Lecture Notes for 6.824, Artificial Itelligence
 Rev. November, 1992

Natural Language Syntax and Semantics

The treatment of natural language syntax and semantics presented here is based on an approach known as categorial grammar. This approach was introduced by a Polish logician Kasimierz Adjuciewicz in the mid 1930's [Adjuciewicz, 1935], rediscovered by Lambek in the 1950's [Lambek, 1958], but largely ignored by linguists until it was discovered yet a third time by Richard Montague [Montague, 1970]. The basic idea behind categorial gramar is that the syntactic categories of English and other natural languages are analogous to the types of the typed λ-calculus. For example, and adjective, like "yellow" modifies a noun phrase like "bird" to give a new noun phrase "yellow bird". This can be treated formally by giving the word "bird" the type N and the word "yellow" the type N $\rightarrow$ N. The phrase "yellow bird" is then a well typed application of the adjectival operator "yellow" to the argument "bird". The phrase "bird yellow" is not well typed under this scheme and is not a grammatical noun phrase of English. Treating syntactic categories as types and assigning types to the individual words works surprisingly well as a theory of gramaticality in English. As we shall see, however, this simple theory is far from providing a complete accout of grammaticality.

In this section we present a very simple version of categorial grammar. This simple treatment only scratches the surface of the rich theories of natural language syntax that have evolved over the past thirty years. However, our simple theory does provide basic insights into the syntax of natural language. In approaching a deep and complex subject one should start by understanding simple theories that are only approximately correct. For example, one should understand properties of ideal gases, or frictionless plains, before attempting to understand the properties of carbon dioxide or tires on wet pavement.

1 Categorial Grammar

In both natural language and the typed λ-calculus there is a notion of a syntactically well formed expression. In the typed λ-calculus, an application $(f\ t_1\ \cdots\ t_n)$ is syntactically well formed if f has type $\tau_1 \times \cdots \times \tau_n \rightarrow \sigma$ and each argument t_i has type τ_i. In English, the sentence "The yellow bird sang" is grammatical (syntactically well formed), while "Yellow the sang bird" is not. In making an analogy with the typed λ-calculus we analyze the grammatical sentence as ((the (Yellow bird)) sang) where the word "bird" has type N, the word "yellow" has type $N \rightarrow N$, the word "the" has type $N \rightarrow NP$, and the word "sang" as type $NP \rightarrow S$. Under this assignment of types, and under the assumption that operator "sang" follows its argument rather than precedes it, the sentence ((The (yellow bird)) sang) is an expression of type S.

English, and every other natural language, has syntactic categories — each phrase in English is either a noun phrase, a verb phrase, an adjectival phrase, and so on. The categories of English are analogous to the types of the typed λ-calculus. We use the symbol S to represent the category of a sentence, the symbol N to represent the category of a noun phrase noun such as "woman", the symbol NP to represent the category of a complete noun phrase such as "every woman" or "Mary", and the symbol VP to represent the category of a verb phrase such as "runs" or "kissed John". An English sentence is an expression of the form $(NP\ VP)$, such as ((Every woman) (kissed (some man))). The fact that an expression of the form $(NP\ VP)$ is a well formed expression of category S can be expressed by the statement that the category NP can be viewed as the arrow type $VP \rightarrow S$. Table 1 gives a list of English categories. Each category is named by a symbol, although many of these symbols are just abbreviations for arrow categories. In the case where a symbol is an abbreviation for an arrow category, the arrow category is shown in separate column.

Unlike the typed λ-calculus, an English operator can appear either before or after its arguments. To distinguish these two kinds of operators we allow the arrow in an arrow category to be written in either direction. An operator that appears before its arguments will have a type of the form $\tau_1 \times \cdots \times$

Category Name	Example Phrases	Arrow Category
S	((the (yellow bird)) flies)	category symbol
N	bird, man, (yellow bird)	category symbol
NP	Mary, (some bird), (every bird)	category symbol
VP	runs, (likes Tom), (flies quickly)	$S \leftarrow NP$
Adj	big, yellow	$N \rightarrow N$
AP	(on (the table))	$N \leftarrow N$
Adv	(into (the house))	$VP \leftarrow VP$

Table 1: Some fundamental categories of English

$\tau_n \rightarrow \sigma$ while an operator that appears after its arguments will have type $\sigma \leftarrow \tau_1 \times \cdots \tau_n$. For example, consider the phrase (big (yellow bird)). In this phrase, "bird" has category N and the adjectives "big" and "yellow" each have category $N \rightarrow N$. So the overall phrase has category N. Now consider the sentences "The yellow bird ate the seed". This sentence can be parsed as a noun phrase followed by a verb phrase ((The yellow bird) (ate the seed)). In this case the verb phrase is viewed as an operator that is applied to the noun phrase. Verb phrases have category $S \leftarrow NP$.[1]

An initial theory of English grammaticality can be constructed by hypothesizing that every word of English associated with a category expression that is either one of the basic category symbols, or a left or right arrow expression built up form other category expressions.

> **Definition** A *lexical entry* is a pair of a word and a category expression.

Intuitively, a lexical entry is an entry in the "internal dictionary" of a native

[1] It is also possible to give a categorial grammar in which VP is a category symbol and NP is the category $VP \rightarrow S$. However, treating VP as $S \leftarrow NP$ allows for quantifier scope ambiguity. There are also other examples of operators that follow rather than procede their arguments. For example, consider the adjectival phrase "on the roof" in "the man on the roof".

English speaker. It contains the category of the word, e.g., that it is a noun. Dictionaries also contain definitions. Definitions give semantic information. The semantics of English is discussed below. A dictionary often contains more than one lexical entry for the same word. In linguistics, the internal dictionary hypothesized to exist for a particular English speaker is called the *lexicon* of that speaker.

> **Definition** A *lexicon* is a set of lexical entries.

We can now define the set of word strings that are grammatical relative to a given lexicon. This definition is similar to the definition of the well-typed expressions of the types λ-calculus except that, in natural languages, operators are allowed to follow their arguments and individual words may be ambiguous in the sense that they have more than one lexical entry in the lexicon. The following definition formally defines a "can be parsed as" relation between word strings and category expressions. If W is a word string and α is a category expression, then we say W can be parsed as α if there exists a parenthesization of the words in W, and an assignment of types to words in accordance with the lexicon, such that the resulting parenthesized expression is well-typed and has type α.

> **Definition** The "can be parsed as" relation between word strings and category expressions is the least relation satisfying the following conditions.
>
> - If there exists a lexical entry for a word w with category expression α, then w can be parsed as α.
> - If the word string W can be written as the concatenation of substrings $U; S_1; \cdots; S_n$ where U can be parsed as $\alpha_1 \times \cdots \times \alpha_n \rightarrow \beta$, and each S_i can be parsed as α_i, then W can be parsed as β.
> - If the word string W can be written as the concatenation of substrings $S_1; \cdots; S_n; U$ where U can be parsed as $\beta \leftarrow \alpha_1 \times \cdots \times \alpha_n$, and each S_i can be parsed as α_i, then W can be parsed as β.

Consider the sentence "Every yellow bird flies quickly". This sentence can be parsed relative to the following assignment of category expressions to words.

- every: $N \rightarrow NP$
- yellow: $N \rightarrow N$
- bird: N
- flies: VP
- quickly: $VP \leftarrow VP$

The parenthesization resulting from the parse is

((Every (yellow bird)) (flies quickly)).

When all category expressions in the lexicon involve only right arrows (or only left arrows), and if every word has only a single lexical entry, then a given word string has at most one parse. However, even if every word has only a single lexical entry, the presence of both left and right arrow categories allow a single word string to have more than one parse. For example, consider the phrase "the old grey man's house". The possessive noun "man's" can be treated as an abbreviation for the word man followed by the "word" poss (for possessive). We treat poss as a word of category $A \leftarrow N$, i.e., it takes a noun phrase argument, which it follows, and produces a phrase of category A, i.e., of category $N \rightarrow N$. The word "the" has category $N \rightarrow NP$, the words "old" and "grey" have category A, and the words "man" and "house" have category N. The word string "the old grey man poss house" has three different parses corresponding to the three following partial parenthesizations.

the old grey ([man poss] house)

the old ([(grey man) poss] house)

the ([(old grey man) poss] house)

Another example of an ambiguous phrase is the sentence "john saw the man with the telescope". In this sentence the phrase "with the telescope" can either specify how john saw the man (with the telescope) or it can specify which man John saw (the man with the telescope). In this case the ambiguity can be attributed to two different lexical entries for the preposition "with".

- with: $NP \to Adv$
- with: $NP \to AP$

The first lexical entry specifies that "with" constructs an adverbial phrase, while the second lexical entry specifies that "with" constructs an adjectival phrase. Table 1 lists Adv is the category expression $VP \leftarrow VP$ and lists AP as the category expression $N \leftarrow N$.

The different parses of the above phrases seem to correspond to different semantic interpretations. It is possible to give a theory of the semantics of English sentences that is similar to the semantics of the typed λ-calculus.

2 Compositional Semantics

A theory of the semantics of English can be constructed from the simple assumption that the meaning of a phrase is constructed by applying the meaning of the operator to the meaning of the arguments. Under this theory, English is compositional in the technical sense that the meaning of a phrase is determined by the meaning of its subphrases. As in the case of syntax, the theory of semantics based on this simple compositional assumption is an idealization — a completely accurate theory of English semantics would be much more elaborate. The simple compositional theory is analogous to an ideal gas or a frictionless plain — it is a useful simplification of reality.

We assume two semantic domains denoted by the type symbols D and B. The type symbol D denotes the set of all objects. The type symbols B denotes the set $\{\mathbf{T},\ \mathbf{F}\}$. Table 2 is a copy of table 1 plus an extra column. The

Category Name	Example Phrases	Arrow Category	Type Expression
S	((the bird) flies)	category symbol	B
N	man, (yellow bird)	category symbol	$D \rightarrow B$
NP	Mary, (every bird)	category symbol	$(D \rightarrow B) \rightarrow B$
VP	runs, (likes Tom)	$S \leftarrow NP$	
Adj	big, yellow	$N \rightarrow N$	
AP	(on (the table))	$N \leftarrow N$	
Adv	(into (the house))	$VP \leftarrow VP$	

Table 2: An association of type expressions with category symbols.

extra column associates each of the basic category symbols with a type expression constructed from D and B. Given an interpretation of each category symbol as a type expression, any category expression can be converted to a type expression by replacing each category symbol by the corresponding type expression and replacing each left arrow by a right arrow (type expressions only have right arrows).

Although each category expression maps to a well defined type expression, it seems that category expressions, and not type expressions, determine grammaticality. For example, the direction of the arrow in a category expression is clearly important in determining grammaticality. The use of category expressions rather than type expressions in determining grammaticality is an example of Chomsky's general principle of the autonomy of syntax — syntactic grammaticality is, to some extent, independent of semantics.

For the remainder of this section I will assume a fixed interpretation of D as a set and of B as the set $\{\mathbf{T}, \mathbf{F}\}$. This assignment of sets to D and B provides a fixed semantic value for each type expression built from D and B. Under the compositional semantics for English, the semantic value of an English phrase is always a member of the semantic value of the type expression assigned to that phrase. To construct the meaning of an English phrase we must assume that each entry in the lexicon associates a word with both a category expression and a semantic value. To specify this we strengthen the definition of a lexical entry.

> **Definition** A *lexical entry* is a triple of a word, a category expression, and a semantic value where the semantic value is an element of the set denoted by the type expression derived from the given category expression.

Given the semantic values in lexical entries, we can define the semantic value of a string of English words. Actually, since a single word string can have several different parses, a single word string can also have several possible different semantic values. The following definition specifies a set of triples (similar to the triples in the lexicon) where each triple consists of a word string, a category expression, and a semantic value.

Definition: We say that a word string W *parses as category* β *with semantic value* v if any of the following conditions hold.

- W consists of a single word and there is a lexical entry for W with category β and value v.
- W can be written as the concatenation of substrings $U; S_1; \cdots; S_n$ where U can be parsed as $\alpha_1 \times \cdots \times \alpha_n \rightarrow \beta$ with value f, each S_i can be parsed as α_i with value v_i, and the value v is the value of f applied to the arguments $v_1, \cdots, v_n$.
- W can be written as the concatenation of substrings $S_1; \cdots; S_n; U$ where U can be parsed as $\beta \leftarrow \alpha_1 \times \cdots \times \alpha_n$ with value f, each S_i can be parsed as α_i with value v_i, and the value v is the value of f applied to the arguments $v_1, \cdots, v_n$.

Consider the sentence ((Some bird) (ate (every seed)). This sentence can be assigned an intuitively correct semantic value by assigning an intuitively plausible value to each word and then constructing the value of the sentence by application. The word "bird" has category N which is associated with the type expression $D \rightarrow B$. According to the semantic constraints on lexical entries, the word "bird" must denote an element of the type $D \rightarrow B$, i.e., it must denote a predicate that maps objects to truth value. We can assume that the word "bird" denotes the predicate that is true of an object just in case that object is a bird. Now consider the phrase (some bird). To assign a meaning to this phrase one must assign a meaning to the word "some". The word "some" has category $N \to NP$. This maps to the type expression

$$((D \to B) \to ((D \to B) \to B))$$

This type expression can be written more clearly by introducing P (for predicate) as an abbreviation for the type expression $D \to B$. The type of "some" can now be written as

$$P \to (P \to B).$$

Any operator of type $P \to (P \to B)$ contains exactly the same information as an operator of type $P \times P \to B$. So the word "some" can be viewed as an operator that takes two predicates and returns a truth value. Let Q and R be two predicates. Intuitively, "some" applied to Q and R should be true if "Some Q is R", i.e., there is some x such that Q is true of x and R is true of x. This leads to the hypothesis that the semantic value of "some" is the operator denoted by the λ-expression

$$(\lambda\ (Q)\ (\lambda\ (R)\ \exists_D x\ (Q\ x) \wedge (R\ x))).$$

Given this semantic value for "some", the phrase (some bird) denotes an operator of type $P \to B$ that takes a predicate and returns true just in case that predicate is true of some bird. Similarly, the semantic value of "every" is the operator denoted by the expression

$$(\lambda\ (Q)\ (\lambda\ (R)\ \forall_D x\ (Q\ x) \to (R\ x))).$$

Given this semantic value for "every", the phrase (every bird) denotes an operator that takes a predicate and returns true just in case that predicate is true of every bird.

Some further comments are in order regarding objects denoted by expressions of category NP. I will use the symbol Q as an abbreviation for the type expression associated with the category symbol NP, i.e., Q is an abbreviation for the type expression $P \to B$ (where P is $D \to B$). Objects of type Q will be called *quantifiers* (hence the symbol Q). The quantifiers $\forall_D$

and $\exists_D$ of higher order logic are both objects of type Q. The propositions $(\forall_D(\lambda\ (x)\ \Phi(x)))$ and $(\exists_D(\lambda\ (x)\ \Phi(x))$ of higher order logic represent the first order propositions $\forall x\ \Phi(x)$ and $\exists x\ \Phi(x)$ respectively. To say that $\Phi(x)$ is true for every bird x, one can write

$$((\text{every bird})\ (\lambda\ (x)\ \Phi(x))).$$

We now return to a consideration of the sentence ((some bird) (ate (every seed))). The word "ate" has category $NP \rightarrow VP$ where VP is an abbreviation for $S \leftarrow NP$. The type expression corresponding to $NP \rightarrow VP$ is $Q \rightarrow (Q \rightarrow B)$. The word "ate" must denote an operator that takes two quantifiers and returns a truth value. Let E be a predicate that is true two objects x and y just in case x ate y. The semantic value of "ate" can be taken to be the object denoted by the following λ-term.

$$(\lambda\ (A)\ (\lambda\ (B)\ (A\ (\lambda\ (y)\ (B(\lambda\ (x)\ (Exy)))))))$$

Given this semantic value of value of "ate", the sentence ((some bird) (ate (every seed))) is true if for each seed there exists a bird, possibly different for each seed, that ate the seed. However, the English sentence seems to be semantically ambiguous. It could also mean that there exists a single bird that ate every seed. The second reading can be generated by a different lexical entry for "ate" whose semantic value is object denoted by

$$(\lambda\ (A)\ (\lambda\ (B)\ (B\ (\lambda\ (y)\ (A(\lambda\ (x)\ (Exy))))))).$$

It is possible to construct a lexicon for a set of English words that accounts for both the syntax and semantics of a large number of English sentences constructed from those words. Categorial grammar and compositional semantics seems to be a useful idealization of the syntax and semantics of natural language. This idealization leads to the (idealized) conjecture that the syntax and semantics of any natural language is determined by the lexicon of that language.

3 Syntactic Complications

Empirically, it seems that categorial grammar and compositional semantics do not provide a full account of the syntax and semantics of natural language. One thing that seems unaccounted for is "number agreement". The verb of a sentence must agree with main noun phrase, if the noun phrase is plural then the verb must be in plural form and if the noun phrase is singular the verb phrase must agree. For example consider the following sentences. Sentences that people judge to be ungrammatical are labeled with $\star$.

The bird is pretty.

$\star$ The bird are pretty.

The birds are pretty.

$\star$ The birds is pretty.

Another example of noun-verb agreement is related to whether the noun is first second or third person. For example, consider the following sentences.

I want to go fishing.

John wants to go fishing.

$\star$ I wants to go fishing.

$\star$ John want to go fishing.

Another thing unaccounted for in categorial grammar is case. In many languages, such as latin, word order is not very important — most operators are free to appear before, after, or between their arguments. For example, a transitive verb can be viewed as an operator that takes two noun phrases as arguments. In Latin, the arguments can appear in either order, and the verb can appear before, after, or between the arguments. Fortunately, one of the argument noun phrases must be in nominative case and the other in the accusative case. Case can be used to determine how the operator is to be applied to the arguments.

A third aspect of grammaticality not accounted for by categorial grammar concerns a phenomenon known as traces. A trace is a theoretical device used to explain "missing arguments". For example, consider the sentence

Mary liked the picture John drew.

The phrase "John drew" acts as an adjectival phrase modifying "picture". However, it seems unreasonable to conjecture lexical entries for "John" and "drew" such that "John drew" can be parsed as AP. It seems more reasonable to assume that the phrase "the picture John drew" has a hidden variable, called a trace, that is the argument of "drew". This phrase can be treated as phonological variant of (the (picture (λ (x) (John (drew x))))) where the λ-predicate (λ (x) (John (drew x))) is an adjectival phrase. The variable x that serves as an argument to "drew" is called a trace because it does not appear in the spoken version of the sentence. Traces account for a wide variety of sentences. For example, consider the following.

Mary liked the picture John gave to Sally.

Mary liked the picture John said was expensive.

Mary liked the picture John saw Sally give to Bill.

There is even phonological evidence that traces are actually present in the phrase structures people assign to sentences such as the above. Consider the following two phrases.

The person I want to see

The person I want to go away

Both of these sentences involve traces. They can be taken to be phonological variants of the following expressions.

The person (λ (x) I want to see x)

The person (λ (x) I want x to go away)

The pair of words "want to" are often pronounced as "wanna". However, it seems that if a trace appears between "want" and "to" this contraction is prevented.

The person I wanna see

$\star$ The person I wanna go to the store

Clearly, the simple version of categorial grammar given above must be modified to account for all of the syntactic features found in human languages. Modern linguistics includes a variety of formal systems for describing grammaticality. All of the formal frameworks attempt to provide a framework in which the variety of human languages can be understood.

4 Semantic Complications

In the idealized theory of compositional semantics, it was assumed that the semantic value of a sentence was always a truth value. This idealized semantical picture must be modified to handle sentences such as the following.

Mary believes John stole the painting.

This sentence can be parsed provided that "believe" is assigned the category $S \rightarrow (S \leftarrow NP)$. In other words, "believe" denotes an operator that takes a sentence and a noun phrase. The difficulty is semantic rather than syntactic. We have assumed that the semantic value of any sentence is a truth value. This implies that "believe" is an operator that takes a truth value and a noun phrase. The only plausible semantic interpretation of "believe" is such that a sentence of the form "Mary believes Φ" is true if and only if Φ is true. This is clearly counter to common sense.

One way of attempting to overcome this semantic difficulty is to introduce the notion of "situations". Intuitively, a situation is a time and place in some actual or possible world. For example, 12 o'clock noon on May 3, 1989, at

Category Name	Example Phrases	Type Expression
S	(Mary runs)	$Sit \to B$
N	(yellow bird)	$D \to (Sit \to B)$
NP	Mary, (some bird)	$(D \to (Sit \to B)) \to (Sit \to B)$

Table 3: Situations in the type expressions.

table four of the Cornell faculty club. Situations can also be "imaginary" or "possible but not actual". Santa's work shop on Christmas eve is a situation.

Rather than say that sentences denote truth values, we can say that sentences denote functions from situations to truth values. For example, if Tweety is a particular bird, then the sentence "Tweety can fly" is true at some times but false at times when Tweety's wings have just been clipped. To incorporate situations into compositional semantics we introduce a new type symbol Sit which denotes the set of all possible situations (actual, possible, and imaginary). We retain the two previous type symbols D and B, but we change the assignment of type expressions to category symbols. More specifically, every occurrence of B in the previous assignment is replaced by $Sit \to B$. The resulting assignment is shown in figure 3. It is useful to compare table 3 with table 2.

Given this assignment of type expressions to category symbols, the lexical entries for most words are virtually identical to the entries given above. For example, the semantic value of "some" is the object denoted by the λ-term

$$(\lambda\ (Q)\ (\lambda\ (R)\ (\exists_I(\lambda\ (x)\ (Rx) \wedge_I (Qx))))).$$

This λ-term is identical to the previous semantic value of "some" except that $\exists_D$ has been replaced by $\exists_I$ and $\wedge$ has been replaced by $\wedge_I$. $\exists_I$ and $\wedge_I$ are the *intensional* versions of existential quantification and Boolean conjunction. The operator $\wedge_I$ takes two objects of type $Sit \to B$ and returns an object of type $Sit \to B$ such that $\Phi \wedge_I \Psi$ is true of a situation s if and only if both Φ and Ψ are true of s. The intensional existential quantifier $\exists_I$ is an operator that takes an object of type $D \to (Sit \to B)$ and returns an object of type $Sit \to B$. Such that the object $(\exists_I(\lambda\ (x)\ \Phi(x)))$ is true when applied to a situation s just in case there is some object d such that $\Phi(d)$ is true when applied to s.

Under this semantics, sentences denote objects of type $Sit \to B$. It is possible to assign a semantic value to the word "believes" such that the sentence "Mary believes Φ" is true of a situation s, even though Φ is false of s. Even the situation based semantics of sentences may not be adequate. The following sentences present semantic difficulties of various kinds.

Mary wants John to go away.

Mary saw John run down the street.

Mary knows John's telephone number.

5 Problems

1. Fill the missing semantic types in the following table. Note that this table is different form the one given in the notes. In the table given below VP is a primitive syntactic type. Each semantic type expression should be an arrow expression built from the basic semantic types D and B.

Category Name	Example Phrases	Arrow Category	Semantic Type Expression
S	((the bird) flies)	category symbol	B
N	man, (yellow bird)	category symbol	$D \to B$
VP	runs, (likes Tom)	category symbol	$D \to B$
NP	Mary, (every bird)	$VP \to S$	
Adj	big, yellow	$N \to N$	
AP	(on (the table))	$N \leftarrow N$	
Adv	(into (the house))	$VP \leftarrow VP$	

2. Under the type system given in problem 1 the word "every" has syntactic

category $N \to NP$. Give a semantic value for the word “every” that is compatible with the semantic type expressions of part a. Your value should be a λ-term of the higher order logic and should give the intuitively correct meaning for the sentence “Every bird flys”.

3. Show the equivalence of categorial and context free grammars.

References

[Adjuciewicz, 1935] Kasimierz Adjuciewicz. Die syntaktische konnexitat. *Studia Philophica*, 1:1–27, 1935. Translated as “Sytactic Connection” in Strolls McCall (ed), *Polish Logic: 1920-1939* (Oxford University Press, 1967).

[Lambek, 1958] J. Lambek. The mathematics of sentence structure. *American Mathematical Monthly*, 65:154–169, 1958.

[Montague, 1970] Richard Montague. Universal grammar. *Theoria*, 36:373–398, 1970. Reprinted in: *Formal Philosophy: Selected Papers of Richard Montague*, ed. by R. H. Thomason, Yale University Press, 1974.