations about the world or conversational norms. Default expectations, for example, license reasoning from *DB. is a chess player* to *DB. is human*, but the addition of the premise *DB was built by IBM* then defeats the inference. Expectations about conversational norms are discussed in the chapter on pragmatics (Potts, this volume).

	Upward Monotonic	Veridical
It's true that S	✓	✓
It's false that S	×	×
It's conceivable that S	✓	×
It is not widely known that S	×	✓

**Table 1:** Environmental monotonicity and veridicality are independent properties. All the logical possibilities occur in natural language.

### Environmental (non-)monotonicity

In linguistic semantics a notion of ‘environmental (non-)monotonicity’ emerges. In this sense non-monotonicity relates to the monotonicity of entailments in syntactic embedding. A sentential operator  $O$  is *upward monotonic* iff whenever a sentence  $S \models S'$ ,  $O(S) \models O(S')$ . A sentential operator  $O$  is *downward monotonic* iff whenever  $S \models S'$ ,  $O(S') \models O(S)$ . For example, if  $S$  is *Spassky will castle on the queen's side*, and  $S'$  is *Spassky will castle*, we have  $S \models S'$ . With the upward monotonic operator *it's certain*, (41a) is valid and (41b) is invalid. On the other hand, with the downward monotonic operator *it's doubtful that*, the pattern reverses; (42b) is valid and (42a) is not.

- (41) a. It's certain that Spassky will castle on the queen's side  
 $\models$  It's certain that Spassky will castle.  
 b. It's certain that Spassky will castle  
 $\not\models$  It's certain that Spassky will castle on the queen's side.

- (42) a. It's doubtful that Spassky will castle on the queen's side  
 $\not\models$  It's doubtful that Spassky will castle.  
 b. It's doubtful that Spassky will castle  
 $\models$  It's doubtful that Spassky will castle on the queen's side.

Veridicality and environmental monotonicity are related in that operators can be both (non-)veridical and (non-)monotonic; however they are independent notions as illustrated by Table 1. Also, non-monotone operators can be neither upward or downward monotonic. *Spassky worried that S* is non-monotonic in this sense: if he worried that Fischer would castle on the queen's side, it doesn't follow that he worried that Fischer would castle, and if he worried that Fischer would castle, it doesn't follow that he worried that Fischer would castle on the queens' side.

### Monotonicity & quantification

Monotonicity is not only a property of proposition-embedding operators; it is more generally something that applies to any *environment* embedding a set-denoting term. If  $S$  is a sentence containing a single occurrence of a set denoting term  $\alpha$  and  $S[\alpha/\beta]$  is the sentence in which  $\alpha$  has been replaced by  $\beta$ ,  $\alpha \supseteq \beta$ , then (1)  $\alpha$  occurs in an *downward monotone environment* iff  $S \models S[\alpha/\beta]$  and (2) it occurs in a *upward monotone environment* iff  $S[\alpha/\beta] \models S$ . In other words, downward monotone environments license inferences to subsets and upward monotone environments license inferences to supersets.<sup>15</sup>



## Notes

<sup>1</sup>Frege’s first-order logic was first motivated as a representation for mathematical statements, but as evident in his philosophy of language and its legacy, this was not its only application.

<sup>2</sup>Note that though essential, the object language/metalanguage distinction can easily become confusing. For example, when we talk about the semantics of the semantic meta-language, which for (1b) would just be the standard semantics for FOPL plus a listing of predicates and constants, we are then effectively treating the semantic meta-language as an object-language for a higher level description. The semantics of FOPL can then be thought of as meta-meta-language relative to the original expression of English being translated.

<sup>3</sup>Strawson’s position was that definite descriptions carry *presuppositions* —see, e.g., Beaver 1997; or Bos 2003 for a computational treatment.

<sup>4</sup>See Champollion et al. 2007 for an interactive tutorial on the Lambda Calculus and ? for a modern application that incorporates probability.

<sup>5</sup>There are several handbook articles on DRT: Beaver and Geurts 2007; van Eijck and Kamp 1997; Kamp et al. 2011 as well as a textbook presentation in Kadmon 2001. A generalization of DRT to deal with (anaphoric) reference to abstract objects like propositions is given in Asher 1993, and a broader generalization to model discourse relationships between segments of text is given in Asher and Lascarides 2003.

<sup>6</sup>We will not set out the semantics of the DRS language in full here, for which the reader may refer to any of the references given above. However, if we assume that satisfaction of other types of conditions is defined relative to a model and an assignment, then for implication the generalized semantics would be as follows:

$M, f \models [x_1, \dots, x_i \mid c_1, \dots, c_m] \Rightarrow [x_{i+1}, \dots, x_j][c_{m+1} \dots c_n]$  iff for every assignment  $g$  just like  $f$  except for the values assigned to the variables  $x_1, \dots, x_i$  such that  $M, g$  satisfies all of the conditions  $c_1, \dots, c_m$ , there exists some further assignment  $h$  differing from  $g$  at most with respect to the variables  $x_{i+1}, \dots, x_j$ , and such that  $M, h$  satisfies all of the conditions  $c_{m+1}, \dots, c_n$ .

Note that implementations of DRT and Heim’s file change semantics differ as regards whether re-assignment is *destructive*. In some versions of the semantics, assignments are considered which may overwrite the values previously given to referents, while in other versions *partial assignment functions* are used, so that instead of talking of one assignment differing from another with respect to a referent, we talk of one assignment extending another with respect to that referent, meaning that whereas the old assignment does not provide a value for the referent, the modified assignment does.

<sup>7</sup>In this semantics all updates are relative to a model  $M$ , so strictly we should write  $\sigma[\phi]_M$  not  $\sigma[\phi]$ , but the  $M$  is omitted for notational simplicity.

<sup>8</sup>Note that the semantics given for negation operates *pointwise*, in the sense that it splits an input context into individual assignments. Given such a negation, the semantics could be presented more simply as a relation between single assignments rather than as an update on sets of assignments, but the latter formulation makes clearer the Stalnakerian intuitions underlying dynamic proposals. For discussion of some alternatives, see Groenendijk and Stokhof 1991b; van Benthem 1996; van Eijck and Visser 2010.

<sup>9</sup>The quantificational dependency between anaphors and antecedents is often described as *bound anaphora*; however, it should be noted that many accounts invoke mechanisms other than binding. Some philosophers and linguists have analyzed pronouns in donkey sentences and discourse anaphora as implicit definite descriptions; *every cat chased its tail* means that *every cat chases that cat’s tail*.

<sup>10</sup>In early versions of DRT, the promotion of proper names is stipulative. However, van der Sandt 1992 showed how promotion could be explained as part of a general treatment of *presuppositions*.

<sup>11</sup>Kamp and Reyle (1993) do not use thematic roles in their presentation of DRT. This side-steps questions about what kind of things thematic roles are but then the entailment problem has to be solved by stipulating axioms. The “neo-

Davidsonian” approach (Dowty, 1989; Parsons, 1990) we use, despite having philosophical problems, handles event entailments perspicuously. As a practical consideration, neo-Davidsonian analysis is adopted in Bos’s (2008) DRT parser Boxer.

<sup>12</sup>More carefully, we want to show that, for the DRSs, whenever there is assignment function  $f$  that satisfies the former DRS with respect to a model  $M$ ,  $f$  also satisfies the latter DRS with respect to  $M$ .

<sup>13</sup>Note that *intension* with an ‘s’ is distinct from *intention* with a ‘t’, the latter of which refers to the goals of individuals.

<sup>14</sup>For discussion of non-veridical contexts and their linguistic significance, see Zwarts (1995); Giannakidou (1999, 2002).

<sup>15</sup>In linguistics, much of the interest in monotonicity centers on ‘polarity’ items; for example, negative polarity items (NPIs) like *any* and *the slightest bit* that occur in downward monotone contexts, and positive polarity items like *already* and British English *rather* that tend not to occur in downward monotone environments. See Fauconnier (1975); Ladusaw (1980, 1997); Krifka (1995); Von Stechow (1999) for discussion of NPI distribution and monotonicity, and Giannakidou (1999) for discussion of the idea that some NPIs are licensed not via downward monotonicity but by non-veridicality.

<sup>16</sup>This is sometimes notated  $\downarrow\text{EVERY}\uparrow$ ; i.e., *every* is left downward monotone in its first argument, and right upward monotone in its second.

<sup>17</sup>There is a large body of work on the semantics and logic of quantification. Notable papers are Barwise and Cooper 1981; Keenan and Stavi 1986; van Benthem 1984; Westerståhl 2007.

<sup>18</sup>Some other computational work builds on only some aspects of the insights of natural logic. For example, sentiment classification and textual entailment systems often build in notions of *polarity reversal*, switching inferences or sentiment evaluations around in the immediate environment of a negative word. See, e.g., Wilson et al. (2005).

## References

- Allen, James F. (1984). ‘Towards a general theory of action and time’. *Artificial Intelligence*, 23(2), 123–154. doi: 10.1016/0004-3702(84)90008-0.
- Androutopoulos, Ion and Prodromos Malakasiotis (2010). ‘A survey of paraphrasing and textual entailment methods’. *Journal of Artificial Intelligence Research*, 38, 135–187. doi: 10.1613/jair.2985.
- Asher, Nicholas (1993). *Reference to abstract objects in discourse*. Springer. ISBN 0792322428.
- Asher, Nicholas and Alex Lascarides (2003). *Logics of conversation*. Cambridge Univ Pr. ISBN 0521650585.
- Asher, Nicholas and Sylvain Pogodalla (2011). ‘A Montagovian treatment of modal subordination’. In Li, Nan and David Lutz, editors, *Proceedings of SALT 20*. LSA Publications.
- Bach, Emmon (1986). ‘The algebra of events’. *Linguistics and Philosophy*, 9(1), 5–16. doi: 10.1007/BF00627432.
- Barker, Chris and Chung-Chieh Shan (2008). ‘Donkey anaphora is in-scope binding’. *Semantics and Pragmatics*, 1(1), 1–46. doi: 10.3765/sp.1.1.
- Barwise, J. and R. Cooper (1981). ‘Generalized quantifiers and natural language’. *Linguistics and philosophy*, 4(2), 159–219. ISSN 0165-0157.
- Beaver, David I. (1997). ‘Presupposition’. In van Benthem, Johan and Alice ter Meulen, editors, *The Handbook of Logic and Language*. Elsevier, Amsterdam.

- Beaver, David I. (2004). ‘The optimization of discourse anaphora’. *Linguistics and Philosophy*, 27(1), 3–56. ISSN 0165-0157.
- Beaver, David I. and Bart Geurts (2007). ‘Discourse Representation Theory’. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Center for the Study of Language and Information, Stanford, CA.
- van Benthem, Johan (1984). ‘Questions about quantifiers’. *Journal of Symbolic Logic*, 49(2), 443–466. ISSN 0022-4812.
- van Benthem, Johan (1996). *Exploring Logical Dynamics*. Center for the Study of Language and Information, Stanford, CA.
- Blackburn, Patrick and Johan Bos (2000). ‘Working with discourse representation theory: An advanced course in computational semantics’.
- Blackburn, Patrick and Johan Bos (2005). *Representation and Inference for Natural Language: A First Course in Computational Semantics*. CSLI Press, Stanford, CA.
- Blackburn, Patrick, Maarten de Rijke, and Yde Venema (2001). *Modal Logic*. Cambridge University Press, Cambridge.
- Bobrow, Daniel, Cleo Condoravdi, Richard Crouch, , Ron Kaplan, Lauri Karttunen, Tracy King, Valeria de Paiva, and Annie Zaenen (2005). ‘A basic logic for textual inference’. In *AAAI Workshop on Inference for Textual Question Answering*.
- Boole, George (1854). *An Investigation of The Laws of Thought on Which are Founded the Mathematical Theories of Logic and Probabilities*. Macmillan, London. Available from Project Gutenberg.
- Bos, Johan (2003). ‘Implementing the binding and accommodation theory for anaphora resolution and pre-supposition projection’. *Computational Linguistics*, 29(2), 179–210.
- Bos, Johan (2008). ‘Wide-coverage semantic analysis with Boxer’. In Bos, Johan and Rodolfo Delmonte, editors, *Semantics in Text Processing (STEP-2008)*, Research in Computational Semantics, 277–286. College Publications.
- Bos, Johan and Katja Markert (2005). ‘Recognising textual entailment with logical inference’. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 628–635. Association for Computational Linguistics.
- Champollion, Lucas, Josh Tauberer, and Maribel Romero (2007). ‘The penn lambda calculator: Pedagogical software for natural language semantics’. In King, T. Holloway and Emily M. Bender, editors, *Proceedings of the Grammar Engineering Across Frameworks (GEAF) Workshop*.
- Church, Alonzo (1932). ‘A set of postulates for the foundation of logic’. *Annals of Mathematics*, 33(2), 346–366.
- Comrie, Bernard (1976). *Aspect*. Cambridge University Press, Cambridge.
- Condoravdi, Cleo, Richard Crouch, Reinhard Stolle, Valeria de Paiva, and Daniel Bobrow (2003). ‘Entailment, intensionality and text understanding’. In *Workshop on Text Meaning, Human Language Technology Conference (HLT-NAACL-2003)*.
- Dagan, Ido, Oren Glickman, and Bernardo Magnini (2006). ‘The PASCAL recognising textual entailment challenge’. *Machine Learning Challenges*, 3944/2006, 177–190. doi: 10.1007/11736790\_9.

- Davidson, Donald (1967). 'The logical form of action sentences'. In Rescher, Nicholas, editor, *The Logic of Decision and Action*, 81–95. University of Pittsburgh Press, Pittsburgh. Reprinted in Davidson 2001, Davidson 2006.
- Dekker, Paul (1994). 'Predicate logic with anaphora'. In Santelmann, Lynn and Mark Harvey, editors, *Proceedings of Semantics and Linguistic Theory (SALT) IV*, 79–95.
- Dorr, Bonnie J. and Mari Broman Olsen (1997). 'Deriving verbal and compositional lexical aspect for NLP applications'. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, 151–158.
- Dowty, David (1989). 'On the semantic content of the notion of "thematic role"'. In Chierchia, Gennaro, Barbara Partee, and Raymond Turner, editors, *Properties, Types and Meaning*, volume 2 of *Semantic Issues*, 69–130. Kluwer, Dordrecht and Norwell, MA.
- Dowty, David (1994). 'The role of negative polarity and concord marking in natural language reasoning'. In Santelmann, Lynn and Mark Harvey, editors, *Proceedings of Semantics and Linguistic Theory (SALT) IV*, 114–144.
- Dowty, David R., Robert E. Wall, and Stanley Peters (1981). *Introduction to Montague Semantics*. Reidel, Dordrecht.
- van Eijck, Jan and Hans Kamp (1997). 'Representing discourse in context'. In van Benthem, Johan and Alice ter Meulen, editors, *Handbook of Logic and Language*, 179–238. Elsevier Science, Amsterdam. Also published by MIT Press: Cambridge, MA.
- van Eijck, Jan and Albert Visser (2010). 'Dynamic semantics'. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Center for the Study of Language and Information.
- Fauconnier, Giles (1975). 'Pragmatic scales and logical structure'. *Linguistic Inquiry*, 6(3), 353–375. ISSN 0024-3892.
- Frege, Gottlob (1879). *Begriffsschrift, a formal language, modeled on that of arithmetic, for pure thought*. Reprinted in van Heijenoort (ed.) 1967.
- Frege, Gottlob (1892). 'Über Sinn und Bedeutung (on sense and reference)'. *Zeitschrift für Philosophie und philosophische Kritik*, 100(1), 25–50. Reprinted in Black (ed.) 1980.
- Geach, Peter T. (1962). *Reference and generality*. Cornell Univ. Press.
- Giannakidou, Anastasia (1999). 'Affective dependencies'. *Linguistics and Philosophy*, 22(4), 367–421. ISSN 0165-0157.
- Giannakidou, Anastasia (2002). 'N-words and negative concord'. *The Linguistics Companion*.
- Groenendijk, Jeroen and Martin Stokhof (1991a). 'Dynamic Predicate Logic'. *Linguistics and Philosophy*, 14, 39–100.
- Groenendijk, Jeroen and Martin Stokhof (1991b). 'Two theories of dynamic semantics'. *Lecture Notes in Computer Science*, 478(Logics in AI), 55–64. doi: 10.1007/BFb0018433.
- de Groote, Philippe (2006). 'Towards a Montagovian account of dynamics'. In *Proceedings of Semantics and Linguistic Theory (SALT) 16*.
- Grosz, Barbara J., Scott Weinstein, and Aaravind K. Joshi (1995). 'Centering: A framework for modeling the local coherence of discourse'. *Computational linguistics*, 21(2), 203–225. ISSN 0891-2017.

- Halvorsen, Per-Kristian and William A. Ladusaw (1979). 'Montague's 'universal grammar': An introduction for the linguist'. *Linguistics and Philosophy*, 3(2), 185–223. doi: 10.1007/BF00126510.
- Heim, Irene (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts at Amherst.
- Hinrichs, Erhard (1986). 'Temporal anaphora in discourses of english'. *Linguistics and Philosophy*, 9(1), 63–82. doi: 10.1007/BF00627435.
- Hirschman, Lynette and Nancy Chinchor (1997). 'MUC-7 coreference task definition'. In *Proceedings of the 7th Message Understanding Conference (MUC-7)*.
- Hoare, C.A.R. (1969). 'An axiomatic basis for computer programming'. *Communications of the ACM*, 12(10), 576–583. doi: 10.1145/363235.363259.
- Hobbs, Jerry R. and Stuart M. Shieber (1987). 'An algorithm for generating quantifier scopings'. *Computational Linguistics*, 13(1-2), 47–63. ISSN 0891-2017.
- Kadmon, Nirit (2001). *Formal Pragmatics: Semantics, Pragmatics, Presupposition and Focus*. Blackwell, Malden, MA and Oxford.
- Kamp, Hans (1981). 'A theory of truth and semantic representation'. In Groenendijk, Jeroen, Theo M.V. Janssen, and Martin Stokhof, editors, *Formal Methods in the Study of Language*, 277–322. Mathematische Centrum, Amsterdam. Reprinted in Portner & Partee (eds.) 2002, Ch. 8.
- Kamp, Hans, Josef van Genabith, and Uwe Reyle (2011). 'Discourse Representation Theory'. In *Handbook of Philosophical Logic*, volume 15, 125–394. Springer, second edition. doi: 10.1007/978-94-007-0485-5.
- Kamp, Hans and Uwe Reyle (1993). *From Discourse to Logic*. Springer, Dordrecht.
- Kamp, J.A.W. (1968). *Tense Logic and the Theory of Linear Order*. Ph.D. thesis, University of California, Los Angeles.
- Karttunen, Lauri (1969). 'Discourse referents'. In *Proceedings of the 1969 conference on Computational linguistics*, 1–38. Association for Computational Linguistics.
- Keenan, Edward L. and Jonathan Stavi (1986). 'A semantic characterization of natural language determiners'. *Linguistics and philosophy*, 9(3), 253–326. ISSN 0165-0157.
- Kratzer, Angelika (1981). 'The notional category of modality'. In Eikmeyer, Hans-Jürgen and Hannes Rieser, editors, *Words, Worlds, and Contexts*, 38–74. de Gruyter, Berlin, New York.
- Krifka, M. (1995). 'The semantics and pragmatics of polarity items'. *Linguistic Analysis*, 25(3-4), 209–257.
- Ladusaw, William A. (1980). 'On the notion affective in the analysis of negative-polarity items'. *Journal of Linguistic Research*, 1(2), 1–16.
- Ladusaw, William A. (1997). 'Negation and polarity items'. In Lappin, Shalom, editor, *The Handbook of Contemporary Semantic Theory*, 321–341. Blackwell, Oxford.
- Lakoff, George (1970). 'Linguistics and natural logic'. *Synthese*, 22(1), 151–271. ISSN 0039-7857.
- Landauer, Thomas K., Peter .W. Foltz, and Darrell Laham (1998). 'An introduction to latent semantic analysis'. *Discourse processes*, 25(2), 259–284. ISSN 0163-853X.
- MacCartney, Bill (2009). *Natural language inference*. Ph.D. thesis, Citeseer.

- MacCartney, Bill and Christopher D. Manning (2007). 'Natural logic for textual inference'. In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, 193–200. Association for Computational Linguistics.
- Mitkov, Ruslan (to appear). 'Coreference resolution'. In Mitkov, Ruslan, editor, *Oxford Handbook of Computational Linguistics*. Oxford University Press, Oxford.
- Moens, Marc and Mark Steedman (1988). 'Temporal ontology and temporal reference'. *Computational Linguistics*, 14(2), 15–28.
- Montague, Richard (1970a). 'English as a formal language'. In Bruno, Visentini, editor, *Linguaggi nella societ a nella tecnica*, 189–223. Edizioni di Comunit a, Milan.
- Montague, Richard (1970b). 'Universal grammar'. *Theoria*, 36(3), 373–398. doi: 10.1111/j.1755-2567.1970.tb00434.x.
- Montague, Richard (1973). 'The proper treatment of quantification in ordinary english'. In Hintikka, Jaakko, Julius M.E. Moravcsik, and Patrick Suppes, editors, *Approaches to Natural Language*, 221–242. Reidel, Dordrecht. Reprinted in Portner & Partee (eds.) 2002, Ch. 1.
- Muskens, Reinhard (1996). 'Combining Montague semantics and discourse representation'. *Linguistics and Philosophy*, 19(2), 143–186. doi: 10.1007/BF00635836.
- Muskens, Reinhard (2010). 'An Analytic Tableau System for Natural Logic'. *Logic, Language and Meaning*, 104–113.
- Nairn, Rowan, Cleo Condoravdi, and Lauri Karttunen (2006). 'Computing relative polarity for textual inference'. In *Proceedings of Inference in Computational Semantics (ICoS) 5*.
- NIST (2007). 'Automatic content extraction & entity translation (ACE) evaluation plan'. Ms.
- Pad o, Sebastian and Ido Dagan (to appear). 'Textual entailment'. In Mitkov, Ruslan, editor, *Oxford Handbook of Computational Linguistics*. Oxford University Press, Oxford.
- Parsons, Terence (1985). 'Underlying events in the logical analysis of english'. In LePore, Ernest and Brian McLaughlin, editors, *Actions and Events: Perspectives on the Philosophy of Donald Davidson*. Blackwell, Oxford.
- Parsons, Terence (1990). *Events in the Semantics of English: A Study in Subatomic Semantics*, volume 19 of *Current Studies in Linguistics*. MIT Press.
- Partee, Barbara (1973). 'Some structural analogies between tenses and pronouns in english'. *The Journal of Philosophy*, 70(18), 601–609.
- Partee, Barbara (1984). 'Nominal and temporal anaphora'. *Linguistics and Philosophy*, 7(3), 243–286. doi: 10.1007/BF00627707.
- Polanyi, Livia and Martin van den Berg (1999). 'Logical structure and discourse anaphora resolution'. In *ACL99 Workshop on Discourse/Dialogue Structure and Reference*.
- Potts, Christopher (to appear). 'Pragmatics'. In Mitkov, Ruslan, editor, *Oxford Handbook of Computational Linguistics*. Oxford University Press, Oxford.
- Pratt, Vaughan R. (1976). 'Semantical consideration on Floyo-Hoare logic'. In *17th Annual Symposium on Foundations of Computer Science (FOCS 1976)*, 109–121. doi: 10.1109/SFCS.1976.27.

- Prior, Arthur N. (1957). *Time and Modality*. Clarendon Press, Oxford. Republished in 2003, New York: Oxford University Press.
- Prior, Arthur N. (1967). *Past, Present, and Future*. Oxford University Press, Oxford.
- Reichenbach, Hans (1947). *Elements of Symbolic Logic*. Macmillan, New York. Republished in 1980, New York: Dover.
- Roberts, Craige (1998). 'The place of centering in a general theory of anaphora resolution'. *Centering theory in discourse*, 359–399.
- Russell, Bertrand (1905). 'On denoting'. *Mind*, 14(4), 479–493. doi: 10.1093/mind/XIV.4.479.
- van der Sandt, Rob A. (1992). 'Presupposition projection as anaphora resolution'. *Journal of Semantics*, 9(4), 333–377. doi: 10.1093/jos/9.4.333.
- Schubert, Lenhart K., Benjamin Van Durme, and Marzieh Bazrafshan (2010). 'Entailment inference in a natural logic-like general reasoner'. In *Common Sense Knowledge Symposium (CSK-10)*, AAAI 2010 Fall Symposium.
- Shan, Chung-Chieh and Chris Barker (2006). 'Explaining crossover and superiority as left-to-right evaluation'. *Linguistics and Philosophy*, 29(1), 91–134. doi: 10.1007/s10988-005-6580-7.
- Smith, Carlota S. (1991). *The Parameter of Aspect*. Kluwer, Dordrecht. Republished in 1997.
- Stalnaker, Robert (1972). 'Pragmatics'. In Davidson, Donald and Gilbert Harman, editors, *Semantics of Natural Language*, 389–408. Reidel, Dordrecht. Reprinted in János Petöfi and Dorothea Franck (eds) 1973, *Präsuppositionen in Philosophie und Linguistik — Presuppositions in Philosophy and Linguistics*, Athanäum Verlag, Frankfurt. Page numbers in text refer to the 1973 version.
- Stalnaker, Robert (1979). 'Assertion'. In Cole, Peter and Jerry Morgan, editors, *Pragmatics*, volume 9 of *Syntax and Semantics*, 315–332. Academic Press, New York.
- Stalnaker, Robert (2002). 'Common ground'. *Linguistics and Philosophy*, 25(5), 701–721. ISSN 0165-0157.
- Strawson, P.F. (1950). 'On referring'. *Mind*, 29(235), 320–344. doi: 10.1093/mind/LIX.235.320.
- Tarski, Alfred (1944). 'The semantic conception of truth and the foundations of semantics'. *Philosophy and Phenomenological Research*, 4(3), 341–376.
- Valencia, V.M.S. (1991). *Studies on natural logic and categorial grammar*. VM Sanchez Valencia.
- Veltman, Frank (1996). 'Defaults in update semantics'. *Journal of Philosophical Logic*, 25(3), 221–261. doi: 10.1007/BF00248150.
- Vendler, Zeno (1957). 'Verbs and times'. *The Philosophical Review*, 66(2), 143–160. doi: 10.2307/2182371.
- Vendler, Zeno (1967). *Linguistics in Philosophy*. Cornell University Press, Ithaca.
- Verhagen, Marc, Robert Gaizauskas, Frank Schilder, Mark Hepple, Jessica Moszkowicz, and James Pustejovsky (2009). 'The TempEval challenge: Identifying temporal relations in text'. *Language Resources and Evaluation*, 43(2), 161–179. doi: 10.1007/s10579-009-9086-z.
- Von Stechow, P. (1999). 'NPI licensing, Strawson entailment, and context dependency'. *Journal of Semantics*, 16(2), 97. ISSN 0167-5133.

- Westerståhl, Dag (2007). 'Quantifiers in formal and natural languages'. *Handbook of philosophical logic*, 223–338.
- Wilson, T., J. Wiebe, and P. Hoffmann (2005). 'Recognizing contextual polarity in phrase-level sentiment analysis'. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 347–354. Association for Computational Linguistics.
- Zwarts, Frans (1995). 'Nonveridical contexts'. *Linguistic Analysis*, 25, 286–312.