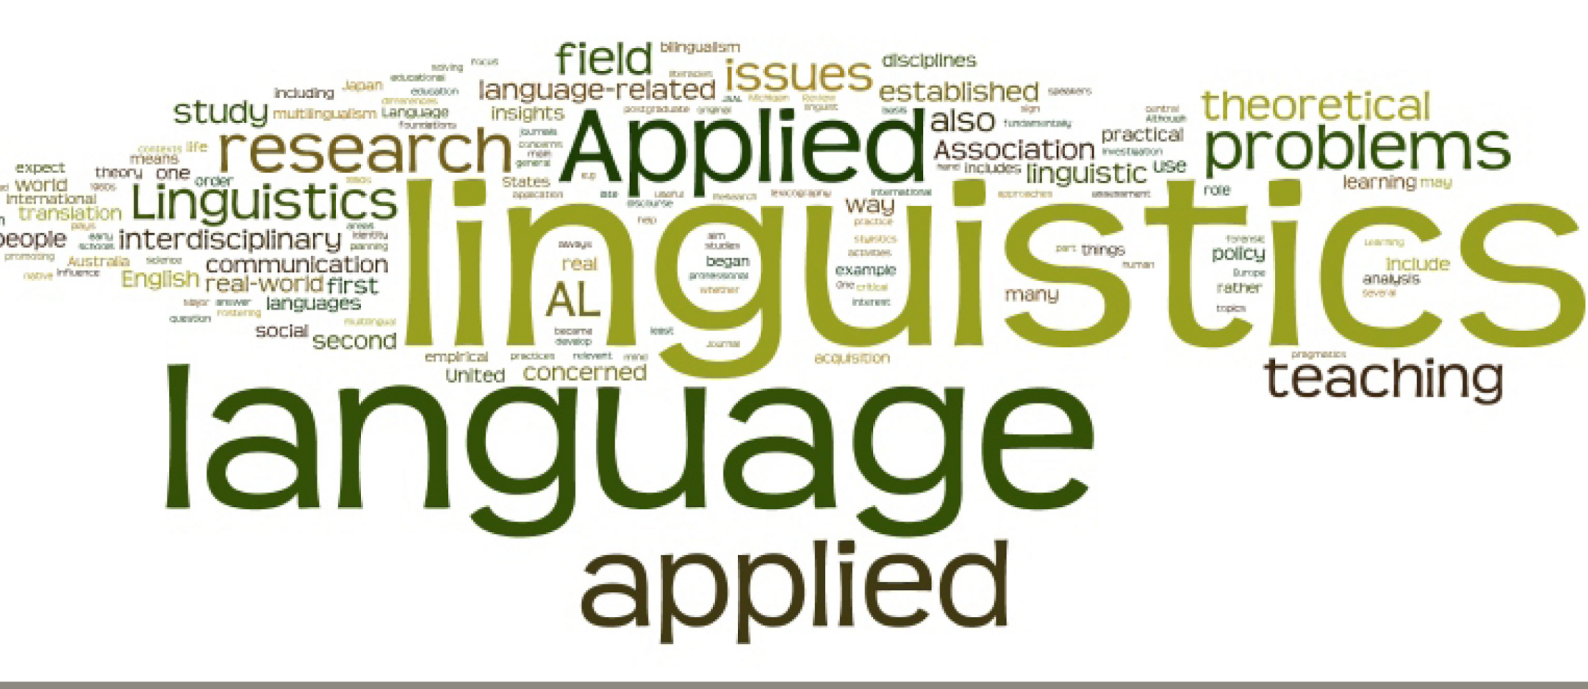

Semantic Theory of Linguistics and Literature

Mario Booth

SEMANTIC THEORY OF LINGUISTICS AND LITERATURE

SEMANTIC THEORY OF LINGUISTICS AND LITERATURE

Mario Booth

BIBLIOTEX
Digital Library

Semantic Theory of Linguistics and Literature
by Mario Booth

www.bibliotex.com

Ebook ISBN: 9781984665294

Published by:
Bibliotex
Canada
Website: www.bibliotex.com

Contents

Chapter 1

Formal Semantics

Formal semantics is the study of grammatical meaning in natural languages using formal tools from logic and theoretical computer science. It is an interdisciplinary field, sometimes regarded as a subfield of both linguistics and philosophy of language. It provides accounts of what linguistic expressions mean and how their meanings are composed from the meanings of their parts. The enterprise of formal semantics can be thought of as that of reverse-engineering the semantic components of natural languages' grammars.

Overview

Formal semantics studies the denotations of natural language expressions. High-level concerns include compositionality, reference, and the nature of meaning. Key topic areas include scope, modality, binding, tense, and aspect. Semantics is distinct from pragmatics, which encompasses aspects of meaning which arise from interaction and communicative intent.

Formal semantics is an interdisciplinary field, often viewed as a subfield of both linguistics and philosophy, while also incorporating work from computer science, mathematical logic, and cognitive psychology. Within philosophy, formal semanticists typically adopt a Platonistic ontology and an externalist view of meaning. Within linguistics, it is more common to view formal semantics as part of the study of

linguistic cognition. As a result, philosophers put more of an emphasis on conceptual issues while linguists are more likely to focus on the syntax-semantics interface and crosslinguistic variation.

Central concepts

Truth conditions

The fundamental question of formal semantics is what you know when you know how to interpret expressions of a language. A common assumption is that knowing the meaning of a sentence requires knowing its truth conditions, or in other words knowing what the world would have to be like for the sentence to be true. For instance, to know the meaning of the English sentence "Nancy smokes" one has to know that it is true when the person Nancy performs the action of smoking.

However, many current approaches to formal semantics posit that there is more to meaning than truth-conditions. In the formal semantic framework of inquisitive semantics, knowing the meaning of a sentence also requires knowing what issues (i.e. questions) it raises.

For instance "Nancy smokes, but does she drink?" conveys the same truth-conditional information as the previous example but also raises an issue of whether Nancy drinks. Other approaches generalize the concept of truth conditionality or treat it as epiphenomenal. For instance in dynamic semantics, knowing the meaning of a sentence amounts to knowing how it updates a context. Pietroski treats meanings as instructions to build concepts.

Compositionality

The Principle of Compositionality is the fundamental assumption in formal semantics. This principle states that the denotation of a complex expression is determined by the denotations of its parts along with their mode of composition. For instance, the denotation of the English sentence "Nancy smokes" is determined by the meaning of "Nancy", the denotation of "smokes", and whatever semantic operations combine the meanings of subjects with the meanings of predicates.

In a simplified semantic analysis, this idea would be formalized by positing that "Nancy" denotes Nancy herself, while "smokes" denotes a function which takes some individual x as an argument and returns the truth value "true" if x indeed smokes. Assuming that the words "Nancy" and "smokes" are semantically composed via function application, this analysis would predict that the sentence as a whole is true if Nancy indeed smokes.

Phenomena

Scope

In formal semantics, the **scope** of a semantic operator is the semantic object to which it applies. For instance, in the sentence "*Paulina doesn't drink beer but she does drink wine*," the proposition that Paulina drinks beer occurs within the scope of negation, but the proposition that Paulina drinks wine does not. Scope can be thought of as the semantic order of

operations. One of the major concerns of research in formal semantics is the relationship between operators' syntactic positions and their semantic scope.

This relationship is not transparent, since the scope of an operator need not directly correspond to its surface position and a single surface form can be semantically ambiguous between different scope construals. Some theories of scope posit a level of syntactic structure called logical form, in which an item's syntactic position corresponds to its semantic scope. Others theories compute scope relations in the semantics itself, using formal tools such as type shifters, monads, and continuations.

Phenomena

Scope ambiguity

The scope of an operator need not correspond directly to the word order of the sentence it occurs in. For instance, some sentences display a *scope ambiguity* in that the relative scopes of two operators can be construed in multiple ways.

- Every hedgehog is friends with a giraffe.

This sentence can be understood in two ways. On the *inverse scope* reading, there is a single giraffe who is very popular in the hedgehog community. On the *surface scope* reading, the sentence can be true even if the hedgehogs are friends with different giraffes.

Split scope

Split scope is the phenomenon where different components of an expressions item's meaning take scope in different places. Negative quantifiers are one category of expression which have been argued to take split scope.

- The company need fire no employees.

On the de re (non-split) reading, this sentence means that there is no employee such that the company needs to fire that employee. This is a non-split scope reading since "no" simply takes scope above the modal "need". On the split scope reading of this sentence, it means that it is not the case that the company needs to fire any employees. On this reading, "no" decomposes into a negation scoping above "need" and an existential quantifier scoping below it.

Indefinites have been argued to have split scope, having separate *existential scope* and *distributive scope*. This fact can be seen in the following example:

- If three relatives of mine die, I will inherit a house.

Among this sentence's reading is one which means "There exists a set of three relatives such that, if those three relatives die, I will inherit a house." On this reading, the indefinite "three relatives of mine" takes existential scope outside the conditional–– it asserts unconditionally that those three relatives do in fact exist. However, it the indefinite takes distributive scope inside the conditional–– the speaker will inherit a house if three relatives die, not if x dies where x can be any of those three relatives.

Definite descriptions have also been argued to have split scope. Definitesare classically considered to presuppose that their referents are unique. For instance, the definite description "the cat" is infelicitous in a context where there are multiple cats which the speaker could have in mind. However, this generalization seems to be contradicted by *Haddock descriptions* such as the following:

- Context: In front of the speaker are numerous hats, one of which contains a rabbit.
- Haddock description: The rabbit in the hat

This noun phrase is felicitous to use in this context, even though there is no unique hat. What seems to license this surprising use of the definite description is the fact that the context contains a unique rabbit-containing hat. To cash out this idea, it has been proposed that the uniqueness presupposition of "the hat" takes scope separately from the rest of the definite's meaning. In other words, a witness set is establishes low in the structure, but is checked for singletonness higher up.

Scope islands

While operators can often take scope above their surface position, there are not entirely free to take scope wherever they want. For instance, as illustrated by Sentence 1 below, quantifiers that originate inside an if-clause usually cannot take scope outside of that "if"-clause. This sentence cannot mean that Beth will inherit one house for each dead relative.

- If every relative of mine dies, I will inherit a house.

This fact parallels the fact that a wh-phrase cannot be extracted from an "if"-clause, as shown in Sentence 2.

- Which relative$_i$ will you inherit a fortune if t_i dies?

Examples of this sort have been used to argue that scope relations are determined by syntactic movement operations.

Aside from their theoretical significance, scope islands are also practically useful since they can be used to construct unambiguous paraphrases of sentences with scope ambiguities.

Exceptional scope

While most operators are unable to scope out of an island, others can.

For instance, the indefinite "a" in the sentence below can take scope outside of its surface position inside an "if"-clause. This sentence can mean that there is a particular relative who must die for the speaker to get a house.

- If a relative of mine dies, I will inherit a house.

Examples of this sort have been used to argue that indefinites do not have standard generalized quantifier denotations. On the choice function approach proposed by Tanya Reinhart, indefinites contribute a variable over choice functions which can be existentially closed at any point higher in the structure. Angelika Kratzer proposed another choice function-based theory, which is similar to Reinhart's except that the choice function variable is leftfree. Recent work such as Charlow

(2020) treats indefinites as denoting sets of individuals which can be type shifted so that they take scope in a manner similar to Karttunen's (1977) alternative-based mechanism for wh-questions.

Formal approaches to scope

The *structural view of scope* is one influential view which posits a close relationship between syntax and semantics. This approach is characterized by the following hypothesis, first formulated by Tanya Reinhart:

- **Hypothesis about scope and domain**: The semantic scope of an operator corresponds to the position of the item which expresses it at some level of syntactic representation.

This view is widely adopted in generative approaches such as that of Heim and Kratzer (1998). In these approaches, the relevant syntactic level is logical form and the syntactic notion which corresponds to semantic scope is typically identified as c-command.

In structural approaches, discrepancies between an expression's surface position and its semantic scope are explained by syntactic movement operations such as quantifier raising. The movement approach is motivated in large part by the fact that quantifier scope seems to obey many of the same restrictions that movement does, e.g. islands.

One prominent alternative to the structural view is the *type shifting view* first proposed by Barbara Partee and Mats Rooth.

This approach uses type shifters to govern scopal relations. Since type shifters are applied during the process of semantic interpretation, this approach allows scopal relations to be partly independent of syntactic structure. The type shifting approach serves as the basis of many recent proposals for exceptional scope, split scope, and other troublesome scope-related phenomena.

Binding (linguistics)

In linguistics, **binding** is the phenomenon in which anaphoric elements such as pronouns are grammatically associated with their antecedents. For instance in the English sentence "Mary saw herself", the anaphor "herself" is bound by its antecedent "Mary". Binding can be licensed or blocked in certain contexts or syntactic configurations, e.g. the pronoun "her" cannot be bound by "Mary" in the English sentence "Mary saw her". While all languages have binding, restrictions on it vary even among closely related languages. Binding has been a major area of research in syntax and semantics since the 1970s, and was a major for the government and binding theory paradigm.

Some basic examples and questions

The following sentences illustrate some basic facts of binding. The words that bear the index ishould be construed as referring to the same person or thing.

- a. **Fred**$_{i}$ is impressed with **himself**$_{i}$. – Indicated reading obligatory

- b. ***Fred**$_i$ is impressed with **him**$_i$. – Indicated reading impossible
- a. ***Susan**$_i$ asked Arthur to help **herself**$_i$. – Indicated reading impossible, sentence ungrammatical
- b. **Susan**$_i$ asked Arthur to help **her**$_i$. – Indicated reading easily possible
- a. **Sue**$_i$ said **she**$_i$ was tired. – Indicated reading easily possible
- b. ***She**$_i$ said **Sue**$_i$ was tired. – Indicated reading impossible
- a. **Fred's**$_i$ friends venerate **him**$_i$. – Indicated reading easily possible
- b. **His**$_i$ friends venerate **Fred**$_i$. – Indicated reading unlikely

These sentences illustrate some aspects of the distribution of reflexive and personal pronouns. In the first pair of sentences, the reflexive pronoun must appear for the indicated reading to be possible. In the second pair, the personal pronoun must appear for the indicated reading to be possible.

The third pair shows that at times a personal pronoun must follow its antecedent, and the fourth pair further illustrates the same point, although the acceptability judgement is not as robust. Based on such data, one sees that reflexive and personal pronouns differ in their distribution and that linear order (of a pronoun in relation to its antecedent or postcedent) is a factor influencing where at least some pronouns can appear.

A theory of binding should be in a position to predict and explain the differences in distribution seen in sentences like these. It should be in a position to answer questions like: What explains where a reflexive pronoun must appear as opposed to

a personal pronoun? When does linear order play a role in determining where pronouns can appear? What other factor (or factors) beyond linear order help predict where pronouns can appear?

Binding domains

The following three subsections consider the binding domains that are relevant for the distribution of pronouns and nouns in English. The discussion follows the outline provided by the traditional binding theory (see below), which divides nominals into three basic categories: reflexive and reciprocal pronouns, personal pronouns, and nouns (common and proper).

Reflexive and reciprocal pronouns ("anaphors")

When one examines the distribution of reflexive pronouns and reciprocal pronouns (which are often subsumed under the general category of "anaphor"), one sees that there are certain domains that are relevant, a "domain" being a syntactic unit that is clause-like. Reflexive and reciprocal pronouns often seek their antecedent close by, in a binding domain that is local, e.g.

- a. **Fred**$_i$ praises **himself**$_i$. – Indicated reading obligatory
- b. ***Fred**$_i$ praises **him**$_i$. – Indicated reading impossible
- a. **The girls**$_i$ like **each other**$_i$. – Indicated reading obligatory
- b. ***Thegirls**$_i$ like **them**$_i$. – Indicated reading impossible

These examples illustrate that there is a domain within which a reflexive or reciprocal pronoun should find its antecedent.

The a-sentences are fine because the reflexive or reciprocal pronoun has its antecedent within the clause. The b-sentences, in contrast, do not allow the indicated reading, a fact illustrating that personal pronouns have a distribution that is different from that of reflexive and reciprocal pronouns. A related observation is that a reflexive and reciprocal pronoun often cannot seek its antecedent in a superordinate clause, e.g.

- a. Susan thinks that **Jill**$_i$ should praise **herself**$_i$. – Indicated reading (almost) obligatory
- b. **Susan**$_i$ thinks that Jill should praise **herself**$_i$. – Indicated reading very unlikely
- a. They asked whether **the girls**$_i$ like **each other**$_i$. – Indicated reading (almost) obligatory
- b. **They**$_i$ asked whether the girls like **each other**$_i$. – Indicated reading very unlikely

When the reflexive or reciprocal pronoun attempts to find an antecedent outside of the immediate clause containing it, it fails. In other words, it can hardly seek its antecedent in the superordinate clause. The binding domain that is relevant is the immediate clause containing it.

Personal pronouns

Personal pronouns have a distribution that is different from reflexive and reciprocal pronouns, a point that is evident with the first two b-sentences in the previous section. The local binding domain that is decisive for the distribution of reflexive and reciprocal pronouns is also decisive for personal pronouns, but in a different way. Personal pronouns seek their antecedent outside of the local binding domain containing them, e.g.

- a. **Fred**$_i$ asked whether Jim mentioned **him**$_i$. – Indicated reading easily possible
- b. *Fred asked whether **Jim**$_i$ mentioned **him**$_i$. – Indicated reading impossible
- a. **Gina**$_i$ hopes that Wilma will mention **her**$_i$. – Indicated reading easily possible
- b. *Gina hopes that **Wilma**$_i$ will mention **her**$_i$. – Indicated reading impossible

In these cases, the pronoun has to look outside of the embedded clause containing it to the matrix clause to find its antecedent. Hence based on such data, the relevant binding domain appears to be the clause. Further data illustrate, however, that the clause is actually not the relevant domain:

- a. **Fred**$_i$ likes the picture of **him**$_i$. – Indicated reading possible
- b. **Gina**$_i$ has heard the rumor about **her**$_i$. – Indicated reading possible

Since the pronouns appear within the same minimal clause containing their antecedents in these cases, one cannot argue that the relevant binding domain is the clause. The most one can say based on such data is that the domain is "clause-like".

Nouns

The distribution of common and proper nouns is unlike that of reflexive, reciprocal, and personal pronouns. The relevant observation in this regard is that a noun is often reluctantly coreferential with another nominal that is within its binding domain or in a superordinate binding domain, e.g.

- a. **Susan**$_i$ admires **herself**$_i$. – Indicated reading obligatory

- b. #**Susan**$_i$ admires **Susan**$_i$. – Indicated reading possible, but special context necessary
- a. **Fred**$_i$ thinks that **he**$_i$ is the best. – Indicated reading easily possible
- b. #**Fred**$_i$ thinks that **Fred**$_i$ is the best. – Indicated reading possible, but special context necessary

The readings indicated in the a-sentences are natural, whereas the b-sentences are very unusual. Indeed, sentences like these b-sentences were judged to be impossible in the traditional binding theory according to Condition C (see below).

Given a contrastive context, however, the b-sentences can work, e.g. *Susan does not admire Jane, but rather Susan*$_i$ *admires Susan*$_i$. One can therefore conclude that nouns are not sensitive to binding domains in the same way that reflexive, reciprocal, and personal pronouns are.

Linear order

The following subsections illustrate the extent to which pure linear order impacts the distribution of pronouns. While linear order is clearly important, it is not the only factor influencing where pronouns can appear.

Linear order is a factor

A simple hypothesis concerning the distribution of many anaphoric elements, of personal pronouns in particular, is that linear order plays a role. In most cases, a pronoun follows its antecedent, and in many cases, the coreferential reading is impossible if the pronoun precedes its antecedent. The

following sentences suggest that pure linear can indeed be important for the distribution of pronouns:

- a. **Jim's**$_i$ grade upsets **him**$_i$. – Indicated reading easily possible
- b. **His**$_i$ grade upsets **Jim**$_i$. – Indicated reading unlikely
- a. **Larry's**$_i$ family avoids **him**$_i$. – Indicated reading easily possible
- b. **His**$_i$ family avoids **Larry**$_i$. – Indicated reading unlikely
- a. We spoke to **Tina's**$_i$ mother about **her**$_i$. – Indicated reading easily possible
- b. We spoke to **her**$_i$ mother about **Tina**$_i$. – Indicated reading unlikely

While the coreferential readings indicated in these b-sentences are possible, they are unlikely. The order presented in the a-sentences is strongly preferred. The following, more extensive data sets further illustrate that linear order is important:

- a. **Sam**$_i$ mentioned twice that **he**$_i$ was hungry. – Indicated reading easily possible
- b. ***He**$_i$ mentioned twice that **Sam**$_i$ was hungry. – Indicated reading impossible
- c. That **Sam**$_i$ was hungry, **he**$_i$ mentioned twice. – Indicated reading possible
- d. That **he**$_i$ was hungry, **Sam**$_i$ mentioned twice. – Indicated reading unlikely
- a. You asked **Fred**$_i$ twice when **he**$_i$ would study. – Indicated reading easily possible
- b. *You asked **him**$_i$ twice when **Fred**$_i$ would study. – Indicated reading impossible
- c. When **Fred**$_i$ would study, you asked **him**$_i$ twice. – Indicated reading possible

- d. When **he**$_i$ would study, you asked **Fred**$_i$ twice. – Indicated reading unlikely

While the acceptability judgements here are nuanced, one can make a strong case that pure linear order is at least in part predictive of when the indicated reading is available. The a- and c-sentences allow the coreferential reading more easily than their b- and d-counterparts.

Linear order is not the only factor

While linear order is an important factor influencing the distribution of pronouns, it is not the only factor. The following sentences are similar to the c- and d-sentences in the previous section insofar as an embedded clause is present.

- a. When **the boys**$_i$ are at home, **they**$_i$ play video games. – Indicated reading easily possible
- b. When **they**$_i$ are at home, **the boys**$_i$ play video games. – Indicated reading possible
- a. If **Susan**$_i$ tries, **she**$_i$ will succeed. – Indicated reading easily possible
- b. If **she**$_i$ tries, **Susan**$_i$ will succeed. – Indicated reading possible

While there may be a mild preference for the order in the a-sentences here, the indicated reading in the b-sentences is also available. Hence linear order is hardly playing a role in such cases. The relevant difference between these sentences and the c- and d-sentences in the previous section is that the embedded clauses here are adjunct clauses, whereas they are argument clauses above. The following examples involve adjunct phrases:

- a. **Rosa**$_i$ found a scratch in Ben's picture of **her**$_i$. – Indicated reading easily possible
- b. ***She**$_i$ found a scratch in Ben's picture of **Rosa**$_i$. – Indicated reading impossible
- c. In Ben's picture of **Rosa**$_i$, **she**$_i$ found a scratch. – Indicated reading unlikely
- d. In Ben's picture of **her**$_i$, **Rosa**$_i$ found a scratch. – Indicated reading possible
- a. **Zelda**$_i$ spent her sweetest hours in **her**$_i$ bed.– Indicated reading easily possible
- b. ***She**$_i$ spent her sweetest hours in **Zelda's**$_i$ bed. – Indicated reading impossible
- c. In **Zelda's**$_i$ bed, **she**$_i$ spent her sweetest hours. – Indicated reading very unlikely
- d. In **her**$_i$ bed, **Zelda**$_i$ spent her sweetest hours. – Indicated reading possible

The fact that the c-sentences marginally allow the indicated reading whereas the b-sentences do not at all allow this reading further demonstrates that linear order is important. But in this regard, the d-sentences are telling, since if linear order were the entire story, one would expect the d-sentences to be less acceptable than they are. The conclusion that one can draw from such data is that there are one or more other factors beyond linear order that are impacting the distribution of pronouns.

Configuration vs. function

Given that linear order is not the only factor influencing the distribution of pronouns, the question is what other factor or factors might also be playing a role. The traditional binding theory (see below) took c-command to be the all important

factor, but the importance of c-command for syntactic theorizing has been extensively criticized in recent years. The primary alternative to c-command is functional rank. These two competing concepts (c-command vs. rank) have been debated extensively and they continue to be debated. C-command is a configurational notion; it is defined over concrete syntactic configurations. Syntactic rank, in contrast, is a functional notion that resides in the lexicon; it is defined over the ranking of the arguments of predicates. Subjects are ranked higher than objects, first objects are ranked higher than second objects, and prepositional objects are ranked lowest. The following two subsections briefly consider these competing notions.

Configuration (c-command)

C-command is a configurational notion that acknowledges the syntactic configuration as primitive. Basic subject-object asymmetries, which are numerous in many languages, are explained by the fact that the subject appears outside of the finite verb phrase (VP) constituent, whereas the object appears inside it. Subjects therefore c-command objects, but not vice versa. C-command is defined as follows:

- **C-command**
- Node A c-commands node B if every node dominating A also dominates B, and neither A nor B dominates the other.

Given the binary division of the clause (S → NP + VP) associated with most phrase structure grammars, this definition sees a typical subject c-commanding everything

inside the verb phrase (VP), whereas everything inside the VP is incapable of c-commanding anything outside of the VP. Some basic binding facts are explained in this manner, e.g.

- a. **Larry**$_i$ promoted **himself**$_i$. – Indicated reading obligatory
- b. ***Himself**$_i$ promoted **Larry**$_i$. – Indicated reading impossible; sentence ungrammatical

Sentence a is fine because the subject *Larry* c-commands the object *himself*, whereas sentence b does not work because the object *Larry* does not c-command the subject *himself*. The assumption has been that within its binding domain, a reflexive pronoun must be c-commanded by its antecedent. While this approach based on c-command makes a correct prediction much of the time, there are other cases where it fails to make the correct prediction, e.g.

- The picture of **himself**$_i$ upsets **Larry**$_i$. – Indicated reading possible

The reading indicated is acceptable in this case, but if c-command were the key notion helping to explain where the reflexive can and must appear, then the reading should be impossible since *himself* is not c-commanded by *Larry*.

As reflexive and personal pronouns occur in complementary distribution, the notion of c-command can also be used to explain where personal pronouns can appear. The assumption is that personal pronouns *cannot* c-command their antecedent, e.g.

- a. When **Alice**$_i$ felt tired, **she**$_i$ lay down. – Indicated reading easily possible

- b When **she**$_i$ felt tired, **Alice**$_i$ lay down. – Indicated reading possible

In both examples, the personal pronoun *she* does not c-command its antecedent *Alice*, resulting in the grammaticality of both sentences despite reversed linear order.

Function (rank)

The alternative to a c-command approach posits a ranking of syntactic functions (SUBJECT > FIRST OBJECT > SECOND OBJECT > PREPOSITIONAL OBJECT). Subject-object asymmetries are addressed in terms of this ranking. Since subjects are ranked higher than objects, an object can have the subject as its antecedent, but not vice versa. With basic cases, this approach makes the same prediction as the c-command approach. The first two sentences from the previous section are repeated here:

- a. **Larry**$_i$ promoted **himself**$_i$. – Indicated reading obligatory
- b. ***Himself**$_i$ promoted **Larry**$_i$. – Indicated reading impossible; sentence ungrammatical

Since the subject outranks the object, sentence a is predictably acceptable, the subject *Larry* outranking the object *himself*. Sentence b, in contrast, is bad because the subject reflexive pronoun *himself* outranks its postcedent*Larry*. In other words, this approach in terms of rank is assuming that within its binding domain, a reflexive pronoun may not outrank its antecedent (or postcedent). Consider the third example sentence from the previous section in this regard:

- The picture of **himself**$_i$ upset **Larry**$_i$. – Indicated reading possible

The approach based on rank does not require a particular configurational relationship to hold between a reflexive pronoun and its antecedent.

In other words, it makes no prediction in this case, and hence does not make an incorrect prediction. The reflexive pronoun *himself*is embedded within the subject noun phrase, which means that it is not the subject and hence does not outrank the object *Larry*.

A theory of binding that acknowledges both linear order and rank can at least begin to predict many of the marginal readings. When both linear order and rank combine, acceptability judgments are robust, e.g.

- a. **Barbara**$_i$ hopes that **she**$_i$will be promoted. – Linear order and rank combine to make the indicated reading easily possible.
- b. ***She**$_i$ hopes that **Barbara**$_i$ will be promoted. – Linear order and rank combine to make the indicated reading impossible.
- a. **Bill's**$_i$ grade upset **him**$_i$. – Linear order alone makes the indicated reading possible; rank is not involved.
- b. **His**$_i$ grade upset **Bill**$_i$. – Linear order alone makes the indicated reading unlikely; rank is not involved.

This ability to address marginal readings is something that an approach combining linear order and rank can accomplish, whereas an approach that acknowledges only c-command cannot do the same.

The traditional binding theory: Conditions A, B, and C

The exploration of binding phenomena got started in the 1970s and interest peaked in the 1980s with Government and Binding Theory, a grammar framework in the tradition of generative syntax that is still prominent today. The theory of binding that became widespread at that time serves now merely as reference point (since it is no longer believed to be correct).

This theory distinguishes between 3 different binding conditions: A, B, and C. The theory classifies nominals according to two features, [±anaphor] and [±pronominal], which are binary. The binding characteristics of a nominal are determined by the values of these features, either plus or minus. Thus, a nominal that is [-anaphor, -pronominal] is an R-expression (referring expression), such as a common noun or a proper name. A nominal that is [-anaphor, +pronominal] is a pronoun, such as *he* or *they*, and a nominal that is [+anaphor, -pronominal] is a reflexive pronoun, such as *himself* or *themselves*. Note that the term *anaphor* here is being used in a specialized sense; it essentially means "reflexive". This meaning is specific to the Government and Binding framework and has not spread beyond this framework.

Based on the classifications according to these two features, three conditions are formulated:

- **Condition A**
- An anaphor (reflexive) must have a local (nearby) antecedent. Thus, $John_i$ *washed* $himself_i$ obeys

Condition A: the antecedent of *himself*, which is *John*, is nearby. In contrast, *$John_i$ *asked Mary to wash* *$himself_i$* is unacceptable, because the reflexive and its antecedent are too far away from each other.

- **Condition B**
- A pronoun can have an antecedent as long as the antecedent is not local or does not c-command the pronoun. Thus*$John_i$ asked Mary to wash him_i* obeys Condition B; *John* is the antecedent of *him*, and *him* is sufficiently far away. On the other hand, *$John_i$ *washed* *him_i* is unacceptable.
- **Condition C**
- An R-expression cannot have an antecedent that c-commands it. Thus *He_i *asked Mary to wash* *$John_i$* is unacceptable.

While the theory of binding that these three conditions represent is no longer held to be valid, as mentioned above, the associations with the three conditions are so firmly anchored in the study of binding that one often refers to, for example, "Condition A effects" or "Condition B effects" when describing binding phenomena.

Modality (natural language)

In linguistics and philosophy, modality is the phenomenon whereby language is used to discuss possible situations. For instance, a modal expression may convey that something is likely, desirable, or permissible. Quintessential modal expressions include modal auxiliaries such as "could", "should", or "must"; modal adverbs such as "possibly" or

"necessarily"; and modal adjectives such as "conceivable" or "probable". However, modal components have been identified in the meanings of countless natural language expressions including counterfactuals, propositional attitudes, evidentials, habituals, and generics.

Modality has been intensely studied from a variety of perspectives. Within linguistics, typological studies have traced crosslinguistic variation in the strategies used to mark modality, with a particular focus on its interaction with Tense–aspect–mood marking. Theoretical linguists have sought to analyze both the propositional content and discourse effects of modal expressions using formal tools derived from modal logic. Within philosophy, linguistic modality is often seen as a window into broader metaphysical notions of necessity and possibility.

Force and flavor

Modal expressions come in different categories called *flavors*. Flavors differ in how the possibilities they discuss relate to reality. For instance, an expression like "might" is said to have epistemic flavor, since it discusses possibilities compatible with some body of knowledge. An expression like "obligatory" is said to have deontic flavor, since it discusses possibilities which are required given the laws or norms obeyed in reality.

- Agatha *must* be the murderer. (expressing epistemic modality)
- Agatha *must* go to jail. (expressing deontic modality)

The sentence in (1) might be spoken by someone who has decided that all of the relevant facts in a particular murder investigation point to the conclusion that Agatha was the murderer, even though it may or may not actually be the case.

The 'must' in this sentence thus expresses epistemic modality, for 'for all we know', Agatha must be the murderer - where 'for all we know' is relative to some knowledge the speakers possess. In contrast, (2) might be spoken by someone who has decided that, according to some standard of conduct, Agatha has committed a vile crime, and therefore the correct course of action is to jail Agatha.

In classic formal approaches to **linguistic modality**, an utterance expressing modality is one that can always roughly be paraphrased to fit the following template:

- (1) According to [a set of rules, wishes, beliefs,...] it is [necessary, possible] that [the main proposition] is the case.

The set of propositions which forms the basis of evaluation is called the **modal base**. The result of the evaluation is called the **modal force**. For example, the utterance in (2) expresses that, according to what the speaker has observed, it is necessary to conclude that John has a rather high income:

- (2) John must be earning a lot of money.

The modal base here is the knowledge of the speaker, the modal force is necessity. By contrast, (3) could be paraphrased as ‘Given his abilities, the strength of his teeth, etc., it is possible for John to open a beer bottle with his teeth’. Here,

the modal base is defined by a subset of John's abilities, the modal force is possibility.

- (3) John can open a beer bottle with his teeth.

Formal semantics

Linguistic modality has been one of the central concerns in formal semantics and philosophical logic. Research in these fields has led to a variety of accounts of the propositional content and conventional discourse effects of modal expressions. The predominant approaches in these fields are based on modal logic. In these approaches, modal expressions such as *must* and *can*are analyzed as quantifiers over a set of possible worlds. In classical modal logic, this set is identified as the set of worlds accessible from the world of evaluation. Since the seminal work of Angelika Kratzer, formal semanticists have adopted a more finely grained notion of this set as determined by two *conversational background functions* called the *modal base* and *ordering source* respectively.

For an epistemic modal like English *must* or *might*, this set is understood to contain exactly those worlds compatible with the knowledge that the speaker has in the actual world. Assume for example that the speaker of sentence (2) above knows that John just bought a new luxury car and has rented a huge apartment. The speaker also knows that John is an honest person with a humble family background and doesn't play the lottery. The set of accessible worlds is then the set of worlds in which all these propositions which the speaker knows about John are true. The notions of **necessity** and **possibility**are then defined along the following lines: A proposition *P* follows

necessarily from the set of accessible worlds, if all accessible worlds are part of *P* (that is, if *p* is true in all of these worlds). Applied to the example in (2) this would mean that in all the worlds which are defined by the speaker's knowledge about John, it is the case that John earns a lot of money (assuming there is no other explanation for John's wealth). In a similar way a proposition *p* is possible according to the set of accessible worlds (i.e. the modal base), if some of these worlds are part of *P*.

Recent work has departed from this picture in a variety of ways. In dynamic semantics, modals are analyzed as *tests* which check whether their prejacent is compatible with (or follows from) the information in the conversational common ground. Probabilistic approaches motivated by gradable modal expressions provide a semantics which appeals to speaker credence in the prejacent. Illocutionary approaches assume a sparser view of modals' propositional content and look to conventional discourse effects to explain some of the nuances of modals' use.

Grammatical expression of modality

Verbal morphology

In many languages modal categories are expressed by verbal morphology – that is, by alterations in the form of the verb. If these verbal markers of modality are obligatory in a language, they are called*mood* markers. Well-known examples of moods in some European languages are referred to as subjunctive, conditional, and indicative as illustrated below with examples

from French, all three with the verb *avoir* 'to have'. As in most Standard European languages, the shape of the verb conveys not only information about modality, but also about other categories such as person and number of the subject.

An example for a non-European language with a similar encoding of modality is Manam. Here, a verb is prefixed by a morpheme which encodes number and person of the subject. These prefixes come in two versions, one *realis* version and one *irrealis* version. Which one is chosen depends on whether the verb refers to an actual past or present event (realis), or merely to a possible or imagined event (irrealis).

Auxiliaries

Modal auxiliary verbs, such as the English words *may, can, must, ought, will, shall, need, dare, might, could, would,* and *should*, are often used to express modality, especially in the Germanic languages.

Ability, desirability, permission, obligation, and probability can all be exemplified by the usage of auxiliary modal verbs in English:

- *Ability*: I **can** ride a bicycle (in the present); I **could** ride a bicycle (in the past)
- *Desirability*: I **should** go; I **ought** to go
- *Permission*: I **may** go
- *Obligation*: I **must** go
- *Likelihood*: He **might** be there; He **may** be there; He **must** be there

Lexical expression

Verbs such as "want," "need," or "belong" can be used to express modality lexically, as can adverbs.

- (5) It *belongs* in a museum!

Other

Complementizers (e.g. Russian) and conjunctions (e.g. Central Pomo) can be used to convey modality.

History

Formal semantics emerged as a major area of research in the early 1970s, with the pioneering work of the philosopher and logician Richard Montague. Montague proposed a formal system now known as Montague grammar which consisted of a novel syntactic formalism for English, a logical system called Intensional Logic, and a set of homomorphic translation rules linking the two. In retrospect, Montague Grammar has been compared to a Rube Goldberg machine, but it was regarded as earth-shattering when first proposed, and many of its fundamental insights survive in the various semantic models which have superseded it.

Montague Grammar was a major advance because it showed that natural languages could be treated as interpretedformal languages. Before Montague, many linguists had doubted that this was possible, and logicians of that era tended to view logic as a replacement for natural language rather than a tool for analyzing it. Montague's work was published during the

Linguistics Wars, and many linguists were initially puzzled by it. While linguists wanted a restrictive theory that could only model phenomena that occur in human languages, Montague sought a flexible framework that characterized the concept of meaning at its most general. At one conference, Montague told Barbara Partee that she was "the only linguist who it is not the case that I can't talk to".

Formal semantics grew into a major subfield of linguistics in the late 1970s and early 1980s, due to the seminal work of Barbara Partee. Partee developed a linguistically plausible system which incorporated the key insights of both Montague Grammar and Transformational grammar. Early research in linguistic formal semantics used Partee's system to achieve a wealth of empirical and conceptual results. Later work by Irene Heim, Angelika Kratzer, Tanya Reinhart, Robert May and others built on Partee's work to further reconcile it with the generative approach to syntax. The resulting framework is known as the *Heim and Kratzer* system, after the authors of the textbook *Semantics in Generative Grammar* which first codified and popularized it. The Heim and Kratzer system differs from earlier approaches in that it incorporates a level of syntactic representation called logical form which undergoes semantic interpretation. Thus, this system often includes syntactic representations and operations which were introduced by translation rules in Montague's system. However, work by others such as Gerald Gazdar proposed models of the syntax-semantics interface which stayed closer to Montague's, providing a system of interpretation in which denotations could be computed on the basis ofsurface structures. These approaches live on in frameworks such as categorial grammar and combinatory categorial grammar.

Cognitive semantics emerged as a reaction against formal semantics, but there have been recently several attempts at reconciling both positions.

Chapter 2

Cognitive Semantics

Cognitive semantics is part of the cognitive linguistics movement. Semantics is the study of linguistic meaning. Cognitive semantics holds that language is part of a more general human cognitive ability, and can therefore only describe the world as people conceive of it. It is implicit that different linguistic communities conceive of simple things and processes in the world differently (different cultures), not necessarily some difference between a person's conceptual world and the real world (wrong beliefs).

The main tenets of cognitive semantics are:

- That grammar manifests a conception of the world held in a culture;
- That knowledge of language is acquired and contextual;
- That the ability to use language draws upon general cognitive resources and not a special language module.

Cognitive semantics has introduced innovations like prototype theory, conceptual metaphors, and frame semantics, and it is the linguistic paradigm/framework that since the 1980s has generated the most studies in lexical semantics. As part of the field of cognitive linguistics, the cognitive semantics approach rejects the traditional separation of linguistics into phonology, morphology, syntax, pragmatics, etc. Instead, it divides semantics into *meaning-construction* and *knowledge*

representation. Therefore, cognitive semantics studies much of the area traditionally devoted to pragmatics as well as semantics.

The techniques native to cognitive semantics are typically used in lexical studies such as those put forth by Leonard Talmy, George Lakoff and Dirk Geeraerts. Some cognitive semantic frameworks, such as that developed by Talmy, take into account syntactic structures as well.

Points of contrast

As a field, semantics is interested in three big questions: what does it mean for units of language, called lexemes, to have "meaning"? What does it mean for sentences to have meaning? Finally, how is it that meaningful units fit together to compose complete sentences? These are the main points of inquiry behind studies into lexical semantics, structural semantics, and theories of compositionality (respectively). In each category, traditional theories seem to be at odds with those accounts provided by cognitive semanticists.

Classic theories in semantics (in the tradition of Alfred Tarski and Donald Davidson) have tended to explain the meaning of parts in terms of *necessary and sufficient conditions*, sentences in terms of *truth-conditions*, and composition in terms of *propositional functions*. Each of these positions is tightly related to the others. According to these traditional theories, the meaning of a particular sentence may be understood as the conditions under which the proposition conveyed by the sentence hold true. For instance, the expression "snow is white" is true if and only if snow is, in fact, white. Lexical

units can be understood as holding meaning either by virtue of set of things they may apply to (called the "extension" of the word), or in terms of the common properties that hold between these things (called its "intension"). The intension provides an interlocutor with the necessary and sufficient conditions that let a thing qualify as a member of some lexical unit's extension. Roughly, propositional functions are those abstract instructions that guide the interpreter in taking the free variables in an open sentence and filling them in, resulting in a correct understanding of the sentence as a whole.

Meanwhile, cognitive semantic theories are typically built on the argument that lexical meaning is conceptual. That is, meaning is not necessarily reference to the entity or relation in some real or possible world. Instead, meaning corresponds with a concept held in the mind based on personal understanding. As a result, semantic facts like "All bachelors are unmarried males" are not treated as special facts about our language practices; rather, these facts are not distinct from encyclopaedic knowledge.

In treating linguistic knowledge as being a piece with everyday knowledge, the question is raised: how can cognitive semantics explain paradigmatically semantic phenomena, like category structure? Set to the challenge, researchers have drawn upon theories from related fields, like cognitive psychology and cognitive anthropology. One proposal is to treat in order to explain category structure in terms of *nodes* in a *knowledge network*. One example of a theory from cognitive science that has made its way into the cognitive semantic mainstream is the theory of prototypes, which cognitive semanticists generally argue is the cause of polysemy.

Cognitive semanticists argue that truth-conditional semantics is unduly limited in its account of full sentence meaning. While they are not on the whole hostile to truth-conditional semantics, they point out that it has limited explanatory power. That is to say, it is limited to indicative sentences, and does not seem to offer any straightforward or intuitive way of treating (say) commands or expressions. By contrast, cognitive semantics seeks to capture the full range of grammatical moods by also making use of the notions of framing and mental spaces.

Another trait of cognitive semantics is the recognition that meaning is not fixed but a matter of construal and conventionalization. The processes of linguistic construal, it is argued, are the same psychological processes involved in the processing of encyclopaedic knowledge and in perception. This view has implications for the problem of compositionality. An account in cognitive semantics called the dynamic construal theory makes the claim that words themselves are without meaning: they have, at best, "default construals," which are really just ways of using words. Along these lines, cognitive semantics argues that compositionality can only be intelligible if pragmatic elements like context and intention are taken into consideration.

The structure of concepts

Cognitive semantics has sought to challenge traditional theories in two ways: first, by providing an account of the meaning of sentences by going beyond truth-conditional accounts; and second, by attempting to go beyond accounts of word meaning that appeal to necessary and sufficient

conditions. It accomplishes both by examining the structure of concepts.

Frame semantics

Frame semantics is a theory of linguisticmeaning developed by Charles J. Fillmore that extends his earlier case grammar. It relates linguisticsemantics to encyclopedic knowledge. The basic idea is that one cannot understand the meaning of a single word without access to all the essential knowledge that relates to that word. For example, one would not be able to understand the word "sell" without knowing anything about the situation of commercial transfer, which also involves, among other things, a seller, a buyer, goods, money, the relation between the money and the goods, the relations between the seller and the goods and the money, the relation between the buyer and the goods and the money and so on. Thus, a word activates, or evokes, a frame of semantic knowledge relating to the specific concept to which it refers (or highlights, in frame semantic terminology).

The idea of the encyclopedic organisation of knowledge itself is old and was discussed by Age of Enlightenment philosophers such as Denis Diderot and GiambattistaVico. Fillmore and other evolutionary and cognitive linguists like John Haiman and Adele Goldberg, however, make an argument against generative grammar and truth-conditional semantics. As is elementary for Lakoffian–Langackerian Cognitive Linguistics, it is claimed that knowledge of *language* is no different from other types of knowledge; therefore there is no grammar in the traditional sense, and language is not an independent cognitive function. Instead, the spreading and survival of linguistic units

is directly comparable to that of other types of units of cultural evolution, like in memetics and other cultural replicator theories.

Use in cognitive linguistics and construction grammar

The theory applies the notion of a semantic frame also used in artificial intelligence, which is a collection of facts that specify "characteristic features, attributes, and functions of a denotatum, and its characteristic interactions with things necessarily or typically associated with it."

A semantic frame can also be defined as a coherent structure of related concepts that are related such that without knowledge of all of them, one does not have complete knowledge of any one; they are in that sense types of gestalt. Frames are based on recurring experiences, therefore the commercial transaction frame is based on recurring experiences of commercial transactions.

Words not only highlight individual concepts, but also specify a certain perspective from which the frame is viewed. For example "sell" views the situation from the perspective of the seller and "buy" from the perspective of the buyer. This, according to Fillmore, explains the observed asymmetries in many lexical relations.

While originally only being applied to lexemes, frame semantics has now been expanded to grammatical constructions and other larger and more complex linguistic units and has more or

less been integrated into construction grammar as the main semantic principle. Semantic frames are also becoming used in information modeling, for example in Gellish, especially in the form of 'definition models' and 'knowledge models'.

Frame semantics has much in common with the semantic principle of profiling from Ronald W. Langacker'scognitive grammar.

The concept of frames has been several times considered in philosophy and psycholinguistics, namely supported by Lawrence W. Barsalou, and more recently by Sebastian Löbner. They are viewed as a cognitive representation of the real world. From a computational linguistics viewpoint, there are semantic models of a sentence. This approach going further than just the lexical aspect is especially studied in SFB 991 in Düsseldorf.

Fillmore: framing

Many pieces of linguistic evidence motivate the frame-semantic project. First, it has been noted that word meaning is an extension of our bodily and cultural experiences. For example, the notion of *restaurant* is associated with a series of concepts, like *food, service, waiters, tables, and eating*. These rich-but-contingent associations cannot be captured by an analysis in terms of necessary and sufficient conditions, yet they still seem to be intimately related to our understanding of "restaurant".

Second, and more seriously, these conditions are not enough to account for asymmetries in the ways that words are used.

According to a semantic feature analysis, there is nothing more to the meanings of "boy" and "girl" than:

- BOY [+MALE], [+YOUNG]
- GIRL [+FEMALE], [+YOUNG]

And there is surely some truth to this proposal. Indeed, cognitive semanticists understand the instances of the concept held by a given certain word may be said to exist in a *schematic relation* with the concept itself. And this is regarded as a legitimate approach to semantic analysis, so far as it goes.

However, linguists have found that language users regularly apply the terms "boy" and "girl" in ways that go beyond mere semantic features. That is, for instance, people tend to be more likely to consider a young female a "girl" (as opposed to "woman"), than they are to consider a borderline-young male a "boy" (as opposed to "man"). This fact suggests that there is a latent frame, made up of cultural attitudes, expectations, and background assumptions, which is part of word meaning. These background assumptions go up and beyond those necessary and sufficient conditions that correspond to a semantic feature account. Frame semantics, then, seeks to account for these puzzling features of lexical items in some systematic way.

Third, cognitive semanticists argue that truth-conditional semantics is incapable of dealing adequately with some aspects of the meanings at the level of the sentence. Take the following:

- You didn't spare me a day at the seaside; you deprived me of one.

In this case, the truth-conditions of the claim expressed by the antecedent in the sentence are not being denied by the proposition expressed after the clause. Instead, what is being denied is the way that the antecedent is framed.

Finally, with the frame-semantic paradigm's analytical tools, the linguist is able to explain a wider range of semantic phenomena than they would be able to with only necessary and sufficient conditions. Some words have the same definitions or intensions, and the same extensions, but have subtly different domains. For example, the lexemes *land* and *ground* are synonyms, yet they naturally contrast with different things—*sea* and *air*, respectively.

As we have seen, the frame semantic account is by no means limited to the study of lexemes—with it, researchers may examine expressions at more complex levels, including the level of the sentence (or, more precisely, the utterance). The notion of framing is regarded as being of the same cast as the pragmatic notion of *background assumptions*. Philosopher of language John Searle explains the latter by asking readers to consider sentences like "The cat is on the mat". For such a sentence to make any sense, the interpreter makes a series of assumptions: i.e., that there is gravity, the cat is parallel to the mat, and the two touch. For the sentence to be intelligible, the speaker supposes that the interpreter has an idealized or default frame in mind.

Langacker: profile and base

An alternate strain of Fillmore's analysis can be found in the work of Ronald Langacker, who makes a distinction between

the notions of *profile* and *base*. The profile is the concept symbolized by the word itself, while the base is the encyclopedic knowledge that the concept presupposes. For example, let the definition of "radius" be "a line segment that joins the center of a circle with any point on its circumference".

If all we know of the concept *radius* is its profile, then we simply know that it is a line segment that is attached to something called the "circumference" in some greater whole called the "circle". That is to say, our understanding is fragmentary until the base concept of *circle*is firmly grasped.

When a single base supports a number of different profiles, then it can be called a "*domain*". For instance, the concept profiles of *arc, center, and circumference* are all in the domain of *circle*, because each uses the concept of *circle* as a base. We are then in a position to characterize the notion of a frame as being either the base of the concept profile, or (more generally) the domain that the profile is a part of.

Categorization and cognition

A major divide in the approaches to cognitive semantics lies in the puzzle surrounding the nature of category structure. As mentioned in the previous section, semantic feature analyses fall short of accounting for the frames that categories may have. An alternative proposal would have to go beyond the minimalistic models given by classical accounts, and explain the richness of detail in meaning that language speakers attribute to categories.

Prototype theories, investigated by Eleanor Rosch, have given some reason to suppose that many natural lexical category structures are graded, i.e., they have prototypical members that are considered to be "better fit" the category than other examples. For instance, robinsare generally viewed as better examples of the category "bird" than, say, penguins. If this view of category structure is the case, then categories can be understood to have central and peripheral members, and not just be evaluated in terms of members and non-members.

In a related vein, George Lakoff, following the later Ludwig Wittgenstein, noted that some categories are only connected to one another by way of *family resemblances*. While some classical categories may exist, i.e., which are structured by necessary and sufficient conditions, there are at least two other kinds: *generative* and *radial*.

*Generative categories*can be formed by taking central cases and applying certain principles to designate category membership. The principle of similarity is one example of a rule that might generate a broader category from given prototypes.

Radial categories are categories motivated by conventions, but not predictable from rules. The concept of "mother", for example, may be explained in terms of a variety of conditions that may or may not be sufficient. Those conditions may include: being married, has always been female, gave birth to the child, supplied half the child's genes, is a caregiver, is married to the genetic father, is one generation older than the child, and is the legal guardian. Any one of the above conditions might not be met: for instance, a "single mother" does not need to be married, and a "surrogate mother" does not

necessarily provide nurturance. When these aspects collectively cluster together, they form a prototypical case of what it means to be a mother, but nevertheless they fail to outline the category crisply. Variations upon the central meaning are established by convention by the community of language users.

For Lakoff, prototype effects can be explained in large part due to the effects of idealized cognitive models. That is, domains are organized with an ideal notion of the world that may or may not fit reality. For example, the word "bachelor" is commonly defined as "unmarried adult male". However, this concept has been created with a particular ideal of what a bachelor is like: an adult, uncelibate, independent, socialized, and promiscuous. Reality might either strain the expectations of the concept, or create false positives. That is, people typically want to widen the meaning of "bachelor" to include exceptions like "a sexually active seventeen-year-old who lives alone and owns his own firm" (not technically an adult but seemingly still a bachelor), and this can be considered a kind of straining of the definition. Moreover, speakers would tend to want to exclude from the concept of *bachelor* certain false positives, such as those adult unmarried males that don't bear much resemblance to the ideal: i.e., the Pope, or Tarzan. Prototype effects may also be explained as a function of either basic-level categorization and typicality, closeness to an ideal, or stereotyping.

So viewed, prototype theory seems to give an account of category structure. However, there are a number of criticisms of this interpretation of the data. Indeed, Rosch and Lakoff, themselves chief advocates of prototype theory, have

emphasized in their later works that the findings of prototype theory do not necessarily tell us anything about category structure. Some theorists in the cognitive semantics tradition have challenged both classical and prototype accounts of category structure by proposing the dynamic construal account, where category structure is always created "on-line"—and so, that categories have no structure outside of the context of use.

Mental spaces

The **mental space** is a theoretical construct proposed by Gilles Fauconnier corresponding to possible worlds in truth-conditional semantics.

The main difference between a mental space and a possible world is that a mental space does not contain a faithful representation of reality, but an idealized cognitive model. Building of mental spaces and establishment of mappings between those mental spaces are the two main processes involved in construction of meaning. It is one of the basic components in Gilles Fauconnier and Mark Turner's blending theory, a theory within cognitive semantics.

Base space and built space

Base space, also known as reality space, presents the interlocutors' shared knowledge of the real world. Space builders are elements within a sentence that establish spaces distinct from, yet related to the base space constructed. Space builders can be expressions like prepositional phrases, adverbs, connectives, and subject-verb combinations that are

followed by an embedded sentence. They require hearers to establish scenarios beyond the present point of time. A built space depicts a situation that only holds true for that space itself, but may or may not be true in reality. The base space and built spaces are occupied by elements that map onto each other. These elements include categories that may refer to specific entities in those categories. According to Fauconnier's Access Principle, specific entities of a category in a space can be described by its counterpart category in another space even if it differs from the specific entity in the other space. An example of a built space can be seen in the example " Mary wants to buy a book". In this case, the built space is not that of reality, but Mary's desire space. Though the book in reality space refers to any book in general, it can still be used to describe the book in Mary's desire space, which may or may not be a specific book...

Foundation and expansion space

'if A then B' sentences create another two spaces called foundation space and expansion space in addition to the base space. The foundation space is a hypothetical space relative to the base space set up by the space builder "if". The expansion space is set up by the space builder "then". If the conditions in the foundation space hold, the expansion space follows.

In traditional semantics, the meaning of a sentence is the situation it represents, and the situation can be described in terms of the possible world that it would be true of. Moreover, sentence meanings may be dependent upon propositional attitudes: those features that are relative to someone's beliefs, desires, and mental states. The role of propositional attitudes

in truth-conditional semantics is controversial. However, by at least one line of argument, truth-conditional semantics seems to be able to capture the meaning of belief-sentences like "Frank believes that the Red Sox will win the next game" by appealing to propositional attitudes. The meaning of the overall proposition is described as a set of abstract conditions, wherein Frank holds a certain propositional attitude, and the attitude is itself a relationship between Frank and a particular proposition; and this proposition is the possible world where the Red Sox win the next game.

Still, many theorists have grown dissatisfied with the inelegance and dubious ontology behind possible-worlds semantics. An alternative can be found in the work of Gilles Fauconnier. For Fauconnier, the meaning of a sentence can be derived from "mental spaces". Mental spaces are cognitive structures entirely in the minds of interlocutors. In his account, there are two kinds of mental space. The *base space* is used to describe reality (as it is understood by both interlocutors). *Space builders* (or *built space*) are those mental spaces that go beyond reality by addressing possible worlds, along with temporal expressions, fictional constructs, games, and so on. Additionally, Fauconnier semantics distinguishes between *roles* and *values*. A semantic role is understood to be description of a category, while values are the instances that make up the category. (In this sense, the role-value distinction is a special case of the type-token distinction.)

Fauconnier argues that curious semantic constructions can be explained handily by the above apparatus. Take the following sentence:

- In 1929, the lady with white hair was blonde.

The semanticist must construct an explanation for the obvious fact that the above sentence is not contradictory. Fauconnier constructs his analysis by observing that there are two mental spaces (the present-space and the 1929-space). His *access principle*supposes that "a value in one space can be described by the role its counterpart in another space has, even if that role is invalid for the value in the first space". So, to use the example above, the value in 1929-space is *the blonde*, while she is being described with the role of *the lady with white hair* in present-day space.

Conceptualization and construal

As we have seen, cognitive semantics gives a treatment of issues in the construction of meaning both at the level of the sentence and the level of the lexeme in terms of the structure of concepts. However, it is not entirely clear what cognitive processes are at work in these accounts. Moreover, it is not clear how we might go about explaining the ways that concepts are actively employed in conversation. It appears to be the case that, if our project is to look at *how* linguistic strings convey different semantic content, we must first catalogue *what* cognitive processes are being used to do it. Researchers can satisfy both requirements by attending to the *construal operations* involved in language processing—that is to say, by investigating the ways that people *structure their experiences* through language.

Language is full of conventions that allow for subtle and nuanced conveyances of experience. To use an example that is

readily at hand, framing is all-pervasive, and it may extend across the full breadth of linguistic data, extending from the most complex utterances, to tone, to word choice, to expressions derived from the composition of morphemes. Another example is *image-schemata*, which are ways that we structure and understand the elements of our experience driven by any given sense.

According to linguists William Croft and D. Alan Cruse, there are four broad cognitive abilities that play an active part in the construction of construals.

They are: attention/salience, judgment/comparison, situatedness, and constitution/gestalt. Each general category contains a number of subprocesses, each of which helps to explain the ways we encode experience into language in some unique way.

In social psychology, a **construal** is how a person perceives, comprehends and interprets their world, particularly the acts of others toward them.

Researchers and theorists within virtually every sub-discipline of psychology have acknowledged the relevance of a subjective construal, especially with regards to the concepts of the false consensus effect and the fundamental attribution error.

There is a difference between self-construal and construal in a social atmosphere. While self-construal is a perception of the self, the latter is a perception of one's surroundings. Construal plays a crucial role when one lacks the knowledge to correctly deal with a situation.

Major theoretical approaches

The concept of construal is not a new one, and the components of construal can be seen in the works of many past psychologists including Kurt Lewin's recognition of the importance of a subjective reality and its effect on one's personal significance; Kurt Koffka's theories of gestalt psychology; Brunswik's emphasis on subjective distinction; Murray's discussion of "beta press"; Kelly's account of personal constructs; Merleau-Ponty's reference to personal situations; and more recent discussions by personality theorists such as Endler and Pervin. Construal used to be viewed as an obstruction in one's perception of the world, but has evolved into a mechanism used to explain how or why a person thinks the way they do.

Cognitive psychologists have been perhaps the most preoccupied with the idea of construal. This is evident in their emphasis on a human's formation of schemas "that help perceivers to resolve ambiguity, fill in the gaps, and generally perceive predictability and coherence." They focus on the idea that we rely on other sources to form our ideas of our surroundings.

Solomon Asch presented an important concept in construal theory when he stated, "that the very meaning of a message can change as a function of the source to which it is attributed." His most classic example is the effect of the phrase "a little rebellion...is a good thing." This statement coming from Thomas Jefferson has a different meaning to the recipient than it does coming from V.I. Lenin. The meaning of the statement

is dependent on not only who says it, but also on how the recipient of the message interprets it.

There are three major sources of construal in human beings: the need to feel good about ourselves, the need to be accurate, and the need to belong. The American social psychologist Leon Festinger was one of the first to acknowledge that these needs may not always coincide (see cognitive dissonance). The Austrian social psychologist Fritz Heider described the concept of construal when he said, "Generally, a person reacts to what he thinks the other person is perceiving, feeling, and thinking, in addition to what the other person may be doing." In other words, a person bases his or her opinions and actions on the opinions and action of everybody else.

For example, take this situation into consideration:

Christopher likes Samantha and wants to ask her to the school prom. He is shy and concerned that Samantha may respond negatively. A social psychologist observes not only Samantha's behavior towards Christopher, but also how Christopher perceives and interprets her behavior toward him. An objective observer may perceive Samantha smiling as friendly, but Christopher may think that she is laughing at something in his appearance, and as a result, he might not invite her.

Contemporary views on construal include the concepts of naïve realism, the accessibility principle, and a focus on the idea of self-construal. Lee Ross's concept of naïve realism is especially important in the context of construal. It is the conviction all of us have that we perceive things how they really are. Essentially, people acknowledge the fact that others experience the effects of construal, but personally think that they form

their own thoughts without being affected by construal. Being blinded by this process often leads people to commit the fundamental attribution error.

Similar to Asch's theory, the accessibility principle suggests that "mental construals are based on the information that is most accessible at the time applies to how we make sense of new information as well as to how we form judgments based on information retrieved from memory." Lastly, self-construal is how a person views the self in comparison to the others. This would suggest that self-construal influences a person's self-esteem. Construal itself is a broad concept in the realm of social psychology and can be applied to many different situations that will be discussed later in this article.

Major empirical evidence

In 1946, Solomon Asch directed one of the earliest known empirical studies of human construal. In this study, Asch focused on the formation of character impressions by asking each participant to study a list of personality traits and make judgments and/or inferences about the possessor of each of these listed traits.

The results of this study demonstrated two different types of phenomena: the primacy effect and the disproportionate effect of certain types of words. For the primacy effect, those personality traits that were listed earlier in the list seemed to have much more influence on the subject's impression of the person with that trait. However, Asch's finding that there was a variability in the effect of categorical terms such as "warmth" and "coldness" hint that those listed traits were "susceptible to

variable interpretation or construal—and the specific meaning attached depended upon the more global impressions adopted by the subjects".

In a study headed by Lee D. Ross, David Green, and Pamela House (1976), eighty Stanford University undergraduates were asked if they were willing to walk around campus for at least thirty minutes while wearing a large sandwich board sign that read "Eat at Joe's" and record the responses of their peers to this novel situation. The subjects were not only asked to answer whether or not they would participate, but they were also asked to estimate other people's responses, and make inferences about the disposition of each group of people based on their agreement and disagreement to participate. Overall, the experimenters found that "those who agreed to participate thought that an average of 62% of their peers would agree"; but, those who disagreed with participating thought that an average 33% of their peers would agree to the job". Furthermore, those who agreed had more extreme inferences about the personal dispositions of those who disagreed, and vice versa. The results indicated that the subjects failed to recognize that their peer's construal or interpretation of the situation may be quite different from the perspective they personally take. (see also false consensus effect)

In 2004, Lee D. Ross developed a theory of a type of construal that he calls "naïve realism." In a simple experiment, Ross took peace proposals created by Israeli negotiators, labeled them as Palestinian proposals, and told Israeli citizens that the ideas in the proposal were the ideas that Palestinians wanted the Israeli to adopt. Then, he took the original proposals and told the Israeli subjects that ideas on the proposal were the ideas

that the Israelis wanted the Palestinians to adopt. The Israeli citizens liked the proposals from the Israelis to the Palestinians more than the proposal from the Palestinians to the Israelis, even though they were the same proposal. Ross stated:

Even when each side recognizes that the other side perceives the issues differently, each thinks that the other side is biased while they themselves are objective and that their own perceptions of reality should provide the basis for settlement."

Self

Hazel Rose Markus and ShinobuKitayama argue that differences between independent and interdependent self-concepts lead to different consequences for a number of cognitive and motivational processes. They argue that the distinctions made regarding independent and interdependent construalsshould be viewed as general tendencies that may emerge when the members of the culture are considered as a whole. Also, "According to Markus and Kitayama (1991), those with an independent self-construal define themselves in terms of internal attributes such as traits, abilities, values, and preferences. In contrast, those with an interdependent self-construal define themselves in terms of their relationships with others." Many who argue these separate views of construal say that both views can strongly affect a person's experience.

The following is a research study about the way in which a person's construal can affect his/her mental health status. Michael S. Christopher and Gemma D. Skillman conducted a study to test the link between self-construal and distress

among African American and Asian American college students. Their research is primarily based on previous assertions that ethnic minorities are more likely to experience distress and express apprehension about the rigors of college. One body of literature has commonly viewed three major minority groups- African Americans, Asian Americans, and Latino Americans- as more likely to display traits of interdependent self-construal. It suggests that on the other hand White Americans were more likely to show traits of individualism or independent self-construal. Current research, however, has begun to move away from this view, and indicates African American students also show a more independent view of self. Therefore, to try and contrast these two views the researchers chose to study African Americans and Asian Americans. In their study they found "African American students reported greater independent construals than did Asian American students, whereas Asian American students reported greater interdependent self-construals than did African American students." In regard to whether self-construal contributed to reports of distress they found that viewing self-construal as independent or interdependent did not predict distress. A person reported to have a more interdependent view of self was more likely to experience distress symptoms. This type of research finding can have major effects on future counseling practices. These researchers encourage counselors to measure self-construals upon intake to help guide treatment.

Markus and Kitayama's self-construal theory postulates that there are two basic ways of conceptualizing the human person and that cultures differ according to which of those they subscribe to. Egoism, individual pride, individual uniqueness, independent thinking, self-expression, self-reliance and self-

enhancement are believed to be facets of one and the same coherent entity named the independent self-construal. On the other hand, altruism, modesty, belief in one's similarity with others, conformity to group norms, self-censorship for the sake of group harmony and cooperation are believed to be facets of the interdependent self-construal. However, the idea that there exists cultures that promote an independent or independent form of self construal has been severely criticized Vignoles et al.'s comparative study of 55 cultural groups shows that there are no such kind of culturally normative independent or interdependent self-contruals as predicted by Markus and Kitayama's theory. At the cultural level of analysis, believing in and valuing uniqueness is postulated by this theory to be positively correlated with egoism, while in reality they are negatively correlated. Experiencing the self as emotionally detached and independent from social contexts was supposed to be positively related to self-reliance, while in reality the two concepts are negatively related. Independent thinking was supposed to be correlated positively with self-reliance, but in reality they are negatively correlated. The same is true for independent thinking and self-consistency across contexts as well as self-reliance and self-expression. Moreover, Japanese culture was found to promote individual uniqueness and independent thinking more than other human cultures in the sample, which directly contradicts Markus and Kitayama.

In the classroom

Rebecca wing-yi Cheng and Shui-fong Lam measured the effects of self-construal in the classroom. They studied the "role of self-construal as a moderator of the social comparison effects in authentic classrooms." With the use of 96 Chinese

seventh grade students they compared independent and interdependent views of self-construal to upward social comparisonand downward social comparison. They noted that "self comparison is commonly used when people are uncertain of their self-evaluation. It allows an individual to gain information about where they stand." The students participated in an Abstract Reasoning Test and reading comprehension task. They manipulated construal by telling the students they were either being compared to others within their school (independent self-construal) or between their school and another school (interdependent self-construal). The results showed that the school children who performed very well experienced negative self-evaluation while those who did not experienced negative self-evaluation. However, those in the interdependent self construal condition always reported positive self-evaluation showcasing a term known as basking-in-reflected-glory. The hope of this study is to encourage classrooms to ensure that interdependent self-construal is being emphasized.

In social marketing

Research drawing on self-construals now shows ways to reduce the intentions of people to binge drink or engage in dangerous driving. An article by Martin, Lee, Weeks and Kaya (2013) suggests that understanding consumer personality and how people view others is important. People were shown ads talking of the harmful effects of binge drinking. People who valued close friends as a sense of who they are, were less likely to want to binge drink after seeing an ad featuring them and a close friend. People who were loners or who did not see close friends important to their sense of who they were reacted

better to ads featuring a person. A similar pattern was shown for ads showing a person driving at dangerous speeds. This suggests ads showing potential harm to citizens from binge drinking or dangerous driving are less effective than ads highlighting a person's close friends.

Mammography

Some researches have believed that construal can have major implications on self-perception of health status. As cited by a Times article "women ages 20–49 should have a physical examination by a health professional every one or two years." Specifically, in relation to breast cancer, women should do monthly self-examinations. However, after about age 40 women should begin mammograms, an effective low-radiation screening method for breast cancer. Although the age and frequency of which women should begin breast exams are highly debated the general consensus is that those over age 50 should be examined annually. Unfortunately, the number of women partaking in regular screenings is still not as high as it should be. Consequently, one study of construal conducted by Gallagher and colleagues looked at the link between message framing and perceptions about breast cancer susceptibility. The research is primarily based on the assumption that "people's responses to framed messages may not always be a simple reflection of the presumed risky nature of screening behaviors, but rather shaped by beliefs about risk." Therefore, in this particular study they "assessed women's illness-detecting v. health-affirming construal of mammography." They found "that among women who have a family history of breast cancer, their construal of mammography moderates their

responses to framed messages. Such that, loss-framed messages are more effective in promoting screening for those with illness-detecting construals, but gain-framed messages are more effective for those with health-affirming construals." Their research shows that the perception of susceptibility to the development of breast cancer was not directly associated with a person's construal of the function of mammography.

Market research

Researchers are trying to establish links between construal, self and economics. In 2007, Liberman and colleagues discussed the links between construal level theory and consumer choice, better decision making, the nature of regret and choice set management.

Chapter 3

Lexical Semantics

Lexical semantics (also known as lexicosemantics), as a subfield of linguisticsemantics, is the study of word meanings. It includes the study of how words structure their meaning, how they act in grammar and compositionality, and the relationships between the distinct senses and uses of a word.

The units of analysis in lexical semantics are lexical units which include not only words but also sub-words or sub-units such as affixes and even compound words and phrases. Lexical units include the catalogue of words in a language, the lexicon. Lexical semantics looks at how the meaning of the lexical units correlates with the structure of the language or syntax. This is referred to as syntax-semantics interface.

The study of lexical semantics looks at:

the classification and decomposition of lexical items

the differences and similarities in lexical semantic structure cross-linguistically

the relationship of lexical meaning to sentence meaning and syntax.

Lexical units, also referred to as syntactic atoms, can stand alone such as in the case of root words or parts of compound words or they necessarily attach to other units such as prefixes and suffixes do. The former are calledfree morphemes and the latter bound morphemes. They fall into a narrow range

of meanings (semantic fields) and can combine with each other to generate new denotations.

Cognitive semantics is the linguistic paradigm/framework that since the 1980s has generated the most studies in lexical semantics, introducing innovations like prototype theory, conceptual metaphors, and frame semantics.

Semantic fields

How lexical items map onto concepts

First proposed by Trier in the 1930s, semantic field theory proposes that a group of words with interrelated meanings can be categorized under a larger conceptual domain. This entire entity is thereby known as a semantic field. The words *boil*, *bake*, *fry*, and *roast*, for example, would fall under the larger semantic category of *cooking*. Semantic field theory asserts that lexical meaning cannot be fully understood by looking at a word in isolation, but by looking at a group of semantically related words. Semantic relations can refer to any relationship in meaning between lexemes, including synonymy *(big* and *large)*,antonymy*(big* and *small)*,hypernymy and hyponymy *(rose* and *flower),* converseness *(buy* and *sell),* and incompatibility. Semantic field theory does not have concrete guidelines that determine the extent of semantic relations between lexemes. The abstract validity of the theory is a subject of debate.

Knowing the meaning of a lexical item therefore means knowing the semantic entailments the word brings with it. However, it is also possible to understand only one word of a semantic field without understanding other related words.

Take, for example, a taxonomy of plants and animals: it is possible to understand the words *rose* and *rabbit* without knowing what a *marigold* or a *muskrat* is. This is applicable to colors as well, such as understanding the word *red* without knowing the meaning of *scarlet,* but understanding *scarlet* without knowing the meaning of *red* may be less likely. A semantic field can thus be very large or very small, depending on the level of contrast being made between lexical items. While cat and dog both fall under the larger semantic field of animal, including the breed of dog, like *German shepherd,* would require contrasts between other breeds of dog (e.g. *corgi*, or *poodle*), thus expanding the semantic field further.

How lexical items map onto events

Event structure is defined as the semantic relation of a verb and its syntactic properties. Event structure has three primary components:

- primitive event type of the lexical item
- event composition rules
- mapping rules to lexical structure

Verbs can belong to one of three types: states, processes, or transitions.

- The door is **closed**.
- The door **closed**.
- John **closed** the door.

(1a) defines the state of the door being closed; there is no opposition in this predicate. (1b) and (1c) both have predicates showing transitions of the door going from being implicitly *open* to *closed*. (1b) gives the intransitive use of the verb close, with no explicit mention of the causer, but (1c) makes explicit mention of the agent involved in the action.

Syntactic basis of event structure: a brief history

Generative semantics in the 1960s

The analysis of these different lexical units had a decisive role in the field of "generative linguistics" during the 1960s. The term *generative* was proposed by Noam Chomsky in his book Syntactic Structures published in 1957. The term *generative linguistics* was based on Chomsky's generative grammar, a linguistic theory that states systematic sets of rules (X' theory) can predict grammatical phrases within a natural language. Generative Linguistics is also known as Government-Binding Theory. Generative linguists of the 1960s, including Noam Chomsky and Ernst von Glasersfeld, believed semantic relations between transitive verbs and intransitive verbswere tied to their independent syntactic organization. This meant that they saw a simple verb phrase as encompassing a more complex syntactic structure.

Lexicalist theories in the 1980s

Lexicalist theories became popular during the 1980s, and emphasized that a word's internal structure was a question of

morphology and not of syntax. Lexicalist theories emphasized that complex words (resulting from compounding and derivation of affixes) have lexical entries that are derived from morphology, rather than resulting from overlapping syntactic and phonological properties, as Generative Linguistics predicts. The distinction between Generative Linguistics and Lexicalist theories can be illustrated by considering the transformation of the word *destroy* to *destruction*:

- **Generative Linguistics theory**: states the transformation of *destroy→destruction* as the nominal, *nom* + *destroy*, combined with phonological rules that produce the output *destruction*. Views this transformation as independent of the morphology.
- **Lexicalist theory**: sees *destroy* and *destruction* as having idiosyncratic lexical entries based on their differences in morphology. Argues that each morpheme contributes specific meaning. States that the formation of the complex word *destruction*is accounted for by a set of *Lexical Rules*, which are different and independent from syntactic rules.

A lexical entry lists the basic properties of either the whole word, or the individual properties of the morphemes that make up the word itself.

The properties of lexical items include their category selection *c-selection*, selectional properties *s-selection*, (also known as semantic selection), phonological properties, and features. The properties of lexical items are idiosyncratic, unpredictable, and contain specific information about the lexical items that they describe.

Lexicalist theories state that a word's meaning is derived from its morphology or a speaker's lexicon, and not its syntax. The degree of morphology's influence on overall grammar remains controversial. Currently, the linguists that perceive one engine driving both morphological items and syntactic items are in the majority.

Micro-syntactic theories: 1990s to the present

By the early 1990s, Chomsky's minimalist framework on language structure led to sophisticated probing techniques for investigating languages. These probing techniques analyzed negative data over prescriptive grammars, and because of Chomsky's proposed Extended Projection Principle in 1986, probing techniques showed where specifiers of a sentence had moved to in order to fulfill the EPP. This allowed syntacticians to hypothesize that lexical items with complex syntactic features (such as ditransitive, inchoative, and causative verbs), could select their own specifier element within a syntax tree construction. (For more on probing techniques, see Suci, G., Gammon, P., &Gamlin, P. (1979)).

This brought the focus back on the syntax-lexical semantics interface; however, syntacticians still sought to understand the relationship between complex verbs and their related syntactic structure, and to what degree the syntax was projected from the lexicon, as the Lexicalist theories argued.

In the mid 1990s, linguists Heidi Harley, Samuel Jay Keyser, and Kenneth Hale addressed some of the implications posed by complex verbs and a lexically-derived syntax. Their proposals indicated that the predicates CAUSE and BECOME, referred to

as subunits within a Verb Phrase, acted as a lexical semantic template. *Predicates* are verbs and state or affirm something about the subject of the sentence or the argument of the sentence.

For example, the predicates *went* and *is here* below affirm the argument of the subject and the state of the subject respectively.

The subunits of Verb Phrases led to the Argument Structure Hypothesis and Verb Phrase Hypothesis, both outlined below. The recursion found under the "umbrella" Verb Phrase, the VP Shell, accommodated binary-branching theory; another critical topic during the 1990s. Current theory recognizes the predicate in Specifier position of a tree in inchoative/anticausative verbs (intransitive), or causative verbs (transitive) is what selects the theta role conjoined with a particular verb.

Hale & Keyser 1990

Kenneth Hale and Samuel Jay Keyser introduced their thesis on lexical argument structure during the early 1990s. They argue that a predicate's argument structure is represented in the syntax, and that the syntactic representation of the predicate is a lexical projection of its arguments. Thus, the structure of a predicate is strictly a lexical representation, where each phrasal head projects its argument onto a phrasal level within the syntax tree. The selection of this phrasal head is based on Chomsky's Empty Category Principle. This lexical projection of the predicate's argument onto the syntactic

structure is the foundation for the Argument Structure Hypothesis. This idea coincides with Chomsky's Projection Principle, because it forces a VP to be selected locally and be selected by a Tense Phrase (TP).

Based on the interaction between lexical properties, locality, and the properties of the EPP (where a phrasal head selects another phrasal element locally),

Hale and Keyser make the claim that the Specifier position or a complement are the only two semantic relations that project a predicate's argument. In 2003, Hale and Keyser put forward this hypothesis and argued that a lexical unit must have one or the other, Specifier or Complement, but cannot have both.

Halle & Marantz 1993

Morris Halle and Alec Marantz introduced the notion of distributed morphology in 1993. This theory views the syntactic structure of words as a result of morphology and semantics, instead of the morpho-semantic interface being predicted by the syntax.

Essentially, the idea that under the Extended Projection Principle there is a local boundary under which a special meaning occurs. This meaning can only occur if a head-projecting morpheme is present within the local domain of the syntactic structure. The following is an example of the tree structure proposed by distributed morphology for the sentence *"John's destroying the city"*. *Destroy* is the root, V-1 represents verbalization, and D represents nominalization.

Ramchand 2008

In her 2008 book, *Verb Meaning and The Lexicon: A First-Phase Syntax*, linguist Gillian Ramchand acknowledges the roles of lexical entries in the selection of complex verbs and their arguments. 'First-Phase' syntax proposes that event structure and event participants are directly represented in the syntax by means of binary branching. This branching ensures that the Specifier is the consistently subject, even when investigating the projection of a complex verb's lexical entry and its corresponding syntactic construction. This generalization is also present in Ramchand's theory that the complement of a head for a complex verb phrase must co-describe the verb's event.

Ramchand also introduced the concept of Homomorphic Unity, which refers to the structural synchronization between the head of a complex verb phrase and its complement. According to Ramchand, Homomorphic Unity is "when two event descriptors are syntactically Merged, the structure of the complement must unify with the structure of the head."

Classification of event types

Intransitive verbs: unaccusative versus unergative

The unaccusative hypothesis was put forward by David Perlmutter in 1987, and describes how two classes of intransitive verbs have two different syntactic structures. These are unaccusative verbs and unergative verbs. These

classes of verbs are defined by Perlmutter only in syntactic terms. They have the following structures underlyingly:

- unaccusative verb: __ [$_{VP}$ V NP]
- unergative verb: NP [$_{VP}$ V]

Transitivity alternations: the inchoative/causative alternation

The change-of-state property of Verb Phrases (VP) is a significant observation for the syntax of lexical semantics because it provides evidence that subunits are embedded in the VP structure, and that the meaning of the entire VP is influenced by this internal grammatical structure. (For example, the VP *the vase broke* carries a change-of-state meaning of the vase becoming broken, and thus has a silent BECOME subunit within its underlying structure.) There are two types of change-of-state predicates: inchoative and causative.

Inchoative verbs are intransitive, meaning that they occur without a direct object, and these verbs express that their subject has undergone a certain change of state. Inchoative verbs are also known as anticausative verbs. Causative verbs are transitive, meaning that they occur with a direct object, and they express that the subject causes a change of state in the object.

Linguist Martin Haspelmath classifies inchoative/causative verb pairs under three main categories: causative, anticausative, and non-directed alternations. Non-directed

alternations are further subdivided into labile, equipollent, and suppletive alternations.

English tends to favourlabile alternations, meaning that the same verb is used in the inchoative and causative forms. This can be seen in the following example: *broke* is an intransitive inchoative verb in (3a) and a transitive causative verb in (3b).

As seen in the underlying tree structure for (3a), the silent subunit BECOME is embedded within the Verb Phrase (VP), resulting in the inchoative change-of-state meaning (y become z). In the underlying tree structure for (3b), the silent subunits CAUS and BECOME are both embedded within the VP, resulting in the causative change-of-state meaning (x cause y become z).

English change of state verbs are often de-adjectival, meaning that they are derived from adjectives.

In example (4a) we start with a stative intransitive adjective, and derive (4b) where we see an intransitive inchoative verb. In (4c) we see a transitive causative verb.

Marked inchoatives

- Some languages (e.g., German, Italian, and French), have multiple morphological classes of inchoative verbs. Generally speaking, these languages separate their inchoative verbs into three classes: verbs that are obligatorily unmarked (they are not marked with a reflexive pronoun, clitic, or affix), verbs that are optionally marked, and verbs that are obligatorily marked. The causative verbs in these languages

remain unmarked. Haspelmath refers to this as the anticausative alternation.

For example, inchoative verbs in Germanare classified into three morphological classes. *Class A* verbs necessarily form inchoatives with the reflexive pronoun *sich*, *Class B* verbs form inchoatives necessarily without the reflexive pronoun, and *Class C* verbs form inchoatives optionally with or without the reflexive pronoun. In example (5), the verb *zerbrach* is an unmarked inchoative verb from *Class B*, which also remains unmarked in its causative form.

In contrast, the verb *öffnete* is a *Class A*verb which necessarily takes the reflexive pronoun *sich* in its inchoative form, but remains unmarked in its causative form.

There has been some debate as to whether the different classes of inchoative verbs are purely based in morphology, or whether the differentiation is derived from the lexical-semantic properties of each individual verb. While this debate is still unresolved in languages such as Italian, French, and Greek, it has been suggested by linguist Florian Schäfer that there are semantic differences between marked and unmarked inchoatives in German. Specifically, that only unmarked inchoative verbs allow an unintentional causer reading (meaning that they can take on an "*x unintentionally caused y*" reading).

Marked causatives

Causative morphemes are present in the verbs of many languages (e.g., Tagalog, Malagasy, Turkish, etc.), usually appearing in the form of an affix on the verb. This can be seen

in the following examples from Tagalog, where the causative prefix *pag-* (realized here as *nag*) attaches to the verb *tumba* to derive a causative transitive verb in (7b), but the prefix does not appear in the inchoative intransitive verb in (7a). Haspelmath refers to this as the causative alternation.

Ditransitive verbs

Kayne's 1981 unambiguous path analysis

Richard Kayne proposed the idea of unambiguous paths as an alternative to c-commanding relationships, which is the type of structure seen in examples (8). The idea of unambiguous paths stated that an antecedent and an anaphor should be connected via an unambiguous path. This means that the line connecting an antecedent and an anaphor cannot be broken by another argument. When applied to ditransitive verbs, this hypothesis introduces the structure in diagram (8a). In this tree structure it can be seen that the same path can be traced from either DP to the verb. Tree diagram (7b) illustrates this structure with an example from English. This analysis was a step toward binary branching trees, which was a theoretical change that was furthered by Larson's VP-shell analysis.

Larson's 1988 "VP-shell" analysis

- Larson posited his Single Complement Hypothesis in which he stated that every complement is introduced with one verb. The Double Object Construction presented in 1988 gave clear evidence of a hierarchical structure using asymmetrical binary branching. Sentences with double objects occur with

ditransitive verbs, as we can see in the following example:

- John sent Mary a package
- John sent a package to Mary.

It appears as if the verb *send* has two objects, or complements (arguments): both *Mary*, the recipient and *parcel*, the theme. The argument structure of ditransitive verb phrases is complex and has undergone different structural hypothesis.

The original structural hypothesis was that of ternary branching seen in (9a) and (9b), but following from Kayne's 1981 analysis, Larson maintained that each complement is introduced by a verb.

Their hypothesis shows that there is a lower verb embedded within a VP shell that combines with an upper verb (can be invisible), thus creating a VP shell (as seen in the tree diagram to the right). Most current theories no longer allow the ternary tree structure of (9a) and (9b), so the theme and the goal/recipient are seen in a hierarchical relationship within a binary branching structure.

Following are examples of Larson's tests to show that the hierarchical (superior) order of any two objects aligns with a linear order, so that the second is governed (c-commanded) by the first. This is in keeping with X'Bar Theory of Phrase Structure Grammar, with Larson's tree structure using the empty Verb to which the V is raised.

Reflexives and reciprocals (anaphors) show this relationship in which they must be c-commanded by their antecedents, such that the (10a) is grammatical but (10b) is not:

- (10) a. I showed Mary herself.
- b. *I showed herself Mary.

A pronoun must have a quantifier as its antecedent:

- (11) a. I gave every worker his paycheck.
- b. *I gave its owner every paycheck.

Question words follow this order:

- (12) a. Who did you give which paycheck?
- b. *Which paycheck did you give who?

The effect of negative polarity means that "any" must have a negative quantifier as an antecedent:

- (13) a. I showed no one anything.
- b. *I showed anyone nothing.

These tests with ditransitive verbs that confirm c-command also confirm the presence of underlying or invisible causative verbs. In ditransitive verbs such as *give someone something*, *send someone something*, *show someone something* etc. there is an underlying causative meaning that is represented in the underlying structure. As seen in example in (9a) above, *John sent Mary a package*, there is the underlying meaning that 'John "caused" Mary to have a package'.

Larson proposed that both sentences in (9a) and (9b) share the same underlying structure and the difference on the surface lies in that the double object construction "John sent Mary a package" is derived by transformation from a NP plus PP construction "John sent a package to Mary".

Beck & Johnson's 2004 double object construction

Beck and Johnson, however, give evidence that the two underlying structures are not the same. In so doing, they also give further evidence of the presence of two VPs where the verb attaches to a causative verb. In examples (14a) and (b), each of the double object constructions are alternated with NP + PP constructions.

- (14) a. Satoshi sent Tubingen the Damron Guide.
- b. Satoshi sent the Damron Guide to Tübingen.

Beck and Johnson show that the object in (15a) has a different relation to the motion verb as it is not able to carry the meaning of HAVING which the possessor (9a) and (15a) can. In (15a), Satoshi is an animate possessor and so is caused to

HAVE kisimen. The PP *for Satoshi* in (15b) is of a benefactive nature and does not necessarily carry this meaning of HAVE either.

- (15) a. Thilo cooked Satoshi kisimen.
- b. Thilo cooked kisimen for Satoshi.

The underlying structures are therefore not the same. The differences lie in the semantics and the syntax of the sentences, in contrast to the transformational theory of Larson.

Further evidence for the structural existence of VP shells with an invisible verbal unit is given in the application of the adjunct or modifier "again". Sentence (16) is ambiguous and looking into the two different meanings reveals a difference in structure.

- (16) Sally opened the door again.

However, in (17a), it is clear that it was Sally who repeated the action of opening the door.

In (17b), the event is in the door being opened and Sally may or may not have opened it previously. To render these two different meanings, "again" attaches to VPs in two different places, and thus describes two events with a purely structural change.

- (17) a. Sally was so kind that she went out of her way to open the door once again.
- b. The doors had just been shut to keep out the bugs but Sally opened the door again.

Chapter 4

Computational Semantics

Computational semantics is the study of how to automate the process of constructing and reasoning with meaning representations of natural language expressions. It consequently plays an important role in natural language processing and computational linguistics.

Some traditional topics of interest are:construction of meaning representations, semantic underspecification, anaphora resolution, presupposition projection, and quantifier scope resolution. Methods employed usually draw from formal semantics or statistical semantics. Computational semantics has points of contact with the areas of lexical semantics (word sense disambiguation and semantic role labeling), discourse semantics, knowledge representation and automated reasoning (in particular, automated theorem proving). Since 1999 there has been an ACL special interest group on computational semantics, SIGSEM.

Semantic analysis (linguistics)

In linguistics, semantic analysis is the process of relating syntactic structures, from the levels of phrases, clauses, sentences and paragraphs to the level of the writing as a whole, to their language-independent meanings. It also involves removing features specific to particular linguistic and cultural contexts, to the extent that such a project is possible. The elements of idiom and figurative speech, being cultural,

are often also converted into relatively invariant meanings in semantic analysis. Semantics, although related to pragmatics, is distinct in that the former deals with word or sentence choice in any given context, while pragmatics considers the unique or particular meaning derived from context or tone. To reiterate in different terms, semantics is about universally coded meaning, and pragmatics, the meaning encoded in words that is then interpreted by an audience.

Semantic analysis can begin with the relationship between individual words. This requires an understanding of lexical hierarchy, including hyponymy and hypernymy, meronomy, polysemy, synonyms, antonyms, and homonyms. It also relates to concepts like connotation (semiotics) and collocation, which is the particular combination of words that can be or frequently are surrounding a single word. This can include idioms, metaphor, and simile, like, "white as a ghost."

With the availability of enough material to analyze, semantic analysis can be used to catalog and trace the style of writing of specific authors.

Underspecification

In theoretical linguistics, underspecification is a phenomenon in which certain featuresare omitted in underlying representations. Restricted underspecification theory holds that features should only be underspecified if their values are predictable. For example, in most dialects of English, all front vowels (/i, ɪ, e, ɛ, æ/) are unrounded. It is not necessary for these phonemes to include the distinctive feature [–round], because all [–back] vowels are [–round] vowels, so the

roundness feature is not distinctive for front vowels. Radical underspecification theory, on the other hand, also allows for traditionally binary features to be specified for only one value, where it is assumed that every segment not specified for that value has the other value. For example, instead of the features [+voice] and [–voice], only [+voice] is specified and voicelessness is taken as the default.

The concept of underspecificationis also used in morphological theory, particularly to refer to cases in which a morpheme does not bear an entire set of feature-values, and is thus compatible with a wide range of potential morphological environments. In this approach to morphology, for example, while the English pronouns *he* vs. *she* are specified for gender, the plural pronoun *they* would be underspecified for gender.

Example of underspecification in phonology

In Tuvan, phonemic vowels are specified with the articulatory features of tongue height, backness, and lip rounding. The archiphoneme |U| is an underspecified high vowel where only the tongue height is specified.

• **phoneme/** • **archiphoneme**	• **height**	• **backness**	• **roundedness**

• /i/	• high	• front	• unrounded
• /ɯ/	• high	• back	• unrounded
• /u/	• high	• back	• rounded
• |U|	• high		•

Whether |U| is pronounced as front or back and whether rounded or unrounded depends on vowel harmony. If |U| occurs following a front unrounded vowel, it will be pronounced as the phoneme /i/; if following a back unrounded vowel, it will be as an/ɯ/; and if following a back rounded vowel, it will be an /u/.

Underspecification in Morphology

Underspecification in morphology uses feature decomposition to create abstract, binary features that allow for the creation of natural classes in relation to morphology.

In German, there are three classes of gender. These are feminine, masculine, and neuter.

Anaphora (linguistics)

In linguistics, **anaphora** (/əˈnæfərə/) is the use of an expression whose interpretation depends upon another expression in context (its antecedent or postcedent). In a narrower sense, anaphora is the use of an expression that depends specifically upon an antecedent expression and thus is contrasted with cataphora, which is the use of an expression that depends upon a postcedent expression. The anaphoric (referring) term is called an **anaphor**. For example, in the sentence*Sally arrived, but nobody saw her*, the pronoun*her* is an anaphor, referring back to the antecedent *Sally*. In the sentence *Before her arrival, nobody saw Sally*, the pronoun *her* refers forward to the postcedent*Sally*, so *her* is now a *cataphor* (and an anaphor in the broader, but not the narrower, sense). Usually, an anaphoric expression is a proform or some other kind of deictic (contextually-dependent) expression. Both anaphora and cataphora are species of endophora, referring to something mentioned elsewhere in a dialog or text.

Anaphora is an important concept for different reasons and on different levels: first, anaphora indicates how discourse is constructed and maintained; second, anaphora binds different syntactical elements together at the level of the sentence; third, anaphora presents a challenge to natural language processing in computational linguistics, since the identification of the reference can be difficult; and fourth, anaphora partially reveals how language is understood and processed, which is relevant to fields of linguistics interested in cognitive psychology.

Nomenclature and examples

The term *anaphora*is actually used in two ways.

In a broad sense, it denotes the act of referring. Any time a given expression (e.g. a proform) refers to another contextual entity, anaphora is present.

In a second, narrower sense, the term *anaphora* denotes the act of referring backwards in a dialog or text, such as referring to the left when an anaphor points to its left toward its antecedent in languages that are written from left to right. Etymologically, *anaphora* derives from Ancient Greekἀναφορά (anaphorá, "a carrying back"), from ἀνά (aná, "up") + φέρω (phérō, "I carry"). In this narrow sense, anaphora stands in contrast to cataphora, which sees the act of referring forward in a dialog or text, or pointing to the right in languages that are written from left to right: Ancient Greek καταφορά (kataphorá, "a downward motion"), from κατά (katá, "downwards") + φέρω (phérō, "I carry"). A proform is a cataphor when it points to its right toward its postcedent. Both effects together are called either anaphora (broad sense) or less ambiguously, along with self-reference they comprise the category of endophora.

Examples of anaphora (in the narrow sense) and cataphoraare given next. Anaphors and cataphors appear in bold, and their antecedents and postcedents are underlined:

- **Anaphora** (in the narrow sense, species of endophora)

- a. Susan dropped the plate. **It** shattered loudly. – The pronoun *it* is an anaphor; it points to the left toward its antecedent *the plate*.
- b. The music stopped, and **that** upset everyone. – The demonstrative pronoun *that* is an anaphor; it points to the left toward its antecedent *The music stopped*.
- c. Fred was angry, and **so** was I. – The adverb *so* is an anaphor; it points to the left toward its antecedent *angry*.
- d. If Sam buys a new bike, I will **do it** as well. – The verb phrase *do it* is an anaphor; it points to the left toward its antecedent *buys a new bike*.
- **Cataphora** (included in the broad sense of anaphora, species of endophora)
- a. Because **he** was very cold, David put on his coat. – The pronoun *he* is a cataphor; it points to the right toward its postcedent*David*.
- b. **His** friends have been criticizing Jim for exaggerating. – The possessive adjective *his* is a cataphor; it points to the right toward its postcedent*Jim*.
- c. Although Sam might **do so**, I shall not buy a new bike. – The verb phrase *do so* is a cataphor; it points to the right toward its postcedent*buy a new bike*.
- d. In **their** free time, the boys play video games. – The possessive adjective *their* is a cataphor; it points to the right toward its postcedent*the boys*.

A further distinction is drawn between endophoric and exophoric reference. Exophoric reference occurs when an expression, an exophor, refers to something that is not directly present in the linguistic context, but is rather present in the situational context. Deictic proforms are stereotypical exophors, e.g.

- **Exophora**

- a. **This** garden hose is better than **that** one. – The demonstrative adjectives *this* and *that* are exophors; they point to entities in the situational context.
- b. Jerry is standing over **there**. – The adverb *there* is an exophor; it points to a location in the situational context.

Exophors cannot be anaphors as they do not substantially refer within the dialog or text, though there is a question of what portions of a conversation or document are accessed by a listener or reader with regard to whether all references to which a term points within that language stream are noticed (i.e., if you hear only a fragment of what someone says using the pronoun *her*, you might never discover who *she* is, though if you heard the rest of what the speaker was saying on the same occasion, you might discover who *she* is, either by anaphoric revelation or by exophoric implication because you realize who *she* must be according to what else is said about *her* even if *her* identity is not explicitly mentioned, as in the case of homophoric reference).

A listener might, for example, realize through listening to other clauses and sentences that *she* is *a Queen* because of some of her attributes or actions mentioned. But which queen? Homophoric reference occurs when a generic phrase obtains a specific meaning through knowledge of its context. For example, the referent of the phrase *the Queen* (using an emphatic definite article, not the less specific*a Queen*, but also not the more specific *Queen Elizabeth*) must be determined by the context of the utterance, which would identify the identity of the queen in question. Until further revealed by additional contextual words, gestures, images or other media, a listener would not even know what monarchy or historical period is being discussed, and even after hearing *her* name is *Elizabeth*

does not know, even if an English-UK Queen Elizabeth becomes indicated, if this queen means *Queen Elizabeth I* or *Queen Elizabeth II* and must await further clues in additional communications.Similarly, in discussing 'The Mayor' (of a city), the Mayor's identity must be understood broadly through the context which the speech references as general 'object' of understanding; is a particular human person meant, a current or future or past office-holder, the office in a strict legal sense, or the office in a general sense which includes activities a mayor might conduct, might even be expected to conduct, while they may not be explicitly defined for this office.

In generative grammar

The term *anaphor*is used in a special way in the generative grammar tradition. Here it denotes what would normally be called a reflexive or reciprocal pronoun, such as *himself* or *each other* in English, and analogous forms in other languages. The use of the term *anaphor* in this narrow sense is unique to generative grammar, and in particular, to the traditional binding theory. This theory investigates the syntactic relationship that can or must hold between a given proform and its antecedent (or postcedent). In this respect, anaphors (reflexive and reciprocal pronouns) behave very differently from, for instance, personal pronouns.

Complement anaphora

In some cases, anaphora may refer not to its usual antecedent, but to its complement set. In the following example a, the anaphoric pronoun *they* refers to the children who are eating

the ice-cream. Contrastingly, example b has *they* seeming to refer to the children who are not eating ice-cream:

- a. Only a few of the children ate their ice-cream. **They** ate the strawberry flavor first. – **They** meaning the children who ate ice-cream
- b. Only a few of the children ate their ice-cream. **They** threw it around the room instead. – **They** meaning either the children who did not eat ice-cream or perhaps the children who did not eat ice-cream and some of those who ate ice-cream but did not finish it or who threw around the ice-cream of those who did not eat it, or even all the children, those who ate ice-cream throwing around part of their ice-cream, the ice-cream of others, the same ice-cream which they may have eaten before or after throwing it, or perhaps only some of the children so that **they** does not mean to be all-inclusive

In its narrower definition, an anaphoric pronoun must refer to some noun (phrase) that has already been introduced into the discourse. In complement anaphora cases, however, the anaphor refers to something that is not yet present in the discourse, since the pronoun's referent has not been formerly introduced, including the case of 'everything but' what has been introduced. The set of ice-cream-eating-children in example b is introduced into the discourse, but then the pronoun *they* refers to the set of non-ice-cream-eating-children, a set which has not been explicitly mentioned.

Both semantic and pragmatics considerations attend this phenomenon, which following discourse representation theory since the early 1980s, such as work by Kamp (1981) and Heim (File Change Semantics, 1982), and generalized quantifier theory, such as work by Barwise and Cooper (1981), was

studied in a series of psycholinguistic experiments in the early 1990s by Moxey and Sanford (1993) and Sanford et al. (1994).In complement anaphora as in the case of the pronoun in example b, this anaphora refers to some sort of complement set (i.e. only to the set of non-ice-cream-eating-children) or to the maximal set (i.e. to all the children, both ice-cream-eating-children and non-ice-cream-eating-children) or some hybrid or variant set, including potentially one of those noted to the right of example b.The various possible referents in complement anaphora are discussed by Corblin (1996), Kibble (1997), and Nouwen (2003). Resolving complement anaphora is of interest in shedding light on brain access to information, calculation, mental modeling, communication.

Anaphora resolution – centering theory

There are many theories that attempt to prove how anaphors are related and trace back to their antecedents, with centering theory (Grosz, Joshi, and Weinstein 1983) being one of them. Taking the computational theory of mind view of language, centering theory gives a computational analysis of underlying antecedents. In their original theory, Grosz, Joshi, & Weinstein (1983) propose that some discourse entities in utterances are more "central" than others, and this degree of centrality imposes constraints on what can be the antecedent.

In the theory, there are different types of centers: forward facing, backwards facing, and preferred.

Forward facing centers

A ranked list of discourse entities in an utterance. The ranking is debated, some focusing on theta relations (Yıldırım et al. 2004) and some providing definitive lists.

Backwards facing center

The highest ranked discourse entity in the previous utterance.

Preferred center

The highest ranked discourse entity in the previous utterance realised in the current utterance.

Presupposition

In the branch of linguistics known as pragmatics, a **presupposition** (or **PSP**) is an implicit assumption about the world or background belief relating to an utterance whose truth is taken for granted in discourse. Examples of presuppositions include:

- *Jane no longer writes fiction.*
- Presupposition: Jane once wrote fiction.
- *Have you stopped eating meat?*
- Presupposition: you had once eaten meat.
- *Have you talked to Hans?*
- Presupposition: Hans exists.

A presupposition must be mutually known or assumed by the speaker and addressee for the utterance to be considered

appropriate in context. It will generally remain a necessary assumption whether the utterance is placed in the form of an assertion, denial, or question, and can be associated with a specific lexical item or grammatical feature (presupposition trigger) in the utterance.

Crucially, negation of an expression does not change its presuppositions: *I want to do it again* and *I don't want to do it again* both presuppose that the subject has done it already one or more times; *My wife is pregnant* and *My wife is not pregnant* both presuppose that the subject has a wife. In this respect, presupposition is distinguished from entailment and implicature. For example, *The president was assassinated* entails that *The president is dead*, but if the expression is negated, the entailment is not necessarily true.

Negation of a sentence containing a presupposition

If presuppositions of a sentence are not consistent with the actual state of affairs, then one of two approaches can be taken. Given the sentences*My wife is pregnant* and *My wife is not pregnant* when one has no wife, then either:

- Both the sentence and its negation are false; or
- Strawson's approach: Both "my wife is pregnant" and "my wife is not pregnant" use a wrong presupposition (i.e. that there exists a referent which can be described with the noun phrase *my wife*) and therefore can not be assigned truth values.

Bertrand Russell tries to solve this dilemma with two interpretations of the negated sentence:

- "There exists exactly one person, who is my wife and who is not pregnant"
- "There does not exist exactly one person, who is my wife and who is pregnant."

For the first phrase, Russell would claim that it is false, whereas the second would be true according to him.

Projection of presuppositions

A presupposition of a part of an utterance is sometimes also a presupposition of the whole utterance, and sometimes not. For instance, the phrase *my wife* triggers the presupposition that I have a wife. The first sentence below carries that presupposition, even though the phrase occurs inside an embedded clause. In the second sentence, however, it does not. John might be mistaken about his belief that I have a wife, or he might be deliberately trying to misinform his audience, and this has an effect on the meaning of the second sentence, but, perhaps surprisingly, not on the first one.

- John thinks that **my wife** is beautiful.
- John said that **my wife** is beautiful.

Thus, this seems to be a property of the main verbs of the sentences, *think* and *say*, respectively. After work by Lauri Karttunen, verbs that allow presuppositions to "pass up" to the whole sentence ("project") are called **holes**, and verbs that block such passing up, or *projection* of presuppositions are

called **plugs**. Some linguistic environments are intermediate between plugs and holes: They block some presuppositions and allow others to project.

These are called**filters**. An example of such an environment are indicative conditionals ("If-then" clauses). A conditional sentence contains an *antecedent* and a *consequent*. The antecedent is the part preceded by the word "if," and the consequent is the part that is (or could be) preceded by "then." If the consequent contains a presupposition trigger, and the triggered presupposition is explicitly stated in the antecedent of the conditional, then the presupposition is blocked. Otherwise, it is allowed to project up to the entire conditional. Here is an example:

- If **I have a wife**, then **my wife** is blonde.

Here, the presupposition triggered by the expression *my wife* (that I have a wife) is blocked, because it is stated in the antecedent of the conditional: That sentence doesn't imply that I have a wife. In the following example, it is not stated in the antecedent, so it is allowed to project, i.e. the sentence *does* imply that I have a wife.

- If it's already 4am, then **my wife** is probably angry.

Hence, conditional sentences act as *filters* for presuppositions that are triggered by expressions in their consequent.

A significant amount of current work in semantics and pragmatics is devoted to a proper understanding of when and how presuppositions project.

Presupposition triggers

A presupposition trigger is a lexical item or linguistic construction which is responsible for the presupposition, and thus "triggers" it. The following is a selection of presuppositional triggers following Stephen C. Levinson's classic textbook on *Pragmatics*, which in turn draws on a list produced by Lauri Karttunen. As is customary, the presuppositional triggers themselves are italicized, and the symbol » stands for 'presupposes'.

Definite descriptions

Definite descriptions are phrases of the form "the X" where X is a noun phrase. The description is said to be *proper* when the phrase applies to exactly one object, and conversely, it is said to be *improper* when either there exist more than one potential referents, as in "the senator from Ohio", or none at all, as in "the king of France". In conventional speech, definite descriptions are implicitly assumed to be proper, hence such phrases trigger the presupposition that the referent is unique and existent.

John saw *the man with two heads.*

»there exists a man with two heads.

Factive verbs

In Western epistemology, there is a tradition originating with Plato of defining knowledge as justified true belief. On this definition, for someone to know X, it is required that X be true.

A linguistic question thus arises regarding the usage of such phrases: does a person who states "John knows X" implicitly claim the truth of X? Steven Pinker explored this question in a popular science format in a 2007 book on language and cognition, using a widely publicized example from a speech by a U.S. president.

A 2003 speech by George W. Bush included the line, "British Intelligence has learned that Saddam Hussein recently sought significant quantities of uranium from Africa." Over the next few years, it became apparent that this intelligence lead was incorrect. But the way the speech was phrased, using a factive verb, implicitly framed the lead as truth rather than hypothesis. There is however a strong alternative view that *factivity thesis*, the proposition that relational predicates having to do with knowledge, such as *knows, learn, remembers,* and *realized*, presuppose the factual truth of their object, is incorrect.

- Martha *regrets* drinking John's home brew.
- Presupposition: Martha did in fact drink John's home brew.
- Frankenstein was *aware* that Dracula was there.
- Presupposition: Dracula was in fact there.
- John *realized* that he was in debt.
- Presupposition: John was in fact in debt.
- It was *odd* how proud he was.
- Presupposition: He was in fact proud.

Some further factive predicates: *know; be sorry that; be proud that; be indifferent that; be glad that; be sad that.*

Implicative verbs

- John *managed* to open the door.
- »John tried to open the door.
- John *forgot* to lock the door.
- »John ought to have locked, or intended to lock, the door.

Some further implicative predicates: *X happened to V*»X didn't plan or intend to *V*; *X avoided Ving*»X was expected to, or usually did, or ought to *V*, etc.

Change of state or continuation of state verbs

With these presupposition triggers, the current unfolding situation is considered presupposed information.

- John *stopped* teasing his wife.
- »John had been teasing his wife.
- Joan *began* teasing her husband.
- »Joan hadn't been teasing her husband.

Some further change of state verbs: *start; finish; carry on; cease; take* (as in *X took Y from Z* » Y was at/in/with Z); *leave; enter; come; go; arrive;* etc.

Iteratives

These types of triggers presuppose the existence of a previous state of affairs.

- The flying saucer came *again*.
- »The flying saucer came before.

- You can't get gobstoppers *anymore*.
- »You once could get gobstoppers.
- Carter *returned* to power.
- »Carter held power before.

Further iteratives: *another time; to come back; restore; repeat; for the nth time.*

Temporal clauses

The situation explained in a clause that begins with a temporal clause constructor is typically considered backgrounded information.

- *Before*Strawson was even born, Frege noticed presuppositions.
- »Strawson was born.
- *While* Chomsky was revolutionizing linguistics, the rest of social science was asleep.
- »Chomsky was revolutionizing linguistics.
- *Since* Churchill died, we've lacked a leader.
- »Churchill died.

Further temporal clause constructors: *after; during; whenever; as* (as in *As John was getting up, he slipped*).

Cleft sentences

Cleft sentence structures highlight particular aspects of a sentence and consider the surrounding information to be backgrounded knowledge. These sentences are typically not spoken to strangers, but rather to addressees who are aware of the ongoing situation.

- *Cleft construction:*It was Henry that kissed Rosie.
- »Someone kissed Rosie.
- *Pseudo-cleft construction:* What John lost was his wallet.
- »John lost something.

Comparisons and contrasts

Comparisons and contrasts may be marked by stress (or by other prosodic means), by particles like "too", or by comparatives constructions.

- Marianne called Adolph a male chauvinist, and then *HE* insulted *HER*.
- »For Marianne to call Adolph a male chauvinist would be to insult him.
- Carol *is a better linguist than*Barbara.
- »Barbara is a linguist.

Counterfactual conditionals

- *If* the notice *had* only *said* 'mine-field' in Welsh as well as in English, we *would* never *have lost* poor Llewellyn.
- »The notice didn't say 'mine-field' in Welsh.

Questions

Questions often presuppose what the assertive part of the question presupposes, but interrogative parts might introduce further presuppositions. There are three different types of

questions: yes/no questions, alternative questions and WH-questions.

- *Is there* a professor of linguistics at MIT?
- »Either there is a professor of linguistics at MIT or there isn't.
- *Is* Newcastle in England *or* in Australia?
- »Newcastle is in England or Newcastle is in Australia.
- *Who is* the professor of linguistics at MIT?
- »Someone is the professor of linguistics at MIT.

Possessive case

- John's children are very noisy.
- »John has children.

Accommodation of presuppositions

A presupposition of a sentence must normally be part of the common ground of the utterance context (the shared knowledge of the interlocutors) in order for the sentence to be felicitous. Sometimes, however, sentences may carry presuppositions that are not part of the common ground and nevertheless be felicitous. For example, I can, upon being introduced to someone, out of the blue explain that *my wife is a dentist,* this without my addressee having ever heard, or having any reason to believe that I have a wife. In order to be able to interpret my utterance, the addressee must assume that I have a wife. This process of an addressee assuming that a presupposition is true, even in the absence of explicit information that it is, is usually called**presupposition accommodation**. We have just

seen that presupposition triggers like *my wife* (definite descriptions) allow for such accommodation. In "Presupposition and Anaphora: Remarks on the Formulation of the Projection Problem", the philosopher Saul Kripke noted that some presupposition triggers do not seem to permit such accommodation. An example of that is the presupposition trigger ***too***. This word triggers the presupposition that, roughly, something parallel to what is stated has happened. For example, if pronounced with emphasis on *John*, the following sentence triggers the presupposition that somebody other than John had dinner in New York last night.

- John had dinner in New York last night, too.

But that presupposition, as stated, is completely trivial, given what we know about New York. Several million people had dinner in New York last night, and that in itself doesn't satisfy the presupposition of the sentence. What is needed for the sentence to be felicitous is really that somebody relevant to the interlocutors had dinner in New York last night, and that this has been mentioned in the previous discourse, or that this information can be recovered from it. Presupposition triggers that disallow accommodation are calledanaphoric presupposition triggers.

Presupposition in critical discourse analysis

Critical discourse analysis (CDA) is a broad study belonging to not one research category. It focuses on identifying presuppositions of an abstract nature from varying

perspectives. CDA is considered critical, not only in the sense of being analytical, but also in the ideological sense. Through the analysis of written texts and verbal speech, Teun A. van Dijk (2003) says CDA studies power imbalances existing in both the conversational and political spectrum. With the purpose of first identifying and then tackling inequality in society, van Dijk describes CDA as a nonconformist piece of work. One notable feature of ideological presuppositions researched in CDA is a concept termed synthetic personalisation

Logical construct

To describe a *presupposition* in the context of propositional calculus and truth-bearers, Belnap defines "A sentence is a *presupposition* of a question if the truth of the sentence is a necessary condition of the question's having some true answer." Then referring to the semantic theory of truth, interpretationsare used to formulate a *presupposition*: "Every interpretation which makes the question truly answerable is an interpretation which makes the presupposed sentence true as well."

A sentence that *expresses a presupposition* in a question may be characterized as follows: the question has some true answer if and only if the sentence is true.

Word-sense disambiguation

Word-sense disambiguation (**WSD**) is an open problem in computational linguistics concerned with identifying which

sense of a wordis used in a sentence. The solution to this issue impacts other computer-related writing, such as discourse, improving relevance of search engines, anaphora resolution, coherence, and inference.

Due to the fact thatnatural language requires reflection of neurological reality, as shaped by the abilities provided by the brain's neural networks, computer science has had a long-term challenge in developing the ability in computers to do natural language processing and machine learning.

Many techniques have been researched, including dictionary-based methods that use the knowledge encoded in lexical resources, supervised machine learning methods in which a classifier is trained for each distinct word on a corpus of manually sense-annotated examples, and completely unsupervised methods that cluster occurrences of words, thereby inducing word senses.

Among these, supervised learning approaches have been the most successful algorithms to date.

Accuracy of current algorithms is difficult to state without a host of caveats. In English, accuracy at the coarse-grained (homograph) level is routinely above 90%, with some methods on particular homographs achieving over 96%. On finer-grained sense distinctions, top accuracies from 59.1% to 69.0% have been reported in evaluation exercises (SemEval-2007, Senseval-2), where the baseline accuracy of the simplest possible algorithm of always choosing the most frequent sense was 51.4% and 57%, respectively.

About word-sense disambiguation

Disambiguation requires two strict inputs: a dictionary to specify the senses which are to be disambiguated and a corpus of language data to be disambiguated (in some methods, a training corpus of language examples is also required). WSD task has two variants: "lexical sample" (disambiguating the occurrences of a small sample of target words which were previously selected) and "all words" task (disambiguation of all the words in a running text). "All words" task is generally consider a more realistic form of evaluation, but the corpus is more expensive to produce because human annotators have to read the definitions for each word in the sequence every time they need to make a tagging judgement, rather than once for a block of instances for the same target word.

History

WSD was first formulated into as a distinct computational task during the early days of machine translation in the 1940s, making it one of the oldest problems in computational linguistics. Warren Weaver first introduced the problem in a computational context in his 1949 memorandum on translation. Later, Bar-Hillel (1960) argued that WSD could not be solved by "electronic computer" because of the need in general to model all world knowledge.

In the 1970s, WSD was a subtask of semantic interpretation systems developed within the field of artificial intelligence, starting with Wilks' preference semantics. However, since WSD

systems were at the time largely rule-based and hand-coded they were prone to a knowledge acquisition bottleneck.

By the 1980s large-scale lexical resources, such as the Oxford Advanced Learner's Dictionary of Current English (OALD), became available: hand-coding was replaced with knowledge automatically extracted from these resources, but disambiguation was still knowledge-based or dictionary-based.

In the 1990s, the statistical revolution advanced computational linguistics, and WSD became a paradigm problem on which to apply supervised machine learning techniques.

The 2000s saw supervised techniques reach a plateau in accuracy, and so attention has shifted to coarser-grained senses, domain adaptation, semi-supervised and unsupervised corpus-based systems, combinations of different methods, and the return of knowledge-based systems via graph-based methods. Still, supervised systems continue to perform best.

Difficulties

Differences between dictionaries

One problem with word sense disambiguation is deciding what the senses are, as different dictionaries and thesauruses will provide different divisions of words into senses. Some researchers have suggested choosing a particular dictionary, and using its set of senses to deal with this issue use. Generally, however, research results using broad distinctions in senses have been much better than those using narrow ones. Most researchers continue to work on fine-grained WSD.

Most research in the field of WSD is performed by using WordNet as a reference sense inventory for English. WordNet is a computational lexicon that encodes concepts as synonym sets (e.g. the concept of car is encoded as { car, auto, automobile, machine, motorcar }). Other resources used for disambiguation purposes include Roget's Thesaurus and Wikipedia. More recently, BabelNet, a multilingual encyclopedic dictionary, has been used for multilingual WSD.

Part-of-speech tagging

In any real test, part-of-speech tagging and sense tagging having been proven to be very closely related with each potentially making constraints to the other. The question whether these tasks should be kept together or decoupled is still not unanimously resolved, but recently scientists incline to test these things separately (e.g. in the Senseval/SemEval competitions parts of speech are provided as input for the text to disambiguate).

Both WSM part-of-speech tagging involve disambiguating or tagging with words. However, algorithms used for one do not tend to work well for the other, mainly because the part of speech of a word is primarily determined by the immediately adjacent one to three words, whereas the sense of a word may be determined by words further away. The success rate for part-of-speech tagging algorithms is at present much higher than that for WSD, state-of-the art being around 96% accuracy or better, as compared to less than 75% accuracy in word sense disambiguation with supervised learning. These figures are typical for English, and may be very different from those for other languages.

Inter-judge variance

Another problem is inter-judgevariance. WSD systems are normally tested by having their results on a task compared against those of a human. However, while it is relatively easy to assign parts of speech to text, training people to tag senses has been proven to be far more difficult.

While users can memorize all of the possible parts of speech a word can take, it is often impossible for individuals to memorize all of the senses a word can take. Moreover, humans do not agree on the task at hand – give a list of senses and sentences, and humans will not always agree on which word belongs in which sense.

As human performance serves as the standard, it is an upper bound for computer performance. Human performance, however, is much better on coarse-grained than fine-grained distinctions, so this again is why research on coarse-grained distinctions has been put to test in recent WSD evaluation exercises.

Pragmatics

Some AI researchers like Douglas Lenat argue that one cannot parse meanings from words without some form of common sense ontology. This linguistic issue is calledpragmatics. As agreed by researchers, to properly identify senses of words one must know common sense facts. Moreover, sometimes the common sense is needed to disambiguate such words like pronouns in case of having anaphoras or cataphoras in the text.

Sense inventory and algorithms' task-dependency

A task-independent sense inventory is not a coherent concept: each task requires its own division of word meaning into senses relevant to the task. Additionally, completely different algorithms might be required by different applications. In machine translation, the problem takes the form of target word selection. The "senses" are words in the target language, which often correspond to significant meaning distinctions in the source language ("bank" could translate to the French "banque"—that is, 'financial bank' or "rive"—that is, 'edge of river'). In information retrieval, a sense inventory is not necessarily required, because it is enough to know that a word is used in the same sense in the query and a retrieved document; what sense that is, is unimportant.

Discreteness of senses

Finally, the very notion of "word sense" is slippery and controversial. Most people can agree in distinctions at the coarse-grainedhomograph level (e.g., pen as writing instrument or enclosure), but go down one level to fine-grainedpolysemy, and disagreements arise. For example, in Senseval-2, which used fine-grained sense distinctions, human annotators agreed in only 85% of word occurrences. Word meaning is in principle infinitely variable and context-sensitive. It does not divide up easily into distinct or discrete sub-meanings. Lexicographers frequently discover in corpora loose and overlapping word meanings, and standard or conventional meanings extended, modulated, and exploited in a bewildering variety of ways. The art of lexicography is to generalize from the corpus to definitions that evoke and explain the full range of meaning of

a word, making it seem like words are well-behaved semantically. However, it is not at all clear if these same meaning distinctions are applicable in computational applications, as the decisions of lexicographers are usually driven by other considerations. In 2009, a task – named lexical substitution – was proposed as a possible solution to the sense discreteness problem. The task consists of providing a substitute for a word in context that preserves the meaning of the original word (potentially, substitutes can be chosen from the full lexicon of the target language, thus overcoming discreteness).

Approaches and methods

There are two main approaches to WSD – deep approaches and shallow approaches.

Deep approaches presume access to a comprehensive body of world knowledge. These approaches are generally not considered to be very successful in practice, mainly because such a body of knowledge does not exist in a computer-readable format, outside very limited domains. Additionally due to the long tradition in computational linguistics, of trying such approaches in terms of coded knowledge and in some cases, it can be hard to distinguish between knowledge involved in linguistic or world knowledge. The first attempt was that by Margaret Masterman and her colleagues, at the Cambridge Language Research Unit in England, in the 1950s. This attempt used as data a punched-card version of Roget's Thesaurus and its numbered "heads", as an indicator of topics and looked for repetitions in text, using a set intersection algorithm. It was not very successful, but had strong

relationships to later work, especially Yarowsky's machine learning optimisation of a thesaurus method in the 1990s.

Shallow approaches don't try to understand the text, but instead consider the surrounding words. These rules can be automatically derived by the computer, using a training corpus of words tagged with their word senses. This approach, while theoretically not as powerful as deep approaches, gives superior results in practice, due to the computer's limited world knowledge.

There are four conventional approaches to WSD:

- Dictionary- and knowledge-based methods: These rely primarily on dictionaries, thesauri, and lexical knowledge bases, without using any corpus evidence.
- Semi-supervised or minimally supervised methods: These make use of a secondary source of knowledge such as a small annotated corpus as seed data in a bootstrapping process, or a word-aligned bilingual corpus.
- Supervised methods: These make use of sense-annotated corpora to train from.
- Unsupervised methods: These eschew (almost) completely external information and work directly from raw unannotated corpora. These methods are also known under the name of word sense discrimination.

Almost all these approaches work by defining a window of *n* content words around each word to be disambiguated in the corpus, and statistically analyzing those *n* surrounding words. Two shallow approaches used to train and then disambiguate

are Naïve Bayes classifiers and decision trees. In recent research, kernel-based methods such as support vector machines have shown superior performance in supervised learning. Graph-based approaches have also gained much attention from the research community, and currently achieve performance close to the state of the art.

Dictionary- and knowledge-based methods

The Lesk algorithm is the seminal dictionary-based method. It is based on the hypothesis that words used together in text are related to each other and that the relation can be observed in the definitions of the words and their senses. Two (or more) words are disambiguated by finding the pair of dictionary senses with the greatest word overlap in their dictionary definitions. For example, when disambiguating the words in "pine cone", the definitions of the appropriate senses both include the words evergreen and tree (at least in one dictionary). A similar approach searches for the shortest path between two words: the second word is iteratively searched among the definitions of every semantic variant of the first word, then among the definitions of every semantic variant of each word in the previous definitions and so on. Finally, the first word is disambiguated by selecting the semantic variant which minimizes the distance from the first to the second word.

An alternative to the use of the definitions is to consider general word-sense relatedness and to compute the semantic similarity of each pair of word senses based on a given lexical knowledge base such as WordNet. Graph-based methods reminiscent of spreading activation research of the early days

of AI research have been applied with some success. More complex graph-based approaches have been shown to perform almost as well as supervised methods or even outperforming them on specific domains.

Recently, it has been reported that simple graph connectivity measures, such as degree, perform state-of-the-art WSD in the presence of a sufficiently rich lexical knowledge base. Also, automatically transferring knowledge in the form of semantic relations from Wikipedia to WordNet has been shown to boost simple knowledge-based methods, enabling them to rival the best supervised systems and even outperform them in a domain-specific setting.

The use of selectional preferences (or selectional restrictions) is also useful, for example, knowing that one typically cooks food, one can disambiguate the word bass in "I am cooking basses" (i.e., it's not a musical instrument).

Supervised methods

Supervised methods are based on the assumption that the context can provide enough evidence on its own to disambiguate words (hence, common sense and reasoning are deemed unnecessary). Probably every machine learning algorithm going has been applied to WSD, including associated techniques such as feature selection, parameter optimization, and ensemble learning. Support Vector Machines and memory-based learninghave been shown to be the most successful approaches, to date, probably because they can cope with the high-dimensionality of the feature space. However, these supervised methods are subject to a new knowledge acquisition

bottleneck since they rely on substantial amounts of manually sense-tagged corpora for training, which are laborious and expensive to create.

Semi-supervised methods

Because of the lack of training data, many word sense disambiguation algorithms use semi-supervised learning, which allows both labeled and unlabeled data. The Yarowsky algorithm was an early example of such an algorithm. It uses the ‘One sense per collocation’ and the ‘One sense per discourse’ properties of human languages for word sense disambiguation. From observation, words tend to exhibit only one sense in most given discourse and in a given collocation.

The bootstrapping approach starts from a small amount of seed data for each word: either manually tagged training examples or a small number of surefire decision rules (e.g., 'play' in the context of 'bass' almost always indicates the musical instrument). The seeds are used to train an initial classifier, using any supervised method. This classifier is then used on the untagged portion of the corpus to extract a larger training set, in which only the most confident classifications are included. The process repeats, each new classifier being trained on a successively larger training corpus, until the whole corpus is consumed, or until a given maximum number of iterations is reached.

Other semi-supervised techniques use large quantities of untagged corpora to provide co-occurrence information that supplements the tagged corpora. These techniques have the

potential to help in the adaptation of supervised models to different domains.

Also, an ambiguous word in one language is often translated into different words in a second language depending on the sense of the word. Word-aligned bilingual corpora have been used to infer cross-lingual sense distinctions, a kind of semi-supervised system.

Unsupervised methods

Unsupervised learning is the greatest challenge for WSD researchers. The underlying assumption is that similar senses occur in similar contexts, and thus senses can be induced from text by clustering word occurrences using some measure of similarity of context, a task referred to as word sense induction or discrimination.

Then, new occurrences of the word can be classified into the closest induced clusters/senses. Performance has been lower than for the other methods described above, but comparisons are difficult since senses induced must be mapped to a known dictionary of word senses. If a mapping to a set of dictionary senses is not desired, cluster-based evaluations (including measures of entropy and purity) can be performed. Alternatively, word sense induction methods can be tested and compared within an application. For instance, it has been shown that word sense induction improves Web search result clustering by increasing the quality of result clusters and the degree diversification of result lists. It is hoped that unsupervised learning will overcome the knowledge acquisition bottleneck because they are not dependent on manual effort.

Representing words considering their context through fixed size dense vectors (word embeddings) has become one of the most fundamental blocks in several NLP systems. Even though most of traditional word embedding techniques conflate words with multiple meanings into a single vector representation, they still can be used to improve WSD. In addition to word embeddings techniques, lexical databases (e.g., WordNet, ConceptNet, BabelNet) can also assist unsupervised systems in mapping words and their senses as dictionaries. Some techniques that combine lexical databases and word embeddingsare presented in AutoExtend and Most Suitable Sense Annotation (MSSA). In AutoExtend, they present a method that decouples an object input representation into its properties, such as words and their word senses. AutoExtend uses a graph structure to map words (e.g. text) and non-word (e.g. synsets in WordNet) objects as nodes and the relationship between nodes as edges. The relations (edges) in AutoExtend can either express the addition or similarity between its nodes. The former captures the intuition behind the offset calculus, while the latter defines the similarity between two nodes. In MSSA, an unsupervised disambiguation system uses the similarity between word senses in a fixed context window to select the most suitable word sense using a pre-trained word embedding model and WordNet. For each context window, MSSA calculates the centroid of each word sense definition by averaging the word vectors of its words in WordNet's glosses (i.e., short defining gloss and one or more usage example) using a pre-trained word embeddings model. These centroids are later used to select the word sense with the highest similarity of a target word to its immediately adjacent neighbors (i.e., predecessor and successor words). After all words are annotated and disambiguated, they can be used as a

training corpus in any standard word embedding technique. In its improved version, MSSA can make use of word sense embeddings to repeat its disambiguation process iteratively.

Other approaches

Other approaches may vary differently in their methods:

- Domain-driven disambiguation;
- Identification of dominant word senses;
- WSD using Cross-Lingual Evidence.
- WSD solution in John Ball's language independent NLU combining Patom Theory [1] and RRG (Role and Reference Grammar)
- Type inference in constraint-based grammars

Other languages

- **Hindi** : Lack of lexical resources in Hindi have hindered the performance of supervised models of WSD, while the unsupervised models suffer due to extensive morphology. A possible solution to this problem is the design of a WSD model by means of parallel corpora. The creation of the Hindi WordNet has paved way for several Supervised methods which have been proven to produce a higher accuracy in disambiguating nouns.

Local impediments and summary

The knowledge acquisition bottleneck is perhaps the major impediment to solving the WSD problem. Unsupervised

methods rely on knowledge about word senses, which is only sparsely formulated in dictionaries and lexical databases. Supervised methods depend crucially on the existence of manually annotated examples for every word sense, a requisite that can so far be met only for a handful of words for testing purposes, as it is done in the Senseval exercises.

One of the most promising trends in WSD research is using the largest corpus ever accessible, the World Wide Web, to acquire lexical information automatically. WSD has been traditionally understood as an intermediate language engineering technology which could improve applications such as information retrieval (IR). In this case, however, the reverse is also true: web search engines implement simple and robust IR techniques that can successfully mine the Web for information to use in WSD. The historic lack of training data has provoked the appearance of some new algorithms and techniques, as described in Automatic acquisition of sense-tagged corpora.

External knowledge sources

Knowledge is a fundamental component of WSD. Knowledge sources provide data which are essential to associate senses with words. They can vary from corpora of texts, either unlabeled or annotated with word senses, to machine-readable dictionaries, thesauri, glossaries, ontologies, etc. They can be classified as follows:

Structured:

- Machine-readable dictionaries (MRDs)
- Ontologies

- Thesauri

Unstructured:

- Collocation resources
- Other resources (such as word frequency lists, stoplists, domain labels, etc.)
- Corpora: raw corpora and sense-annotated corpora

Evaluation

Comparing and evaluating different WSD systems is extremely difficult, because of the different test sets, sense inventories, and knowledge resources adopted. Before the organization of specific evaluation campaigns most systems were assessed on in-house, often small-scale, data sets. In order to test one's algorithm, developers should spend their time to annotate all word occurrences. And comparing methods even on the same corpus is not eligible if there is different sense inventories.

In order to define common evaluation datasets and procedures, public evaluation campaigns have been organized. Senseval (now renamed SemEval) is an international word sense disambiguation competition, held every three years since 1998: Senseval-1 (1998), Senseval-2 (2001), Senseval-3 (2004), and its successor, SemEval (2007). The objective of the competition is to organize different lectures, preparing and hand-annotating corpus for testing systems, perform a comparative evaluation of WSD systems in several kinds of tasks, including all-words and lexical sample WSD for different languages, and, more recently, new tasks such as semantic role labeling, gloss WSD, lexical substitution, etc. The systems submitted for

evaluation to these competitions usually integrate different techniques and often combine supervised and knowledge-based methods (especially for avoiding bad performance in lack of training examples).

In recent years , the WSD evaluation task choices had grown and the criterion for evaluating WSD has changed drastically depending on the variant of the WSD evaluation task. Below enumerates the variety of WSD tasks:

Task design choices

As technology evolves, the Word Sense Disambiguation (WSD) tasks grows in different flavors towards various research directions and for more languages:

- Classic monolingual WSD evaluation tasks use WordNet as the sense inventory and are largely based on supervised/semi-supervised classification with the manually sense annotated corpora:
- Classic English WSD uses the Princeton WordNet as it sense inventory and the primary classification input is normally based on the SemCor corpus.
- Classical WSD for other languages uses their respective WordNet as sense inventories and sense annotated corpora tagged in their respective languages. Often researchers will also tapped on the SemCor corpus and aligned bitexts with English as its source language
- Cross-lingual WSD evaluation task is also focused on WSD across 2 or more languages simultaneously. Unlike the Multilingual WSD tasks, instead of

providing manually sense-annotated examples for each sense of a polysemous noun, the sense inventory is built up on the basis of parallel corpora, e.g. Europarl corpus.

- Multilingual WSD evaluation tasks focused on WSD across 2 or more languages simultaneously, using their respective WordNets as its sense inventories or BabelNet as multilingual sense inventory. It evolved from the Translation WSD evaluation tasks that took place in Senseval-2. A popular approach is to carry out monolingual WSD and then map the source language senses into the corresponding target word translations.
- Word Sense Induction and Disambiguation task is a combined task evaluation where the sense inventory is first induced from a fixed training set data, consisting of polysemous words and the sentence that they occurred in, then WSD is performed on a different testing data set.

Software

- Babelfy, a unified state-of-the-art system for multilingual Word Sense Disambiguation and Entity Linking
- BabelNet API, a Java API for knowledge-based multilingual Word Sense Disambiguation in 6 different languages using the BabelNetsemantic network

- WordNet::SenseRelate, a project that includes free, open source systems for word sense disambiguation and lexical sample sense disambiguation
- UKB: Graph Base WSD, a collection of programs for performing graph-based Word Sense Disambiguation and lexical similarity/relatedness using a pre-existing Lexical Knowledge Base
- pyWSD, python implementations of Word Sense Disambiguation (WSD) technologies

Semantic role labeling

In natural language processing, **semantic role labeling** (also called **shallow semantic parsing** or **slot-filling**) is the process that assigns labels to words or phrases in a sentence that indicates their semantic role in the sentence, such as that of an agent, goal, or result.

It serves to find the meaning of the sentence. To do this, it detects the arguments associated with the predicate or verb of a sentence and how they are classified into their specific roles. A common example is the sentence "Mary sold the book to John."

The agent is "Mary," the predicate is "sold" (or rather, "to sell,") the theme is "the book," and the recipient is "John." Another example is how "the book belongs to me" would need two labels such as "possessed" and "possessor" and "the book was sold to John" would need two other labels such as theme and recipient, despite these two clauses being similar to "subject" and "object" functions.

History

In 1968, the first idea for semantic role labeling was proposed by Charles J. Fillmore. His proposal led to the FrameNetproject which produced the first major computational lexicon that systematically described many predicates and their corresponding roles. Daniel Gildea (Currently at University of Rochester, previously University of California, Berkeley / International Computer Science Institute) and Daniel Jurafsky (currently teaching at Stanford University, but previously working at University of Colorado and UC Berkeley) developed the first automatic semantic role labeling system based on FrameNet. The PropBank corpus added manually created semantic role annotations to the Penn Treebank corpus of Wall Street Journal texts. Many automatic semantic role labeling systems have used PropBank as a training dataset to learn how to annotate new sentences automatically.

Uses

Semantic role labeling is mostly used for machines to understand the roles of words within sentences. This benefits applications similar to Natural Language Processing programs that need to understand not just the words of languages, but how they can be used in varying sentences. A better understanding of semantic role labeling could lead to advancements in question answering, information extraction, automatic text summarization, text data mining, and speech recognition.

Chapter 5

Semantics (Psychology)

Semantics within psychology is the study of how meaning is stored in the mind. Semantic memory is a type of long-term declarative memory that refers to facts or ideas which are not immediately drawn from personal experience. It was first theorized in 1972 by W. Donaldson and EndelTulving. Tulving employs the word semantic to describe a system of memory that involves “words and verbal symbols, their meanings and referents, the relations between them, and the rules, formulas, or algorithms for influencing them”.

Ideasthesia

Ideasthesia is a psychological phenomenon in which activation of concepts evokes sensory experiences. For example, in synesthesia, activation of a concept of a letter (e.g., that of the letter *A*) evokes sensory-like experiences (e.g., of red color).

Psychosemantics

In the 1960s, psychosemantic studies became popular after Charles E. Osgood's massive cross-cultural studies using his semantic differential (SD) method that used thousands of nouns and adjective bipolar scales. A specific form of the SD, Projective Semantics method uses only most common and neutral nouns that correspond to the 7 groups (factors) of adjective-scales most consistently found in cross-cultural studies (Evaluation, Potency, Activity as found by Osgood, and

Reality, Organization, Complexity, Limitation as found in other studies). In this method, seven groups of bipolar adjective scales corresponded to seven types of nouns so the method was thought to have the object-scale symmetry (OSS) between the scales and nouns for evaluation using these scales. For example, the nouns corresponding to the listed 7 factors would be: Beauty, Power, Motion, Life, Work, Chaos, Law. Beauty was expected to be assessed unequivocally as “very good” on adjectives of Evaluation-related scales, Life as “very real” on Reality-related scales, etc. However, deviations in this symmetric and very basic matrix might show underlying biases of two types: scales-related bias and objects-related bias. This OSS design meant to increase the sensitivity of the SD method to any semantic biases in responses of people within the same culture and educational background.

Semantic memory

Semantic memory is one of the two types of explicit memory (or declarative memory) (our memory of facts or events that is explicitly stored and retrieved). Semantic memory refers to general world knowledge that we have accumulated throughout our lives. This general knowledge (facts, ideas, meaning and concepts) is intertwined in experience and dependent on culture. Semantic memory is distinct from episodic memory, which is our memory of experiences and specific events that occur during our lives, from which we can recreate at any given point. For instance, semantic memory might contain information about what a cat is, whereas episodic memory might contain a specific memory of petting a particular cat. We can learn about new concepts by applying our knowledge

learned from things in the past. The counterpart to declarative or explicit memory is nondeclarative memory or implicit memory.

History

The idea of semantic memory was first introduced following a conference in 1972 between EndelTulving, of the University of Toronto, and W. Donaldson on the role of organization in human memory. Tulving constructed a proposal to distinguish between episodic memory and what he termed semantic memory. He was mainly influenced by the ideas of Reiff and Scheers, who in 1959 made the distinction between two primary forms of memory. One form was entitled "remembrances", the other "memoria". The remembrance concept dealt with memories that contained experiences of an autobiographic index, whereas the memoria concept dealt with those memories that did not reference experiences having an autobiographic index.

Semantic memory reflects our knowledge of the world around us, hence the term 'general knowledge' is often used. It holds generic information that is more than likely acquired across various contexts and is used across different situations. According to Madigan in his book titled *Memory*, semantic memory is the sum of all knowledge one has obtained—whether it be vocabulary, understanding of math, or all the facts one knows. In his book titled "Episodic and Semantic Memory", EndelTulving adopted the term "semantic" from linguists to refer to a system of memory for "words and verbal symbols, their meanings and referents, the relations between them, and the rules, formulas, or algorithms for influencing them." The

use of semantic memory is quite different from that of episodic memory. Semantic memory refers to general facts and meanings one shares with others whereas episodic memory refers to unique and concrete personal experiences. Tulving's proposal of this distinction between semantic and episodic memory was widely accepted, primarily because it allowed the separate conceptualization of knowledge of the world. Tulving discusses conceptions of episodic and semantic memory in his book titled Elements of Episodic Memory, in which he states that several factors differentiate between episodic memory and semantic memory in ways that include

- the characteristics of their operations,
- the kind of information they process,
- their application to the real world as well as the memory laboratory.

Before Tulving's proposal, this area of human memory had been neglected by experimental psychologists. Since Tulving's inception of these distinctions, several experimenters have conducted tests to determine the validity of his hypothesized differences between episodic and semantic memory.

Recent research has focused on the idea that when people access a word's meaning, sensorimotor information that is used to perceive and act on the concrete object the word suggests is automatically activated. In the theory of grounded cognition, the meaning of a particular word is grounded in the sensorimotor systems. For example, when one thinks of a pear, knowledge of grasping, chewing, sights, sounds, and tastes used to encode episodic experiences of a pear are recalled through sensorimotor simulation. A grounded simulation

approach refers to context-specific re-activations that integrate the important features of episodic experience into a current depiction. Such research has challenged previously utilized amodal views. The brain encodes multiple inputs such as words and pictures to integrate and create a larger conceptual idea by using amodal views (also known as amodal perception). Instead of being representations in modality-specific systems, semantic memory representations had previously been viewed as redescriptions of modality-specific states. Some accounts of category-specific semantic deficits that are amodal remain even though researchers are beginning to find support for theories in which knowledge is tied to modality-specific brain regions. This research defines a clear link between episodic experiences and semantic memory. The concept that semantic representations are grounded across modality-specific brain regions can be supported by the fact that episodic and semantic memory appear to function in different yet mutually dependent ways. The distinction between semantic and episodic memory has become a part of the broader scientific discourse. For example, it has been speculated that semantic memory captures the stable aspects of our personality while episodes of illness may have a more episodic nature.

Empirical evidence

Jacoby and Dallas (1981)

This study was not created to solely provide evidence for the distinction of semantic and episodic memory stores. However, they did use the experimental dissociation method which provides evidence for Tulving's hypothesis.

- Part one

Subjects were presented with 60 words (one at a time) and were asked different questions.

- Some questions asked were to cause the subject to pay attention to the visual **appearance**: Is the word typed in bold letters?
- Some questions caused the participants to pay attention to the **sound** of the word: Does the word rhyme with ball?
- Some questions caused the subjects to pay attention to the **meaning** of the word: Does the word refer to a form of communication?
- Half of the questions were "no" answers and the other half "yes"
- Part Two

In the second phase of the experiment, 60 "old words" seen in stage one and "20 new words" not shown in stage one were presented to the subjects one at a time.

The subjects were given one of two tasks:

- *Perceptual Identification task (semantic)*: The words were flashed on a video-screen for 35ms and the subjects were required to say what the word was.
- *Episodic Recognition Task*: Subjects were presented with each word and had to decide whether they had seen the word in the previous stage of the experiment.
- Results:

- The percentages correct in the Semantic task (perceptual identification) did not change with the encoding conditions of appearance, sound, or meaning.
- The percentages for the episodic task increased from the appearance condition (.50), to the sound condition (.63), to the meaning condition (.86). – The effect was also greater for the "yes" encoding words than the "no" encoding words. (see stage one)
- Conclusion:

It displays a strong distinction of performance of episodic and semantic tasks, thus supporting Tulving's hypothesis.

Models

The essence of semantic memory is that its contents are not tied to any particular instance of experience, as in episodic memory. Instead, what is stored in semantic memory is the "gist" of experience, an abstract structure that applies to a wide variety of experiential objects and delineates categorical and functional relationships between such objects. Thus, a complete theory of semantic memory must account not only for the representational structure of such "gists", but also for how they can be extracted from experience. Numerous models of semantic memory have been proposed; they are summarized below.

Network models

Networks of various sorts play an integral part in many theories of semantic memory. Generally speaking, a network is

composed of a set of nodes connected by links. The nodes may represent concepts, words, perceptual features, or nothing at all. The links may be weighted such that some are stronger than others or, equivalently, have a length such that some links take longer to traverse than others. All these features of networks have been employed in models of semantic memory, examples of which are found below.

Teachable Language Comprehender (TLC)

One of the first examples of a network model of semantic memory is the Teachable Language Comprehender (TLC). In this model, each node is a word, representing a concept (like "Bird"). With each node is stored a set of properties (like "can fly" or "has wings") as well as pointers (i.e., links) to other nodes (like "Chicken").

A node is directly linked to those nodes of which it is either a subclass or superclass (i.e., "Bird" would be connected to both "Chicken" and "Animal"). Thus, TLC is a hierarchical knowledge representation in that high-level nodes representing large categories are connected (directly or indirectly, via the nodes of subclasses) to many instances of those categories, whereas nodes representing specific instances are at a lower level, connected only to their superclasses.

Furthermore, properties are stored at the highest category level to which they apply. For example, "is yellow" would be stored with "Canary", "has wings" would be stored with "Bird" (one level up), and "can move" would be stored with "Animal" (another level up). Nodes may also store negations of the properties of their superordinate nodes (i.e., "NOT-can fly"

would be stored with "penguin"). This provides an economy of representation in that properties are only stored at the category level at which they become essential, that is, at which point they become critical features (see below).

Processing in TLC is a form of spreading activation. That is, when a node becomes active, that activation spreads to other nodes via the links between them. In that case, the time to answer the question "Is a chicken a bird?" is a function of how far the activation between the nodes for "Chicken" and "Bird" must spread, i.e., the number of links between the nodes "Chicken" and "Bird".

The original version of TLC did not put weights on the links between nodes. This version performed comparably to humans in many tasks, but failed to predict that people would respond faster to questions regarding more typical category instances than those involving less typical instances.

Collins and Quillian later updated TLC to include weighted connections to account for this effect. This updated TLC is capable of explaining both the familiarity effect and the typicality effect. Its biggest advantage is that it clearly explains priming: you are more likely to retrieve information from memory if related information (the "prime") has been presented a short time before. There are still a number of memory phenomena for which TLC has no account, including why people are able to respond quickly to obviously false questions (like "is a chicken a meteor?"), when the relevant nodes are very far apart in the network.

Semantic networks

TLC is an instance of a more general class of models known as semantic networks. In a semantic network, each node is to be interpreted as representing a specific concept, word, or feature. That is, each node is a symbol. Semantic networks generally do not employ distributed representations for concepts, as may be found in a neural network. The defining feature of a semantic network is that its links are almost always directed (that is, they only point in one direction, from a base to a target) and the links come in many different types, each one standing for a particular relationship that can hold between any two nodes. Processing in a semantic network often takes the form of spreading activation (see above).

Semantic networks see the most use in models of discourse and logicalcomprehension, as well as in Artificial Intelligence. In these models, the nodes correspond to words or word stems and the links represent syntactic relations between them. For an example of a computational implementation of semantic networks in knowledge representation, see Cravo and Martins (1993).

Feature models

Feature models view semantic categories as being composed of relatively unstructured sets of features. The semantic feature-comparison model, proposed by Smith, Shoben, and Rips (1974), describes memory as being composed of feature lists for different concepts. According to this view, the relations between categories would not be directly retrieved, they would be indirectly computed. For example, subjects might verify a

sentence by comparing the feature sets that represent its subject and predicate concepts. Such computational feature-comparison models include the ones proposed by Meyer (1970), Rips (1975), Smith, et al. (1974).

Early work in perceptual and conceptual categorization assumed that categories had critical features and that category membership could be determined by logical rules for the combination of features. More recent theories have accepted that categories may have an ill-defined or "fuzzy" structure and have proposed probabilistic or global similarity models for the verification of category membership.

Associative models

The "association"—a relationship between two pieces of information—is a fundamental concept in psychology, and associations at various levels of mental representation are essential to models of memory and cognition in general. The set of associations among a collection of items in memory is equivalent to the links between nodes in a network, where each node corresponds to a unique item in memory. Indeed, neural networks and semantic networks may be characterized as associative models of cognition. However, associations are often more clearly represented as an $N \times N$ matrix, where N is the number of items in memory. Thus, each cell of the matrix corresponds to the strength of the association between the row item and the column item.

Learning of associations is generally believed to be a Hebbian process; that is, whenever two items in memory are simultaneously active, the association between them grows

stronger, and the more likely either item is to activate the other. See below for specific operationalizations of associative models.

Search of Associative Memory (SAM)

A standard model of memory that employs association in this manner is the Search of Associative Memory (SAM) model. Though SAM was originally designed to model episodic memory, its mechanisms are sufficient to support some semantic memory representations, as well. The SAM model contains a short-term store (STS) and long-term store (LTS), where STS is a briefly activated subset of the information in the LTS. The STS has limited capacity and affects the retrieval process by limiting the amount of information that can be sampled and limiting the time the sampled subset is in an active mode.

The retrieval process in LTS is cue dependent and probabilistic, meaning that a cue initiates the retrieval process and the selected information from memory is random. The probability of being sampled is dependent on the strength of association between the cue and the item being retrieved, with stronger associations being sampled and finally one is chosen. The buffer size is defined as r, and not a fixed number, and as items are rehearsed in the buffer the associative strengths grow linearly as a function of the total time inside the buffer. In SAM, when any two items simultaneously occupy a working memory buffer, the strength of their association is incremented. Thus, items that co-occur more often are more strongly associated. Items in SAM are also associated with a specific context, where the strength of that association

determined by how long each item is present in a given context. In SAM, then, memories consist of a set of associations between items in memory and between items and contexts. The presence of a set of items and/or a context is more likely to evoke, then, some subset of the items in memory. The degree to which items evoke one another—either by virtue of their shared context or their co-occurrence—is an indication of the items' semantic relatedness.

In an updated version of SAM, pre-existing semantic associations are accounted for using a semantic matrix. During the experiment, semantic associations remain fixed showing the assumption that semantic associations are not significantly impacted by the episodic experience of one experiment. The two measures used to measure semantic relatedness in this model are the Latent semantic analysis (LSA) and the Word association spaces (WAS). The LSA method states that similarity between words is reflected through their co-occurrence in a local context. WAS was developed by analyzing a database of free association norms. In WAS, "words that have similar associative structures are placed in similar regions of space."

ACT-R: a production system model

The ACT (Adaptive Control of Thought) (and later ACT-R (Adaptive Control of Thought-Rational)) theory of cognition represents declarative memory (of which semantic memory is a part) with "chunks", which consist of a label, a set of defined relationships to other chunks (i.e., "this is a _", or "this has a _"), and any number of chunk-specific properties. Chunks, then, can be mapped as a semantic network, given that each

node is a chunk with its unique properties, and each link is the chunk's relationship to another chunk. In ACT, a chunk's activation decreases as a function of the time since the chunk was created and increases with the number of times the chunk has been retrieved from memory.

Chunks can also receive activation from Gaussiannoise, and from their similarity to other chunks. For example, if "chicken" is used as a retrieval cue, "canary" will receive activation by virtue of its similarity to the cue (i.e., both are birds, etc.). When retrieving items from memory, ACT looks at the most active chunk in memory; if it is above threshold, it is retrieved, otherwise an "error of omission" has occurred, i.e., the item has been forgotten. There is, additionally, a retrieval latency, which varies inversely with the amount by which the activation of the retrieved chunk exceeds the retrieval threshold. This latency is used in measuring the response time of the ACT model, to compare it to human performance.

While ACT is a model of cognition in general, and not memory in particular, it nonetheless posits certain features of the structure of memory, as described above. In particular, ACT models memory as a set of related symbolic chunks which may be accessed by retrieval cues. While the model of memory employed in ACT is similar in some ways to a semantic network, the processing involved is more akin to an associative model.

Statistical models

Some models characterize the acquisition of semantic information as a form of statistical inference from a set of

discrete experiences, distributed across a number of "contexts". Though these models differ in specifics, they generally employ an (Item × Context) matrix where each cell represents the number of times an item in memory has occurred in a given context. Semantic information is gleaned by performing a statistical analysis of this matrix.

Many of these models bear similarity to the algorithms used in search engines (for example, see Griffiths, *et al.*, 2007 and Anderson, 1990), though it is not yet clear whether they really use the same computational mechanisms.

Other statistical models of semantic memory

The success of LSA and HAL gave birth to a whole field of statistical models of language. A more up-to-date list of such models may be found under the topic Measures of semantic relatedness.

Location of semantic memory in the brain

The cognitive neuroscience of semantic memory is a somewhat controversial issue with two dominant views.

On the one hand, many researchers and clinicians believe that semantic memory is stored by the same brain systems involved in episodic memory, that is, the medial temporal lobes (MTL), including the hippocampal formation. In this system, the hippocampal formation "encodes" memories, or makes it

possible for memories to form at all, and the neocortex stores memories after the initial encoding process is completed.

Recently, new evidence has been presented in support of a more precise interpretation of this hypothesis. The hippocampal formation includes, among other structures: the hippocampus itself, the entorhinal cortex, and the perirhinal cortex.

These latter two make up the "parahippocampal cortices". Amnesics with damage to the hippocampus but some spared parahippocampal cortex were able to demonstrate some degree of intact semantic memory despite a total loss of episodic memory. This strongly suggests that encoding of information leading to semantic memory does not have its physiological basis in the hippocampus.

Other researchers believe the hippocampus is only involved in episodic memory and spatial cognition. This then raises the question where semantic memory may be located. Some believe semantic memory lives in temporal cortex. Others believe that semantic knowledge is widely distributed across all brain areas. To illustrate this latter view, consider your knowledge of dogs. Researchers holding the 'distributed semantic knowledge' view believe that your knowledge of the sound a dog makes exists in your auditory cortex, whilst your ability to recognize and imagine the visual features of a dog resides in your visual cortex. Recent evidence supports the idea that the temporal pole bilaterally is the convergence zone for unimodal semantic representations into a multimodal representation. These regions are particularly vulnerable to damage in semantic dementia, which is characterised by a global semantic deficit.

Neural correlates and biological workings

The hippocampal areas are important to semantic memory's involvement with declarative memory. The left inferior prefrontal cortex (PFC) and the left posterior temporal areas are other areas involved in semantic memory use. Temporal lobe damage affecting the lateral and medial cortexes have been related to semantic impairments. Damage to different areas of the brain affect semantic memory differently.

Neuroimaging evidence suggests that left hippocampal areas show an increase in activity during semantic memory tasks. During semantic retrieval, two regions in the right middle frontal gyrus and the area of the right inferior temporal gyrus similarly show an increase in activity. Damage to areas involved in semantic memory result in various deficits, depending on the area and type of damage. For instance, Lambon Ralph, Lowe, & Rogers (2007) found that category-specific impairments can occur where patients have different knowledge deficits for one semantic category over another, depending on location and type of damage. Category-specific impairments might indicate that knowledge may rely differentially upon sensory and motor properties encoded in separate areas (Farah and McClelland, 1991).

Category-specific impairments can involve cortical regions where living and nonliving things are represented and where feature and conceptual relationships are represented. Depending on the damage to the semantic system, one type might be favored over the other. In many cases, there is a point

where one domain is better than the other (i.e. - representation of living and nonliving things over feature and conceptual relationships or vice versa)

Different diseases and disorders can affect the biological workings of semantic memory. A variety of studies have been done in an attempt to determine the effects on varying aspects of semantic memory.

For example, Lambon, Lowe, & Rogers (2007) studied the different effects semantic dementia and herpes simplex virus encephalitis have on semantic memory. They found that semantic dementia has a more generalized semantic impairment. Additionally, deficits in semantic memory as a result of herpes simplex virus encephalitis tend to have more category-specific impairments. Other disorders that affect semantic memory - such as Alzheimer's disease - has been observed clinically as errors in naming, recognizing, or describing objects. Whereas researchers have attributed such impairment to degradation of semantic knowledge (Koenig et al. 2007).

Various neural imaging and research points to semantic memory and episodic memory resulting from distinct areas in the brain. Still other research suggests that both semantic memory and episodic memory are part of a singular declarative memory system, yet represent different sectors and parts within the greater whole.

Different areas within the brain are activated depending on whether semantic or episodic memory is accessed. Certain experts are still arguing whether or not the two types of memory are from distinct systems or whether the neural

imaging makes it appear that way as a result of the activation of different mental processes during retrieval.

Disorders

Category specific semantic impairments

Category specific semantic impairments are a neuropsychological occurrence in which an individual ability to identify certain categories of objects is selectively impaired while other categories remain undamaged. This condition can result in brain damage which can be widespread, patchy, or localized to a specific part of the brain. Research suggests that the temporal lobe, more specifically the structural description system might be responsible for category specific impairments of semantic memory disorders.

Impairment categories

Category specific semantic deficits tend to fall into two different categories, each of which can be spared or emphasized depending on the individual's specific deficit. The first category consists of animate objects with "animals" being the most common deficit. The second category consists of inanimate objects with two subcategories of "fruits and vegetables" (biological inanimate objects) and "artifacts" being the most common deficits. The type of deficit, however, does not indicate a lack of conceptual knowledge associated with that category. This is because the visual system used to identify and describe the structure of objects functions independently of an individual's conceptual knowledge base.

Most of the time, these two categories are consistent with case-study data. However, there are a few exceptions to the rule as is the case with most neuropsychological conditions. Things like food, body parts, and musical instruments have been shown to defy the animate/inanimate or biological/non-biological categorical division. For example, it has been shown that musical instruments tend to be impaired in patients with damage to the living things category despite the fact that musical instruments fall in the non-biological/inanimate category. However, there are also cases of biological impairment where musical instrument performance is at a normal level. Similarly, food has been shown to be impaired in those with biological category impairments. The category of food specifically can present some irregularities though because it can be natural, but it can also be highly processed. This can be seen in a case study of an individual who had impairments for vegetables and animals, while their category for food remained intact. These findings are all based on individual case studies, so although they are the most reliable source of information, they are also full of inconsistencies because every brain and every instance of brain damage is unique in its own way.

Theories

When looking at category specific semantic deficits, it is important to consider how semantic information is stored in the brain. Theories on this subject tend to fall into two different groups based on their underlying principles. Theories based on the "correlated structure principle", which states that conceptual knowledge organization in the brain is a reflection of how often an object's properties occur, assume that the

brain reflects the statistical relation of object properties and how they relate to each other. Theories based on the "neural structure principle", which states that the conceptual knowledge organization in the brain is controlled by representational limits imposed by the brain itself, assume that organization is internal. These theories assume that natural selective pressures have caused neural circuits specific to certain domains to be formed, and that these are dedicated to problem-solving and survival. Animals, plants, and tools are all examples of specific circuits that would be formed based on this theory.

The role of modality

Modality refers to a semantic category of meaning which has to do with necessity and probability expressed through language. In linguistics, certain expressions are said to have modal meanings. A few examples of this include conditionals, auxiliaries, adverbs, and nouns. when looking at category specific semantic deficits, there is another kind of modality that looks at word relationships which is much more relevant to these disorders and impairments.

For category specific impairments, there are modality-specific theories which all rest on a few general predictions. These theories state that damage to the visual modality will result in a deficit of biological objects while damage to the functional modality will result in a deficit of non-biological objects (artifacts). Modality-based theories also assume that if there is damage to modality-specific knowledge, then all the categories that fall under it will be damaged. In this case, damage to the visual modality would result in a deficit for all biological

objects with no deficits restricted to the more specific categories. In other words, there would be no category specific semantic deficits for just "animals" or just "fruits and vegetables".

Category specific semantic deficit causes

Semantic Dementia

Semantic Dementia is a semantic memory disorder that causes patients to lose the ability to match words or images to their meanings. However, it is fairly rare for patients with semantic dementia to develop category specific impairments, though there have been documented cases of it occurring. Typically, a more generalized semantic impairment results form dimmed semantic representations in the brain.

Alzheimer's disease is a subcategory of semantic dementia which can cause similar symptoms. The main difference between the two being that Alzheimer's is categorized by atrophy to both sides of the brain while semantic dementia is categorized by loss of brain tissue in the front portion of the left temporal lobe. With Alzheimer's disease in particular, interactions with semantic memory produce different patterns in deficits between patients and categories over time which is caused by distorted representations in the brain. For example, in the initial onset of Alzheimer's disease, patients have mild difficulty with the artifacts category. As the disease progresses, the category specific semantic deficits progress as well, and patients see a more concrete deficit with natural categories. In other words, the deficit tends to be worse with living things as opposed to non-living things.

Herpes Simplex Virus Encephalitis

Herpes Simplex Virus Encephalitis (HSVE) is a neurological disorder which causes inflammation of the brain. It is caused by the herpes simplex virus type 1. Early symptoms include headache, fever, and drowsiness, but over time symptoms including diminished ability to speak, memory loss, and aphasia will develop. HSVE can also cause category specific semantic deficits to occur. When this does happen, patients typically have damage temporal lobe damage that affects the medial and lateral cortex as well as the frontal lobe. Studies have also shown that patients with HSVE have a much higher incidence of category specific semantic deficits than those with semantic dementia, though both cause a disruption of flow through the temporal lobe.

Brain lesions

A brain lesion refers to any abnormal tissue in or on the brain. Most often, this is caused by a trauma or infection. In one particular case study, a patient underwent surgery to remove an aneurysm, and the surgeon had to clip the anterior communicating artery which resulted in basal forebrain and fornix lesions. Before surgery, this patient was completely independent and had no semantic memory issues. However, after the operation and the lesions occurred, the patient reported difficulty with naming and identifying objects, recognition tasks, and comprehension. For this particular case, the patient had a much more significant amount of trouble with objects in the living category which could be seen in the drawings of animals which the patient was asked to do and in the data from the matching and identification tasks.

Every lesion is different, but in this case study researchers suggested that the semantic deficits presented themselves as a result of disconnection of the temporal lobe. This would lead to the conclusion that any type of lesion in the temporal lobe, depending on severity and location, has the potential to cause semantic deficits.

These results give us a baseline for the differences in semantic knowledge across gender for healthy subjects. When looking at category specific semantic deficits, we can compare the data to the table above to see if the results line up. Experimental data tells us that men with category specific semantic deficits are mainly impaired with fruits and vegetables while women with category specific semantic deficits are mainly impaired with animals and artifacts. This leads to the conclusion that there are significant gender differences when it comes to category specific semantic deficits, and that the patient will tend to be impaired in categories that had less existing knowledge to begin with.

Modality specific impairments

Semantic memory is also discussed in reference to modality. Different components represent information from different sensorimotor channels. Modality specific impairments are divided into separate subsystems on the basis of input modality. Examples of different input modalities include visual, auditory and tactile input. Modality specific impairments are also divided into subsystems based on the type of information. Visual vs. verbal and perceptual vs. functional information are examples of information types. Modality specificity can account for category specific impairments in semantic memory

disorders. Damage to visual semantics primarily impairs knowledge of living things, and damage to functional semantics primarily impairs knowledge of nonliving things.

Semantic refractory access and semantic storage disorders

Semantic memory disorders fall into two groups. Semantic refractory access disorders are contrasted with semantic storage disorders according to four factors. Temporal factors, response consistency, frequency and semantic relatedness are the four factors used to differentiate between semantic refractory access and semantic storage disorders. A key feature of semantic refractory access disorders is temporal distortions. Decreases in response time to certain stimuli are noted when compared to natural response times. Response consistency is the next factor. In access disorders you see inconsistencies in comprehending and responding to stimuli that have been presented many times.

Temporal factors impact response consistency. In storage disorders, you do not see an inconsistent response to specific items like you do in refractory access disorders. Stimulus frequency determines performance at all stages of cognition. Extreme word frequency effects are common in semantic storage disorders while in semantic refractory access disorders word frequency effects are minimal. The comparison of 'close' and 'distant' groups tests semantic relatedness. 'Close' groupings have words that are related because they are drawn from the same category. For example, a listing of clothing types would be a 'close' grouping. 'Distant' groupings contain words with broad categorical differences. Non-related words would

fall into this group. Comparing close and distant groups shows that in access disorders semantic relatedness had a negative effect. This is not observed in semantic storage disorders. Category specific and modality specific impairments are important components in access and storage disorders of semantic memory.

Present and future research

Semantic memory has had a comeback in interest in the past 15 years, due in part to the development of functional neuroimaging methods such as positron emission tomography (PET) and functional magnetic resonance imaging (fMRI), which have been used to address some of the central questions about our understanding of semantic memory.

Positron emission tomography (PET) and functional magnetic resonance (fMRI) allow cognitive neuroscientists to explore different hypotheses concerning the neural network organization of semantic memory. By using these neuroimaging techniques researchers can observe the brain activity of participants while they perform cognitive tasks. These tasks can include, but are not limited to, naming objects, deciding if two stimuli belong in the same object category, or matching pictures to their written or spoken names.

Rather than any one brain region playing a dedicated and privileged role in the representation or retrieval of all sorts of semantic knowledge, semantic memory is a collection of functionally and anatomically distinct systems, where each attribute-specific system is tied to a sensorimotormodality (i.e. vision) and even more specifically to a property within that

modality (i.e. color). Neuroimaging studies also suggest a distinction between semantic processing and sensorimotor processing.

A new idea that is still at the early stages of development is that semantic memory, like perception, can be subdivided into types of visual information—color, size, form, and motion. Thompson-Schill (2003) found that the left or bilateral ventral temporal cortex appears to be involved in retrieval of knowledge of color and form, the left lateral temporal cortex in knowledge of motion, and the parietal cortex in knowledge of size.

Neuroimaging studies suggest a large, distributed network of semantic representations that are organized minimally by attribute, and perhaps additionally by category. These networks include "extensive regions of ventral (form and color knowledge) and lateral (motion knowledge) temporal cortex, parietal cortex (size knowledge), and premotor cortex (manipulation knowledge). Other areas, such as more anterior regions of temporal cortex, may be involved in the representation of nonperceptual (e.g. verbal) conceptual knowledge, perhaps in some categorically-organized fashion." It is suggested that within the temperoparietal network, the anterior temporal lobe is relatively more important for semantic processing, and posterior language regions are relatively more important for lexical retrieval.

Prototype theory

Prototype theory is a theory of categorization in cognitive science, particularly in psychology and cognitive linguistics, in

which there is a graded degree of belonging to a conceptual category, and some members are more central than others. It emerged in 1971 with the work of psychologist Eleanor Rosch, and it has been described as a "Copernican revolution" in the theory of categorization for its departure from the traditional Aristotelian categories. It has been criticized by those that still endorse the traditional theory of categories, like linguist Eugenio Coseriu and other proponents of the structural semanticsparadigm.

In this prototype theory, any given concept in any given language has a real world example that best represents this concept. For example: when asked to give an example of the concept *furniture*, a *couch* is more frequently cited than, say, a *wardrobe*. Prototype theory has also been applied in linguistics, as part of the mapping from phonological structure to semantics.

In formulating prototype theory, Rosch drew in part from previous insights in particular the formulation of a category model based on family resemblance by Wittgenstein (1953), and by Roger Brown's *How shall a thing be called?* (1958).

Overview and terminology

The term *prototype*, as defined in psychologist Eleanor Rosch's study "Natural Categories", was initially defined as denoting a stimulus, which takes a salient position in the formation of a category, due to the fact that it is the first stimulus to be associated with that category. Rosch later defined it as the most central member of a category.

Rosch and others developed prototype theory as a response to, and radical departure from, the classical theory of concepts, which defines concepts by necessary and sufficient conditions. Necessary conditions refers to the set of features every instance of a concept must present, and sufficient conditions are those that no other entity possesses. Rather than defining concepts by features, the prototype theory defines categories based on either a specific artifact of that category or by a set of entities within the category that represent a prototypical member. The prototype of a category can be understood in lay terms by the object or member of a class most often associated with that class. The prototype is the center of the class, with all other members moving progressively further from the prototype, which leads to the gradation of categories. Every member of the class is not equally central in human cognition. As in the example of *furniture* above, *couch* is more central than *wardrobe*. Contrary to the classical view, prototypes and gradations lead to an understanding of category membership not as an all-or-nothing approach, but as more of a web of interlocking categories which overlap.

In Cognitive linguistics it has been argued that linguistic categories also have a prototype structure, like categories of common words in a language.

Categories

Basic level categories

The other notion related to prototypes is that of a *basic level* in cognitive categorization. Basic categories are relatively

homogeneous in terms of sensory-motor affordances — a chair is associated with bending of one's knees, a fruit with picking it up and putting it in your mouth, etc. At the subordinate level (e.g. [dentist's chairs], [kitchen chairs] etc.) few significant features can be added to that of the basic level; whereas at the superordinate level, these conceptual similarities are hard to pinpoint. A picture of a chair is easy to draw (or visualize), but drawing furniture would be more difficult.

Linguist Eleanor Rosch defines the basic level as that level that has the highest degree of cue validity. Thus, a category like [animal] may have a prototypical member, but no cognitive visual representation. On the other hand, basic categories in [animal], i.e. [dog], [bird], [fish], are full of informational content and can easily be categorized in terms of Gestalt and semantic features.

Clearly semantic models based on attribute-value pairs fail to identify privileged levels in the hierarchy. Functionally, it is thought that basic level categories are a decomposition of the world into maximally informative categories. Thus, they

- maximize the number of attributes shared by members of the category, and
- minimize the number of attributes shared with other categories

However, the notion of Basic Level is problematic, e.g. whereas dog as a basic category is a species, bird or fish are at a higher level, etc. Similarly, the notion of frequency is very closely tied to the basic level, but is hard to pinpoint.

More problems arise when the notion of a prototype is applied to lexical categories other than the noun. Verbs, for example, seem to defy a clear prototype: [to run] is hard to split up in more or less central members.

In her 1975 paper, Rosch asked 200 American college students to rate, on a scale of 1 to 7, whether they regarded certain items as good examples of the category *furniture*. These items ranged from chair and sofa, ranked number 1, to a love seat (number 10), to a lamp (number 31), all the way to a telephone, ranked number 60.

While one may differ from this list in terms of cultural specifics, the point is that such a graded categorization is likely to be present in all cultures. Further evidence that some members of a category are more privileged than others came from experiments involving:

- 1. *Response Times*: in which queries involving prototypical members (e.g. *is a robin a bird*) elicited faster response times than for non-prototypical members.
- 2. *Priming*: When primed with the higher-level (superordinate) category, subjects were faster in identifying if two words are the same. Thus, after flashing *furniture*, the equivalence of *chair-chair*is detected more rapidly than *stove-stove*.
- 3. *Exemplars*: When asked to name a few exemplars, the more prototypical items came up more frequently.

Subsequent to Rosch's work, prototype effects have been investigated widely in areas such as colour cognition, and also

for more abstract notions: subjects may be asked, e.g. "to what degree is this narrative an instance of telling a lie?". Similar work has been done on actions (verbs like look, kill, speak, walk [Pulman:83]), adjectives like "tall", etc.

Another aspect in which Prototype Theory departs from traditional Aristotelian categorization is that there do not appear to be natural kind categories (bird, dog) vs. artifacts (toys, vehicles).

A common comparison is the use of prototype or the use of exemplars in category classification. Medin, Altom, and Murphy found that using a mixture of prototype and exemplar information, participants were more accurately able to judge categories. Participants who were presented with prototype values classified based on similarity to stored prototypes and stored exemplars, whereas participants who only had experience with exemplar only relied on the similarity to stored exemplars. Smith and Minda looked at the use of prototypes and exemplars in dot-pattern category learning. They found that participants used more prototypes than they used exemplars, with the prototypes being the center of the category, and exemplars surrounding it.

Distance between concepts

The notion of prototypes is related to Wittgenstein's (later) discomfort with the traditional notion of category. This influential theory has resulted in a view of semantic components more as *possible* rather than necessary contributors to the meaning of texts. His discussion on the category *game* is particularly incisive:

Consider for example the proceedings that we call 'games'. I mean board games, card games, ball games, Olympic games, and so on. What is common to them all? Don't say, "There must be something common, or they would not be called 'games'"--but look and see whether there is anything common to all. For if you look at them you will not see something common to all, but similarities, relationships, and a whole series of them at that. To repeat: don't think, but look! Look for example at board games, with their multifarious relationships. Now pass to card games; here you find many correspondences with the first group, but many common features drop out, and others appear. When we pass next to ball games, much that is common is retained, but much is lost. Are they all 'amusing'? Compare chess with noughts and crosses. Or is there always winning and losing, or competition between players? Think of patience. In ball games there is winning and losing; but when a child throws his ball at the wall and catches it again, this feature has disappeared. Look at the parts played by skill and luck; and at the difference between skill in chess and skill in tennis. Think now of games like ring-a-ring-a-roses; here is the element of amusement, but how many other characteristic features have disappeared! And we can go through the many, many other groups of games in the same way; can see how similarities crop up and disappear. And the result of this examination is: we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail.

Wittgenstein's theory of family resemblance describes the phenomenon when people group concepts based on a series of overlapping features, rather than by one feature which exists throughout all members of the category. For example,

basketball and baseball share the use of a ball, and baseball and chess share the feature of a winner, etc, rather than one defining feature of "games". Therefore, there is a distance between focal, or prototypical members of the category, and those that continue outwards from them, linked by shared features.

Recently, Peter Gärdenfors has elaborated a possible partial explanation of prototype theory in terms of multi-dimensional feature spaces called conceptual spaces, where a category is defined in terms of a conceptual distance. More central members of a category are "between" the peripheral members. He postulates that most *natural* categories exhibit a convexity in conceptual space, in that if x and y are elements of a category, and if z is *between* x and y, then z is also likely to belong to the category.

Combining categories

Within language we find instances of combined categories, such as *tall man* or *small elephant*. Combining categories was a problem for extensional semantics, where the semantics of a word such as *red* is to be defined as the set of objects having this property. This does not apply as well to modifiers such as *small*; a *small mouse* is very different from a *small elephant*.

These combinations pose a lesser problem in terms of prototype theory. In situations involving adjectives (e.g. *tall*), one encounters the question of whether or not the prototype of [tall] is a 6 foot tall man, or a 400-foot skyscraper. The solution emerges by contextualizing the notion of prototype in terms of the object being modified. This extends even more

radically in compounds such as *red wine* or *red hair* which are hardly *red* in the prototypical sense, but the red indicates merely a shift from the prototypical colour of wine or hair respectively. The addition of *red* shifts the prototype from the one of *hair* to that of *red hair.* The prototype is changed by additional specific information, and combines features from the prototype of *red* and *wine*.

Critique

Prototype theory has been criticized by those that still endorse the classic theory of categories, like linguist Eugenio Coseriu and other proponents of the structural semanticsparadigm.

Exemplar theory

Douglas L. Medin and Marguerite M. Schaffer showed by experiment that a context theory of classification which derives concepts purely from exemplars (cf. exemplar theory) worked better than a class of theories that included prototype theory.

Graded categorization

Linguists, including Stephen Laurence writing with Eric Margolis, have suggested problems with the prototype theory. In their 1999 paper, they raise several issues. One of which is that prototype theory does not intrinsically guarantee graded categorization. When subjects were asked to rank how well certain members exemplify the category, they rated some members above others. For example robins were seen as being "birdier" than ostriches, but when asked whether these categories are "all-or-nothing" or have fuzzier boundaries, the

subjects stated that they were defined, "all-or-nothing" categories. Laurence and Margolis concluded that "prototype structure has no implication for whether subjects represent a category as being graded" (p. 33).

Compound concepts

Daniel Osherson and Edward Smith raised the issue of *pet fish* for which the prototype might be a guppy kept in a bowl in someone's house. The prototype for *pet* might be a dog or cat, and the prototype for *fish* might be trout or salmon. However, the features of these prototypes do not present in the prototype for *pet fish*, therefore this prototype must be generated from something other than its constituent parts.

Antonio Lieto and Gian Luca Pozzato have proposed a typicality-based compositional logic (TCL) that is able to account for both complex human-like concept combinations (like the PET-FISH problem) and conceptual blending. Thus, their framework shows how concepts expressed as prototypes can account for the phenomenon of prototypical compositionality in concept combination.

Chapter 6

Stubs

Semantic Folding

Semantic folding theory describes a procedure for encoding the semantics of natural language text in a semantically grounded binary representation. This approach provides a framework for modelling how language data is processed by the neocortex.

Theory

Semantic folding theory draws inspiration from Douglas R. Hofstadter's *Analogy as the Core of Cognition* which suggests that the brain makes sense of the world by identifying and applying analogies. The theory hypothesises that semantic data must therefore be introduced to the neocortex in such a form as to allow the application of a similarity measure and offers, as a solution, the sparsebinary vector employing a two-dimensional topographic semantic space as a distributional reference frame. The theory builds on the computational theory of the human cortex known as hierarchical temporal memory (HTM), and positions itself as a complementary theory for the representation of language semantics.

A particular strength claimed by this approach is that the resulting binary representation enables complex semantic operations to be performed simply and efficiently at the most basic computational level.

Two-dimensional semantic space

Analogous to the structure of the neocortex, Semantic Folding theory posits the implementation of a semantic space as a two-dimensional grid. This grid is populated by context-vectors in such a way as to place similar context-vectors closer to each other, for instance, by using competitive learning principles. This vector space model is presented in the theory as an equivalence to the well known word space model described in the information retrieval literature.

Given a semantic space (implemented as described above) a word-vector can be obtained for any given word *Y* by employing the following algorithm:

- **For each** position *X* in the semantic map (where X represents cartesian coordinates)
- **if** the word *Y* is contained in the context-vector at position *X*
- then add 1 to the corresponding position in the word-vector for *Y*
- **else**
- add 0 to the corresponding position in the word-vector for *Y*

The result of this process will be a word-vector containing all the contexts in which the word Y appears and will therefore be representative of the semantics of that word in the semantic space.

It can be seen that the resulting word-vector is also in a sparse distributed representation (SDR) format [Schütze, 1993] &

[Sahlgreen, 2006]. Some properties of word-SDRs that are of particular interest with respect to computational semantics are:

- highnoise resistance: As a result of similar contexts being placed closer together in the underlying map, word-SDRs are highly tolerant of false or shifted "bits".
- boolean logic: It is possible to manipulate word-SDRs in a meaningful way using boolean (OR, AND, exclusive-OR) and/or arithmetical (SUBtract) functions .
- sub-sampling: Word-SDRs can be sub-sampled to a high degree without any appreciable loss of semantic information.
- topological two-dimensional representation: The SDR representation maintains the topological distribution of the underlying map therefore words with similar meanings will have similar word-vectors. This suggests that a variety of measures can be applied to the calculation of semantic similarity, from a simple overlap of vector elements, to a range of distance measures such as:Euclidean distance, Hamming distance, Jaccard distance, cosine similarity, Levenshtein distance, Sørensen-Dice index, etc.

Semantic spaces

Semantic spaces in the natural language domain aim to create representations of natural language that are capable of capturing meaning. The original motivation for semantic spaces stems from two core challenges of natural language:

Vocabulary mismatch (the fact that the same meaning can be expressed in many ways) and ambiguity of natural language (the fact that the same term can have several meanings).

The application of semantic spaces in natural language processing (NLP) aims at overcoming limitations of rule-based or model-based approaches operating on the keyword level. The main drawback with these approaches is their brittleness, and the large manual effort required to create either rule-based NLP systems or training corpora for model learning. Rule-based and machine learning-based models are fixed on the keyword level and break down if the vocabulary differs from that defined in the rules or from the training material used for the statistical models.

Research in semantic spaces dates back more than 20 years. In 1996, two papers were published that raised a lot of attention around the general idea of creating semantic spaces: latent semantic analysis from Microsoftand Hyperspace Analogue to Language from the University of California. However, their adoption was limited by the large computational effort required to construct and use those semantic spaces. A breakthrough with regard to the accuracy of modelling associative relations between words (e.g. "spider-web", "lighter-cigarette", as opposed to synonymous relations such as "whale-dolphin", "astronaut-driver") was achieved by explicit semantic analysis (ESA) in 2007. ESA was a novel (non-machine learning) based approach that represented words in the form of vectors with 100,000 dimensions (where each dimension represents an Article in Wikipedia). However practical applications of the approach are limited due to the large number of required dimensions in the vectors.

More recently, advances in neural networking techniques in combination with other new approaches (tensors) led to a host of new recent developments: Word2vec from Google and GloVe from Stanford University.

Semantic folding represents a novel, biologically inspired approach to semantic spaces where each word is represented as a sparse binary vector with 16,000 dimensions (a semantic fingerprint) in a 2D semantic map (the semantic universe). Sparse binary representation are advantageous in terms of computational efficiency, and allow for the storage of very large numbers of possible patterns.

Visualization

The topological distribution over a two-dimensional grid (outlined above) lends itself to a bitmap type visualization of the semantics of any word or text, where each active semantic feature can be displayed as e.g. a pixel. As can be seen in the images shown here, this representation allows for a direct visual comparison of the semantics of two (or more) linguistic items.

Image 1 clearly demonstrates that the two disparate terms "dog" and "car" have, as expected, very obviously different semantics.

Image 2 shows that only one of the meaning contexts of "jaguar", that of "Jaguar" the car, overlaps with the meaning of Porsche (indicating partial similarity). Other meaning contexts of "jaguar" e.g. "jaguar" the animal clearly have different non-overlapping contexts. The visualization of semantic similarity

using Semantic Folding bears a strong resemblance to the fMRI images produced in a research study conducted by A.G. Huth et al., where it is claimed that words are grouped in the brain by meaning.

Semantic holism

Semantic holism is a theory in the philosophy of language to the effect that a certain part of language, be it a term or a complete sentence, can only be understood through its relations to a (previously understood) larger segment of language. There is substantial controversy, however, as to exactly what the larger segment of language in question consists of. In recent years, the debate surrounding semantic holism, which is one among the many forms of holism that are debated and discussed in contemporary philosophy, has tended to centre on the view that the "whole" in question consists of an entire language.

Background

Since the *use* of a linguistic expression is only possible if the speaker who uses it understands its *meaning*, one of the central problems for analytic philosophers has always been the question of meaning. What is it? Where does it come from? How is it communicated? And, among these questions, what is the smallest unit of meaning, the smallest fragment of language with which it is possible to communicate something? At the end of the 19th and beginning of the 20th century, GottlobFrege and his followers abandoned the view, common at the time, that a word gets its meaning in isolation,

independently from all the rest of the words in a language. Frege, as an alternative, formulated his famous context principle, according to which it is only within the context of an entire sentence that a word acquires its meaning. In the 1950s, the agreement that seemed to have been reached regarding the primacy of sentences in semantic questions began to unravel with the collapse of the movement of logical positivism and the powerful influence exercised by the later Ludwig Wittgenstein. Wittgenstein wrote in the *Philosophical Investigations* that "comprehending a proposition means comprehending a language". About the same time or shortly after, W. V. O. Quine wrote that "the unit of measure of empirical meaning is all of science in its globality"; and Donald Davidson, in 1967, put it even more sharply by saying that "a sentence (and therefore *a word*) has meaning only in the context of a (whole) language".

Problems

If semantic holism is interpreted as the thesis that any linguistic expression *E* (a word, a phrase or sentence) of some natural language *L* cannot be understood in isolation and that there are inevitably many ties between the expressions of *L*, it follows that to understand *E* one must understand a set *K* of expressions to which *E* is related. If, in addition, no limits are placed on the size of *K* (as in the cases of Davidson, Quine and, perhaps, Wittgenstein), then *K* coincides with the "whole" of *L*.

The many and substantial problems with this position have been described by Michael Dummett, Jerry Fodor, Ernest Lepore and others. In the first place, it is impossible to

understand how a speaker of *L* can acquire knowledge of (learn) the meaning of *E*, for any expression *E* of the language. Given the limits of our cognitive abilities, we will never be able to master the whole of the English (or Italian or German) language, even on the assumption that languages are static and immutable entities (which is false).

Therefore, if one must understand all of a natural language *L* to understand the single word or expression *E*, then language learning is simply impossible.

Semantic holism, in this sense, also fails to explain how two speakers can mean the same thing when using the same linguistic expression, and therefore how communication is even possible between them. Given a sentence *P*, since Fred and Mary have each mastered different parts of the English language and *P*is related to the sentences in each part differently, the result is that *P* means one thing for Fred and something else for Mary.

Moreover, if a sentence *P* derives its meaning from the relations it entertains with the totality of sentences of a language, as soon as the vocabulary of an individual changes by the addition or elimination of a sentence *P′*, the totality of relations changes, and therefore also the meaning of *P*. As this is a very common phenomenon, the result is that *P* has two different meanings in two different moments during the life of the same person. Consequently, if I accept the truth of a sentence and then reject it later on, the meaning of what I rejected and what I accepted are completely different, and therefore I cannot change my opinions regarding the same sentences.

Holism of mental content

These sorts of counterintuitive consequences of semantic holism also affect another form of holism, often identified with but, in fact, distinct from semantic holism: the holism of mental content. This is the thesis that the meaning of a particular propositional attitude (thought, desire, belief) acquires its content by virtue of the role that it plays within the web that connects it to all the other propositional attitudes of an individual.

Since there is a very tight relationship between the content of a mental state *M* and the sentence *P*, which expresses it and makes it publicly communicable, the tendency in recent discussion is to consider the term "content" to apply indifferently both to linguistic expressions and to mental states, regardless of the extremely controversial question of which category (the mental or the linguistic) has priority over the other and which, instead, possesses only a *derived* meaning.

So, it would seem that semantic holism ties the philosopher's hands. By making it impossible to explain language learning and to provide a unique and consistent description of the meanings of linguistic expressions, it blocks off any possibility of formulating a theory of meaning; and, by making it impossible to individuate the exact contents of any propositional attitude—given the necessity of considering a potentially infinite and continuously evolving set of mental states—it blocks off the possibility of formulating a theory of the mind.

Confirmation holism

The key to answering this question lies in going back to Quine and his attack on logical positivism. The logical positivists, who dominated the philosophical scene for almost the entire first half of the twentieth century, maintained that genuine knowledge consisted in all and only such knowledge as was capable of manifesting a strict relationship with empirical experience.

Therefore, they believed, the only linguistic expressions (manifestations of knowledge) that had meaning were those that either directly referred to observable entities, or that could be reduced to a vocabulary that directly referred to such entities. A sentence S contained knowledge only if it possessed a meaning, and it possessed a meaning only if it was possible to refer to a set of experiences that could, at least potentially, verify it and to another set that could potentially falsify it. Underlying all this, there is an implicit and powerful connection between epistemological and semantic questions. This connection carries over into the work of Quine in *Two Dogmas of Empiricism*.

Quine's holistic argument against the neo-positivists set out to demolish the assumption that every sentence of a language is bound univocally to its own set of potential verifiers and falsifiers and the result was that the epistemological value of every sentence must depend on the entire language. Since the epistemological value of every sentence, for Quine just as for the positivists, was the meaning of that sentence, then the meaning of every sentence must depend on every other. As Quine states it:

- All of our so-called knowledge or convictions, from questions of geography and history to the most profound laws of atomic physics or even mathematics and logic, are an edifice made by man that touches experience only at the margins. Or, to change images, science in its globality is like a force field whose limit points are experiences...a particular experience is never tied to any proposition inside the field except indirectly, for the needs of equilibrium which affect the field in its globality.

For Quine then (although Fodor and Lepore have maintained the contrary), and for many of his followers, confirmation holism and semantic holism are inextricably linked. Since confirmation holism is widely accepted among philosophers, a serious question for them has been to determine whether and how the two holisms can be distinguished or how the undesirable consequences of *unbuttoned holism*, as Michael Dummett has called it, can be limited.

Moderate holism

Numerous philosophers of language have taken the latter avenue, abandoning the early Quinean holism in favour of what Michael Dummett has labelled *semantic molecularism*. These philosophers generally deny that the meaning of an expression *E* depends on the meanings of the words of the entire language *L* of which it is part and sustain, instead, that the meaning of *E* depends on some subset of *L*. These positions, notwithstanding the fact that many of their proponents continue to call themselves holists, are actually intermediate between holism and atomism.

Dummett, for example, after rejecting Quinean holism (holism *tout court* in his sense), takes precisely this approach.

But those who would opt for some version of moderate holism need to make the distinction between the parts of a language that are "constitutive" of the meaning of an expression *E* and those that are not without falling into the extraordinarily problematic analytic/synthetic distinction.

Fodor and Lepore (1992) present several arguments to demonstrate that this is impossible.

Holism and compositionality

- The relationship between compositionality and semantic holism has also been of interest to many philosophers of language. On the surface it would seem that these two ideas are in complete and irremediable contradiction. Compositionality is the principle that states that the meaning of a complex expression depends on the meaning of its parts and on its mode of composition. As stated before, holism, on the other hand, is the thesis that the meanings of expressions of a language are determined by their relations with the other expressions of the language as a whole. Peter Pagin, in an essay called *Are Compositionality and Holism Compatible* identifies three points of incompatibility between these two hypotheses. The first consists in the simple observation that while, for holism, the meaning of the whole would seem to *precede* that of its parts in terms of priority, for compositionality, the reverse is

true, the meaning of the parts *precedes* that of the whole. The second incoherence consists in the fact that a necessity to attribute "strange" meanings to the components of larger expressions would apparently result from any attempt to reconcile compositionality and holism. Pagin takes a specific holistic theory of meaning – inferential role semantics, the theory according to which the meaning of an expression is determined by the inferences that it involves – as his paradigm of holism. If we interpret this theory holistically, the result will be that *every* accepted inference that involves some expression will enter into the meaning of that expression. Suppose, for example, that Fred believes that "Brown cows are dangerous". That is, he accepts the inference from "brown cows" to "dangerous." This entails that this inference is now part of the meaning of "brown cow." According to compositionality then, "cow implies dangerous" and "brown implies dangerous" are both true because they are the constituents of the expression "brown cow." But is this really an inevitable consequence of the acceptance of the holism of inferential role semantics? To see why it's not assume the existence of a relation of inference *I* between two expressions *x* and *y* and that the relation applies just in case *F* accepts the inference from *x* to *y*. Suppose that in the extension of *I*, there are the following pairs of expressions ("The sky is blue and leaves are green", "the sky is blue") and ("brown cow", "dangerous").

Holism and externalism

Since the concept of semantic holism, as explained above, is often used to refer to not just theories of meaning in natural languages but also to theories of mental content such as the hypothesis of a language of thought, the question often arises as to how to reconcile the idea of semantic holism (in the sense of the meanings of expressions in mental languages) with the phenomenon called externalism in philosophy of mind.

Externalism is the thesis that the propositional attitudes of an individual are determined, at least in part, by her relations with her environment (both social and natural). Hilary Putnam formulated the thesis of the *natural* externalism of mental states in his *The Meaning of "Meaning"*. In it, he described his famous thought experiment involving Twin Earths: two individuals, Calvin and Carvin, live, respectively, on the real earth (E) of our everyday experience and on an exact copy (E') with the only difference being that on E "water" stands for the substance while on E' it stands for some substance macroscopically identical to water but which is actually composed of XYZ. According to Putnam, only Calvin has genuine experiences that involve water, so only his term "water" really refers to water.

Tyler Burge, in *Individualism and the Mental*, describes a different thought experiment that led to the notion of the *social* externalism of mental contents. In Burge's experiment, a person named Jeffray believes that he has arthritis in his thighs and we can correctly attribute to him the (mistaken) belief that he has arthritis in his thighs because he is ignorant of the fact that arthritis is a disease of the articulation of the

joints. In another society, there is an individual named Goodfrey who also believes that he has arthritis in the thighs. But in the case of Goodfrey the belief is correct because in the counterfactual society in which he lives "arthritis" is defined as a disease that can include the thighs.

The question then arises of the possibility of reconciling externalism with holism. The one seems to be saying that meanings are determined by the external relations (with society or the world), while the other suggests that meaning is determined by the relation of words (or beliefs) to all the other words (or beliefs). Frederik Stjernfelt identifies at least three possible ways to reconcile them and then points out some objections.

The first approach is to insist that there is no conflict because holists do not mean the phrase "determine beliefs" in the sense of individuation but rather of attribution. But the problem with this is that if one is not a "realist" about mental states, then all we are left with is the attributions themselves and, if these are holistic, then we really have a form of hidden *constitutive holism* rather than a genuine *attributive holism*. But if one is a "realist" about mental states, then why not say that we can actually individuate them and therefore that instrumentalist attributions are just a short-term strategy?

Another approach is to say that externalism is valid only for certain beliefs and that holism only suggests that beliefs are determined only *in part* by their relations with other beliefs. In this way, it is possible to say that externalism applies only to those beliefs not determined by their relations with other beliefs (or for the part of a belief that is not determined by its

relations with other parts of other beliefs), and holism is valid to the extent that beliefs (or parts of beliefs) are not determined externally. The problem here is that the whole scheme is based on the idea that certain relations are constitutive (i.e. necessary) for the determination of the beliefs and others are not. Thus, we have reintroduced the idea of an analytic/synthetic distinction with all of the problems that that carries with it.

A third possibility is to insist that there are two distinct *types* of belief: those determined holistically and those determined externally. Perhaps the external beliefs are those that are determined by their relations with the external world through observation and the holistic ones are the theoretical statements. But this implies the abandonment of a central pillar of holism: the idea that there can be no one to one correspondence between behavior and beliefs. There will be cases in which the beliefs that are determined externally correspond one to one with perceptual states of the subject.

One last proposal is to carefully distinguish between so-called narrow content states and broad content states. The first would be determined in a holistic manner and the second non-holistically and externalistically. But how to distinguish between the two notions of content while providing a justification of the possibility of formulating an idea of narrow content that does not depend on a prior notion of broad content?

These are some of the problems and questions that have still to be resolved by those who would adopt a position of "holistic externalism" or "externalist holism".

Semantic heterogeneity

Semantic heterogeneity is whendatabase schema or datasets for the same domain are developed by independent parties, resulting in differences in meaning and interpretation of data values. Beyond structured data, the problem of semantic heterogeneity is compounded due to the flexibility of semi-structured data and various tagging methods applied to documents or unstructured data. Semantic heterogeneity is one of the more important sources of differences in heterogeneous datasets.

Yet, for multiple data sources to interoperate with one another, it is essential to reconcile these semantic differences. Decomposing the various sources of semantic heterogeneities provides a basis for understanding how to map and transform data to overcome these differences.

Classification

One of the first known classification schemes applied to data semantics is from William Kent more than two decades ago. Kent's approach dealt more with structural mapping issues than differences in meaning, which he pointed to data dictionaries as potentially solving.

One of the most comprehensive classifications is from Pluempitiwiriyawej and Hammer, "Classification Scheme for Semantic and Schematic Heterogeneities in XML Data Sources". They classify heterogeneities into three broad classes:

- *Structural* conflicts arise when the schema of the sources representing related or overlapping data exhibit discrepancies. Structural conflicts can be detected when comparing the underlying schema. The class of structural conflicts includes generalization conflicts, aggregation conflicts, internal path discrepancy, missing items, element ordering, constraint and type mismatch, and naming conflicts between the element types and attribute names.
- *Domain* conflicts arise when the semantics of the data sources that will be integrated exhibit discrepancies. Domain conflicts can be detected by looking at the information contained in the schema and using knowledge about the underlying data domains. The class of domain conflicts includes schematic discrepancy, scale or unit, precision, and data representation conflicts.
- *Data* conflicts refer to discrepancies among similar or related data values across multiple sources. Data conflicts can only be detected by comparing the underlying sources. The class of data conflicts includes ID-value, missing data, incorrect spelling, and naming conflicts between the element contents and the attribute values.

Moreover, mismatches or conflicts can occur between set elements (a "population" mismatch) or attributes (a "description" mismatch).

Relevant applications

Besides data interoperability, relevant areas in information technology that depend on reconciling semantic heterogeneities include data mapping, semantic integration, and enterprise information integration, among many others. From the conceptual to actual data, there are differences in perspective, vocabularies, measures and conventions once any two data sources are brought together. Explicit attention to these semantic heterogeneities is one means to get the information to integrate or interoperate.

A mere twenty years ago, information technology systems expressed and stored data in a multitude of formats and systems. The Internet and Web protocols have done much to overcome these sources of differences. While there is a large number of categories of semantic heterogeneity, these categories are also patterned and can be anticipated and corrected. These patterned sources inform what kind of work must be done to overcome semantic differences where they still reside.

Semantic integration

Semantic integration is the process of interrelating information from diverse sources, for example calendars and to do lists, email archives, presence information (physical, psychological, and social), documents of all sorts, contacts (including social graphs), search results, and advertising and marketing relevance derived from them. In this regard, semantics focuses on the organization of and action upon information by acting

as an intermediary between heterogeneous data sources, which may conflict not only by structure but also context or value.

Applications and methods

In enterprise application integration (EAI), semantic integration can facilitate or even automate the communication between computer systems using metadata publishing. Metadata publishing potentially offers the ability to automatically linkontologies. One approach to (semi-)automated ontology mapping requires the definition of a semantic distance or its inverse, semantic similarity and appropriate rules. Other approaches include so-called *lexical methods*, as well as methodologies that rely on exploiting the structures of the ontologies. For explicitly stating similarity/equality, there exist special properties or relationships in most ontology languages. OWL, for example has "owl:equivalentClass", "owl:equivalentProperty" and "owl:sameAs".

Eventually system designs may see the advent of composable architectures where published semantic-based interfaces are joined together to enable new and meaningful capabilities. These could predominately be described by means of design-time declarative specifications, that could ultimately be rendered and executed at run-time.

Semantic integration can also be used to facilitate design-time activities of interface design and mapping. In this model, semantics are only explicitly applied to design and the run-time systems work at the syntax level. This "early semantic

binding" approach can improve overall system performance while retaining the benefits of semantic driven design.

Semantic integration situations

From the industry use case, it has been observed that the semantic mappings were performed only within the scope of the ontology class or the datatype property.

These identified semantic integrations are (1) integration of ontology class instances into another ontology class without any constraint, (2) integration of selected instances in one ontology class into another ontology class by the range constraint of the property value and (3) integration of ontology class instances into another ontology class with the value transformation of the instance property. Each of them requires a particular mapping relationship, which is respectively: (1) equivalent or subsumption mapping relationship, (2) conditional mapping relationship that constraints the value of property (data range) and (3) transformation mapping relationship that transforms the value of property (unit transformation). Each identified mapping relationship can be defined as either (1) direct mapping type, (2) data range mapping type or (3) unit transformation mapping type.

KG vs. RDB approaches

In the case of integrating supplemental data source,

- KG(Knowledge graph) formally represents the meaning involved in information by describing

concepts, relationships between things, and categories of things. These embedded semantics with the data offer significant advantages such as reasoning over data and dealing with heterogeneous data sources. The rules can be applied on KG more efficiently using graph query. For example, the graph query does the data inference through the connected relations, instead of repeated full search of the tables in relational database. KG facilitates the integration of new heterogeneous data by just adding new relationships between existing information and new entities. This facilitation is emphasized for the integration with existing popular linked open data source such as Wikidata.org.

- SQL query is tightly coupled and rigidly constrained by datatype within the specific database and can join tables and extract data from tables, and the result is generally a table, and a query can join tables by any columns which match by datatype. SPARQL query is the standard query language and protocol for Linked Open Data on the web and loosely coupled with the database so that it facilitates the reusability and can extract data through the relations free from the datatype, and not only extract but also generate additional knowledge graph with more sophisticated operations(logic: transitive/symmetric/inverseOf/functional). The inference based query (query on the existing asserted facts without the generation of new facts by logic) can be fast comparing to the reasoning based query (query on the existing plus the generated/discovered facts based on logic).

- The information integration of heterogeneous data sources in traditional database is intricate, which requires the redesign of the database table such as changing the structure and/or addition of new data. In the case of semantic query, SPARQL query reflects the relationships between entities in a way that aligned with human's understanding of the domain, so the semantic intention of the query can be seen on the query itself. Unlike SPARQL, SQL query, which reflects the specific structure of the database and derived from matching the relevant primary and foreign keys of tables, loses the semantics of the query by missing the relationships between entities. Below is the example that compares SPARQL and SQL queries for medications that treats "TB of vertebra".

Examples

The Pacific Symposium on Biocomputing has been a venue for the popularization of the ontology mapping task in the biomedical domain, and a number of papers on the subject can be found in its proceedings.

Semantic Interoperability Community of Practice

Semantic Interoperability Community of Practice (SICoP) is a group of people who seek to make the Semantic Web

operational in their respective settings by achieving "semantic interoperability" and "semantic data integration".

SICoP seeks to enable Semantic Interoperability, specifically the "operationalizing" of relevant technologies and approaches, through online conversation, meetings, tutorials, conferences, pilot projects, and other activities aimed at developing and disseminating best practices.

The individuals making up this Community of Practice are from various settings, however, the SICoP claims neither formal nor implied endorsement by any organization.

Semantic lexicon

A semantic lexicon is a digital dictionary of words labeled with semantic classes so associations can be drawn between words that have not previously been encountered. Semantic lexicons are built upon semantic networks, which represent the semantic relations between words. The difference between a semantic lexicon and a semantic network is that a semantic lexicon has definitions for each word, or a "gloss".

Structure

Semantic lexicons are made up of lexical entries. These entries are not orthographic, but semantic, eliminating issues of homonymy and polysemy. These lexical entries are interconnected with semantic relations, such as hyperonymy, hyponymy, meronymy, or troponymy. Synonymous entries are grouped together in what the Princeton WordNet calls "synsets"

Most semantic lexicons are made up of four different "sub-nets": nouns, verbs, adjectives, and adverbs, though some researchers have taken steps to add an "artificial node" interconnecting the sub-nets.

Nouns

Nouns are ordered into a taxonomy, structured into a hierarchy where the broadest and most encompassing noun is located at the top, such as "thing", with the nouns becoming more and more specific the further they are from the top. The very top noun in a semantic lexicon is called a *unique beginner*. The most specific nouns (those that do not have any subordinates), are *terminal nodes*.

Semantic lexicons also distinguish between types, where a type of something has characteristics of a thing such as a *Rhodesian Ridgeback* being a type of dog, and instances, where something is an example of said thing, such as *Dave Grohl* is an instance of a *musician*. Instances are always terminal nodes because they are solitary and don't have other words or ontological categories belonging to them.

Semantic lexicons also address meronymy, which is a "part-to-whole" relationship, such as keys are part of a laptop. The necessary attributes that define a specific entry are also necessarily present in that entry's hyponym. So, if a *computer* has *keys*, and a *laptop* is a type of *computer*, then a *laptop* must have *keys*. However, there are many instances where this distinction can become vague. A good example of this is the item *chair*. Most would define a chair as having legs and a seat (as in the part one sits on). However, there are some very

“artistic” and “modern” chairs in overpriced boutiques that do not have legs at all. Beanbags also do not have legs, but few would argue that they aren't chairs. Questions like this are the core questions that drive research and work in the fields of taxonomy and ontology.

Verbs

Verb synsetsare arranged much like their noun counterparts: the more general and encompassing verbs are near the top of the hierarchy while troponyms (verbs that describe a more specific way of doing something) are grouped beneath. Verb specificity moves along a vector, with the verbs becoming more and more specific in reference to a certain quality. For example. The set "walk / run / sprint" becomes more specific in terms of the speed, and "dislike / hate / abhor" becomes more specific in terms of the intensity of the emotion.

The ontological groupings and separations of verbs is far more arguable than their noun counterparts. It is widely accepted that a *dog* is a type of *animal* and that a *stool* is a type of *chair*, but it can be argued that *abhor* is on the same emotional plane as *hate* (that they are synonyms and not super/subordinates). It can also be argued that *love* and *adore* are synonyms, or that one is more specific than the other. Thus, the relations between verbs are not as agreed-upon as that of nouns.

Another attribute of verb synset relations is that there are also ordered into verb pairs. In these pairs, one verb necessarily entails the other in the way that *massacre* entails *kill*, and *know* entails *believe*. These verb pairs can be troponyms and

their superordinates, as is the case in the first example, or they can be in completely different ontological categories, as in the case in the second example.

Adjectives

Adjective synset relations are very similar to verb synset relations. They are not quite as neatly hierarchical as the noun synsetrelations, and they have fewer tiers and more terminal nodes. However, there are generally less terminal nodes per ontological category in adjective synset relations than that of verbs. Adjectives in semantic lexicons are organized in word pairs as well, with the difference being that their word pairs are antonyms instead of entailments. More generic polar adjectives such as *hot* and *cold*, or *happy* and *sad* are paired. Then other adjectives that are semantically similar are linked to each of these words. *Hot*is linked to *warm*, *heated*, *sizzling*, and *sweltering*, while *cold* is linked to *cool*, *chilly*, *freezing*, and *nippy*. These semantically similar adjectives are considered *indirect antonyms* to the opposite polar adjective (i.e. *nippy* is an indirect antonym to *hot*). Adjectives that are derived from a verb or a noun are also directly linked to said verb or noun across sub-nets. For example, *enjoyable*is linked to the semantically similar adjectives *agreeable*, and *pleasant*, as well as to its origin verb, *enjoy*.

Adverbs

There are very few adverbs accounted for in semantic lexicons. This is because most adverbs are taken directly from their adjective counterparts, in both meaning and form, and changed only morphologically (i.e. *happily* is derived from *happy*, and

luckily is derived from *lucky*, which is derived from *luck*). The only adverbs that are accounted for specifically are ones without these connections, such as *really*, *mostly*, and *hardly*.

Challenges facing semantic lexicons

The effects of the Princeton WordNet project extend far past English, though most research in the field revolves around the English language. Creating a semantic lexicon for other languages has proved to be very useful for Natural Language Processing applications. One of the main focuses of research in semantic lexicons is linking lexicons of different languages to aid in machine translation. The most common approach is to attempt to create a shared ontology that serves as a "middleman" of sorts between semantic lexicons of two different languages. This is an extremely challenging and as-of-yet unsolved issue in the Machine Translation field. One issue arises from the fact that no two languages are word-for-word translations of each other. That is, every language has some sort of structural or syntactic difference from every other. In addition, languages often have words that don't translate easily into other languages, and certainly not with an exact word-to-word match. Proposals have been madeto create a set framework for wordnets. Research has shown that every known human language has some sort of concept resembling synonymy, hyponymy, meronymy, and antonymy. However, every idea so far proposed has been met with criticism for using a pattern that works best for English and less for other languages.

Another obstacle in the field is that no solid guidelines exist for semantic lexicon framework and contents. Each lexicon

project in each different language has had a slightly (or not so slightly) different approach to their wordnet. There is not even an agreed-upon definition of what a "word" is. Orthographically, they are defined as a string of letters with spaces on either side, but semantically it becomes a very debated subject. For example, though it is not difficult to define *dog* or *rod* as words, but what about *guard dog* or *lightning rod*? The latter two examples would be considered orthographically separate words, though semantically they make up one concept: one is a type of dog and one is a type of rod. In addition to these confusions, wordnets are also idiosyncratic, in that they do not consistently label items. They are redundant, in that they often have several words assigned to each meaning (synsets). They are also open-ended, in that they often focus on and extend into terminology and domain-specific vocabulary.

Other names

- wordnet
- computational lexicon

List of semantic lexicons

- WordNet
- EuroWordNet
- Multilingual Central Repository
- Global Wordnet
- MindNet