

Introduction to Linguistics

Marcus Kracht
Department of Linguistics, UCLA
3125 Campbell Hall
450 Hilgard Avenue
Los Angeles, CA 90095–1543
kracht@humnet.ucla.edu

Lecture 1: Introduction

Before we start. The notes to the lectures reflect the content of the lectures more accurately than the book by Fromkin et. al. You will find the notes much more technical than the book. Do not panic! You are not required to master these technicalities right away. The technical character is basically due to my desire to be as explicit and detailed as possible. For some of you this might actually be helpful. If you are not among them you may concentrate on the book instead. However, linguistics is getting increasingly formal and mathematical, and you are well advised to get used to this style of doing science. So, even if you do not understand right away what I am saying, do try again and again. And keep asking questions. New words and technical terms that are used for the first time are typed in bold-face. If you are supposed to know what they mean, a definition will be given right away. The definition is valid throughout the entire course, but be aware of the fact that other people might define things differently. (This applies when you read other books, for example.) If you are not given a definition elsewhere, be cautious.

Language is a means to communicate, it is a semiotic system. Its units are **signs**. Principally, a sign is a pair consisting — in the words of Ferdinand de Saussure — of a **signifier** and a **signified**. We prefer to call the signifier the **exponent** and the signified the **meaning**. For example, in English the string dog is a signifier, and its signified is the set of all dogs. Sign systems are ubiquitous: clocks, road signs, pictograms — they all are parts of sign systems. Language differs from them only in the complexity. This explains why language signs have much more internal structure than ordinary signs. For notice that language allows to express virtually every thought that we have, and the number of signs that we can produce is literally endless. Although one may find it debatable whether or not language is actually infinite, it is clear that we are able to understand utterances that we have never heard before. Every year, hundreds of thousands of books appear, and clearly each of them is new. If it were the same as a previously published book this would be considered a breach of copyright! However, no native speaker of the language experiences trouble understanding them (apart from technical books).

It might be far fetched, though, to speak of an entire book as a sign. But it certainly is. Linguists mostly study only signs that consist of just one a sentence. And this is what we shall do here, too. However, texts are certainly more than

a sequence of sentences, and the study of **discourse** (which includes texts and dialogs) is certainly a very vital one. Unfortunately, even sentences are so complicated that it will take all our time to study them. The methods, however, shall be useful for discourse analysis as well.

For a linguist, language signs are constituted of four different levels, not just two: **phonology**, **morphology**, **syntax** and **semantics**. **Semantics** deals with the meanings (what is signified), while the other three are all concerned with the exponent. At the lowest level we find that everything is composed from a small set of sounds, or — when we write — of letters. (Chinese is exceptional in that the alphabet consists of around 50,000 ‘letters’, but each sign stands for a syllable — a sequence of sounds, not just a single one.) With some exceptions (for example tone and intonation) every utterance can be seen as a sequence of sounds. For example, **dog** consists of three letters (and three sounds): **d**, **o** and **g**. In order not to confuse sounds with letters we denote the sounds by enclosing them in square brackets. So, the sounds that make up [dog] are [d], [o] and [g], in that order. What is important to note here is that sounds by themselves in general have no meaning. The decomposition into sounds has no counterpart in the semantics. Just as every sign can be decomposed into sounds, it can also be decomposed into words. In written language we can spot words as minimal parts of text enclosed by blanks. In spoken language the definition of word becomes very tricky. The part of linguistics that deals with how words are put together into sentences is called **syntax**. On the other hand, words are not the smallest meaningful units of language. For example, **dogs** is the plural of **dog** and as such it is formed by a regular process, and if we only know the meaning of **dog** we also know the meaning of **dogs**. Thus, we can decompose **dogs** into two parts: **dog** and **s**. The minimal parts of speech that bear meaning are called **morphemes**. Often, it is tacitly assumed that a morpheme is a part of a word; bigger chunks are called **idioms**. Idioms are **kick the bucket**, **keeps taps on someone**, and so on. The reason for this division is that while idioms are intransparent as far as their meaning is concerned (if you die you do not literally kick a bucket), syntactically they often behave as if they are made from words (for example, they inflect: **John kicked the bucket**).

So, a word such as ‘**dogs**’ has four manifestations: its meaning, its sound structure, its morphological structure and its syntactic structure. The levels of manifestation are also called **strata**. (Some use the term **level of representation**.) We use the following notation: the sign is given by enclosing the string in brack-

ets: ‘dog’. $[\text{dog}]_P$ denotes its phonological structure, $[\text{dog}]_M$ its morphological structure, $[\text{dog}]_L$ its syntactic structure and $[\text{dog}]_S$ its semantical structure. We also agree to use typewriter font for symbols in print, and that for the most part we analyse language as written language, unless otherwise indicated. With that in mind, we have $[\text{dog}]_P = \text{dog}$. The latter is a string composed from three symbols, d, o and g. So, ‘dog’ refers to the sign whose exponent is written here dog. We shall agree on the following.

Definition 1 *A sign is a quadruple $\langle \pi, \mu, \lambda, \sigma \rangle$, where π is its **exponent** (or **phonological structure**), μ its **morphological structure**, λ its **syntactic structure** and σ its **meaning** (or **semantic structure**).*

We write signs vertically, in the following way.

$$(1) \quad \begin{bmatrix} \sigma \\ \lambda \\ \mu \\ \pi \end{bmatrix}$$

This definition should not be taken as saying something deep. It merely fixes the notion of a linguistic sign, saying that it consists of nothing more (and nothing less) than four things: its phonological structure, its morphological structure, its syntactic structure and its semantic structure. Moreover, in the literature there are numerous *different* definitions of signs. You should not worry too much here: the present definition is valid throughout this book only. Other definitions have other merits.

The power of language to generate so many signs comes from the fact that it has rules by which complex signs are made from simpler ones.

(2) Cars are cheaper this year.

In (2), we have a sentence composed from 5 words. The meaning of each word is enough to understand the meaning of (2). Exactly how this is possible is one question that linguistics has to answer. We shall illustrate the technical approach. We assume that there is a binary operation \bullet , called **merge**, which takes two signs and forms a new sign. \bullet operates on each of the strata (or levels of manifestation) independently. This means that there are four distinct operations, \textcircled{P} , \textcircled{M} , \textcircled{L} , and

Ⓢ, which simultaneously work together as follows.

$$(3) \quad \begin{bmatrix} \sigma_1 \\ \lambda_1 \\ \mu_1 \\ \pi_1 \end{bmatrix} \bullet \begin{bmatrix} \sigma_2 \\ \lambda_2 \\ \mu_2 \\ \pi_2 \end{bmatrix} = \begin{bmatrix} \sigma_1 \textcircled{\text{S}} \sigma_2 \\ \lambda_1 \textcircled{\text{L}} \lambda_2 \\ \mu_1 \textcircled{\text{M}} \mu_2 \\ \pi_1 \textcircled{\text{P}} \pi_2 \end{bmatrix}$$

Definition 2 A *language* is a set of signs. A *grammar* consists of a set of signs (called *lexicon*) together with a finite set of functions that operate on signs.

Typically, though not necessarily, the grammars that linguists design for natural languages consist in the lexicon plus a single binary operation \bullet of merge. There may also be additional operations (such as movement), but let's assume for simplicity that this is not so. Such a grammar is said to **generate** the following language (= set of signs) L :

- ① Each member of the lexicon is in L .
- ② If S and S' are in L , then so is $S \bullet S'$.
- ③ Nothing else is in L .

(Can you guess what a general definition would look like?) We shall now give a glimpse of how the various representations look like and what these operations are. It will take the entire course (and much more) to understand the precise consequences of Definitions 1 and 2 and the idea that operations are defined on each stratum *independently*. But it is a very useful one in that it forces us to be clear and concise. Everything has to be written into one of the representations in order to have an effect on the way in which signs combine and what the effect of combination is.

For example, $\textcircled{\text{P}}$ is typically concatenation, with a blank added. Let us represent strings by \vec{x}, \vec{y} etc., and concatenation by \frown . So,

$$(4) \quad \text{d} \frown \text{a} \frown \text{x} \text{y} = \text{d} \frown \text{a} \text{c} \text{x} \text{y}$$

$$(5) \quad \text{a} \text{d} \frown \square \frown \text{x} \text{y} = \text{a} \text{d} \frown \text{ } \text{x} \text{y}$$

Notice that visually, \square ('blank') is not represented at the end of a word. However, it is a symbol (on a typewriter you have to press space) to get it. So, $\text{x} \frown \square$ is *not*

the same as \mathbf{x} ! Now we have

$$(6) \quad \vec{x} \oplus \vec{y} := \vec{x} \square \vec{y}$$

For example, ‘this year’ is composed from the words ‘this’ and ‘year’. And we have

$$(7) \quad \text{this year} = [\text{this year}]_P = [\text{this}]_P \oplus [\text{year}]_P = \text{this} \square \text{year}$$

This, however, is valid only for words and only for written language. The composition of smaller units is different. No blank is inserted. For example, the word ‘car’ (more exactly the morpheme) $[\text{car}]_M$ and the plural morpheme ‘s’ compose to give *cars*, not *car s*. Moreover, the plural of *man* is *men*, so it is not all formed by adding *s*. We shall see below how this is dealt with.

Morphology does not get to see the individual makeup of its units. In fact, the difference between ‘car’ and ‘cat’ is morphologically speaking as great as that between ‘car’ and ‘moon’. Also, both are subject to the same morphological rules and behave in the same way, for example form the plural by adding ‘s’. That makes them belong to the same noun class. Still, they are counted as different morphemes. This is because they are manifested differently (the sound structure is different). This is why we distinguish between a morpheme and its **morphological structure**. The latter is only the portion that is needed on the morphological stratum to get everything right.

Definition 3 A *morpheme* is an indecomposable sign.

A morpheme can only be defined relative to a grammar. If we have only \bullet , then S is a morpheme if there are no S' and S'' with $S = S' \bullet S''$. (If you suspect that essentially the lexicon may consist in all and only the morphemes, you are right. Though the lexicon may contain more elements, it cannot contain less.) A word is something that is enclosed by blanks and/or punctuation marks. So the punctuation marks make the morpheme a word. To morphology, ‘car’ is known as a noun that takes an s-plural. We write

$$(8) \quad \left[\begin{array}{l} : n \\ : s\text{-pl} \end{array} \right]$$

to say that the item is of morphological category ‘n’ (nominal) and that it has inflectional category ‘s-pl’ (which will take care of the fact that its plural will be formed by adding ‘s’).

To the syntactic stratum the item *cars* is known only as a plural noun despite the fact that it consists of two morphs. Also, syntax is not interested in knowing how the plural was formed. The syntactic representation therefore is the following.

$$(9) \quad \left[\begin{array}{l} : \text{ N} \\ : \text{ pl} \end{array} \right]$$

This says that we have an object of category N whose number is plural. We shall return to the details of the notation later during the course. Now, for the merge on the syntactic stratum let us look again at ‘this year’. The second part, ‘year’ is a noun, the first a determiner. The entire complex has the category of a determiner phrase (DP). Both are singular. Hence, we have that in syntax

$$(10) \quad \left[\begin{array}{l} : \text{ D} \\ : \text{ sg} \end{array} \right] \oplus \left[\begin{array}{l} : \text{ N} \\ : \text{ sg} \end{array} \right] = \left[\begin{array}{l} : \text{ DP} \\ : \text{ sg} \end{array} \right]$$

This tells us very little about the action of \oplus . In fact, large parts of syntactic theory are consumed by finding out what merge does in syntax!

Semantical representations are too complex to be explained here (they require a course in model–theory or logic). We shall therefore not say much here. Fortunately, most of what we shall have to say here will be clear even without further knowledge of the structures. Suffice it to say, for example, that the meaning of ‘car’ is the set of all cars (though this is a massive simplification this is good enough for present purposes); it is clearly different from the meaning of ‘cat’, which is the set of all cats. What is the meaning of ‘cars’? It turns out that it is the set of all sets of cars that have at least two members. The operation of forming the plural takes a set A and produces the set of all subsets of A that have at least two members. So:

$$(11) \quad [s]_S : \{\clubsuit, \heartsuit, \diamond, \clubsuit\} \mapsto \\ \{ \{\clubsuit, \heartsuit\}, \{\clubsuit, \diamond\}, \{\clubsuit, \clubsuit\}, \{\heartsuit, \diamond\}, \{\heartsuit, \clubsuit\}, \{\diamond, \clubsuit\}, \\ \{\clubsuit, \heartsuit, \diamond\}, \{\clubsuit, \heartsuit, \clubsuit\}, \{\clubsuit, \diamond, \clubsuit\}, \{\heartsuit, \diamond, \clubsuit\}, \{\heartsuit, \diamond, \clubsuit\}, \\ \{\clubsuit, \heartsuit, \diamond, \clubsuit\} \}$$

With this defined we can simply say that \oplus is function application.

$$(12) \quad M \oplus N := \begin{cases} M(N) & \text{if defined,} \\ N(M) & \text{otherwise.} \end{cases}$$

The function is $[s]_S$ and the argument is $[car]_S$, which is the set of all cars. By definition, what we get is the set of all sets of cars that have at least two members in it. Our typographical convention is the following. For a given word, say ‘cat’ the semantics is denoted by sans–serife font plus an added prime: *cat’*.

Here is a synopsis of the merge of ‘this’ and ‘year’.

$$(13) \quad \left[\begin{array}{c} \text{this}' \\ \left[\begin{array}{c} \text{D} \\ \text{sg} \end{array} \right] \\ \left[\begin{array}{c} \text{n} \\ \text{abl} \end{array} \right] \\ \text{this} \end{array} \right] \bullet \left[\begin{array}{c} \text{year}' \\ \left[\begin{array}{c} \text{N} \\ \text{sg} \end{array} \right] \\ \left[\begin{array}{c} \text{n} \\ \text{s-pl} \end{array} \right] \\ \text{year} \end{array} \right] = \left[\begin{array}{c} \text{this' (year')} \\ \left[\begin{array}{c} \text{DP} \\ \text{sg} \end{array} \right] \\ \left[\begin{array}{c} \text{np} \\ \star \end{array} \right] \\ \text{this year} \end{array} \right]$$

(Here, ‘abl’ stands for ‘ablaut’. What it means is that the distinction between singular and plural is signaled only by the vowel. In this case it changes from [ɪ] to [i:]. ★ means: no value.) One may ask why it is at all necessary to distinguish morphological from syntactic representation. Some linguists sharply divide between lexical and syntactical operations. Lexical operations are those that operate on units below the level of words. So, the operation that combines ‘car’ and plural is a lexical operation. The signs should have no manifestation on the syntactical stratum, and so by definition, then, they should not be called signs. However, this would make the definition unnecessarily complicated. Moreover, linguists are not unanimous in rejecting syntactic representations for morphemes, since it poses more problems than it solves (this will be quite obvious for so-called polysynthetic languages). We shall not attempt to solve the problem here. Opinions are quite diverse and most linguists do accept that there is a separate level of morphology.

A last issue that is of extreme importance in linguistics is that of deep and surface structure. Let us start with phonology. The sound corresponding to the letter *l* differs from environment to environment (see Page 525 of Fromkin et. al.). The ‘*l*’ in *slight* is different from the ‘*l*’ in *listen*. If we pronounce *listen* using the ‘*l*’ sound of *slight* we get a markedly different result (it sounds a bit like Russian accent). So, one letter has different realizations, and the difference is recognized by the speakers. However, the difference between these sounds is redundant in the language. In fact, in written language they are represented by

just one symbol. Thus, one distinguishes a **phone** (= sound) from a **phoneme** (= set of sounds). While phones are language independent, phonemes are not. For example, the letter p has two distinct realizations, an aspirated and an unaspirated one. It is aspirated in *pot* but unaspirated in *spit*. Hindi recognizes two distinct phonemes here. A similar distinction exists in all other strata, though we shall only use the distinction between **morph** and **morpheme**. A morpheme is a set of morphs. For example, the plural morpheme contains a number of morphs. One of them consists in the letter s, another in the letters en (which are appended, as in *ox/oxen*), a third is zero (*fish/fish*). And some more. The morphs of a morpheme are called allomorphs of each other. If a morpheme has several allomorphs, how do we make sure that the correct kind of morph is applied in combination? For example, why is the plural of *car* not *caren* or *car*? The answer lies in the morphological representation. Indeed, we have proposed that morphological representations contain information about word classes. This means that for nouns it contains information about the kind of plural morph that is allowed to attach to it. If one looks carefully at the setup presented above, the distinction between deep and surface stratum is however nonexistent. There is no distinction between morpheme and morph. Thus, either there are no morphs or there are no morphemes. Both options are theoretically possible.

Some notes. The idea of stratification is implicit in many syntactic theories. There are differences in how the strata look like and how many there are. **Transformational grammar** recognizes all four of the strata (they have been called **Logical Form** (for the semantical stratum) **S-structure** (for syntax) and **Phonetic Form** (for phonological stratum). Morphology has sometimes been considered a separate, lexical stratum, although some theories (for example **Distributed Morphology**) try to integrate it into the overall framework. **Lexical Functional Grammar (LFG)** distinguishes **c(onstituent)-structure** (= syntax), **a(rgument)-structure**, **f(unctional)-structure** and **m(orphological)-structure**.

There has also been **Stratificational Grammar**, which basically investigated the stratal architecture of language. The difference with the present setup is that Stratificational Grammar assumes independent units at all strata. For example, a morpheme is a citizen of the morphological stratum. The morpheme 'car' is different from the morpheme 'cat', for example. Moreover, the lexeme 'car' is once again different from the morpheme 'car', and so on. This multiplies the linguistic ontology beyond need. Here we have defined a morpheme to be a sign of some sort, and so it has just a manifestation on all strata rather than belonging

to any of them. That means that our representation shows no difference on the morphological stratum, only on the semantical and the phonological stratum.

Lecture 2: Phonetics

We begin with phonology and phonetics. It is important to understand the difference between phonetics and phonology. Phonetics is the study of actual sounds of human languages, their production and their perception. It is relevant to linguistics for the simple reason that the sounds are the primary physical manifestation of language. Phonology on the other hand is the study of sound *systems*. The difference is very important though often enough it is not evident whether a phenomenon is phonetic in nature or phonological. English, for example, has a basic sound [t]. While from a phonological point of view there is only one *phoneme* [t], there are infinitely many actual sounds that realize this phoneme. So, while there are infinitely many different sounds for any given language there are only finitely many phonemes, and the upper limit is around 120. English has 40 (see Table 7). The difference can be illustrated also with music. There is a continuum of pitches, but the piano has only 88 keys, so you can produce only 88 different pitches. The chords of the piano are given (unlike a synthesizer), so that the basic sound colour cannot be altered. But you can still manipulate the loudness, for example. Sheet music reflects this state of affairs in the same way as written language. The musical sounds are described by discrete signs, the keys. Returning now to language: the difference between various different realizations of the letter t, for example, are negligible in English and often enough we cannot even tell the difference between them. Still, if we recorded the sounds and mapped them out in a spectrogram we could actually see the difference. (Spectrograms are one important instrument in phonetics because they visualize sounds so that you can see what you often cannot hear.) Other languages cut the sound continuum in a different way. Not all realizations of t in English are acceptable in French (all the aspirated ones), for example. This means that if we think of the sounds as forming a 'space' the so-called basic sounds of a language occupy some region of that space. These regions vary from one language to another.

Languages are written in alphabets, and many use the Latin alphabet. It turns out that not only is the Latin alphabet not always suitable for other languages, orthographies are often not a reliable source for pronunciation. English is a case in point. To illustrate the problems, let us look at the following tables (taken from Florian Coulmas: *Writing Systems*, Cambridge University Press, 2003). Table 1 concerns the values of the letter x in different languages: As one can see, the correspondence between letters and sounds is not at all uniform. On the other

Table 1: The letter x in various languages

Language	Value
Albanian	[dʒ]
Basque	[x]
English	[gz]
French	[gz]
German	[ks]
Portuguese	[ʃ]
Spanish	[ç]
Pinyin of Mandarin	[ɕ]

Table 2: The sound [ə] in English

Letter	Example
a	about
e	believe
i	compatible
o	oblige
u	circus

hand, even in one and the same language the correspondence can be nonuniform. Table 2 lists ways to represent [ə] in English by letters. Basically any of the vowel letters can represent [ə]. This mismatch has various reasons, a particular one being language change (in time or space). For one, if language changes in time, so does its material side. The good side about a stable writing system is that we can (in principle) read older texts even if we do not know how to pronounce them. Second, language with strong dialectal variation often fix writing according to one of the dialects. Once again this means that documents are understood across dialects, even though they are read out differently.

The disadvantage for the linguist is that the standard orthographies have to be learned (if you study many different languages this can be a big impediment) and second they do not reveal what is nevertheless important: the sound quality. For that reason one has agreed on a special alphabet, the so-called **International**

Phonetic Alphabet (IPA). In principle this alphabet is designed to give an accurate written transcription of sounds, one that is uniform for all languages. Since the IPA is an international standard, it is vital that one understands how it works (and can read or write using it). The complete set of symbols is rather complex, but luckily one does not have to know all of it.

The Analysis of Speech Sounds

First of all, the continuum of speech is broken up into a sequence of discrete units, which we referred to as sounds. Thus we are claiming that language utterances are sequences of sounds. Right away we mention that there is an exception. **Intonation** and **stress** do not comply with this. The sentences below are distinct only in intonation (falling pitch versus falling and rising pitch).

- (14) You spoke with the manager.
(15) You spoke with the manager?

Also, the word **protest** has two different pronunciations; when it is a noun the stress is on the first syllable, when it is a verb it is on the second. Stress and intonation obviously affect the way in which the sounds are produced (changing loudness and / or pitch), but in terms of decomposition of an utterance into segments intonation and stress have to be taken apart. We shall return to stress later. Suffice it to say that in **IPA** stress is marked not on the vowel but on the syllable (by a ['] before the stressed syllable), since it is thought to be a property of the syllable. **tone** is considered to be a suprasegmental, too. It does not play a role in European languages, but for example in language of South East Asia (including Chinese and Vietnamese), in languages of Africa and Native American languages. We shall not deal with tone.

Sounds are produced in the vocal tract. Air is flowing through the mouth and nose and the characteristics of the sounds are manipulated by several so-called **articulators**. A rough picture is that the mouth aperture is changed by moving the jaw, and that the shape of the cavity can be manipulated by the tongue in many ways. The parts of the body that are involved in shaping the sound, the **articulators**, can be **active** (in which case they move) or **passive**. The articulators are as follows: **oral cavity, upper lip, lower lip, upper teeth, alveolar ridge** (the section of the mouth just behind the upper teeth stretching to the 'corner'),

Table 3: IPA consonant column labels

	Articulators involved
bilabial	the two lips, both active and passive
labiodental	active lower lip to passive upper teeth
dental	active tongue tip/blade to passive upper teeth
alveolar	active tongue tip/blade to passive front part of alveolar ridge
postalveolar	active tongue blade to passive behind alveolar
retroflex	active tongue tip raised or curled to passive postalveolar (difference between postalveolar and retroflex: blade vs. tip)
palatal	tongue blade/body to hard palate behind entire alveolar ridge
velar	active body of tongue to passive soft palate (sometimes to back of soft palate)
uvular	active body of tongue to passive (or active) uvula
pharyngeal	active body/root of tongue to passive pharynx
glottal	active body/root of tongue to passive pharynx

tongue tip, **tongue blade** (the flexible part of the tongue), **tongue body**, **tongue root**, **epiglottis** (the leaf-like appendage to the tongue in the pharynx), **pharynx** (the back vertical space of the vocal tract, between uvula and larynx), **hard palate** (upper part of the mouth just above the tongue body in normal position), **soft palate** or **velum** (the soft part of the mouth above the tongue, just behind the hard palate), **uvula** (the hanging part of the soft palate), and **larynx** (the part housing the vocal chords). For most articulators it is clear whether they can be active or passive, so we shall not concern ourselves with this definition.

It is evident that the **vocal chords** play a major role in sounds (they are responsible for the distinction between **voiced** and **unvoiced**), and the sides of the tongue are also used (in sounds known as **laterals**). Table 3 gives some definitions of phonetic features in terms of articulators for consonants. Column labels here refers to what defines the **place of articulation** as opposed to the **manner of articulation**. The **degree of constriction** is roughly the distance of the active articulator to the passive articulator. The degree of constriction plays less of a role in consonants,

Table 4: Constriction degrees for consonants

stop	active and passive articulators touch and hold-to-seal (permitting no flow of air out of the mouth)
trill	active articulator vibrates as air flows around it
tap/flap	active and passive articulators touch but don't hold (includes quick touch and fast sliding)
fricative	active and passive articulators form a small constriction, creating a narrow gap causing noise as air passes through it
approximant	active and passive articulators form a large constriction, allowing almost free flow of air through the vocal tract

though it does vary say between full contact [d] and ‘close encounter’ [z], and it certainly varies during the articulation (for example in affricates [dz] where the tongue retreats in a slower fashion than with [d]). The **manner of articulation** combines the degree of constriction together with the way it changes in time. Table 4 gives an overview of the main terms used in the IPA and Table 5 identifies the row labels of the IPA chart. Vowels differ from consonants in that there is no constriction of air flow. The notions of active and passive articulator apply, but here we find at least four degrees of constriction (**close**, **close–mid**, **open–mid** and **open**), corresponding to the height of the tongue body (plus degree of mouth aperture). There is a second dimension for the horizontal position of the tongue body. The combination of these two parameters is often given in the form of a two dimensional trapezoid, which shows with more accuracy the position of the tongue. There is a third dimension, which defines the rounding (**round** versus **unrounded**, which is usually not marked). We add a fourth dimension, **nasal** versus **nonnasal**, depending on whether the air flows partly through the nose.

Naming the Sounds

The way to name a sound is by stringing together the attributes that define it. However, there is a distinction between vowels and consonants. First we describe the names of consonants. For example, [p] is described as a voiceless, bilabial

Table 5: IPA consonant row labels

plosive	a pulmonic-egressive, oral stop
nasal	a pulmonic-egressive stop with a nasal flow; not a plosive, because not oral
fricative	a sound with fricative constriction degree; implies that airflow is central
lateral fricative	a fricative in which the airflow is lateral
approximant	a sound with approximant constriction degree; implies that the airflow is central
lateral approximant	an approximant in which the airflow is lateral

Table 6: IPA vowel row and column labels

close	compared with other vowels, overall height of tongue is greatest; tongue is closest to roof of mouth (also: high)
open	compared with other vowels, overall height of mouth is least; mouth is most open (also: low)
close-mid, open-mid	intermediate positions (also: mid / uppermid / lowermid)
front	compared with other vowels, tongue is overall forward
central	intermediate position
back	compared with other vowels, tongue is overall back (near pharynx)
rounded	lips are constricted inward and protruded forward

stop, [m] is called a (voiced) bilabial nasal. The rules are as follows:

(16) **voicing place manner**

Sometimes other features are added. If we want to describe [p^h] we say that it is a voiceless bilabial aspirated stop. The additional specification ‘aspirated’ is a manner attribute, so it is put after the place description (but before the attribute ‘stop’, since the latter is a noun). For example, the sequence ‘voiced retroflex fricative’ refers to [ʒ], as can be seen from the IPA chart.

Vowels on the other hand are always described as ‘vowels’, and all the other features are attributes. We have for example the description of [y] as ‘high front rounded vowel’. This shows that the sequence is

(17) **height place lip–attitude [nasality] vowel**

Nasality is optional. If nothing is said, the vowel is not nasal.

On Strict Transcription

Since IPA tries to symbolize a sound with precision, there is a tension between accuracy and usefulness. As we shall see later, the way a phoneme is realized changes from environment to environment. Some of these changes are so small that one needs a trained ear to even hear them. The question is whether we want the difference to show up in the notation. At first glance the answer would be negative. But two problems arise: (a) linguists sometimes *do* want to represent the difference and there should be a way to do that, and (b) a contrast that speakers of one language do not even hear might turn out to be distinctive and relevant in another. (An example is the difference between English [d] (alveolar) and a sound where the tongue is put between the teeth (dental). Some languages in India distinguish these sounds, though I hardly hear a difference.) Thus, on the one hand we need an alphabet that is highly flexible on the other we do not want to use it always in full glory. This motivates using various systems of notation, which differ mainly in accuracy. Table 7 gives you a list of English speech sounds and a phonetic symbol that is exact insofar that knowing the IPA would tell an English speaker exactly what sound is meant by what symbol. This is called **broad transcription**. The dangers of broad transcription are that a symbol like [p] does not reveal exact details of which sounds fall under it, it merely tells us

that we have a voiceless bilabial stop. Since French broad transcription might use the same symbol [p] for that we might be tempted to conclude that they are the same. But they are not.

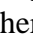
Thus in addition to broad transcription there exists strict or narrow transcription, which consists in adding more information (say, whether [p] is pronounced with aspiration or not). Clearly, the precision of the IPA is limited. Moreover, the more primitive symbols it has the harder it is to memorize. Therefore, IPA is based on a set of a hundred or so primitive symbols, and a number of diacritics by which the characteristics of the sound can be narrowed down.

Table 7: The Sounds of English

	Phonetic Symbol	Word illustrating it
1	p	<u>p</u> ope
2	b	<u>b</u> ar <u>b</u> er
3	m	<u>m</u> um
4	f	<u>f</u> ife
5	v	<u>v</u> ital, li <u>v</u> e
6	t	<u>t</u> aunt
7	d	<u>d</u> eed
8	n	<u>n</u> un
9	r	<u>r</u> are
10	θ	<u>th</u> ousand <u>th</u>
11	ð	<u>th</u> is, breathe
12	s	<u>s</u> ource, fu <u>ss</u>
13	z	<u>z</u> anies
14	ʃ	<u>sh</u> ush
15	ʒ	mea <u>ʒ</u> ure
16	l	<u>l</u> ul
17	tʃ	<u>ch</u> urch
18	dʒ	<u>judge</u>
19	j	<u>y</u> oke
20	k	<u>c</u> ook
21	g	<u>g</u> ag
22	ŋ	<u>sing</u> ing
23	w	<u>w</u> e
24	h	<u>h</u> e
25	i	<u>e</u> asy
26	ɪ	<u>i</u> mitate
27	e	<u>a</u> ble
28	ɛ	<u>e</u> dge
29	æ	<u>b</u> attle, att <u>a</u> ck
30	ɑ	<u>f</u> ather
31	ɔ	<u>f</u> ought
32	o	<u>r</u> oad
33	ʊ	<u>b</u> ook, sh <u>ou</u> ld
34	u	<u>f</u> ood
35	ə	<u>a</u> roma
36	ʌ	<u>b</u> ut
37	ɜː (or ɜ˞ or ɪ)	<u>b</u> ird
38	aɪ	<u>r</u> ide
39	aʊ	<u>h</u> ouse
40	ɔɪ	<u>b</u> oy

Phonology I: Features and Phonemes

Distinctiveness

There is a continuum of sounds but there is only a very limited set of distinctions that we look out for. It is the same with letters: although you can write them in many different ways (there are hundreds of different fonts for example, but whether you write the letter 'a' like this: a or like this: , it does not matter), most differences do not matter at all. Thus, some phonetic contrasts are relevant others are not. The question is: what do we mean by relevance? The answer is: if the contrast makes a difference in meaning it is relevant. The easiest test is to find to words that mean different things but differ only in one sound. These are called **minimal pairs**. Here are some examples of minimal pairs.

(18) hat : cat

(19) cat : cap

(20) cap : cup

(21) flight : fright

(22) flight : plight

We see from the first pair that the change from [h] to [k] may induces a change in meaning. Thus the contrast is relevant. Likewise, (19) shows that the contrast [p]:[t] is relevant (from which we deduce that the contrast labial:dental is relevant, though for other sounds it need not make a difference). Finally, (20) shows that the contrast [æ]:[ʊ] is relevant. Many of the contrasts between the 40 or so basic sounds of English can be demonstrated to be relevant by just choosing two words that are minimally different in that one has one sound and the other has the other sound. (Although this would require to establish $(40 \times 39)/2 = 780$ minimal pairs, one is usually content with far less.) Let us note also that in English certain sounds just do not exist. For example, retroflex consonants, lateral fricatives are not used at all by English speakers. Thus we may say that English uses only *some* of the available sounds, and other languages use others (there are languages that have retroflex consonants, for example many languages spoken in India). Additionally, the set of English sounds is divided into 40 groups, each corresponding to one letter in Table 7. These groups are called **phonemes** (and correspond to the 40 letters used in the broad transcription). The letter 'l' for example pretty much corresponds to a phoneme of English, which in turn is realized by many distinct

sounds. The IPA actually allows to represent the different sounds to some degree:

- (23) file ['faɪ̯] slight ['sl̩laɪ̯t] wealth ['weɪ̯lθ̩] listen ['lɪsən̩]
 fool ['fuɪ̯] flight ['fl̩laɪ̯t] health ['heɪ̯lθ̩] lose ['luːz]
 all ['aɪ̯] plow ['pl̩laʊ̯] filthy ['fɪ̯lθ̩i] allow [ə'laʊ̯]

The phoneme therefore contains the ‘sounds’ [ɪ̯], [l̩], [l̩] and [l]. (In fact, since the symbols are again only approximations, they are themselves not sounds but sets of sounds. But let’s ignore that point of detail here.) The following picture emerges. Utterances are strings of sounds, which the hearer (subconsciously) represents as sequences of phonemes:

- (24) sounds → σ_1 σ_2 σ_3 σ_4 ...
 phonemes → p_1 p_2 p_3 p_4 ...

The transition from sounds to phonemes is akin to the transition from narrow ((25)) to broad ((26)) transcription:

- (25) ['dð̩s ɪz ə fə'neɪ̯t̩k̩ t̩s̩h̩.ɛən'skɹɪp̩ʃ̩n̩]

- (26) /ð̩s ɪz ə fəʊnɛdɪk trənskɹɪpʃən/

- (27) this is a phonetic transcription

The conversion to phonemic representation means that a lot of information about the actual sound structure is lost, but what is lost is immaterial to the message itself.

We mention right away that the different sounds of a phoneme do not always occur in the same environment. If one sound σ can always be exchanged by σ' of the same phoneme, then σ and σ' are said to be in **free variation**. If however σ and σ' are not in free variation, we say that the realization of the phoneme as either σ or σ' is **conditioned by the context**.

Table 8 gives a list of the phonemes of American English. The slanted brackets denote phonemes not sounds, but the sounds are nevertheless given in IPA.

Some Concerns in Defining a Phoneme

So while a phone is just a sound, a concrete, linearly indecomposable sound (with the exception of certain diphthongs and affricates), a phoneme on the other hand

Table 8: Phonemes of English

Consonants

		bila- bial	labio- dental	den- tal	alve- olar	palato- alveolar	pala- tal	velar	glot- tal
stops	vl	/p/			/t/	/tʃ/		/k/	
	vd	/b/			/d/	/dʒ/		/g/	
fricatives	vl		/f/	/θ/	/s/	/ʃ/			/h/
	vd		/v/	/ð/	/z/	/ʒ/			
nasals		/m/			/n/			/ŋ/	
approximants	lat				/l/				
	cnt	/w/			/r/		/j/		

vl = voiceless, vd = voiced, lat = lateral, cnt = central

Vowels and Diphthongs

	front unrounded	central unrounded	back unrounded	back rounded	diphthongs
upper high	/i/			/u/	/aɪ/, /aʊ/,
lower high	/ɪ/			/ʊ/	/oɪ/
upper mid	/e/	/ə/		/o/	syllabic
lower mid	/ɛ/		/ʌ/		consonant
low	/æ/		/ɑ/		ɾ

is a set of sounds. Recall that in the book a phoneme is defined to be a basic speech sound. It is claimed, for example, that in Maasai [p], [b] and [β] are in complementary distribution. Nevertheless Maasai is said to have a phoneme /p/, whose feature specification is that of [p]. This means among other that it can only be pronounced as [p]. By contrast we define the following. Let $\hat{\ } \text{ denote concatenation.}$

Definition 4 (Provisional) *A phoneme is a set of phones (= speech sounds). In a language L , two sounds a and b belong to the same phoneme if and only if for all strings of sounds \vec{x} and \vec{y} : if both $\vec{x}\hat{a}\vec{y}$ and $\vec{x}\hat{b}\vec{y}$ belong to L , they have the same meaning. a and b are **allophones** if and only if they belong to the same phoneme.*

So, if a and b belong to the same phoneme, then either in a given word (or text) containing a one cannot substitute b for a , or one can but the result has the same meaning; and in a text containing b somewhere either one cannot substitute a for b or one can and the result has the same meaning. Take the sounds [t] and [ɾ] in (American) English (see Page 529). They are in complementary distribution, that is, in a context $\vec{x}_\vec{y}$ at most one of them can appear. So, we have ['dɛrə] but not ['dɛtə] (the context is 'de__ə); we have ['tæn] but not ['ɾæn] (the context is '__æn). (Notice that to pronounce data by ['dɛtə] or even ['dɛt^hə] is actually not illegitimate; this is the British pronunciation, and it is understood though not said. The meaning attributed to this string is just the same. The complications arising from the distinction between how something is pronounced correctly and how much variation is tolerated shall not be dealt with here.) On the other hand, if we change the position of the tongue slightly (producing, say [t̚] in place of [t]), the resulting string is judged to be the same. Hence it also means the same. We say that [t] and [t̚] are in **free variation**. So, two sounds are allophones if and only if they are (i) in complementary distribution or (ii) in free variation.

Definition 5 *If L is a language, and p a specific sound then $|p|_L$ denotes the phoneme containing p in L .*

The definition of a phoneme as *one of the basic speech sounds of a language* is different from ours. So it needs comment why we do it differently. First, it needs to be established what a basic speech sound is. For example, in Maasai [p] and [β] are in complementary distribution. By our definition, the sounds instantiating either [p] or [β] all belong to the same Maasai phoneme, which we denote by

$/p/_{\text{Maasai}}$. But is $/p/_{\text{Maasai}}$ a basic speech sound? How can we know? It seems that Fromkin et al. do not believe that it is. They take instead the phoneme to be [p], and assume that the context distorts the realization. Now look at English. The sound [p] is sometimes pronounced with aspiration and sometimes not. The two realizations of p, [p] and [p^h] do not belong to the same phoneme in Hindi. If this is the case it is difficult to support the idea that $/p/_{\text{English}}$ can be basic. If we look carefully at the definition above, it involves also the notion of meaning. Indeed, if we assume that a word, say car, has endlessly many realizations, the only way to tell that we produced something that is not a realization of car is to establish that it does not mean what a realization of car means. Part of the problem derives from the notation [p], which suggests that it is clear what we mean. But it is known that the more distinctions a language makes in some dimension, the narrower defined the basic speech sounds are. English, for example, has only two bilabial stops, which we may write [p] and [b]. Sanskrit (and many languages spoken in India today) had four: [p], [p^h], [b] and [b^h]. There is every reason to believe that the sounds that pass for a 'p' in English is narrower defined than in Sanskrit (or Hindi or Thai, which are similar in this respect). Thus, to be perfect we should write [p]_{English}, [p]_{Sanskrit} and so on. Indeed, the crucial parameter that distinguishes all these sounds, the Voice Onset Time (Page 660) is a *continuous parameter*. The distinction that is binary on the abstract level turns out to be based on a continuum which is sliced up in a somewhat arbitrary way. The discussion also has to do with the problem of narrow versus wide transcription. When we write [p] we mean something different for English than for Hindi, because it would be incorrect to transcribe for a Hindi speaker the sound that realizes p in paḷ by [p]; we should use [p^h] instead.

Features

By definition, any set of sounds can constitute a phoneme. However, it turns out that phonemes are constituted by classes of sounds that have certain properties in common. These are defined by **features**. Features are phonetic, and supposed to be not subject to cross-language variation. What exactly is a feature? The actual features found in the literature take a (more or less) articulatory standpoint. Take any sound realizing English b. It is produced by closing the lips, thereby obstructing the air flow ('bilabial') and then releasing it, and at the same time letting the vocal cords vibrate ('voiced'). If the vocal cords do not vibrate we get

the sound corresponding to p. We can analyse the sound as a motor program that is executed on demand. Its execution is not totally fixed, so variation is possible (as it occurs with all kinds of movements that we perform). Second, the motor program directs various parts of the vocal tracts, some of which are independent from each other. We may see this as a music score which has various parts for different 'instruments'. The score for the voicing feature is one of them. The value '+' tells us that the cords have to vibrate during the production of the corresponding sound, while '-' tells us that they should not. We have to be a bit cautious, though. It is known, for example, that b is not pronounced with immediate voicing. Rather, the voicing is delayed by a certain onset time. This onset time varies from language to language. Hence, the actual realization of a feature is different across languages, a fact that is rather awkward for the idea that phonemes are defined by recourse to *phonetic* features. The latter should namely be language independent. The problem just mentioned can of course be resolved by making finer distinctions with the features. But the question remains: just how much detail do the phonetic features need to give?

We shall use the following notation. There is a set of so-called **attributes** and a set of so-called **values**. A pair [: val] consisting of an attribute and a value is called a **feature**. We treat +voiced as a notational alternative of [: +]. An attribute is associated with a **value range**. For phonology, we may assume the following set of attributes:

(28) , , , , , , ...

and we may assume the following set of values:

(29) bilabial, labiodental, plosive, approximant, high, mid, +, -, ...

The range of is obviously different from that of , since 'dental' is a value of the former and not of the latter. A set of features is called an **attribute value structure (AVS)**. You have seen AVSs already in the first lecture. The notation is as follows. The attributes and values are arranged vertically, the rows just having the attribute paired with its value, separated by a colon:

(30)
$$\left[\begin{array}{l} : \text{dental} \\ : \text{fricative} \\ : + \end{array} \right]$$

Notice that the following are also legitimate AVSs:

$$(31) \quad \left[\begin{array}{l} : \text{dental} \\ : \text{dental} \\ : \text{fricative} \\ : + \end{array} \right] \quad \left[\begin{array}{l} : \text{dental} \\ : \text{uvular} \\ : + \end{array} \right]$$

The first is identical to (30) in the sense that it specifies the same object (the features are read conjunctively). The second however does not specify any sound, since the values given to the same feature are incompatible. (Features must have one and only one value.) We say the second AVS is **inconsistent**. Notice that AVSs are not sounds, they are just representations thereof, and they may specify the sounds only partly.

The following is however illegitimate because it gives a value to [d] that is outside of its value range.

$$(32) \quad \left[\begin{array}{l} : \text{fricative} \\ : + \end{array} \right]$$

On Page 584 – 585 you will find a number of arguments why features must exist. We add one very important reason here. Suppose that you have a rule

$$(33) \quad \left[\begin{array}{l} : + \\ : + \end{array} \right] \rightarrow \left[\begin{array}{l} : + \\ : - \end{array} \right] / _ \#$$

It says that a consonant becomes unvoiced at the end of a word. The part before the arrow specifies the situation before the rule applies; the part to the right and before the slash show us how it looks after the application of the rule. The part after the slash shows in what context the rule may be applied. The underscore shows where the left part of the rule must be situated (and where the right part will be substituted in its place). Here, it says: it must occur right before #, the end of the word.

Suppose that you have to say this without features. It is not enough to say that the voiced consonants are transformed into the voiceless ones; we need to know which voiceless consonant will replace which voiced consonant. The tie between [p] and [b], between [k] and [g] and so on needs to be established. Because features have an independent motivation the correspondence is specified uniformly for all sounds ('voice' refers to the fact whether or not the vocal cords vibrate). As

will be noted throughout this course, some rules are not really specific to one language but a whole group of them (final devoicing is a case in point). This seems to be contradictory, because the rules are stated using phonemes, and phonemes are language dependent, as we have seen. However, this need not be what is in fact going on. The contrast between voiced and voiceless may be present in a language, even though the exact specification of what counts, say, as a voiced bilabial stop as opposed to a voiceless bilabial stop, may be open to variation. Important is that the contrast exists and is one of voicing.

For example, Hungarian, Turkish and Finnish both have a rule called vowel harmony. Modulo some difficulties all rules agree that words cannot both contain a back vowel and a front vowel. On the other hand, the front close–mid rounded vowel of Finnish (written *ö*) is pronounced with more lip rounding than the Hungarian one (also written *ö*). Nevertheless, both languages systematically oppose *ö* with *o*, which differs in the position of the tongue body (close–mid back rounded vowel). The situation is complicated through the fact that Hungarian long and short vowels do not only contrast in length but also in a feature that is called **tension**. Finnish *ö* is tense even when short, while in Hungarian it is **lax** (which means less rounded and less close). However, even if short and long vowels behave in this way, and even if back and front vowels are different across these languages, there is good reason to believe that the contrast is between ‘front’ and ‘back’, no matter *what else is involved*. Thus, among the many parameters that define a sound language decide to systematically encode only a limited set (which is phonologically relevant and on which the rules operate) even though one still needs to fill in details as for the exact nature of the sounds. Precisely this is the task of **realization rules**. These are the rules that make the transition from phonemes to sounds. They will be discussed in the next lecture.

Binarism

There is a never concluded debate on the thesis of **binarism** of features. **Binarism** is the thesis that features have just two values: + and –. Although any feature system can be reconstructed using binary valued features, the two systems are not equivalent, since they define different natural classes. Suppose we fix the set of attributes and values for all languages in the same way. On the basis of this classification we can define the following.

Definition 6 *A natural class of sounds is a set of sounds that can be specified by a single AVS.*

Consider, by way of example, the sounds [p], [t] and [k]. They are distinct only in the place of articulation (bilabial versus alveolar versus velar). The only natural classes are: the empty one or the one containing all three. If we assume a division into binary features, either [p] and [t] or [p] and [k] or [t] and [k] must form a natural class in addition. This is so since binary features can cut a set only into two parts. If your set has three members, you can single out a given member by two cuts and only sometimes by one. So you need two binary features to distinguish the three from each other. But which ones do we take? In the present case we could take either [+labial], [+dental] or [+velar]. The choice between these three is not easy and hotly disputed. It depends on the way the rules of the language can be simplified which classification is used. But if that is so, the idea becomes problematic as a foundational tool. It is perhaps better not to enforce binarism.

In structuralism, the following distinction has been made: a distinction is **equipollent** or **privative**. To begin with the latter: the distinction between *a* and *b* is **privative** if (i) *a* has something that *b* does not have or (ii) *b* has something that *a* does not have. In case that (i) obtains, we call *a* **marked** (in opposition to *b*) and in case that (ii) obtains we call *b* **marked**. An equipollent distinction is one that is not of this kind. (So, neither *a* nor *b* can be said to be marked.) We have suggested above that the distinctions between speech sounds is always equipollent; for example, [p] and [b] are distinct because the one has the feature [–voiced] the other has the feature [+voiced]. We may however say the following: if in an AVS the feature [+voiced] is missing it will be realized as [p], that is, as if it were [–voiced]. So, by default, labials are voiceless. One also says that [b] is marked in contrast to [p].

Phonology II: Realization Rules and Representations

Determining Natural Classes

Let us start with a simple example to show what is meant by a natural class. Sanskrit had the following obstruents and nasals

(34)

p	p ^h	b	b ^h	m
t	t ^h	d	d ^h	n
ʈ	ʈ ^h	ɖ	ɖ ^h	ɳ
c	c ^h , ʃ	ɟ	ɟ ^h	ɲ
k	k ^h	g	g ^h	ŋ

(By the way, if you read the sounds as they appear here, this is exactly the way they are ordered in Sanskrit. The alphabet is much more logically arranged than the Latin alphabet!) To describe these sounds we use the following features and values:

(35) :bilab(ial), dental, retroflex, velar, palatal

(36) () :+, -

(37) () :+, -

(38) :+, -

Here is how the phonemes from the first row are to be represented:

(39)

p	p ^h	b
⌈ :bilab ⌋	⌈ :bilab ⌋	⌈ :bilab ⌋
⌈ :- ⌋	⌈ :+ ⌋	⌈ :- ⌋
⌈ :- ⌋	⌈ :- ⌋	⌈ :- ⌋
⌈ :- ⌋	⌈ :- ⌋	⌈ :+ ⌋
b ^h	m	
⌈ :bilab ⌋	⌈ :bilab ⌋	
⌈ :+ ⌋	⌈ :- ⌋	
⌈ :- ⌋	⌈ :+ ⌋	
⌈ :+ ⌋	⌈ :+ ⌋	

Let us establish the natural classes. First, each feature (a single pair of an attribute and its value) defines a natural class:

(40)	[: bilab]	{p, p ^h , b, b ^h , m}
	[: dental]	{t, t ^h , d, d ^h , n}
	[: retroflex]	{t, t ^h , d, d ^h , n}
	[: palatal]	{c, c ^c , ʃ, ʃ ^h , ɲ}
	[: velar]	{k, k ^h , g, g ^h , ŋ}
	[: +]	{p ^h , b ^h , t ^h , d ^h , t ^h , d ^h , c ^h , ʃ ^h , k ^h , g ^h }
	[: -]	{p, b, m, t, d, n, t, d, n, c, ʃ, ɲ, k, g, ŋ}
	[: +]	{m, n, ɲ, ɲ, ŋ}
	[: -]	{p, p ^h , b, b ^h , t, t ^h , d, d ^h , t, t ^h , d, d ^h , c, c ^h , ʃ, ʃ ^h , k, k ^h , g, g ^h }
	[: +]	{b, b ^b , m, d, d ^d , n, d, d ^h , n, ʃ, ʃ ^h , ɲ, g, g ^h , ŋ}
	[: -]	{p, p ^p , t, t ^h , t, t ^h , c, c ^h , k, k ^h }

All other classes are intersections of the ones above. For example, the class of phonemes that are both retroflex and voiced can be formed by looking up the class of retroflex phonemes, the class of voiced phonemes and then taking the intersection:

$$(41) \quad \{t, t^h, d, d^h, n\} \cap \{b, b^b, m, d, d^d, n, d, d^h, n, g, g^h, \eta\} = \{d, d^h, \eta\}$$

Basically, there are at most $6 \times 3 \times 3 \times 3 = 162$ (!) different natural classes. How did I get that number? For each attribute you can either give a value, or leave the value undecided. That gives 6 choices for place, 3 for nasality, 3 for voice, and three for aspiratedness. In fact, nasality does not go together with aspiratedness or with being voiceless, so some combinations do not exist. All the phonemes constitute a natural class of their own. This is so since the system is set up this way: each phoneme has a unique characteristic set of features. Obviously, things have to be this way, since the representation has to be able to represent each phoneme by itself. Now, if 162 classes strikes you as a large number, look at this. There are 25 phonemes, so there are $2^{25} = 33,554,432$ different sets of phonemes (if you cannot be bothered about the maths here, just believe me)! But how do we know that a set of phonemes is not natural. First method: try all possibilities. This might be a little slow, but you will soon find some shortcuts. Second method. You have to find a description that fits all and only the sounds in your set. It has to be of the form ‘has this feature, this feature and this feature’ — so no disjunction, no negation. You take two sounds and look at the attributes on which

they differ. Obviously, these ones you cannot use for the description. After you have established the set of attributes (and values) on which all agree, determine the set that is described by this combination. If it is your set, that set is natural. Otherwise not. Take the set {m, p^h, d}. The first is nasal, but the others are not. So the description cannot involve nasality. The second is voiceless, the others are voicedness. The description cannot involve voicing. Similarly for aspiratedness and place. It means that the smallest natural class that contains this set is — the entire set of them. (Yes, the entire set of sounds is a natural class. Why? Well, no condition is also a condition. Technically, it corresponds to the empty AVS, which is denoted by []. Nothing is in there, so any phoneme fits that description.)

Now, why are natural classes important and how do we use them? Let us look at a phenomenon of Sanskrit (and not only Sanskrit) called **sandhi**. Sanskrit words may end in the following of the above: p, m, t, n, ṭ, k, and ŋ. This consonant changes depending on the initial phoneme of the following word. Sometimes the initial phoneme also changes. An example is /tat ja ri:ram/, which becomes /tac c^hari:ram/. We shall concentrate here on the more common effect that the last phoneme changes. The books give you the following look-up table:

	word ends in:						
	k	ṭ	t	p	ŋ	n	m̐
p, p ^h	k	ṭ	t	p	ŋ	n	m̐
b, b ^h	g	ḍ	d	b	ŋ	n	m̐
t, t ^h	k	ṭ	t	p	ŋ	n	m̐
d, d ^h	g	ḍ	d	b	ŋ	n	m̐
ṭ, ṭ ^h	k	ṭ	ṭ	p	ŋ	m̐ʂ	m̐
ḍ, ḍ ^h	g	ḍ	ḍ	b	ŋ	ŋ	m̐
c, c ^h	k	ṭ	c	p	ŋ	m̐j	m̐
ʃ, ʃ ^h	g	ḍ	ʃ	b	ŋ	ŋ	m̐
k, k ^h	k	ṭ	t	p	ŋ	n	m̐
g, g ^h	g	ḍ	d	b	ŋ	n	m̐
n, m	ŋ	ŋ	n	m	ŋ	n	m̐

[ʂ] is a voiceless retroflex fricative, [j] is a voiceless palatal fricative. There is one symbol that needs explanation. m̐ denotes a nasalisation of the preceding vowel (thus it is not a phoneme in the strict sense — see below on a similar issue concerning English).

Suppose we want to write rules that capture Sandhi. The lookup table presents

individual rules of the form

$$(43) \quad k \rightarrow g / ___\#b$$

which says that /k/ changes to /g/ when preceding a word boundary followed by /b/. If you calculate, you will see that this table codes 154 rules (and I haven't given you the whole table). Obviously this is unsatisfactory. Using some tricks we can compress these rules somewhat. But we certainly want to see if there is a more systematic representation of this.

There is. Notice for example the behaviour of /k/, /t/ and /p/. If the consonant of the following word is voiced, they become voiced, too. If the consonant is voiceless, they remain voiceless. This can be encoded into a single rule. Observe that the last consonant is a stop; and the first consonant is a stop, too. Using our representations, we can capture the content of all of these rules as follows.

$$(44) \quad \left[\begin{array}{c} :- \\ :- \end{array} \right] \rightarrow \left[\begin{array}{c} :+ \\ :- \end{array} \right] / ___\# \left[\begin{array}{c} :+ \\ :- \end{array} \right]$$

Read this as follows: if the phoneme you look at has features that are not mentioned in the rule, they will be left as is. This is the way we achieve generality.

Next, if the consonant of the following word is a nasal, the preceding consonant becomes a nasal. The choice of the nasal is completely determined by the place of articulation, which is not changed.

$$(45) \quad \left[\begin{array}{c} :- \end{array} \right] \rightarrow \left[\begin{array}{c} :+ \\ :- \\ :+ \end{array} \right] / ___\# \left[\begin{array}{c} :+ \end{array} \right]$$

The reason we specified voicing and aspiratedness in the result is that we want the rule to apply to all obstruents. But if they are voiceless and we change only nasality, we get an voiceless nasal, which does not exist.

/ŋ/ never changes. /m/ becomes that mysterious m̥ . Only /n/ and /t/ show more action. Basically, they adapt to the place of articulation of the following consonant provided it is palatal or retroflex. (These are the next door neighbours, phonetically speaking.) However, if the following consonant is voiceless, the nasal changes to a sibilant, and the nasalisation is thrown onto the preceding vowel.

Neutralization of Contrast

There are also cases where the phonological rules actually obliterate a phonological contrast (on such case is stop nasalization in Korean). We discuss here a phenomenon called *final devoicing*. In Russian and German, stops become devoiced at the end of a syllable. It is such rules that cannot be formulated in the same way as above, namely as rules of specialization. This is so since they involve two sounds that are not allophones, for example [p] and [p̥] in Korean or [k] and [k̥] in German. We shall illustrate the German phenomenon, which actually is rather widespread. The contrast between voiced and voiceless is phonemic:

- (46) Kasse ['kasə] (*cashier*) : Gasse ['gasə]
 Daten ['da:tən] (*data*) : Taten ['ta:tən] (*deeds*)
 Peter ['pe:tə̃] (*Peter*) : Beter ['be:tə̃] (*praying person*)

Now look at the following words: Rad (*wheel*) and Rat (*advice*). They are pronounced alike: ['ʁa:t]. This is because at the end of the syllable (and so at the end of the word), voiced stops become unvoiced:

- (47) [+stop] → [+stop
-voiced] / ___#

(Here, # symbolizes the word boundary.) So how do we know that the sound that underlies Rad is [d] and not [t]? It is because when we form the genitive the [d] actually reappears:

- (48) (des) Rades ['ʁa:dəs]

The genitive of Rat on the other hand is pronounced with [t]:

- (49) (des) Rates ['ʁa:təs]

This is because the genitive adds an [s] (plus an often optional epenthetic schwa) and this schwa suddenly makes the [t] and [d] nonfinal, so that the rule of devoicing cannot apply.

Phonology: Deep and Surface

The fact that rules change representations has led linguists to posit two distinct sublevels. One is the level of **deep phonological representations** and the other

is that of **surface phonological representation**. The deep representation is more abstract and more regular. For German, it contains the information about voicing no matter what environment the consonant is in. The surface representation however contains the phonological description of the sounds that actually appear; so it will contain only voiced stops at the end of a syllable. The two representations are linked by rules that have the power of changing the representation. In terms of IPA-symbols, we may picture the change as follows.

$$(50) \quad \begin{array}{c} ['\text{ba:d}] \\ \downarrow \text{ (Final Devoicing)} \\ ['\text{ba:t}] \end{array}$$

However, what we should rather be thinking of is this:

$$(51) \quad \begin{array}{c} \left[\begin{array}{l} -\text{vowel} \\ +\text{approximant} \\ +\text{velar} \\ +\text{voiced} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{open} \\ +\text{front} \\ +\text{long} \end{array} \right] \left[\begin{array}{l} -\text{vowel} \\ +\text{stop} \\ +\text{labiodental} \\ +\text{voiced} \end{array} \right] \# \\ \downarrow \\ \left[\begin{array}{l} -\text{vowel} \\ +\text{approximant} \\ +\text{velar} \\ +\text{voiced} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{open} \\ +\text{front} \\ +\text{long} \end{array} \right] \left[\begin{array}{l} -\text{vowel} \\ +\text{stop} \\ +\text{labiodental} \\ -\text{voiced} \end{array} \right] \# \end{array}$$

(I mention in passing another option: we can deny that the voicing contrast at the end of a syllable is phonological — a voiced stop like [b] is just realized (= put into sound) in two different ways, depending on the environment. This means that the burden is on the phonology-to-phonetics mapping. However, the evidence seems to be that the syllable final [b] is pronounced just like [p], and so it simply is transformed into [p].)

We may in fact view all rules proposed above as rules that go from deep to surface phonological mapping. Some of the rules just *add* features while some of the *change* features. The view that emerges is that deep phonological structure contains the minimum specification necessary to be able to figure how the object sounds, while preserving the highest degree of abstraction and regularity. For example, strings are formed at deep phonological level by concatenation, while on the surface this might not be so. We have seen that effect with final devoicing.

While on the deep level the plural ending (or a suffix like *chen*) are simply added, the fact that a stop might find itself at the end of the word (or syllable) may make it change it to something else. The picture is thus this: the word *Rad* is stored as a sequence of three phonemes, and no word boundary exists because we might decide to add something. However, when we form the singular nominative, suddenly a word boundary gets added, and this is the moment the rule of final devoicing can take effect.

The setup is not without problems. Look at the English word *mouse*. Its plural is *mice* (the root vowel changes). How is this change accounted for? Is there a phonological rule that says that the root vowel is changed? (We consider the diphthong for simplicity to be a sequence of two vowels of which the second is relevant here.) The answer to the second question is negative. Not because such a rule could not be written, but because it would be incredibly special: it would say that the sequence [maʊss#] (with the second 's' coming from the plural!) is to be changed into [maɪs#]. Moreover, we expect that phonological rules can be grounded in the articulatory and perceptive quality of the sounds. There is nothing that suggests why the proposed change is motivated in terms of difficulty of pronunciation. We could equally well expect the form [maʊsəz], which is phonologically well-formed. It just is no plural of the word [maʊs] though English speakers are free to change that. (Indeed, English did possess more irregular plurals. The plural of [bʊk] was once [be:k], the plural of ['tunge] was ['tungan], and many more. These have been superseded by regular formations.) So, if it is not phonology that causes the change, something else must do that. One approach is to simply list the singular [maʊs] and the plural [maɪs], and no root form. Then [maɪs] is not analyzable. Another solution is to say that the root has *two* forms; in the singular we find [maʊs], in the plural [maɪs]. The plural has actually *no* exponent, like the singular. Plural as opposed to singular is signaled just by the vowel. The last approach has the advantage that it handles the change not in phonology.

There is — finally — a third approach. Here the quality of the second half of the diphthong is indeterminate between *ʊ* and *ɪ*. It is specified as 'lower high'. If you consult Table 8 you see that there are exactly *two* vowels that fit this description: *ʊ* and *ɪ*. So, it is a natural class. Hence we write the representation as

follows.

$$(52) \quad \begin{array}{c} \text{m} \qquad \text{a} \qquad \text{ʊ/ɪ} \qquad \text{s} \\ \left[\begin{array}{l} +\text{stop} \\ +\text{nasal} \\ +\text{bilab} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{low} \\ +\text{back} \\ -\text{rounded} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{lower high} \end{array} \right] \left[\begin{array}{l} +\text{fric} \\ +\text{alveol} \\ -\text{voiced} \end{array} \right] \end{array}$$

This representation leaves enough freedom to fill in either front or back so as to get the plural or the singular form. Notice however that it does not represent a phoneme. In early phonological theory one called these objects **archiphonemes**.

We shall briefly comment on this solution. First, whether or not one wants to use two stems or employ an underspecified sound is a matter of generality. The latter solution is only useful when it covers a number of different cases; in English we have at least *this:these*, *woman:women*, *foot:feet*. In German this change is actually much more widespread (we shall return to that phenomenon). So there is a basis for arguing that we find an archiphoneme here. Second, we still need to identify the representation of the plural. Obviously, the plural is not a sound in itself, it is something that makes another sound become one or another. For example, we can propose a notation $\leftarrow [+ \text{ front}]$ which says the following: go leftward and attach yourself to the first possible sound. So, the plural of mouse becomes represented as follows:

$$(53) \quad \begin{array}{c} \text{m} \qquad \text{a} \qquad \text{ʊ/ɪ} \qquad \text{s} \\ \left[\begin{array}{l} +\text{stop} \\ +\text{nasal} \\ +\text{bilab} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{low} \\ +\text{back} \\ -\text{rounded} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{lower high} \end{array} \right] \left[\begin{array}{l} +\text{fric} \\ +\text{alveol} \\ -\text{voiced} \end{array} \right] \leftarrow [+ \text{ front}] \end{array}$$

By convention, this is

$$(54) \quad \begin{array}{c} \text{m} \qquad \text{a} \qquad \text{ɪ} \qquad \text{s} \\ \left[\begin{array}{l} +\text{stop} \\ +\text{nasal} \\ +\text{bilab} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{low} \\ +\text{back} \\ -\text{rounded} \end{array} \right] \left[\begin{array}{l} +\text{vowel} \\ +\text{lower high} \\ +\text{front} \end{array} \right] \left[\begin{array}{l} +\text{fric} \\ +\text{alveol} \\ -\text{voiced} \end{array} \right] \end{array}$$

Note that the singular has the representation $\leftarrow [+ \text{ back}]$. You may have wondered about the fact that the so-called **root** of a noun like *cat* was nondistinct from its

singular form. This is actually not so. The root does not have a word boundary while the plural does (nothing can attach itself after the plural suffix). Moreover, there is nothing wrong with signs being empty.

This, however, leaves us with a more complex picture of a phonological representation. It contains not only items that define sounds but also funny symbols that act on sounds somewhere else in the representation.

Phonology III: Syllable Structure, Stress

Utterances are not mere strings of sounds. They are structured. An important unit is the *syllable*. Words consist of one or several syllables. Syllables in English begin typically with some consonants, then comes a vowel or a diphthong and then some consonants again. The first set of consonants is the **onset**, the group of vowels the **nucleus** and the second group of consonants the **coda**. The combination of nucleus and coda is called **rhyme**. So, syllables have the following structure:

(55) [onset [nucleus coda]]

For example, *strength* is a word consisting of a single syllable:

(56) [str [ε ŋθ]]

Thus, the onset consists of three consonants: [s], [t] and [r], the nucleus consists just of [ε], and the coda has [ŋ] and [θ].

The division into syllables is clearly felt by any speaker, although there sometimes is hesitation as to exactly how to divide a word into syllables. Consider the word *atmosphere*. Is the *s* part of the second syllable or part of the third? The answer is not straightforward. In particular the stridents (that is, the sounds [s], [ʃ]) enjoy a special status. Some claim that they are **extrasyllabic** (not part of any syllable at all), some maintain that they are **ambisyllabic** (they belong to *both* syllables). We shall not go into that here. The existence of rhymes can be attested by looking at verses (which also explains the terminology): words that rhyme do not need to end in the same syllable, they only need to end in the same rhyme: *fun – run – spun – shun*. Also, the coda is the domain of a rule that affects many languages: For example, in English and Hungarian, within the coda the obstruents must either all be voiced or unvoiced; in German and Russian, all obstruents in coda must be voiceless. Germanic verse in the Middle Ages used a rhyming technique where the onsets of the rhyming words had to be the same. (This is also called alliteration. It allowed to rhyme two words of the same stem; German had a lot of Umlaut and ablaut, that is to say, it had a lot of root vowel change making it impossible to use the same word to rhyme with itself (say *run – ran*.) It is worthwhile to remain with the notion of the **domain** of a rule. Many phonological constraints are seen as conditions that concern two adjacent sounds. When these sounds come into contact, they undergo change to a smaller or greater extent, for some sound combinations are better pronounceable than others. We have

Table 9: The Sonoricity Hierarchy

dark vowels [a], [o]	> mid vowels [æ], [œ]	> high vowels [i], [y]
> r-sounds [r]	> nasals; laterals [m], [n]; [l]	> vd. fricatives [z], [ʒ]
> vd. plosives [b], [d]	> vl. fricatives [s], [ʃ]	> vl. plosives [p], [t]

discussed sandhi at length in Lecture 4. For example, the Latin word *in* ‘in’ is a verbal prefix, which changes in writing (and therefore in pronunciation) to *im* when it precedes a labial (*impedire*). Somewhat more radical is the change from [ml] to [mpl] to avoid the awkward combination [ml] (the word *templum* derives from *temlom*, with *tem* being the root, meaning ‘to cut’). There is an influential theory in phonology, **autosegmental phonology**, which assumes that phonological features are organized on different scores (**tiers**) and can spread to adjacent segments independently from each other. Think for example of the feature [±voiced]. The condition on the coda in English is expressed by saying that the feature [±voiced] spreads along the coda. Clearly, we cannot allow the feature to spread indiscriminately, otherwise the total utterance is affected. Rather, the spreading is blocked by certain constituent boundaries; these can be the coda, onset, nucleus, rhyme, syllable, foot or the word. To put that on its head: the fact that features are blocked indicates that we are facing a constituent boundary. So, voicing harmony indicates that English has a coda.

The nucleus is the element that bears the stress. Typically it is a vowel, but this need not be so. Consider the standard pronunciation of *beaten*: [ˈbi:tʰn]. For my ears the division is into two syllables: [bi:] and [tʰn]. (In German this is certainly so the verb *retten* is pronounced [ˈʁɛtʰn]. The [n] must therefore occupy the nucleus of the second syllable.) In general, phonologists have posited the following conditions on syllable structure. Sounds are divided as follows. The sounds are aligned into a so-called **sonority hierarchy**, which is shown in Table 9 (vd. = voiced, vl. = voiceless). The syllable is organized as follows.

Syllable Structure

Within a syllable the sonority strictly increases and then decreases again.

This means that a syllable must contain at least one sound which is at least as sonorous as all the others in the syllable. It is called the **sonority peak** and is found in the nucleus. Thus, in the onset consonants must be organized such that the sonority rises, while in the coda it is the reverse. The conditions say nothing about the nucleus. In fact, some diphthongs are increasing ([iə] as in the British English pronunciation of *here*) others are decreasing ([aɪ], [oɪ]). This explains why the phonotactic conditions are opposite at the beginning of the syllable than at the end. You can end a syllable in [rt], but you cannot begin it that way. You can start a syllable by [tr], but you cannot end it that way (if you to make up words with [tr], automatically, [r] or even [tr] will be counted as part of the following syllable).

A moment's reflection shows why the position of stridents is problematic: the sequence [ts] is the only legitimate onset according to the sonority hierarchy, [st] is ruled out. Unfortunately, both are attested in English, with [ts] only occurring in non-native words (eg *tse-tse*, *tsunami*). There are phonologists who even believe that [s] is part of its own little structure here ('extrasyllabic').

Representing the Syllabification

The division of words into syllables is called **syllabification**. In written language the syllable boundaries are not marked, so words are not explicitly syllabified. We only see where words begin and end. The question is: does this hold also for the representations that we need to assume, for example, in the brain of a speaker? The answer is that we need almost no information about the syllabification. Moreover, the status of sounds at the boundary change as soon as new material comes in. Let us look at the plural *cats* of *cat*. The plural marker is *s*, and it is added at the end. Let's suppose the division into syllables was already given in the lexicon. Then we would have something like this: †*cat*†, where † marks the syllable boundary. Then the plural will be †*cat*†*s*†, with the plural 's' forming a syllable of its own. This is not the desired result, although the sonority hierarchy would predict exactly that. Let us look harder. We have seen that in German coda consonants devoice. The word *Rad* is pronounced [ˈʁa:t] as if written *Rat*. Suppose

we had added the syllable boundary: †Rad†. Now add the genitive es (which would also be †es†, by the way...). Then we get †Rat†es†, which by the rules of German would have to be pronounced [ˈʁa:t ʔəs], with an inserted glottal stop, because German (like English) does not have syllables beginning with a vowel (like French) and prevents that by inserting the glottal stop. In actual fact the pronunciation is [ˈʁa:dəs]. There are two indicators why the d is now at the beginning of the second syllable: (1) there is no glottal stop, (2) it is pronounced voiced.

We notice right away a consequence of this: syllables are *not* the same as morphemes. And morphemes neither necessarily are syllables or sequences thereof, nor do syllable boundaries constitute morpheme boundaries. Morphemes can be as small as a single phoneme (like the English plural), a phonemic feature (like the plural of mouse), they can just be a stress shift (nominalisation of the verb protest ([proˈtɛst] into protest [ˈprotɛst]) or they can even be phonemically zero. For example, in English, you can turn a noun into a verb (to google, to initial, to cable, to TA). The representation does not change at all neither in writing nor in speaking. You just verb it...

Of course it may be suggested that the syllabification *is* explicitly given and changed as more things come in. But this position unnecessarily complicates matters. It is known that syllabification is to a large extent predictable. So it is enough to just insert syllable boundary markers in the lexicon where absolutely necessary and leave the insertion of the other boundary markers to be determined later.

Syllable Structure

Languages differ in what types of syllables they allow. Thus, not only do they use different sounds they also restrict the possible combinations of these sounds in particular ways. Finnish allows (with few exceptions) only one consonant at the onset of a syllable. Moreover, Finnish words preferably end in a vowel, a nasal or 's'. Japanese syllables are preferably CV (consonant plus vowel). This has effects when these languages adopt new words. Finns for example call the East German car Trabant simply rabantti (with a vowel at the end and a long [t!]). The onset tr is simplified to just the r. There are also plenty of loanwords: koulu 'school' has lost the 's', rahti 'freight' has lost the 'f', and so on. Notice that it is always the last consonant that wins.

English too has constraints on the structure of syllables. Some are more strict than others. We notice right away that English does allow the onset to contain several consonants, similarly with the coda. However, some sequences are banned. Onsets are not to contain a nasal except in first place (exception: [sm] and [sn]). There are some loanwords that break the rule: *mnemonic*. The sequence sonorant–obstruent is also banned ([mp], [rk], [rp], [lt] and so on). Stridents are not found other than in first place; exceptions are [ts] and [ks]. The cluster [ps] is reduced to [s]. It also makes a difference whether the syllable constitutes a word or is peripheral to the word. Typically, inside a word syllables have to be simpler than at the boundary.

Syllabification is based on expectations concerning syllable structure that derive from the well–formedness conditions of syllables. However, these leave room. ['pleito] (Plato) can be syllabified ['plei.to] or ['pleit.o]. In both cases the syllables we get are well formed. However, if a choice exists, then preference is given to creating either **open syllables** (no coda) or **syllables with onset**. Another strategy that has been proposed is

Maximise Onset

Try to put as many consonants into the onset as possible.

For example, take the words *restless* and *restricted*. (From now on we denote the syllable boundary by a dot (.), as is common practice.) We do not find *re.stless* nor *res.tless*, since there is no onset of the form *stl* or *tl*. Hence, the only possibilities are *rest.less* or *restl.ess*. Here, we find that there is no coda of the form *stl*. Hence, no choice is possible. Now look at *restricted*. There are onsets of the form *str* (*strive*), so the principle of maximal onsets mandates that this be the only syllabification. Notice that *res.tric.ted* and *rest.ric.ted* are also possible since both coda and onset are legitimate. Indeed, maximal onsets work towards making the preceding syllable open, and to have syllables with onset.

In the ideal case all consonants are onset consonants, so syllables are open. Occasionally this strategy breaks down. For example, *suburb* is syllabified *sub.urb*. This reflects the composition of the word (from Latin *sub* ‘under’ and *urbs* ‘city’, so it means something like the lower city; cf. English *downtown* which has a different meaning!). One can speculate about this case.

Stress

Syllables are not the largest phonological unit. They are themselves organised into larger units. A group of two, sometimes three syllables is called a **foot**. A foot contains one syllable that is more prominent than the others. Feet are organized in higher units, where again one is more prominent than the others, and so on. Prominence is marked by **stress**. There are various ways to give prominence to a syllable. Ancient Greek is said to have marked prominence by pitch (the stressed syllable was about a fifth higher ($3/2$ of the frequency of an unstressed syllable)). Other languages (like German) use **loudness**. Other languages use combination of the two (Swedish). Within a given word there is one syllable that is the most prominent. In IPA it is marked by a preceding [ˈ]. We say that it carries **primary stress**. Languages differ with respect to the placement of primary stress. Finnish and Hungarian place the stress on the first syllable, French on the last. Latin put the stress on the last but one (**penultimate**), if it was long (that is to say, had a long vowel or was closed); otherwise, if was a syllable that preceded it (the **antepenultimate**) then that syllable got primary stress. Thus we had *pe.re.gri.nus* ('foreign') with stress on the penultimate (*gri*) since the vowel was long, but *in.fe.ri.or* with stress on the antepenultimate *fe* since the *i* in the penultimate was short. (Obviously, monosyllabic words had the stress on the last syllable.) Sanskrit was said to have free stress, that is to say, stress was free to fall anywhere in the word.

Typically, within a foot the syllables like to follow in a specific pattern. If the foot has two syllables, this is either an unstressed followed by a stressed syllable (**iambic metre**), or vice versa (**trochaic metre**). Sometimes a foot carries three syllables (a stressed followed by two unstressed ones, a **dactylus**). So, if the word has more than three syllables, there will be a syllable that is more prominent than its neighbours but not carrying main stress. You may try this with the word *antepenultimate*. You will find that the first syllable is more prominent than the second but less than the fourth. We say that it carries **secondary stress**: [ˌantɛpɛnˈultimɛt]. Or [əˌsɪməˈleɪʃn]. The so-called **metrical stress** theory tries to account for stress as follows. The syllables are each represented by a cross (×). This is a Layer 0 stress. Then, in a sequence of cycles, syllables get assigned more crosses. The more crosses, the heavier the weight of the syllable. The number of crosses is believed to correspond to the absolute weight of a syllable. So, a word that has a syllable of weight 3 (three crosses) is less prominent than one with a

Notice that the stress is governed by a number of heterogeneous factors. The first is the **weight** of the syllable; this decides about Layer 1 stress. Then there is the position of the main stress (which in English must to a large extent be learned — equivalently, it must explicitly be given in the representation). Third, it depends on the way in which the word is embedded into larger units (so syntactic criteria play a role here). Also, morphological formation rules can change the location of the main stress! For example, the suffix (a)tion attracts stress ([kɒm'beɪn] and [kɒmbɪ'neɪʃn]) so does the suffix ee (as in employee), but ment does not ([ˈgavəʃn] and [ˈgavəʃnment]). The suffix al does move the accent without attracting it ([ˈæŋkədɔt] versus [æŋkədɔtəl]).

Finally, we mention a problem concerning the representations that keeps coming up. It is said that a certain syllable cannot receive stress because it has a vowel that cannot be stressed (for example, schwa: [ə]). On the other hand, we can also say that a vowel is schwa because it is unstressed. Take, for example, the pair [ˈɪələʊz] and [ɪəliˈzeɪʃn]. When the stress shifts, the realisation of *i* changes. So, is it rather the stress that changes and makes the vowel change quality or does the vowel change and make the stress shift? Often, these problems find no satisfactory answer. In this particular example it seems that the stress shift is first, and it induces the vowel change. It is known that unstressed vowels undergo reduction in time. The reason why French stress is always on the last syllable is because it inherited the stress pattern from Latin, but the syllables following the stressed syllable got eventually lost. Historically, the sound was the same, they differed only in stress (these words come from medieval Latin together with their stress pattern; and they became partly Anglicized). However, due to English rules of pronunciation the stress pattern induced a change in vowel quality.

Phonology IV: Rules, Constraints and Optimality Theory

A Problem Concerning the Underlying Form

Let us look at the formation of the past tense.

	[t]		[d]		[əd]	
(62)	licked	[likt]	bugged	[bʌgd]	mended	[mɛndəd]
	squished	[skwɪʃt]	leaned	[liːnd]	parted	[pɑːrtəd]
	kept	[kɛpt]	buzzed	[bʌzɪd]	feasted	[fiːstəd]
	laughed	[lɑːft]	played	[pleɪd]	batted	[bætəd]

The past tense is either formed by adding a [t], or by adding a [d], or by adding [əd]. We ask: what is the source of this difference? Surely, it is possible to say that the past has three different forms and that depending on the verb a different form must be chosen. This, however, misses one important point, namely that the choice of the form is determined solely by the phonological form of the verb. The facts can be summarized as follows.

- ① [d] is found if the last sound of the verb is voiced but unequal to [d].
- ② [t] is found if the last sound of the verb is voiceless but unequal to [t].
- ③ [əd] is found if the verb ends in [d] or [t].

We mention here that it is required that the verb is *regular*. Thus, *run* and *catch* are of course not covered by this rule.

It seems that the choices can be accounted for solely by applying some general principles. First, notice that in a coda, with the exception of sonorants ([l], [m], [n]), all consonants agree in voicing:

Voice Agreement Principle

Adjacent obstruent sequences must either be both [+ voice] or both [–voice] at the end of a word.

(Notice that although the rule talks only about two obstruents at a time, it actually forces an entire sequence of obstruents to agree in the voice–feature. Second, the principle is less general than possible: this applies to all sequences contained in the same coda.) This goes half way in explaining the choice of suffix–form. We know why we see [d] after voiced consonants. But we do not know why it is that we get [lɪnd] rather than [lɪnt], because either of them is legitimate according to this principle. Furthermore, we do not know why we find the inserted schwa. The latter can be explained as follows: suppose there was no schwa. Then the present and the past forms would sound alike (mendd would be [mɛnd]). Languages try to avoid homonymy (although they never completely manage), and English employs the strategy to insert schwa also in the plural. We find *busses* (or *buses*), not *buss*.

We can explain the various surface forms as variants of one underlying form, which is formed by applying some rules.

Analysis 1. We assume that the underlying form is [d]. There is a rule that de-voices [d] right after a voiceless obstruent. There is a second rule which inserts a schwa right after [d]. For the purpose of the definition of the rules, two consonants are called **similar** if they differ at most in the voicing feature (for example, [t] is similar to both [t] and [d], but nothing else).

$$(63a) \quad \left[\begin{array}{c} +\text{voice} \\ +\text{obstruent} \end{array} \right] \rightarrow \left[\begin{array}{c} -\text{voice} \\ +\text{obstruent} \end{array} \right] \quad /[-\text{voice}]_ \#$$

$$(63b) \quad \emptyset \rightarrow [ə] \quad /C_C' \\ (C \text{ and } C' \text{ similar})$$

The symbol \emptyset denotes an empty string. The rule effectively says that it is legal to *insert* schwa anywhere between similar consonants. Notice that schwa insertion has to take place before the first rule. This gives three possible outcomes:

$$(64) \quad \begin{array}{l} \text{root} \quad /b_{\Delta}gd/ \quad /lɪkd/ \quad /mɛndd/ \quad /startd/ \\ (63b) \quad /b_{\Delta}gd/ \quad /lɪkd/ \quad /mɛndəd/ \quad /startəd/ \\ (63a) \quad /b_{\Delta}gd/ \quad /lɪkt/ \quad /mɛndəd/ \quad /startəd/ \end{array}$$

Notice that sometimes the rules do not apply. However, this means that the form is left *unchanged* by the rule. You may figure out for yourselves in which cases

dissimilar consonants.

$$(69a) \quad \begin{bmatrix} +\text{voice} \\ +\text{obstruent} \end{bmatrix} \rightarrow \begin{bmatrix} -\text{voice} \\ +\text{obstruent} \end{bmatrix} \quad /[-\text{voice}]_ \#$$

$$(69b) \quad [\text{ə}] \rightarrow \emptyset \quad /C_C' \\ (C \text{ and } C' \text{ dissimilar})$$

	root	/b Δ gəð/	/lɪkəð/	/mɛndəð/	/startəð/
(70)	(69b)	/b Δ gd/	/lɪkd/	/mɛndəð/	/startəð/
	(69a)	/b Δ gd/	/lɪkt/	/mɛndəð/	/startəð/

If we schedule (69a) before (69b), this would be the outcome:

	root	/b Δ gəð/	/lɪkəð/	/mɛndəð/	/startəð/
(71)	(69a)	/b Δ gəð/	/lɪkəð/	/mɛndəð/	/startəð/
	(69b)	/b Δ gd/	/lɪkt/	/mɛndəð/	/startət/

In principle, there are many more analyses. We can assume the underlying form to be anything we like (say, even /xx/). The principal difference between them is solely the extent to which the rules that transform them can be motivated language internally as well as language externally. And this is also the criterion that will make us choose one analysis over the others. Let's look carefully. First, let us go back to the voice agreement principle. It says only that adjacent obstruents agree in voicing, it does not claim that obstruents must agree with the preceding vowel, since we do actually find forms like [kæt]. Analysis 2 incorporates the wrong version of the Voice Agreement Principle. Rule (66a) repairs some of the forms without need, while Analysis 1 repairs the forms if and only if they do not conform to the Voice Agreement Principle. Now look at Analysis 3: it does not conflict with the Voice Agreement Principle. However, it proposes to eliminate schwa in certain forms such as [lɪked]. There is however no reason why this form is bad. So, it repairs forms that are actually totally well-formed.

Let us summarize: an analysis is preferred over another if it proposes laws of change that are widely attested (schwa insertion is one of them, and final devoicing is another). Also, an analysis is dispreferred if its rules change representations that are actually well-formed. Thus, rules of the kind discussed here are seen as *repair* strategies that explain why a form sometimes does not appear in the way expected. What we are looking at here is, by the way, the mapping from deep phonological form to surface phonological form. The last bit of evidence that makes us go for Analysis 1 is the following principle:

Not-Too-Similar Principle

No English word contains a sequence of subsequent similar obstruents.

Namely, Rule (63a) devoices an obstruent in coda if the preceding consonant is voiceless. And in that case, the Voice Agreement Principle is violated. After the rule has applied, the offending part is gone. Rule (63b) applies if there are two subsequent similar consonants, precisely when the Not-Too-Similar Principle is violated. After application of the rule the offending part is gone.

Which Repair for Which Problem?

The idea that we are pursuing is that deep-to-surface mappings institutionalize a repair of impossible situations. Thus, every rule is motivated from the fact that the input structure violates a constraint of the language and the output structure removes that offending part. Unfortunately, this is not all of the story. Look at the condition that onsets may not branch. This constraint exists in many languages, for example in Japanese and Finnish. But the two languages apply different repair strategies. While Japanese likes to insert vowels, Finnish likes to cut the onset to the last vowel. Both repair strategies are attested in other languages. However, we could imagine a language that simplifies an onset cluster to its first element: *trabant* becomes not *rabantti* but *tabantti* in Finnish, Germanic *hrengas* becomes not *rengas*, but *hengas* and so on. This latter strategy is however not attested. So, we find that among the infinite possibilities to avoid forbidden cluster combinations, only some get used at all, while others are completely disfavoured. Among those that are in principle available, certain languages choose one and not the other, but in other languages it is the opposite.

Some general strategies can actually be motivated to a certain degree. Look at schwa insertion as opposed to schwa deletion. While the insertion of a schwa is inevitably going to improve the structure (because languages all agree in that CV is a syllable...) the deletion of a schwa can in principle produce clusters that are illegitimate. Moreover, it is highly unlikely that deleting a schwa will make matters better. Thus, we would expect that there is a bias towards the repair by insertion of schwa. Yet, all these arguments have to be taken with care. For example, if a rule changes stress, this can be a motivating factor in reducing or eliminating a vowel.

Optimality Theory

Several ways to look at the regularities of language have been proposed:

- ☞ The **generative approach** proposes **representations** and **rules**. The rules shall generate all and only the legitimate representations of the language.
- ☞ The **descriptive** or **model–theoretic** approach proposes only **representations** and conditions that a representation has to satisfy in order to belong to a given language.


Note that generative linguists do not always propose that rules are real, that is, in the head of a speaker and that derivations take place in time. They would say that the rules are a way to systematize the data. If that is so, however, it is not clear why we should not adopt a purely descriptive account, characterizing all only the legal representations rather than pretending that they have been derived in a particular way. The more so since many arguments drawn in favour of a given analysis comes from data on child development and language learning. To interpret the data coming from learning we need to have a theory of the internal knowledge of language ('language faculty'). This knowledge may either consist in representations and conditions on the representations, or in fact in representations plus a set of rules.

The discussion concerning the problem whether we should have rules or not will probably go on forever. **Optimality Theory (OT)** adds a new turn to the issue. OT tries to do away with rules (though we shall see that this is an illusion). Also, rather than saying exactly which representations are legitimate it simply proposes a list a desiderata for an optimal result. If a result is not optimal, still it may be accepted if it is the best possible among its kin. Thus, to begin, OT must posit two levels: **underlying representation (UR)** and **surface representation (SR)**. We start with the UR [bætɔd] (battɔd). Which now is the SR? OT assumes that we generate all possible competitors and rank them according to how many and how often they violate a given constraint.


Here is how it can work in the present situation. We shall assume that the SR deviates in the least possible way from the UR. To that effect, we assign each segment a slot in a grid:

(72) b æ t d
 • • • •


We may subsequently insert or delete material and we may change the representations of the segments, but we shall track our segments through the derivation. In the present circumstances we shall require, for example, that they do not change order with respect to each other. (Although this sometimes happens, it is called metathesis.) Here is an example where a segment changes:



Here is an example where one segment is dropped:



And here is an example where one segment is added:



Now, we look at pairs $\langle I, O \rangle$ of representations. Consider the following constraints on such pairs.

Recover the Morpheme

At least one segment of any given morpheme must be preserved in the SR.

This principle is carefully phrased. If the underlying morpheme is empty, there need not be anything in the surface. But if it is underlyingly not empty, then we must see one of its segments.

Recover Obstruency

If an underlying segment is an obstruent, it must be an obstruent in the SR.

The last principle says that the segment that is underlyingly hosting an obstruent must first of all survive; it cannot be deleted. Second, the phoneme that it hosts must be an obstruent.

Recover Voicing

If a segment is voiced in UR, it must be voiced in SR.

Syllable Equivalence

The UR must contain the same number of syllables than the SR.

Recover Adjacency

Segments that are adjacent in the UR must be adjacent in the SR.

These principles are too restrictive in conjunction. The idea is that one does not have to satisfy all of them; but the more the better. An immediate idea is to do something like linear programming: for each constraint there is a certain penalty, which is 'awarded' on violation. For each violation, the corresponding penalty is added. (In school it's basically the same: bad behaviour is punished, depending on your behaviour you heap up more or less punishment.) Here however the violation of a more valuable constraint cannot be made up for. No matter how often someone else violates a lesser valued constraint, if you violate a higher valued constraint, you lose.

To make this more precise, for a pair $\langle I, O \rangle$ of underlying representation (I) and a surface representation (O), we take note of which principles are violated. Each language defines a partial linear order on the constraints, such as the ones given above. It says in effect, given two constraints C and C' , whether C is more valuable than C' , or C' is more valuable than C , or whether they are equally valuable.

The name of the game is now this. Given an UR I :

- ① Suppose that $\pi = \langle I, O \rangle$ and $\pi' = \langle I, O' \rangle$ are such that for all constraints that π violates there is a more valuable constraint that π' violates. Then O is **optimal with respect to O'** .

- ② Suppose that $\langle I, O \rangle$ and $\pi' = \langle I, O' \rangle$ are such that C is the most valuable constraint that π violates; and that C is also the most valuable constraint that π' violates. Then if π' violates C more often than π , O is **optimal with respect to O'** .
- ③ O is **optimal** if it is optimal with respect to every other pair with same UR.

So, given I , if we want to know which SR corresponds to it, we must find an O which is optimal. Notice that this O need not be unique. (OT uses the following talk. O and O' are called **candidates** and they get **ranked**. However, candidacy is always relative to the UR.)

Let's apply this for [bætd]. We rank the constraints NTS (Not-Too-Similar), RObs (Recover Obstruency), RMorph (Recover the Morpheme) and RAdj (Recover Adjacency) as follows:

- (76) NTS, RObs, RMorph > RAdj

The first three are ranked equal, but more than the fourth. Other principles are left out of consideration.

/bætd/	NTS	RObs	RMorph	RAdj
[bætd]	★			
[bænd]		★		
[bæt]			★	
[bætəd]				★

The forms (a), (b), (c) all violate constraints that are higher ranked than RAdj. (d) violates the latter only. Hence it is optimal among the four. (Notice that we have counted a violation of Recover Obstruency for (c), even though one obstruent was dropped. This will concern us below.) Note that the optimal candidate still violates some constraint.

We now turn to the form /sæpd/. Assume that the ranking is (with VAg = Voice Agreement, RVoice = Recover Voicing)

- (78) VAg, RMorph, RObs > RAdj > RVoice

This yields the following ranking among the candidates:

	/sæpd/	VAgr	RMorph	RObs	RAdj	RVoice
(79)	[sæpd]	★				
	[sæp]		★			
	[sæmp]			★		
	[sæpəd]				★	
	[sæpt]					★

Some Conceptual Problems with OT

First notice that OT has no rules but its constraints are not well-formedness conditions on representations either. They talk about the relation between an UR and an SR. They tell us in effect that certain repair strategies are better than others, something that well-formedness conditions do not do. This has consequences. Consider the form [bænd]. We could consider it to have been derived by changing [t] to [n]:

(80)	b	æ	t	d
	•	•	•	•
	↓	↓	↓	↓
	•	•	•	•
	b	æ	n	t

Or we could think of it as being derived by deleting [t] and inserting [n]:

(81)	b	æ	t	d
	•	•	•	•
	↓	↓		↓
	•	•	•	•
	b	æ	n	d

More fanciful derivations can be conceived. Which one is correct if that can be said at all? And how do we count violations of rules? The first one violates the principle that obstruents must be recoverable. The second does not; it does not even violate adjacency! Of course, we may forbid that obstruents be deleted. But note that languages do delete obstruents (Finnish does so to simplify onset clusters). Thus obstruents can be deleted, but we count that also as a violation of

the principle of recoverability of obstruents. Then the second pair violates that principle. Now again: what is the punishment associated with that pair? How does it get ranked? Maybe we want to say that of all possibilities takes the most favourable one for the candidate.

The problem is a very subtle one: how do we actually measure violation? The string [cat] — how many constraints must we violate how often to get it from [startd]? For example, we may drop [s], then [d] and change the place of articulation of the first sound to [+velar]. Or may we drop [s] and [d] in one step, and then change [t] to [k] — or we do all in one step. How many times have we violated **Rule C**? The idea of violating a constraint 3 times as opposed to once makes little sense unless we assume that we apply certain rules.

Morphology I: Basic Facts

Word Classes

Morphology is the study of the minimal meaningful units of language. More precisely, it studies the structure of words. For everything that is larger than a word is the domain of syntax. Thus within morphology one considers the structure of words. The first to notice is that words come in different classes. For example, there are verbs (to imagine) and there are nouns (a car), there are adverbs (slowly) and adjectives (red). Intuitively, one is inclined to divide them according to their meaning: verbs denote activities, nouns denote things adverbs denote ways of performing an activities and adjectives denote properties. However, language has its own mind. The noun trip denotes an activity, yet it is a noun. Thus, the semantic criterion is misleading. From a morphological point of view, the three are distinct in the following way. Verbs take the endings s, ed, and ing, nouns only take the ending s. Adjectives and adverbs on the other hand do not change. (They can be distinguished by other criteria, though.)

(82) We imagine.

(83) He imagines.

(84) We are imagining.

(85) He imagined.

This distinction is made solely on the basis on morphological change. The criterion is at times not so easy to use. There are verbs that are not formed in the way above. For example, the verb run has no form *runned. Still, we classify it as a verb. For example, the English nouns take a subset of the endings that the verb takes. The word veto is both a noun and a verb, but this analysis predicts that it is a verb. Therefore, more criteria must be used. One is that of taking a context and looking which words fit into it.

(86) The governor ____ the bill.

If you fill the gap by a word, it is certainly a verb (more exactly a transitive verb, one that takes a direct object). On the other hand, if it can fill the gap in the next example it is a noun:

(87) The ____ vetoed the bill.

When we say ‘fill the gap’ we do not mean however that what we get is meaningful sentence when we put in that word; we only mean that it is grammatically (= syntactically) well-formed. We can fill in *cat*, but that stretches our imagination a bit. When we fill in *democracy* we have to stretch it even further, and so on. Adjectives can fill the position between the determiner (*the*) and the noun:

(88) The ____ governor vetoed the bill.

Finally, adverbs (*slowly*, *surprisingly*) can fill the slot just before the main verb.

(89) The governor ____ vetoed the bill.

Another test for word classes in the combinability with affixes. (**Affixes** are parts that are not really words by themselves, but get glued onto words in some way.) Table 10 shows a few English affixes and lists the word classes to which it can be applied. We see that the list of affixes is heterogeneous, and that affixes do not always attach to all members of a class with equal ease (*anti-house*, for example, is yet to be found in English). Still, the test reveals a lot about the division into different word classes.

Morphological Formation

Words are formed from simpler words, using various processes. This makes it possible to create very large words. Those words or parts thereof that are not composed and must therefore be drawn from the lexicon is called **root**. Roots are ‘main’ words, those that carry meaning. (This is a somewhat hazy definition. It becomes clearer only through examples.) Affixes are not roots. Inflectional endings are also not roots. An example of a root is *cat*, which as have explained, is form identical with the singular. However, the latter also has a word boundary marker at the right and (so it looks more like (*cat#*, but this detail is often generously ignored). In other languages, roots are clearly distinct from every form you get to see on paper. Latin *deus* ‘god’ has two parts: the root *de*, and the nominative ending *us*. This can be clearly seen if we add the other forms as well: genitive *dei*, dative *deo*, accusative *deum*, and so on. However, dictionaries avoid using roots. Instead, you find the words by their citation form, which in Latin is the nominative singular. So, you find the root in the dictionary under *deus* not

Table 10: English Affixes and Word Classes

Affix	Attaches to	Forming	Examples
anti-	nouns	nouns	anti-matter, ant-aircraft
	adjectives	adjectives	anti-democratic
un-	adjective	adjectives	un-happy, un-lucky
	verbs	verbs	un-bridle, un-lock
re-	verbs	verbs	re-establish, re-assure
dis-	verbs	verbs	dis-enfranchise, dis-own
	adjectives	adjectives	dis-ingenuous, dis-honest
-ment	verbs	nouns	establish-ment, amaze-ment
-ize	nouns	verbs	burglar-ize
	adjective	verbs	steril-ize, Islamic-ize
-ism	nouns	nouns	Lenin-ism, gangster-ism
	adjectives	nouns	real-ism, American-ism
-ful	nouns	adjectives	care-ful, soul-ful
-ly	adjectives	adverbs	careful-ly, nice-ly
-er	adjectives	adjectives	nic-er, angry-er

under de. (Just an aside: verbs are cited in their infinitival form; this need not be so. Hungarian dictionaries often list them in their 3rd singular form. This is because the 3rd singular reveals more about the inflection than the infinitive. This saves memory!)

There are several distinct ways in which words get formed; moreover, languages differ greatly in the extent to which they make use of them. The most important ones are

- ① **compounding**: two words, neither an affix, become one by juxtaposition. Each of them is otherwise found independently. Examples are *goalkeeper*, *whistleblower* (verb + noun compound), *hotbed* (adjective + noun).
- ② **derivation**: only one of the parts is a word; the other is only found in combination, and it acts by changing the word class of the host. Examples are the affixes which we have discussed above (*anti*, *dis*, *ment*).
- ③ **inflection**: one part is an independent word, the other is not. It does however not change the category, it adds some detail to the category (inflection of verbs by person, number, tense ...).

Compounding

In English, compounds are disfavoured over genuine constructions involving several words. It has to be said, though, that the spelling does not really tell you whether you are dealing with one or two words. For example, although one writes *rest room*, the stress pattern sounds as if one is dealing with only one word. Fortunately, there are languages where the difference can be told immediately. German is such a language.

- (90) *Regierung – s – vorlage*
 government proposal
- (91) *Schwein – e – stall*
 pig sty
- (92) *Räd – er – werk*
 wheel work = mechanism

The compound *Regierungsentwurf* not only contains the two words *Regierung* ('government') and *Vorlage* ('proposal'), it also contains an 's' (called 'Fugen s' = 'gap s'). To make German worthwhile for students, what gets inserted is not always an 's' but sometimes 'e'; *Schweinestall* is composed from *Schwein* and *Stall*. *Räderwerk* is composed from *Rad* and *Werk*. *Schweine* and *Räder* sound exactly like the plural, while *Regierungs* is not like any case form of *Regierung*. In none of these cases can the compound be mistaken for a multiword construction.

The meaning of compounds often enough is not determined from the meaning of its parts. It is characteristic of compounds that they often juxtapose some words and leave it open as to what the whole means (take *money laundry* or *coin washer*, which are generally not places where you launder money or wash coins). This is more true of noun+noun compounds than of verb+noun / noun+verb compounds, though.

Derivation

English has a fair amount of derivational affixes. Table 10 shows some affixes. We said that derivation changes the category of the word; this is not necessarily so.

Thus it might be hard to distinguish derivation from inflection in that case. However, derivation is optional (while inflection is not), and can be iterated (inflection cannot be iterated). And, inflectional affixes are typically outside derivational affixes. To give an example: you can form *republican* from *republic*, then you can add *anti*: *antirepublican* and finally form the plural: *antirepublicans*. You could have added the 's' to *republic*, but then you could not go on with the derivation. There is no word **republicsan*. Whether or not a word is formed by derivation is not always clear. For example, is *reside* formed by affixation? Actually, it once was, but nowadays it is not. This is because we do not have a verb **side*. Thus, derivation may form words that initially are perceived as complex, but later lose their transparent structure. This maybe because they start to sound different or because the base form gets lost. Nobody would guess that the word *nest* once was a complex word **nizdo* (here the star means: this form is reconstructed), derived from the words **ni* ('down') and **sed* ('sit').

Inflection

To fit a word into a syntactic construction, it has to undergo some changes. In English, the verb has to get an 's' suffix if the subject is third person singular. The addition of the 's' does not change the category of the verb; it makes it more specific, however. Likewise, the addition of past tense. Adding inflection thus makes the word more specific in category, narrowing down the contexts in which it can occur. Inflection is not optional; you must choose an inflectional ending. In Latin, adjectives agree in gender, number and case with the noun they modify:

(93)	<i>discipul-us</i> student-nom.sg	<i>secund-us</i> second-masc.nom.sg
(94)	<i>discipul-orum</i> student-gen.pl	<i>secund-orum</i> second-masc.gen.pl
(95)	<i>puell-arum</i> girl-gen.pl	<i>secund-arum</i> second-fem.gen.pl
(96)	<i>poet-arum</i> poet-gen.pl	<i>secund-orum</i> second-masc.gen.pl

The last example was chosen on purpose: form identity is not required. It is actually true that the forms of adjectives resemble those of nouns. The word *poeta*

belongs a form class of nouns that are mostly feminine, this is why adjectives show this form class if agreeing with a feminine noun (this has historic reasons). But the form class contains some masculine nouns, and to agree with them adjectives show a different form class. The latter actually is identical to that which contains more masculine nouns. This also explains why we have not added gender specifications to the nouns; unlike adjectives, nouns cannot be decomposed into gender and a genderless root.

The morphological characteristic of inflection is that it is harder to identify an actual affix (morph).

This way of indicating structure is actually not helpful, because if we have several occurrences of the same word it does not tell us which occurrence is going together with which other. Luckily, there are good ways to represent the structure. Here are namely a few observations about constituents. Let us think of a sentence as a sequence of words, like this:

(102) **this villa costs a fortune**
 w_1 w_2 w_3 w_4 w_5

The constituents can be named as follows:

(103) $\{w_1w_2w_3w_4w_5, w_1w_2, w_1, w_2, w_3w_4w_5, w_3, w_4w_5, w_4, w_5\}$

Let $w_1w_2w_3 \dots w_n$ be a sequence of words constituting a sentence English. Then a constituent of that sentence has the form $w_iw_{i+1}w_{i+2} \dots w_j$, for example $w_2w_3, w_5w_6w_7$ but not $w_2w_4w_7$. It always involves a *continuous stretch of words*.

Continuity of Constituents

Constituents are continuous parts of the sentence.

Non-Crossing

Given two constituents that share a word, one must be completely inside the other.

Words are Constituents

Every occurrence of a word forms its own constituent.

Here is a useful terminology. A constituent C is an **immediate constituent** of another constituent D if C is properly contained in D , but there is no constituent D' such that C is properly contained in D' and D' is properly contained in D . Our sentence (97) has only two immediate subconstituents: **this villa** and **costs a fortune**. The latter has the immediate constituents **costs** and **a fortune**. It is not hard to see that it is enough to establish for each constituents its immediate subconstituents. Notice also that we can extend the notion of precedence to constituents. A constituent C **precedes** a constituent D if all words of C precede all words of D . So, **this villa** precedes **a fortune** because **this** precedes both **a** and **fortune** and **villa** precedes both **a** and **fortune**.

This coincidence opens the way to a few alternative representations. One is by enclosing constituents in brackets:

(104) [[**this**] [**villa**]] [[**costs**] [**a**] [**fortune**]]]]]




Figure 1: An Unlabelled Syntactic Tree

Typically, the brackets around single words are omitted, though. This gives the slightly more legible

(105) [[this villa] [[costs [a fortune]]]]

Another one is to draw a tree, with each constituent represented by a node, and drawing lines as given in Figure 1. Each node is connected by a line to its immediate subconstituents. These lines go down; a line going up is consequently from a constituent to the constituent that it is an immediate part of. So, 8 and 9 are the immediate constituents of 7, 6 and 7 are the immediate subconstituents of 3, and so on. It follows that 6, 7 and 8 are subconstituents of 3, which is to say that 3 consists of costs, a and fortune.

Definition 7 A *tree* is a pair $\langle T, < \rangle$ where T is a set and $<$ is a set of pairs of elements from T such that (a) if $x < y$ and $x < z$ then either $y < z$, or $y = z$ or $z < y$ and (b) there is an element r such that $x < r$ for all $x \neq r$ (r is called the *root*).

The tree in Figure 1 consists of the nodes 1, 2, 3, ..., 9 (we ignore the stuff below them; this concerns wordforms and is of no interest to syntax; we add these things for orientation only). The relation $<$ is as follows: $2 < 1$, $3 < 1$, $4 < 2$, $5 < 2$,

$6 < 1, 7 < 1, 8 < 1, 9 < 1; 4 < 2, 5 < 2; 6 < 3, 7 < 3, 8 < 3, 9 < 3; 8 < 7, 9 < 7;$
no other relations hold.

Notice also that the order in which elements follow each other is also reflected in the tree, simply by placing the earlier constituents to the left of the later ones.

Categories

A **context** is a pair of strings. We denote them usually as follows:

(106) \vec{x} ___ \vec{y}

An example of a context is

(107) this ___ a fortune

Given a sentence, every substring is uniquely identified by its context: the part that precedes it and the part that follows it. The missing part in (107) is *villa costs*. So, every substring can be denoted by its left context and its right context. Substituting a substring by another is extracting the context and then putting back the replacing string. Replacing *villa costs* by *misses* gives

(108) this misses a fortune

Not a good sentence, but grammatical. (Does this show, by the way that *villa costs* is a constituent ...?)

We agree to call arbitrary strings of English words **constituents** just in case they occur as constituents in some sentence. It may happen that a sequence occurs in one sentence as a constituent, but not in another. Here is an example. The string *you called* is a constituent in (109) but not in (110):

(109) This is the man you called.

(110) This is a song for you called "Sweet Georgia Brown".

To know why try to substitute *you called* by *you called* and *Mary liked*. So, first of all we need to talk about sequences that occur *as constituents*. If they do, we call the occurrence a *constituent occurrence*.

Definition 8 A *category* (of a given language) is a set Δ of constituents such that any constituent of a given member of Δ can be replaced by any other member of Δ in that context preserving grammaticality (and the occurrence will continue to be a constituent occurrence in the new sentence).

The definition of a category is actually similar to that of a phoneme; phonemes were defined to be substitution classes up to meaning preservation. Here we define the classes up to grammaticality (more or less, since we have the caveat about ‘constituent occurrences’ because we are dealing with the substitution of strings for strings, not just of an item for another item). We give an example. The intransitive verbs in the 3rd personal singular form a category:

(111) {falls, runs, talks, ...}

By this definition, however, *talks* and *talk* are not members of the same category, for in (112) we cannot replace *talk* by *talks*. The result is simply ungrammatical.

(112) Mary and Paul talk.

It seems that the number of categories of any given language must be enormous if not infinite. It is certainly true that the number of categories is large, but it has a highly regular structure which we shall unravel in part for English.

Now that we have defined the constituents, let us go back to our tree in Figure 1. Each of the nodes in that tree is a constituent, hence belongs to some category. Omitting some detail, the categories are given in Figure 2. (We call S a sentence, D a determiner, DP a determiner phrase, NP a noun phrase, VP a verb-phrase, V verb.)

Context Free Grammars

Now look at Figure 2. If the labelling is accurate, the following should follow: any sequence of a determiner phrase followed by a verb phrase is a sentence. Why is this so? Look at the tree. If DP is the category to which *this villa* belongs, and if VP is the category to which *costs a fortune* belongs, then we are entitled to




Figure 2: A Labelled Syntactic Tree

substitute *any* DP for *this villa*, and *any* VP for *costs a fortune*:

$$(113) \quad \left\{ \begin{array}{l} \text{this villa} \\ \text{a car} \\ \text{it} \\ \text{tomorrow's sunshine} \end{array} \right\} \left\{ \begin{array}{l} \text{costs a fortune} \\ \text{walks} \\ \text{catches the bus} \\ \text{sings praise to the lord} \end{array} \right\}$$

None of the sentences you get are ungrammatical, so this actually works. We state the fact that a DP followed by a VP forms a sentence in the following way:

$$(114) \quad S \rightarrow \text{DP VP}$$

We have seen earlier statements of the form ‘something on the left’ \rightarrow ‘something on the right’, with a slash after which some conditions to the context were added. The condition on context is absent; this is why the rule is called **context free**. Because it can be used no matter what the context is. However, notice one difference, namely that the thing on the left is always a single symbol, denoting a category, while the thing on the right is a sequence of symbols, which may each be either a category or a word. A **grammar** is a set of such rules, together with a special symbol, called **start symbol** (usually S). Consider by way of the example

the following grammar. The start symbol is S, the rules are

- (115a) $S \rightarrow DP \ VP$
 (115b) $DP \rightarrow D \ NP$
 (115c) $D \rightarrow a \mid \text{this}$
 (115d) $NP \rightarrow \text{villa} \mid \text{fortune}$
 (115e) $VP \rightarrow V \ DP$
 (115f) $V \rightarrow \text{costs}$

Here, the vertical stroke ‘|’ is a disjunction. It means ‘can be either the one or the other’. For example, the notation $D \rightarrow a \mid \text{this}$ is a shorthand for two rules: $D \rightarrow a$ and $D \rightarrow \text{this}$.

This grammar says that (97) is of category S. How does it do that? It says that **this** is a determiner (D; Rule (115c)), and **villa** is a noun phrase (NP; Rule (115d)). By Rule (115b) we know that the two together are a DP. Similarly, it tells us that **a fortune** is a DP, that **costs a fortune** is a VP. Finally, using Rule (115a) we get that the whole is an S. Now, the role of the start symbol is the following. The grammar is said to **generate** the category of its start symbol. This is called the **language** defined by that grammar. Ideally, a grammar for English should generate exactly those sentences that are proper English. In the following sequence, each line is obtained by the previous by doing one replacement according to the rules.

- (116) S
 DP VP
 D NP VP
 D NP V DP
 D NP V D NP
 this NP V D NP
 this villa V D NP
 this villa costs D NP
 this villa costs a NP
 this villa costs a fortune

Syntax II: Argument Structure

Lexical Categories

There are several tens of thousands of categories in a language, maybe even millions. Thus the number of rules that we have to write is far too large to be written one by one. Thus, while in phonology the desire for general rules could still be dismissed as a perhaps needless endeavor, here it becomes absolutely central. We shall put to use our notation of attribute value structures (AVSs). We start off with a few general purpose rules and then refine them as we go along.

First, words fall into roughly two handful of so-called **lexical** or **major categories**. The ones we shall be using are: **noun** (N), verb (V), adjective (A), **adverb** (Adv), **preposition** (P), **complementizer** (C), **determiner** (D), and **tense** (T). Not all classes have single words in them, but most of them do:

	N	car, house, storm, insight
	V	run, tick, hasten, crawl
	A	greedy, raw, shiny, cheerful
(117)	Adv	very, steadily, allegedly, down
	P	in, for, about, below
	C	that, which, because, while
	D	a, the, this, those

Our first attribute is $\langle \text{CAT} \rangle$, and it has the values just displayed (so far, they are N, V, A, Adv, P, C, D, T).

Subject and Object

The next distinction we want to make is that between a **word** and a **phrase**. We have made that distinction earlier, when we called certain words determiners and certain constituents DPs (= determiner phrases). The distinction is intimately connected with the notion of argument structure. The argument structure tells us for each word what kinds of constituents it combines with to form a phrase. Take the verbs *run* and *accuse*. The difference between the two is that the first is happy to combine with just one DP, say *the sailors* to form a sentence, while

the latter is not:

- (118) The sailors run.
 (119) *The sailors accuse.

To make the latter into the sentence, you need to supply two more DPs, one for the one who is accused and one for what he is accused of:

- (120) The sailors accuse the captain of treason.

We say, *run* takes **one argument**, *accuse* takes **three arguments**. It means that in order to get a grammatical sentence, *run* only needs to combine with one DP, *accuse* needs three.

As for verbs, they always have one argument in a sentence, and this is the **subject**. The subject is that argument which is in the nominative case. You can recognize it by the form of pronoun that you have to use. The subject is *she*, *he*, *it*, *I*, *you*, *they* and not *her*, *him*, *me*, *us* and *them*. (The form *you* is both nominative and accusative, so we cannot use it for our test.) If instead you have to use the accusative forms, you are looking at the (**direct**) **object**.

- (121) They run./*Them run.
 (122) They accuse the captain of treason.
 (123) *Them accuse the captain of treason.
 (124) The sailors accuse us of treason.
 (125) *The sailors accuse we of treason.

English puts the subject before the verb, and the object right after it. So, this actually is another diagnostic for subject and object. The third argument is *of treason*. This is formed by using a preposition (*of*), so it is called a PP (= preposition phrase). There are as many preposition phrases as there are prepositions. We shall deal with them a little later.

A verb is called **intransitive** if it has no direct object. Otherwise it is called **transitive**. Hence, *run* is intransitive, *accuse* is transitive. Verbs may also be both; *break* and *eat* can occur both with and without a direct object. Then it is both transitive and intransitive; we say that it is **used** transitively if there is a direct object (as in (126)), and intransitively otherwise (as in (127)).

- (126) The children are eating the cake.
 (127) The children are eating.

The distinction between transitive and intransitive tells us whether or not the verb needs a direct object as an argument.

We observe that intransitive verbs are distributionally equivalent with the combination of transitive verb+direct object:

$$(128) \quad \left\{ \begin{array}{l} \text{The child} \\ \text{Napoleon} \\ \text{My neighbour's dog} \end{array} \right\} \left\{ \begin{array}{l} \text{ran.} \\ \text{lost the battle of Waterloo.} \\ \text{ate the cake.} \\ \text{is beautiful.} \end{array} \right\}$$

This tells us two things: transitive verbs plus (direct) objects form a constituent, and this constituent can be replaced by an intransitive verb.

Linguists have therefore proposed the following scheme. Constituents have two attributes: a major category and a **projection level**. The levels are 0, 1, 2. They are also called **bar levels** (because they used to be denoted by overstrike bars). Other notation is: D^0 for determiner, 0 level; D^1 , \overline{D} or D' for determiner first projection; and D^2 , $\overline{\overline{D}}$ or D'' for determiner level 2, or determiner phrase. Words are uniformly assigned level 0, and phrases are level 2. So, what we used to call a determiner (D) is now a D^0 , and the representation of it and a determiner phrase (DP) is like this:

$$(129) \quad D^0 = \begin{bmatrix} :D \\ :0 \end{bmatrix} \quad DP = \begin{bmatrix} :D \\ :2 \end{bmatrix}$$

The following rules are proposed for English (to the right you find a more user-friendly notation *of the same rule*):

$$(130) \quad \begin{bmatrix} :V \\ :2 \end{bmatrix} \rightarrow \begin{bmatrix} :D \\ :2 \end{bmatrix} \left[\begin{bmatrix} :V \\ :1 \end{bmatrix} \right] \quad VP \rightarrow DP \quad V'$$

$$(131) \quad \begin{bmatrix} :V \\ :1 \end{bmatrix} \rightarrow \begin{bmatrix} :V \\ :0 \end{bmatrix} \quad V' \rightarrow V^0$$

$$(132) \quad \begin{bmatrix} :V \\ :1 \end{bmatrix} \rightarrow \begin{bmatrix} :V \\ :0 \end{bmatrix} \left[\begin{bmatrix} :D \\ :2 \end{bmatrix} \right] \quad V' \rightarrow V' \quad DP$$

The start symbol is momentarily VP. Now, the rule (130) says that a sentence (= VP) is something that begins with a determiner phrase and then has a level 1 V. (131) is for intransitive verbs, (132) is for transitive verbs. Thus, we further add

an attribute with values + and –, to prevent an intransitive from taking a direct object:

$$(133) \quad \begin{bmatrix} :V \\ :1 \end{bmatrix} \rightarrow \begin{bmatrix} :V \\ :0 \\ :- \end{bmatrix}$$

$$(134) \quad \begin{bmatrix} :V \\ :1 \end{bmatrix} \rightarrow \begin{bmatrix} :V \\ :0 \\ :+ \end{bmatrix} \begin{bmatrix} :D \\ :2 \end{bmatrix}$$

$$(135) \quad \begin{bmatrix} :V \\ :0 \\ :- \end{bmatrix} \rightarrow \text{sit} \mid \text{walk} \mid \text{talk} \mid \dots$$

$$(136) \quad \begin{bmatrix} :V \\ :0 \\ :+ \end{bmatrix} \rightarrow \text{take} \mid \text{see} \mid \text{eat} \mid \dots$$

Notice that we did not specify transitivity for V^1 . This is because we have just said that a V^1 is a transitive verb plus direct object or an intransitive verb. Thus we say that a V^1 is intransitive, and there is simply no transitive V^1 . (Cases like *They call him an idiot.* seem like involving a verb with two direct objects. This is not so, but to make our case would take us too far afield.) Thus, in the above rules with can simply drop mentioning transitivity for V^1 s.

We should mention here the fact that some verbs want sentences as arguments, not DPs. The sentence must be opened with a complementizer, usually *that*.

(137) John believes that the earth is flat.

(138) John knows that two plus two is four.

These sentences are also called objects. They can be replaced by *this* or *that*, showing that they are either subjects or objects. They follow the verb, so they must be objects. In the following examples they are subjects:

(139) That the earth is flat is believed by John.

(140) That the cat was eating is food annoyed the child.

Oblique Arguments and Adjuncts

Verbs can have other arguments besides the subject and the object, too. These are called **oblique**. Let us look at the verb *accuse* again. In addition to a direct object it also wants a PP expressing subject matter. This PP must be opened by the preposition *of*. There are verbs that require a PP with *on* (*count*), others with *about* (*think*), and so on. Let us add a new attribute whose value can be any of the prepositions of English. Then *accuse* will get the following syntactic category:

$$(141) \quad \begin{bmatrix} :V \\ :0 \\ :+ \\ :of \end{bmatrix}$$

This tells us that *accuse* wants a subject (because all verbs do), a direct object (because it is transitive) and a PP opened by *of* (because this is how we defined the meaning of $:of$). To get the positioning of the phrases right, we propose to add the following rules:

$$(142) \quad \begin{bmatrix} :V \\ :1 \\ :on \end{bmatrix} \rightarrow \begin{bmatrix} :V \\ :1 \\ :on \end{bmatrix} \left[\begin{bmatrix} :P \\ :2 \\ :on \end{bmatrix} \right]$$

$$(143) \quad \begin{bmatrix} :P \\ :2 \\ :on \end{bmatrix} \rightarrow \begin{bmatrix} :P \\ :1 \\ :on \end{bmatrix}$$

$$(144) \quad \begin{bmatrix} :P \\ :1 \\ :on \end{bmatrix} \rightarrow \begin{bmatrix} :P \\ :0 \\ :on \end{bmatrix} \left[\begin{bmatrix} :D \\ :2 \end{bmatrix} \right]$$

$$(145) \quad \begin{bmatrix} :P \\ :0 \\ :on \end{bmatrix} \rightarrow \text{on}$$

Notice that we have used the feature [$: on$] also for prepositions; this makes sure that the right preposition appears at last in the structure!

The rules say that the PP is to be found to right of the direct object, if there is

one. This is generally the case:

- (146) They found the gold in the river.
 (147) *They found in the river the gold.
 (148) The pilot flew the airplane to Alaska.
 (149) *The pilot flew to Alaska the airplane.

Indeed, we can observe that the above rules hold for arbitrary prepositional phrases, so we simply replace *on* by a placeholder, say α . The rules then look like this. We show the Rule (142) only, which becomes (150).

$$(150) \quad \left[\begin{array}{c} :V \\ :1 \\ :\alpha \end{array} \right] \rightarrow \left[\begin{array}{c} :V \\ :1 \\ :\alpha \end{array} \right] \left[\begin{array}{c} :P \\ :2 \\ :\alpha \end{array} \right]$$

The idea is that in this rule schema, α may be instantiated to any appropriate value. In this case the appropriate values are the English prepositions. If we choose α , every occurrence of α must be replaced by the same value. For example, the following is *not* a correct instance of (150):

$$(151) \quad \left[\begin{array}{c} :V \\ :1 \\ :about \end{array} \right] \rightarrow \left[\begin{array}{c} :V \\ :1 \\ :on \end{array} \right] \left[\begin{array}{c} :P \\ :2 \\ :about \end{array} \right]$$

Notice that the PP does not pass directly to P^0+DP , there is an intermediate level 1 projection. (The reason is not apparent from the data given so far, and some syntacticians dispute whether things are this way. However, for our purposes it makes the syntax more homogeneous.) Notice also that the level does not decrease. A verb plus PP has the same level as the verb itself, namely 1. Therefore the PP is called an adjunct. Adjuncts do not change the level of the projection, but direct arguments do.

X-Bar Syntax

The structure of PPs looks similar to that of transitive verb phrases, except that the subject is generally missing. This similarity exists across categories. We take as an example NPs. There exist not only nouns as such, but nouns too take arguments.

These arguments are mostly optional, which is to say that nouns can be used also without these arguments.

- (152) the counting of the horses
 (153) some president of the republic
 (154) a trip to the Philippines
 (155) this talk about the recent events

In English, nouns cannot take direct arguments. The verb count is transitive (count the horses), but the gerund wants the former object in a PP opened by of.

We can account for this in the same way as we did for verbs. We allow nouns to additionally have a feature [: of], or [: in], and so on. And if they do, they can (but need not) add a PP opened by the corresponding preposition.

$$(156) \begin{bmatrix} :N \\ :1 \\ :of \end{bmatrix} \rightarrow \begin{bmatrix} :N \\ :0 \\ :of \end{bmatrix} \left(\begin{bmatrix} :P \\ :2 \\ :of \end{bmatrix} \right)$$

$$(157) \begin{bmatrix} :N \\ :1 \\ :of \end{bmatrix} \rightarrow \begin{bmatrix} :N \\ :0 \\ :of \end{bmatrix}$$

These two rules are abbreviated as follows.

$$(158) \begin{bmatrix} :N \\ :1 \\ :of \end{bmatrix} \rightarrow \begin{bmatrix} :N \\ :0 \\ :of \end{bmatrix} \left(\left(\begin{bmatrix} :P \\ :2 \\ :of \end{bmatrix} \right) \right)$$

The brackets around an item say that it is optional. For completeness, let us note that we have the following rules.

$$(159) \begin{bmatrix} :N \\ :0 \\ :of \end{bmatrix} \rightarrow \text{counting} \mid \text{brother} \mid \text{election} \mid \text{chief} \mid \dots$$

$$(160) \begin{bmatrix} :N \\ :0 \\ :about \end{bmatrix} \rightarrow \text{doubt} \mid \text{question} \mid \text{talk} \mid \text{rumour} \mid \dots$$

The Global Structure of Sentences

The start symbol is CP, for all languages. Constituents that are CPs are often called **clauses**. Sentences have a global structure. In English, it looks like this:

(166) [XP [C⁰ [YP [T⁰ VP]]]]

This means the following. The leftmost element is the specifier of CP. Then follows the complementizer (C⁰), and then the complement of the complementizer. It in turn has the following structure. First comes the specifier of T(ense). Then the T⁰ and then the VP. We will see later why we need all these parts in the sentence. The largest CP, which contains the entire sentence is called the **matrix clause**.

Syntax III: Local Dependencies and Constraints: Selection, Agreement and Case Marking

Grammatical Features

Nouns come in two varieties: there are singular and plural nouns. Singular and plural are called **numbers**. We have already mentioned a few ways in which the singular and plural of nouns are formed. For syntax the different ways of forming the plural are not relevant; only the fact whether a noun is singular or plural is relevant. A way to represent the number of a noun is by adding an attribute $\text{DP}[\text{number}]$, with values *sing* and *pl* (in other languages there will be more...). So we have that the **syntactic representation** of 'mouse' and 'mice' is

$$(167) \quad \text{mouse} : \begin{bmatrix} :N \\ :0 \\ :sing \end{bmatrix}, \text{mice} : \begin{bmatrix} :N \\ :0 \\ :sing \end{bmatrix}$$

This we phrase for written English by using the following rules:

$$(168) \quad \begin{bmatrix} :N \\ :0 \\ :sing \end{bmatrix} \rightarrow \text{mouse}$$

$$(169) \quad \begin{bmatrix} :N \\ :0 \\ :pl \end{bmatrix} \rightarrow \text{mice}$$

An alternative notation, which is used elsewhere and should now be self-explanatory, is $\text{DP}[\text{sing}]$ and $\text{DP}[\text{pl}]$. (The attribute DP is omitted, because *sing* and *pl* are only values of that attribute not of any other). We shall use this type of notation without further warning. It is to be thought of as abbreviatory only.

Also noun phrases in English, like nouns, can be both singular and plural. Whether or not an NP is singular or plural can be seen by looking at the head noun.

(170) *fearful warrior* (singular)

(171) *fearful warriors* (plural)

(172) *brother of the shopkeeper* (singular)

(173) *brothers of the shopkeeper* (plural)

In the first example the head noun is *warrior*, in the second it is *warriors*. This indicates that in the first example the NP is singular, and plural in the second. Thus, the number feature is passed up, so to speak, to the phrase from the head. Notice that the determiner is different depending on whether the NP is singular or plural.

- (174) *this fearful warrior*
 (175) **these fearful warrior*
 (176) **this fearful warriors*
 (177) *these fearful warriors*

We say that as with nouns, determiners have a singular and a plural form and that the determiner **agrees** with the NP complement in number. This will be put into the rules as follows. (The letter α may be instantiated to any legitimate value, in this case *sing* or *pl*. Notice that within a single rule, each occurrence of α must be replaced by the same value.)

$$(178) \quad \begin{bmatrix} :D \\ :1 \\ :\alpha \end{bmatrix} \rightarrow \begin{bmatrix} :D \\ :0 \\ :\alpha \end{bmatrix} \left[\begin{bmatrix} :N \\ :2 \\ :\alpha \end{bmatrix} \right]$$

$$(179) \quad \begin{bmatrix} :D \\ :0 \\ :sing \end{bmatrix} \rightarrow \text{this} \mid \text{the} \mid \text{a} \mid \dots$$

$$(180) \quad \begin{bmatrix} :D \\ :0 \\ :pl \end{bmatrix} \rightarrow \text{these} \mid \text{the} \mid \emptyset \mid \dots$$

Here, \emptyset is the empty string. It is needed for indefinite plurals (the plural of a *car* is *cars*).

The choice of the determiner determines a feature of DPs that is sometimes syntactically relevant: definiteness. DPs are said to be **definite** if, roughly speaking, they refer to a specific entity given by the context or if they are uniquely described by the DP itself. Otherwise they are called **indefinite**. The determiner of definite DPs is for example *this*, *that* and *the*. The determiner of indefinite DPs is for example *a*, *some*.

In other languages, nouns have more features that are syntactically relevant. The most common ones are **case** and **gender**. Latin, for example, has three gen-

ders, called **masculine**, **feminine** and **neuter**. The noun *homo* ('man') is masculine, *luna* ('moon') is feminine, and *mare* ('sea') is neuter. The adjectives have different forms in each case (notice that the adjective likes to follow the noun, but it does not have to in Latin):

- | | | |
|-------|--------------------|------------|
| (181) | <i>homo ruber</i> | = red man |
| (182) | <i>luna rubra</i> | = red moon |
| (183) | <i>mare rubrum</i> | = red sea |

Nouns have many different declension classes (morphology deals with them, syntax does not), and there are rules of thumb as to what declension class nouns have which gender, but they can fail. Nouns ending in *a* are generally feminine, but there are exceptions (*nauta* 'the seafarer', *agricola* 'the farmer' are masculine). Similarly, adjectives have many declension paradigms, so the forms can vary. But for each adjective there are forms to go with a masculine noun, forms to go with a feminine noun, and forms to go with a neuter noun. We therefore say that the adjective **agrees** in gender with the noun. It is implemented by installing a new attribute `gender`, whose values are 'm(asculine)', 'f(eminine)' and 'n(euter)'. Latin nouns also have different **cases**. There are five cases in Latin: nominative (for subjects), accusative (for direct objects), dative, genitive and ablative. Just as verbs select PPs with a particular preposition in English they can also select a DP with a particular case. If it is accusative the verb is transitive; but it can be dative (*placere*+DP[dat] 'to please someone'), ablative (*frui*+DP[abl] 'to enjoy something'), and genitive (*meminisse*+DP[gen] 'to remember someone/something'). There exist verbs that take several DPs with various cases. For example, *inferre* 'to inflict' (with perfect *intuli*) wants both a direct object and a dative DP.

- (184) *Caesar Gallis bellum intulit.*
 Caesar Gauls-dat war-acc inflict.upon-perf
 Caesar inflicted war on the Gauls.

This is just like English *inflict*, that wants a direct object and a PP[on].

For us, a **grammatical feature** is anything that defines a syntactic category. You may think of it as the syntactic analogue of a phonemic feature. However, beware that elsewhere the usage is a little different. Grammatical category and projection level are typically discarded, so that grammatical feature refers more

to the kinds of things we have introduced in this section above: number, gender, definiteness, case. A fifth one is *person*. (This list is not exhaustive.)

The position of subject can be filled with DPs like *the mouse*, *a car*, but also by so-called **pronouns**. Pronouns are distinct from nouns in that they express little more than that they stand in for a DP. However, they have many different forms, depending on the grammatical feature. In addition they show a distinction in **person**. Across languages, there is a pretty universal system of three persons: 1st, 2nd, and 3rd. First person means: includes speaker. Second person means: includes hearer, and third person means: includes neither speaker nor hearer. So, *I* is first person (it includes me and no one else); it is also singular, because it refers to just one thing. The plural *we* is used when one refers to several people, including speaker. By contrast, *you* is used for individuals or groups including the hearer (there is no distinction between singular and plural). The third person pronouns distinguish also gender in the singular (*he*, *she*, *it*), but not in the plural (*they*). Moreover, as we explained above, pronouns distinguish nominative from accusative. Thus, there are more morphological distinctions in the pronominal system in English than there is in the ordinary DPs.

More on Case

Cases have many different functions. One function is to indicate the nature of the argument. A verb has a subject, and the case of the subject is referred to as **nominative**. The direct object has a case that is referred to as **accusative**. Some people believe that English has cases (because pronouns still reflect a distinction between subject and object: *she* : *her*, *he* : *him*, and so on). On the other hand, this is confined to the pronouns, and nouns show no distinction in case whatsoever. This is why we say that English has *no case*. Strictly speaking, it means only that there is no distinction in case. (One might say: there is one case and only one. This is useful. For example, there is a famous principle of syntactic theory which states that nouns need case. If there is no case, this principle fails.) Chinese is another example of a language that has no cases. These languages make no distinction between subject and object in form; nevertheless, one can tell the difference: the subject precedes the verb, and the object follows it (both in English and in Chinese).

Verbs can have more than two arguments, and many more adjuncts. To dis-

tinguish between them some kind of marking is needed. In English this is done by means of prepositions. For example, there often is an argument towards which the action is directed or for which it is performed (the ‘goal’) and it is given by a PP opened by *to* (= PP[*to*], for example *talk to someone*). The goal is also called **indirect object**. Latin has a case for this, the **dative**. There is from a global viewpoint not much of a difference whether the goal is encoded by a case or by a PP. Languages can choose which way to go.

Another important case is the **genitive**. It marks possession. English has basically two ways to mark possession (apart from obvious ones like *which belongs to*). One is the so-called **Anglo-Saxon genitive**, formed by adding an *'s* (*my neighbour's car*). The other is a PP opened by *of* (*the car of my neighbour*). The genitive is used a lot in English. Nouns that have arguments that are not PPs put them in the genitive:

- (185) *the election of the chancellor*
- (186) *the peak of the mountain*
- (187) *Napoleon's destruction of the city*

Notice that two of these nouns have been obtained from transitive verbs. The rule in English is that the noun can take any of the arguments that the verb used to take (though they are now optional). However, the subject and the object must now appear in the genitive. The PPs on the other hand are taken over as is. For example, *destroy* is transitive, so the noun *destruction* can take two genitives, one for the subject and one for the object. The verb *talk* takes a subject, an indirect object and subject matter, expressed by a PP headed by *about*. The latter two are inherited as is by the noun *talk*, while the subject is put in the genitive. Alternatively, it can be expressed by a PP[*by*].

- (188) *John talked to his boss about the recent layoffs.*
- (189) *John's talk to his boss about the recent layoffs*
- (190) *talk by John to his boss about the recent layoffs*

Subject–Verb Agreement

English displays a phenomenon called **subject–verb agreement**. This means that the form of the verb depends on the grammatical features of the subject. The

agreement system is very rudimentary; the only contrast that exists is that between singular 3rd and the rest:

(191) She runs.

(192) They run.

Notice that since the verb does agree in person with the subject it has to make a choice for DPs that are not pronominal. It turns out that the choice it makes is that ordinary DPs trigger 3rd agreement:

(193) The sailor runs.

This applies even when the DP actually refers to the speaker! So, agreement in person is (at least in English) not only a matter of what is actually talked about, but it is also a syntactic phenomenon. There are rules which have to be learned.

Other languages have more elaborate agreement systems. Let us look at Hungarian. The verbal root is *lát* 'to see'.

(194)	Én látok.	Mi látunk.
	I see	We see
	Te láatsz.	Ti látatok.
	You(sg) see	You(pl) see
	Ő lát.	Ők látuk.
	He/she/it sees	They see

Hungarian has no distinction whatsoever in gender (not even in the pronominal system; *ő* must be rendered by 'he', or 'she', or 'it', depending on what is talked about). However, it does distinguish whether the direct object is definite or not. Look at this (translation is actually word by word):

(195) Én látok egy madarat.

I see a bird

(196) Én látom a madarat.

I see the bird

The subject is the same in both sentences, but the object is indefinite in the first (a bird) and definite in the second (the bird). When the object is indefinite, the form

látok is used, to be glossed roughly as ‘I see’, while if the object is definite, then the form látom is used, to be glossed ‘I see it’. Hungarian additionally has a form to be used when subject is first person singular and the direct object is 2nd person:

(197) Én látlak.

(198) I see.you

(The dot is used to say that what is rendered in English by two words is expressed in Hungarian by just one.)

Who Agrees with Whom in What?

Agreement is pervasive in some languages, and absent in others. Chinese has no agreement whatsoever, English has next to none. The most common type of agreement is that of verbs with their subjects. Some languages even have the verb agree with the direct object (Hungarian, Mordvin (a language spoken in Russia but related to Hungarian), Potawatomi (an American Indian language)). Other languages have the verb agree in addition with the indirect object (Georgian). Agreement is typically in person and number, but often also in gender. Above we have seen that definiteness can also come into the picture. Adjectives sometimes agree with the nouns they modify (Latin, German, Finnish), sometimes not (Hungarian). There is no general pattern here. This is one of the things that one has to accept as it is.

Syntax IV: Movement and Non-Local Dependencies

Movement

We have learned that in English the transitive verb requires its direct object immediately to its right. This rule has a number of exceptions. The first sentence below displays a phenomenon known as **topicalisation**, the second is a simple question using a question word.

- (199) Air pilots Harry admires.
- (200) Which country have you visited?

We could of course give up the idea that the direct object is to the right of the verb, but the facts are quite complex. For example, no matter what kind of constituent the question word replaces (subject, object, indirect object and so on), it is at the first place even if it is not the subject.

- (201) Alice has visited Madrid in spring to learn Spanish.
- (202) What has Alice visited in spring to learn Spanish?
- (203) Who has visited Madrid in spring to learn Spanish?
- (204) When has Alice visited Madrid to learn Spanish?
- (205) Why has Alice visited Madrid in spring?

We see that the sentences involving question words differ from (201) in that the question word is in first place and the verb in second place. There is a way to arrive at a question in the following way. First, insert the question word where it ought to belong according to our previous rules. Next, take it out and put in first position. Now move the auxiliary (has) into second place:

- (206) Alice has visited what in spring to learn Spanish?
- (207) What Alice has visited__in spring to learn Spanish?
- (208) What has Alice__visited__to learn Spanish?

(The underscore just helps you to see where the word came from. It is typically neither visible nor audible.) The good side about this proposal is that it is actually simple. Very little needs to be done to save the original approach. However, now we have a two-stage approach to syntactic structure: first we generate representations with a context free grammar and then we mix them up using certain rules

that move constituents. But in principle this is what we have done with phonological representations, too. First we have generated deep representations and then we have changed them according to certain rules. Thus, we say that the context free grammar generates **deep syntactic representations**, but that the rules just considered operate on them to give a final output, the **surface syntactic representation**. The rules are also referred to as **(syntactic) transformations**.

Wh–Movement

Let us investigate the properties of the so-called **Wh–Movement**. This is the transformation which is responsible to put the question word in front of the sentence. Question words are also referred to as **wh–words**, since they all start with **wh** (*who, what, where, why, etc.*). At first blush one would think that syntactic transformations operate on strings; but this is not so. Suppose the original sentence was not (201) but

(209) Alice has visited which famous city in Mexico to wait
for her visa?

Then the output we expect on this account is (210). But it is ungrammatical. Instead, only (211) is grammatical.

(210) *Which has Alice__visited__famous city in Mexico to wait
for her visa?

(211) Which famous city in Mexico has Alice__visited__to wait
for her visa?

It is the entire DP that contains the question word that goes along with it. There is no way to define that on the basis of the string, instead it is defined on the basis of the tree. To see how, let us note that the sentence (201) has the following structure. (Some brackets have been omitted to enhance legibility.)

(212) Alice [has [[visited [which famous city in Mexico][to wait
for her visa]]]]?

So, *which famous city in Mexico* is a constituent. Moreover, it is the object of the verb *visited*. The word *which* is a determiner, and the smallest phrase that contains it is the one that has to move.

Wh-Movement I.

Only *phrases* can be moved by Wh-Movement. What moves is the least phrase containing a given wh-word. It moves to the beginning of a clause (= CP).

This specification is imprecise at various points. First, what happens if there are several wh-words? In English what happens is that only one of them moves and the others stay in place; the choice of the one to move is a bit delicate, so we shall not deal with that question here. In other languages (Rumanian, Bulgarian, Hungarian are examples) all of them move. Second, what happens if the wh-word finds itself inside a sentence that is inside another sentence? Let us take a look.

(213) Mary thinks you ought to see what city? (deep structure)

Here the wh-phrase moves to end of the higher sentence (and notice that something strange happens to the verb too):

(214) What city does Mary think you ought to see?

However, some verbs dislike being passed over. In that case the wh-phrase ducks under; it goes to the left end of the lower sentence.

(215) *What city does Mary wonder you have seen?

(216) Mary wonders what city you have seen.

So, let us add another qualification.

Wh-Movement II

The wh-phrase moves to the beginning of the leftmost phrase possible.

We shall see further below that this is not a good way of putting things, since it refers to linear order and not hierarchical structure.

Verb Second

Many languages display a phenomenon called **Verb Second** or **V2**. German is among them. Unlike English, the verb is not always in second place. Here is a

pair of sentences with word-to-word translation.

(217) Hans geht in die Oper.

Hans goes into the opera

(218) Der Lehrer ist erfreut, weil Hans in die Oper geht.

the teacher is pleased, because Hans into the opera goes

The main verb is *geht* ('goes'). In the first example it is in second place, in the second example it is at the end of the sentence. Notice that in the second example there is a CP which is opened by *weil* ('because'). It is called **subordinate**, because it does not display the same kind of order as a typical clause. Now one may suspect that the verb simply occupies a different place in subordinate clauses. However, if we look at an auxiliary plus a main verb, matters start to become more complex.

(219) Hans will in die Oper gehen.

Hans wants into the opera go

(220) Der Lehrer ist erfreut, weil Hans in die Oper gehen will.

the teacher is pleased, because Hans into the opera go wants

Only the auxiliary (*will*) is found in the second place in the main clause. More facts can be adduced to show that show the following. The verb is at the end of the clause in deep structure. In a subordinate clause it stays there. Otherwise it moves to second position.

Now, what exactly is 'second position'? It cannot be the second word in the sentence. In the next example it is in the fifth place.

(221) Der Frosch im Teich ist kein Prinz.

the frog in the pond is no prince

Obviously, it is not the second word, it is the second *constituent* of the sentence. Once again we find that the operation of movement is not described in terms of strings but in terms of the structure.

How Movement Works

We have to begin with a very important definition. Recall that a tree is a pair $\langle T, < \rangle$, where T is a set, the set of **nodes** and $<$ the relation 'is properly dominated

by'. x and y are said to be **comparable** if $x = y$ or $x < y$ or $x > y$. (Thus, x can be found from y by either following the lines upwards or following them downwards.)

Definition 9 Let $\langle T, < \rangle$ be a tree and $x, y \in T$ be nodes. x **c-commands** y if and only if x and y are not comparable and the mother of x dominates y .

This can be phrased differently as follows. x c-commands y if x is sister to z and y is below z (or identical to it). Thus, c-command is with any sister and its heirs. Recall the tree from Lecture 8, repeated here as Figure 3. Here is a complete list of which nodes c-command which other nodes:

	c-commands
1	–
2	3, 6, 7, 8, 9
3	2, 4, 5
4	5
5	4
6	7, 8, 9
7	6
8	9
9	8

The relation of c-command is inherited by the strings that correspond to the nodes. For example, *this villa* c-commands *costs a fortune* and its sub-constituents.

Now, movement is such that the constituent that is moved is moved to a place that (i) is empty, and (ii) c-commands its original place. To make this proposal work we assume that instead of lexical items we can also find \emptyset below a terminal node. Trees 4 and 5 describe a single movement step, for the movement that is otherwise denoted using strings by

(223) Simon likes which cat?

(224) Which cat Simon likes__?

We have also added the labellings. Note that they do not exactly conform to our previously defined X-bar grammar. We have simplified the real tree a little bit, omitting some nodes. What is important here is only the mechanics of movement.




Figure 3: A Syntactic Tree




Figure 4: Before Movement




Figure 5: After Movement

The constituent consisting of *which cat* (which contains the node dominating these words plus everything below it) is moved into a c-commanding position, and replace the node dominating \emptyset . In place of the earlier constituent a node dominating \emptyset is put instead. (You may also think of the two constituents changing place: one which just dominates \emptyset , the other dominating *which cat*.)

For the whole story to go through it is assumed that the X-bar grammar produces a number of empty nodes that can be used for elements to move into. (You do not move into nodes that are already occupied.) One such nodes is typically the specifier of CP. The positions that an element moves into must also match in label (category). For example, the C^0 -position (which we have not shown above) is also often empty, but a phrase cannot go there, otherwise the labels do not match. It is believed that the verb moves there in the relevant examples from (201). You can check that X-bar syntax places no restriction on the category of the specifier except that it must be a phrase. Thus, wh-phrases can go there. The solution has an advantage worth mentioning. CP has only one specifier, and therefore only one of the wh-phrases can go there. The others have to stay in place.

Now, why do we also need the condition of c-command? Well, it is conceivable so far that the constituent that moves puts itself into any position that is freely

available. For example, take this deep structure

(225)

which salesperson wonders he should promote which product

The surface structure is

(226)

Which salesperson wonders which product he should
promote?

This shows first of all that the position right in front of *wonders* has an empty slot in it (we know it already: specifier of CP). But why does *which salesperson* not move there? That would give

(227)

*Wonders which salesperson he should promote
which product?

As we have said, if the CP is filled, the other stays in place, so this should be grammatical but it is not. The reason it is not is that movement has to be upwards, and into c-commanding position.

Syntax V: Binding

Pronouns

In this chapter we shall look at a phenomenon that is most intimately connected with c-command, namely **binding**. Binding is, as we shall see, as much of a semantic phenomenon as a syntactic one, but we shall ignore its semantic aspects as much as we can. What is at issue is that there are several different kinds of DPs: ordinary DPs, **names** and **pronouns**. Pronouns can be either **reflexive** (like *myself*, *yourself* etc.) or not (*he*, *she*, etc.). In English, the reflexive pronouns exist only in the accusative. There is no **hissself*, for example. (The pronoun *her* can both be genitive and accusative, so *herself* is not a good text case. On the other hand, it is *ourselves* and not **usselves*. English does not do us the favour of being completely consistent.) There additionally are **demonstratives** (*this*, *that*), **relative pronouns** (*who*, *what*, *which*), which are used to open a relative clause.

(228) I could not find the CD [which you told me about].

(229) There is a new film [in which no one ever talks].

The enclosed constituents are called **relative clauses**. They are used to modify nouns, for example. In that they are like adjectives, but they follow the noun in English rather than preceding it.

(230) *I could not find the which you told me about CDs.

(231) *There is a new in which no one ever talks film.

The relative clause is opened by a relative pronoun, *which*, together with the preposition that takes the pronoun as its complement. As with questions we like to think of the relative clauses as being derived from a structure in which the relative pronoun is in the place where the verb expects it.

(232) the CD [\emptyset you told me about which]

(233) a new film [\emptyset no one ever talks in which]

The position of specifier of CP is vacant, and the relative pronoun wants to move there. (The position of C is also empty, but we like to think that there is a silent C that sits there. Anyway, phrases can never go there.) Sometimes the relative

pronoun goes alone (being a DP, hence a phrase, it can do that), sometimes it drags the P along. The latter is known as **Pied-Piping**, from the fairy tale of the piper who promised to empty the city of the rats ...

Different languages have different systems. Latin does not distinguish reflexive and irreflexive pronouns in the 1st and 2nd person. It has however two 3rd pronouns, *is* ('he'), and *se* ('himself'). The reflexive exists in all cases but the nominative. It is the same both in the singular and the plural.

(234)	nom	–
	gen	sui
	dat	sibi
	acc	se
	abl	se

There is another set of pronouns called *possessive*. They are used to denote possession. They are just like adjectives. For example,

(235) equus suus

(236) horse-nom his-nom

(237) equum suum

(238) horse-acc his-acc

Binding

These different kinds of expressions each have their distinct behaviour. Look at the following three sentences.

(239) John votes for John in the election.

(240) John votes for himself in the election.

(241) John votes for him in the election.

There is an understanding that the two occurrences of the word *John* in (239) point to two different people. If they do not, we have to use (240) instead. Moreover, if we use (240) there is no hesitation as to the fact that there is an individual named *John* which casts a vote for the same individual named *John*. If we use (241), finally, we simply cannot mean the same person, *John*, by the word *him*. *John* cannot be self-voting in (241), he votes for someone else.

There are several theories that one may try to understand the distribution of full DPs, pronouns and reflexives. First of all, however, let us notice that a reflexive pronoun expects the thing it refers to to be the same as something else in the sentence. The expression that denotes this is called the **antecedent** of the pronoun. In (240), for example, the antecedent of *himself* is *John*. To express that some constituent refers to the same object we give them little numbers, called **indices**, like this.

- (242) *John₁ votes for John₁ in the election.
 (243) John₁ votes for himself₁ in the election.
 (244) *John₁ votes for him₁ in the election.

We have already assigned grammaticality judgments. Needless to say, any other number (112, 7, 56 or 34) would have done equally well, so that (242) is the same as

- (245) John₃₄ votes for John₃₄ in the election.

as far as syntax is concerned. (The books distribute letters *i*, *j* and *k* in place of concrete numbers, but this not a good idea.) In (239) the two occurrences of *John* are supposed to point to the same individual. If they are supposed to point to different individuals, we write different numbers:

- (246) John₁ votes for John₂ in the election.

The numbers are devices to tell us whether some constituent names the same individual as some other constituent, or whether it names a different one.

Pronouns seem to encourage a difference between subject and object in (240), and similarly with names (241). However, things are tricky. In (247) the pronoun *his* can be taken to refer not only to someone different from *John*, but also to *John* himself. And similarly in (248). (Just an aside: the reflexives cannot be used in genitive, they only have accusative forms. This may be the reason why we do not find them here, but the theory we are going to outline here tries a different line of argumentation.)

- (247) His lack of knowledge worried John.
 (248) He looks at himself in the mirror.

So, it is not really the case that when we have a pronoun that it must be used to talk about a different person than the others in the sentence (= that it must have a different index).

The conditions that regulate the distribution of these expressions are as follows. First, we define the notion of binding.

Definition 10 *A constituent X binds another constituent Y if X c-commands Y and X and Y have the same index.*

Binding is an interesting mixture between semantical conditions (carrying the same index, hence talking about the same individual) and purely structural ones (c-command). A note of warning is in order. Constituents are strings, but talk here as if they are nodes in a tree. This confusion is harmless. What we mean to say is this: suppose that x and y are nodes and the corresponding constituents are X and Y . Then if x c-commands y and has the same index, then x binds y , and X binds Y . So, X binds Y if there are nodes x and y such that x binds y and X is the constituent of x and Y the constituent of y . Now we are ready to say what the conditions on indexing are.

Principle A.

A reflexive pronoun must be bound by a constituent of the same CP (or DP).

Principle B.

A pronoun must not be bound by a constituent inside the same CP.

Principle C.

A name must not be bound.

For example, (239) can be used successfully to talk about two individuals named John. So, the expression alone is not sufficient to rule out a sentence. But the rules do tell us sometimes what the possible meanings are.

To understand the role of c-command, we need to look at the structure of some sentences. The subject c-commands the object of the same verb, since the subject precedes the V' , which contains the object. However, the object does not c-command the subject. So the following is illegitimate no matter what indices we assign:

(249) *Himself voted for John.

The subject also c-commands all other arguments. This accounts for the correctness of the following.

(250) The queen was never quite sure about herself.

(251) The students were all talking to themselves.

Notice that the principles not only say that these sentences are fine, they also tell us about the assignment of indices. It claims that (252) is fine but (253) is not.

(252) The queen₁ was never quite sure about herself₁.

(253) *The queen₁ was never quite sure about herself₂.

This is because the reflexive (*herself*) must be bound inside the sentence; this means there must be some antecedent c-commands it. In (252) it is the queen, but in (253) there is no such constituent. Thus, the overall generalization seems to be good.

Now we look at pronouns. The situation where a pronoun should not be used in the same sentence is when actually a reflexive would be appropriate according to Principle A. For example, the following sentences are ruled out if meant to be talking about the same individual(s) in the same sentence (that is, if *her* refers to the queen and *them* refers to the students).

(254) *The queen₁ was never quite sure about her₁.

(255) *The students₁ were all talking to them₁.

By the same token, if we use numbers we can also write:

(256) *The queen₁ was never quite sure about herself₂.

(257) *The students₁ were all talking to themselves₂.

Here, different numbers mean that the expressions are meant to refer to different individuals or groups.

Notice that the contrast between a reflexive and a nonreflexive pronoun only matters when the antecedent is c-commanding the pronoun. We repeat (247) below:

(258) His_{1,2} lack of knowledge worried John₁.

Here, John is the antecedent. We can take the sentence to mean that John is worried about his own lack of knowledge, or that he is worried about someone else's lack of knowledge. In none of the cases would a reflexive pronoun be appropriate.

Let us now change the position of the two:

(259) *John's₁ lack of knowledge worried himself₁.

Here, John is not c-commanding the pronoun. But the reflexive must be bound by something that c-commands it inside the clause. This can only be John's lack of knowledge. But if we think that, we would have to say *itself* rather than *himself*. Next, why is (260) fine?

(260) John's₁ lack of knowledge worried him₁.

Even though it has the same index as John, it is not c-commanded, hence not bound by it. Having the same index nevertheless means that they refer to the same thing (John), but for syntax binding takes place only if c-command holds in addition. Hence, the next sentence is also fine for syntax:

(261) John's₁ lack of knowledge worried John₁.

Admittedly, we would prefer (262) over (261), but it is agreed that this is not a syntactic issue.

(262) His₁ lack of knowledge worried John₁.

Movement Again

We have argued earlier that an element can move only into a position that c-commands the earlier position. Now that we have binding theory, we can do the following. We declare that every time a constituent moves it leaves behind a silent element (called **trace**). The constituent in the new position and the trace will be coindexed. Now we add the following.

Condition on Traces.

All traces must be bound.

If we assume this much it follows that movement must inevitably be into a position that c-commands the original position. It does something else, too. It ensures that whatever thing we choose to interpret the moved element by, it will be used to interpret the trace.

(263) Air pilots₁ Harry admires *t*₁.

One may wonder why this is a good theory. It postulates empty elements (traces), so how can we be sure that they exist? We do not see them, we do not hear them, so we might as well assume that they do not exist. Furthermore, (263) is not the sentence we shall see in print or hear, so this actually adds structure that is seemingly superfluous.

Opinions on this diverge. What is however agreed is that empty elements are very useful. We have used them occasionally to save our neck. When nouns get transformed into verbs, no change is involved. Other languages are not so liberal, so it is not the nature of nouns and verbs that allows this. It is, we assume, an empty element which English has (and other languages do not) which can be attached to nouns to give a verb. At another occasion we have smuggled in empty complementizers. You will no doubt find more occasions on which we have made use of empty elements.

Binding And Agreement

Pronouns distinguish not only case and number but also gender in English. Crucially, when a pronoun becomes bound it must agree in number and gender with its binder (not in case though).

(264) John voted for himself.

(265) *John voted for herself.

(266) *John voted for themselves.

(267) The committee members voted for themselves.

(268) Mary voted for herself.

The fact that agreement holds between binder and pronoun means that certain indexations are not licit. Here is an example.

(269) *Her₁ lack of knowledge worried John₁.

(270) *Their disinterest in syntax₁ bothered John₁.

Also, we have said that a pronoun cannot be bound inside the same CP. But it can be bound outside of it:

(271) John₁ told his boss₂ that he_{1/2} looked good.

In the previous example, both indices are licit, since binding by John or by boss is compatible with agreement. However, in the following sentence one of the options is gone.

(272) John₁ told Mary₂ that she*_{1/2} looked good.

For if the pronoun is bound by John, it must agree with it in gender, which it does not. So it can only be bound by Mary.

Morphology II: Similarities and Dissimilarities to Syntax. Representational Issues

Kinds of Morphological Processes

In syntax words follow each other like pegs on a line. In morphology this is not always so. We shall review a few ways in which two morphemes can be composed. The general term for grammatical morphemes is **affix**. Generally, an affix is a string that puts itself somewhere in the string. By the terminology set up below, it is either a prefix, a suffix or an infix. Some writers use it in a more general sense, but we shall not. A morpheme is not always an affix, however. A morpheme need not be a piece (= string) that we add somewhere; it may be several pieces (transfix, circumfix), or simply a certain kind of change effected on the string in question. The general term for all of these is **morphological change**. We shall give a few examples of what kinds of morphological changes there are in the languages of the world.

Suffixes and Prefixes

A **suffix** is a string that is added at the end, a **prefix** is a string that is added at the beginning. English has mostly suffixes (derivational ones like *ation*, *ize*, *ee*; inflectional ones like *s* and *d*). But it also has prefixes: *de*, *re*, *un* are prefixes. It is generally agreed that — if we use an analogy with syntax here — the affix is the head, and that it either expects the string on its right (prefix) or on its left (suffix).

If there were only suffixes and prefixes, morphology would look like syntax. Unfortunately, this is not the case.

Circumfixes

A **circumfix** consists of a part that is being added before the word and another that is added thereafter. It is thus a combination of prefix and suffix. The German perfect is a circumfix. It consists in the prefix *ge* and a suffix which changes from

stem to stem (usually it is en or t).

	Infinitive	Root	Perfect
	sehen	seh	gesehen
(273)	back	back	gebacken
	filmen	film	gefilmt
	hausen	haus	gehaust

The two parts (the prefix and the suffix) are historically of different origin (the prefix *ge* did not exist originally, and in English you do not have it). In present day German, however, there is not sense in taking the two separate.

The superlative in Hungarian also is a circumfix.

(274)	nagy	‘great’	legnagyobb	‘greatest’
	fehér	‘white’	legfehérebb	‘whitest’

Here as in the German case, the circumfix has two identifiable parts. The suffix is found also in the comparative:

(275)	nagy	‘great’	nagyobb	‘greater’
	fehér	‘white’	fehérebb	‘whiter’

The same situation is found in Rumanian. We have *frumos* (‘beautiful’) in the positive, *mai frumos* (‘more beautiful’) in the comparative, and *cel mai frumos* (‘most beautiful’). However, whether or not the superlative can be seen as being formed from the comparative by adding a prefix depends on the possibility to decompose the meaning of the superlative in such a way that it is derived from the meaning of the comparative. To my knowledge this has not been proposed.

Infixes

Infixes insert themselves inside the string. Look at the following data from Chrau (a Vietnamese language).

(276)	võh	‘know’	vanõh	‘wise’
(277)	cãh	‘remember’	canãh	‘left over’

The string *an* is inserted after the first consonant! The string is cut into two pieces and the nominaliser inserts itself right there.

(278)	võh	→	v	+	õh	→	v	+	an	+	õh	→	vanõh
-------	-----	---	---	---	----	---	---	---	----	---	----	---	-------

Transfixes

A **transfix** is an even more complex entity. We give an example from Egyptian Arabic. Roots have three consonants, for example *ktb* ‘to write’ and *drs* ‘to study’. Words are formed by adding some material in front (prefixation), some material after (suffixation) and some material in between (infixation). Moreover, all these typically happen *at the same time*. Let’s look at the following list.

	[katab]	‘he wrote’	[daras]	‘he studied’
	[ʔamal]	‘he did’	[naʔal]	‘he copied’
	[baktib]	‘I write’	[badris]	‘I study’
	[baʔmil]	‘I do’	[banʔil]	‘I copy’
(279)	[iktib]	‘write!’	[idris]	‘study!’
	[iʔmil]	‘do!’	[inʔil]	‘copy!’
	[kaatib]	‘writer’	[daaris]	‘studier’
	[ʔaamil]	‘doer’	[naaʔil]	‘copier’
	[maktuub]	‘written’	[madruus]	‘studied’
	[maʔmuu]	‘done’	[manʔuul]	‘copied’

It is a puzzle to find out how the forms get arrived at. It is as you might find not always the same pattern. One thing however is already apparent: the vowels get changed from one form to the other. This explains why the vowel is not thought of as being part of the root.

Other Kinds of Changes

Reduplication is the phenomenon where a string is copied and added. For example, if *abc* is a string, then its (re)duplication is *abcabc*. Austronesian languages like to use reduplication or sometimes triplication (*abcabcabc*) for various purposes (to form the plural, to intensify a verb, to derive nouns and so on). The following is from Indonesian.

	orang	‘man’	orang-orang	‘child’
(280)	anak	‘man’	anak-anak	‘children’
	mata	‘eye’	mata-mata	‘spy’

Reduplication need not copy the full word. For example, in Latin some verbs form the perfect in the following way. The first consonant together with the next vowel

is duplicated and inserted:

(281)	pendit	'he hangs'	pependit	'he has hung'
	tendit	'he stretches'	tetendit	'he has stretched'
	currit	'he runs'	cucurrit	'he has run'
	spondet	'he promises'	spopondit	'he has promised'

The last example shows that the s is exempt from reduplication.

The Difference Between Morphology and Syntax

There is no way to predict whether some piece of meaning is expressed by a morpheme or by a separate lexeme or by both. There exist huge differences across languages. Some languages pack all kinds of meanings into the verb (Inuit, Mohawk), some keep everything separate (Chinese). Most languages lie in between. Morphology is more or less important. However, even within one language itself the means to express something change. Adjectives in English have three forms: **positive** (normal form), **comparative** (form of simple comparison) and **superlative** (form of absolute comparison). Now look at the way they get formed.

	positive	comparative	superlative
(282)	high	higher	highest
	fast	faster	fastest
	common	more common	most common
	good	better	best
	bad	worse	worst

The first set of adjectives take suffixes (zero for the positive, **er** for the comparative and **est** for the superlative). The second set of adjectives take a separate word, which is added in front (**more** and **most**). The third set is irregular. The adjective **good** changes the root in the comparative and superlative before adding the suffix (**bett** in the comparative and **b** in the superlative), while **bad** does not allow an analysis into root and suffix in the comparative and superlative. (We may define a suffix, but it will be a one-time-only suffix, since there is no other adjective like **bad**. It therefore makes not much sense to define a separate comparative suffix for **worse**. However, **worst** is a debatable case.)

It is not far fetched to subsume the difference between PPs and case marked DPs under this heading. Finnish has a case to express movement into a location, and a case for movement to a location:

- (283) Jussi menee taloon.
Jussi goes house-into = Jussi goes into the house.
- (284) Jussi menee talolle.
Jussi goes house-to = Jussi goes to the house.

When it comes to different concepts, however, like ‘under’ and ‘over’ Finnish runs out of cases and starts to use adpositions:

- (285) Jussi menee talon alle.
Jussi goes house-gen under = Jussi goes under the house.
- (286) Jussi menee talon yli.
Jussi goes house-gen over = Jussi goes over the house.

Hungarian has a case for being on top of, but the situation is quite analogous to Finnish.

Thus, the cutoff point between morphology and syntax is arbitrary. However, the two may show to behave differently in a given language so that the choice between morphological and syntactical means of expression has further consequences. We have seen, for example, that the comparative morpheme *er* is a suffix — so it is added after the word. However, the comparative lexeme *more* wants the adjective to its right. The latter is attributable to the general structure of English phrases. The complement is always to the right. Morphemes are exempt from this rule. They can be on the other side, and generally this is what happens. For example, verb+noun compounds in English are formed by placing the verb after the noun: *goalkeeper*, *eggwarmer*, *lifesaver* and so on. If these were two words, we should have *keeper goal*, *warmer egg*, and *saver life*. The reason why we do not get that is interesting in itself. English used to be a language where the verb follows the object (as is the case in German). It then changed into a language where the verb is to the left of the object. This change affected only the syntax, not the morphology. French forms compounds the other way around (*casse-noix* cracker-nut = ‘nutcracker’, *garde-voie* = guard-way = ‘gatekeeper’). This is because when French started to form compounds, verbs already preceded their objects.

The Latin Perfect — A Cabinet of Horrors

This section shall illustrate that the same meaning can be signaled by very different means; sometimes these are added independently. To start, there is a large group of verbs that form the perfect stem by adding *v*. (3rd person does not signal gender; we use ‘he’ instead of the longer ‘he/she/it’. The ending of the 3rd singular perfect is *it*. The ending in the present is *at*, *et* or *it*, depending on the verb.)

	<i>amat</i>	‘he loves’	<i>amavit</i>	‘he has loved’
(287)	<i>delet</i>	‘he destroys’	<i>delevit</i>	‘he has destroyed’
	<i>audit</i>	‘he hears’	<i>audivit</i>	‘he has heard’

We also find a group of verbs that form the perfect by adding *u*.

(288)	<i>vetit</i>	‘he forbids’	<i>vetuit</i>	‘he has forbidden’
	<i>habet</i>	‘he has’	<i>habuit</i>	‘he has had’

The difference is due to modern spelling. The two letters were originally not distinct and denoted the same sound, namely [u], which became a bilabial approximant in between vowels. The difference is thus accounted for by phonemic rules.

Other verbs add an *s*. The combination *gs* is written *x*.

	<i>regit</i>	‘he reigns’	<i>rexit</i>	‘he has reigned’
(289)	<i>augit</i>	‘he fosters’	<i>auxit</i>	‘he has fostered’
	<i>carpit</i>	‘he plucks’	<i>carpsit</i>	‘he has plucked’

There are verbs where the last consonant is lost before the *s*-suffix.

There are verbs where the perfect is signaled by lengthening the vowel (lengthening was not written, we add it here for illustration):

(290)	<i>iuvit</i>	‘he helps’	<i>iūvit</i>	‘he has helped’
	<i>lavat</i>	‘he washes’	<i>lāvit</i>	‘he has washed’

Sometimes this lengthening is induced by a loss of a nasal:

(291)	<i>rumpit</i>	‘he breaks’	<i>rūpit</i>	‘he has broken’
	<i>fundit</i>	‘he pours’	<i>fūdit</i>	‘he has poured’

Finally, there are verbs which signal perfect by reduplication:

- (292) *currit* 'he runs' *cucurr̄it* 'he has run'
 tendit 'he stretches' *tetendit* 'he has stretched'

Now, certain verbs use a combination of these. The verb *frangere* ('to break') uses loss of nasal, accompanied by vowel lengthening, plus ablaut:

- (293) *frangit* 'he breaks' *frēgit* 'he has broken'

The lengthening of the vowel could be attributed to the loss of the nasal (in which case it is called **compensatory lengthening**). However, the next verb shows that this need not be the case.

The verb *tangere* uses a combination of reduplication, loss of nasal and ablaut:

- (294) *tangit* 'he touches' *tetigit* 'he has touched'

The root is *tang*. The nasal is dropped, yielding *tag*. Reduplication yields *tetag* (actually, it should give **tatag*.) Finally, the vowel changes to give *tetig*. The change of vowel is quite common in the reduplicating verbs since the reduplication takes away the stress from the root. We have *'tangit* and *'tetigit*. The reason is that the root vowel is actually still short (so no compensatory lengthening). If it were long, we would have to pronounce it *te'tīgit*.

Semantics I: Basic Remarks on Representation

So far we have been concerned only with the *form* of language expressions, and not with their meaning. Ultimately, however, language is designed to allow us to communicate to each other how the world is like, to make each other do or believe something, and so on. The part of linguistics that deals mainly with the question of what is meant by saying something is called **semantics**. Meaning is not just some aspect of the form in which expressions are put by the language. If I tell you that Paul has pestered at least three squirrels this morning you can conclude that he has pestered at least two squirrels this morning. That follows not by virtue of the form the words **three** and **two** have, but in virtue of what these words mean. Likewise, if I get told that half of my students do not like syntax, and I have 180 students, then I conclude that 90 students do not like syntax. This reasoning works independently of the language in which the sentences are phrased. The same information can be conveyed in English, French, Swahili and so on. The form will be much different, but the message will be the same.

We introduce some bits of terminology. First, a **statement** is a sentence which can be said to be either true or false. In distinction to a statement, a question is not true or false; it is a request for information. Likewise, a command is a request on the part of the speaker that something be done. We shall be concerned here exclusively with statements.

Statements express **propositions**. We think of proposition as language independent. A sentence is a faithful translation of another sentence if both express the same proposition (if speaking about statements). So, the French sentence

(295) Marc a vu la Tour Eiffel.

is translated into English by

(296) Marc has seen the Eiffeltower.

just because they express the same proposition. A principal tool of semantics is to study the logical relations that hold between sentences in order to establish their meanings. Let us look at the following reasoning.

(297)
$$\begin{array}{l} \text{Peter gets married.} \\ \text{Sue gets married.} \\ \hline \therefore \text{Peter and Sue get married.} \end{array}$$

This seems to be true at first blush. However, there is a certain subtlety that might get overlooked. Let's change the wording somewhat.

$$(298) \quad \frac{\begin{array}{l} \text{Peter gets married.} \\ \text{Sue gets married.} \end{array}}{\therefore \text{Peter marries Sue.}}$$

There is a clear sense in which the second reasoning does not go through. Suppose Peter marries Joan and Sue marries Alec. Then the first two sentences are true, but the conclusion fails. Thus (298) does not go through. However, (297) still seems to be OK. And this is because the conclusion is true even in the case that Peter does not marry Sue. It is the task of semantics to explain this.

Definition 11 Let A_1, \dots, A_n and B be propositions. We say that B is a **logical consequence** of A_1 to A_n if whenever A_1, A_2, \dots, A_n are true, so is B . A_1 to A_n are called the **premises** and B the **conclusion**. We write $A_1, \dots, A_n \models B$ in this case; and $A_1, \dots, A_n \not\models B$ otherwise.

An alternative notation has been used above:

$$(299) \quad \frac{\begin{array}{c} A_1 \\ \vdots \\ A_n \end{array}}{\therefore B} \quad \text{or} \quad \frac{A_1 \dots A_n}{\therefore B}$$

We shall use this sometimes, with or without \therefore . This definition talks about propositions, not about statements. However, we still need to use some language to denote the statements that we want to talk about. Often enough linguists are content in using the statements in place of the proposition that they denote. So we write

$$(300) \quad \begin{array}{l} \text{Peter gets married., Sue gets married.} \\ \models \text{Peter and Sue get married.} \end{array}$$

Also the following is true

$$(301) \quad \text{Peter marries Sue.} \models \text{Peter gets married.}$$

But this may get us into trouble sometimes. It so turns out that sentences do not always express just one proposition, they can express various propositions, sometimes depending on the context. The easiest example is provided by pronouns. Suppose I say

(302) I love Sachertorte.

Then the proposition expressed by this sentence is different from the proposition expressed by the same sentence uttered by Arnold Schwarzenegger. This is so since I can truly utter (302) at the same time when he can truly deny (302), and vice versa. Moreover, (302) uttered by me is roughly equivalent to

(303) Marcus Kracht loves Sachertorte.

This of course would not be so if we speak of an utterance by Arnold Schwarzenegger, in which case the utterance expressed the same proposition as

(304) Arnold Schwarzenegger loves Sachertorte.

Another case is the following. Often we use a seemingly weaker sentence to express a stronger one. For example, we say

(305) Peter and Sue got married.

when we want to say that Peter married Sue. There is nothing wrong in doing so, as (305) is true on that occasion, too. But in fact what we intended to convey was that Peter married Sue, and we expect our interlocutor to understand that. Semantics does not deal with the latter problem; it does not investigate what we actually *intend* to say by saying something. This is left to **pragmatics**. (Some would also relegate the Sachertorte-example to pragmatics, but that view is not shared by everyone.) Still, even if such things are excluded, language is not always clear beyond doubt. Let's look at the following sentence.

(306) John is a bright UCLA student.

Suppose it is agreed that in order to be a UCLA student one has to be particularly bright in comparison to other students. Now, is (306) saying that John is bright even in comparison to other UCLA students, so that he is far brighter than average? Or does it only say that John is bright as a student, and that in addition he is a UCLA student? This is not clear beyond doubt.

Here is another case.

- (307) Six companies sent three representatives to the exhibition.

How many representatives got sent? Three or eighteen? We might think it is the latter, but in other sentences intuitions are less clear.

- (308) Six students visited three universities.

The example (306) may just as well involve in total six students and in total three universities.

It is agreed that the difference between (307) and (308) is not due to the particular meanings of the expressions. Rather, it is the way the world normally is that makes us choose one interpretation over another. For example, we expect that one is employed by just one company. So, if a company sends out a representative, it is by default a representative only of this company and not of another. We can imagine things to be otherwise; if the companies strike a deal by which they all pool together and send a three employees to represent them all, (307) would also be true. However, our expectations are different, and this decides that we understand this as saying that there were 18 representatives.

There are two ways of making clear what one is saying. The first is to use a sentence that is clearer on the point. For example, the following sentences make clear on the point in question what is meant.

- (309) Six companies sent three representatives each to the exhibition.

- (310) Six companies sent in total three representatives to the exhibition.

The other option is to use a formal language that has well-defined meanings and into which propositions are rendered. It is the latter approach that is more widespread, since the nuances in meaning are sometimes very difficult to express in natural language. Often enough, a mixture of the two are used. In the previous section we have used indices to make precise which DPs are taken to refer to the same individual. We shall do the same here. For example, we shall use the following notation.

- (311) $\text{run}'(x)$

Here, x is a variable which can be filled by an individual, say john' and it will be true or false. So,

(312) $\text{run}'(\text{john}')$

is a proposition, and it is either true or false. This looks like a funny way of saying the same thing, but the frequent use of brackets and other symbols will actually do the job of making meanings more precise. What it does not do, however, is explain in detail what it means to *run*. It will not help in setting the boundary between *walking* and *running*, even though semantics has a job to do there as well. It is felt, though, that this part of the job falls into the lexicon and is therefore left to lexicographers, while semantics is (mostly) concerned with explaining how the meanings of complex expressions are made from the meanings of simple expressions.

Semantics II: Compositionality

Truth Values

The thesis of compositionality says that the meaning of a complex expression is a function of the meaning of its parts and the way they have been put together. Thus, in order to understand the meaning of a complex expression we should not need to know how it has been phrased, if the two expressions are synonymous. In fact, we have made compositionality a design criterion of our representations. We said that when two constituents are merged together, the meaning of the complex expression is arrived at by applying a function to the meaning of its parts. Rather than study this in its abstract form, we shall see how it works in practice.

We shall use the numbers 0 and 1 for specifying the truth of a sentence. The numbers 0 and 1 are therefore also called **truth values** (they are also numbers, of course). 0 represents the ‘false’ and 1 represents the ‘true’. A given sentence is either false or true. We imagine that there is a ghost which tells us for every possible statement whether or not it is true or false. (In mathematical jargon we call this ghost a **function**.) We represent this ghost by a letter, say g . g needs a string and returns 0 or 1 instead of ‘yes’ or ‘no’. Now, g must abide by the conventions of the language. For example, a sentence of the form S and T is true if and only if both S and T are true. So,

$$(313) \quad g(S \text{ and } T) = 1 \text{ if and only if } g(S) = 1 \text{ and } g(T) = 1$$

This can be rephrased as follows. We introduce a symbol \cap , which has exactly the following meaning. It is an operation which needs two truth values and returns a truth value. Its table is the following.

$$(314) \quad \begin{array}{c|cc} \cap & 0 & 1 \\ \hline 0 & 0 & 0 \\ 1 & 0 & 1 \end{array}$$

Now we can write

$$(315) \quad g(S \text{ and } T) = g(S) \cap g(T)$$

This suggests that the meaning of **and** is actually \cap . Notice however that \cap is a formal symbol which has the meaning that we have given it; **and** on the other

hand is a word of English whose meaning we want to describe. Thus, to say that **and** means \cap is actually to say something meaningful! In the same way we can say that the meaning of **or** is the following function

$$(316) \quad \begin{array}{c|cc} \cup & 0 & 1 \\ \hline 0 & 0 & 1 \\ 1 & 1 & 1 \end{array}$$

and that the meaning of **if... then** is the function

$$(317) \quad \begin{array}{c|cc} \rightarrow & 0 & 1 \\ \hline 0 & 1 & 1 \\ 1 & 0 & 1 \end{array}$$

This takes a while to understand. We claim that a sentence of the form **if** S **then** T is true if either S is false or T is true. Often, people understand this as saying that S must be true exactly when T is true. (And they phrase things accordingly.) But this need not be so. For example, suppose I order a book at the bookstore and they say to me

(318) If the book arrives next week, we shall notify you.

Then if the book indeed arrives and they do not notify me they have issued a false promise. On the other hand, if it does not arrive and still they notify me (that it hasn't arrived yet), that is still alright. To give another example: in mathematics there are lots of theorems which say: 'if the Riemann hypothesis is true then...'. The point is that nobody really knows if Riemann's hypothesis actually true. (If you have a proof, you will be famous in no time!) But the theorems will remain true no matter which way the hypothesis is eventually decided. To see why this is so notice that we have said that

$$(319) \quad g(\text{if } S \text{ then } T) = g(S) \rightarrow g(T)$$

This means that **if** S **then** T is true if either S is false (that is, $g(S) = 0$) or T is true (that is $g(T) = 1$).

Let us see how this works in practice.

(320) Pete talks and John talks or John walks.

The question is: does this imply that Pete talks? Our intuition tells us that this depends on which way we read this sentence. We may read this as being formed

from the sentences `Pete talks and John talks or John walks` using `and`; or we can read it as being formed from `Pete talks and John talks and John walks` using `or`. The string is the same in both cases.

- (321) $g(\text{Pete talks and John talks or John walks})$
 $=g(\text{Pete talks}) \cap (g(\text{John talks or John walks}))$
 $=g(\text{Pete talks}) \cap (g(\text{John talks}) \cup g(\text{John walks}))$
- (322) $g(\text{Pete talks and John talks or John walks})$
 $=g(\text{Pete talks and John talks}) \cup g(\text{John walks})$
 $=(g(\text{Pete talks}) \cap g(\text{John talks})) \cup g(\text{John walks})$

The results are really different. Suppose for example that Pete talks, that John does not talk and that John walks. Then we have

- (323) $g(\text{Pete talks}) \cap (g(\text{John talks}) \cup g(\text{John walks}))$
 $=0 \cap (1 \cup 1)$
 $=0$
- (324) $(g(\text{Pete talks}) \cap g(\text{John talks})) \cup g(\text{John walks})$
 $=(0 \cap 1) \cup 1$
 $=1$

So, the different interpretations can give different results in truth!

The True Picture

Now, we have started out by assuming that there is a ghost g telling us for each sentence whether or not it is true. This ghost is in a predicament with respect to (320), for it may be that it is both true and false. Does this mean that our picture of compositionality is actually wrong? The short answer is: no! To understand why this is so, we need to actually look closer into the syntax and semantics of `and`. First, we have assumed that syntactically everything is binary branching. So the structure of `Pete talks and John walks` is

- (325) $[[\text{Pete talks}]_{\text{CP}} [\text{and } [\text{John walks}]_{\text{CP}}]]_{\text{CP}}$

Thus, `and` forms a constituent together with the right hand CP (the latter is also called **conjunct**). And the two together form a CP with the left hand CP (also

called **conjunct**). The argument in favour of this is that it is legitimate to say

(326) Pete talks. And John walks.

So, we want to have a semantics that gives meanings to all constituents involved. We shall first give the semantics of the expression **and** T . This is a function which, given some truth value x will return the value $x \cap g(T)$. In mathematics, one writes $x \mapsto x \cap g(T)$. There is a more elegant notation saying exactly the same:

(327) $g(\text{and } T) = \lambda x.x \cap g(T)$

The notation $\lambda x.\dots$ is similar to the set notation $\{x|\dots\}$. Given an expression, it yields a function. For example, given the term x^2 it gives the function $\lambda x.x^2$. In ordinary maths one writes this as $y = x^2$, but this latter notation (which is intended to say exactly the same) is actually not useful for the purpose at hand. To avoid the shortcomings of this ‘usual’ notation, the λ got introduced. In practice, if you know what the answer is in each case, say it is x^2 , then $\lambda x.x^2$ is a ghost (or machine, or function, whatever you find more instructive) that gives you the answer n^2 when you give it the number n . For example, you give it the number 3, it gives you back 9; you give it 7, it answers 49. And so on. But why is $\lambda x.x^2$ different from x^2 ? The difference can be appreciated as follows. I can say: *Let* $x^2 < 16$. Because depending on what x is, x^2 is smaller than 16 or it isn’t. It is like saying *Let* $x < 4$. But about the function $f(x) = x^2$ I cannot simply say: *Let* $f < 16$. It is not a function that is smaller than some number, only numbers can be. So, there is an appreciable difference between numbers and functions from numbers to numbers. It is this difference that we shall concentrate on.

By our conventions we write $\lambda x.x^2$ instead of f . This is *not* a number, it is a ghost. You call it, you give it a number, and it gives you back a number. In and of itself, it is not a number. So why use the notation with the λ ’s? The reason is that the notation is so useful because it can be iterated. We can write $\lambda y.\lambda x.(x - y)$. What is this? It is a ghost that, when given a number m , calls another ghost (so call it a second order ghost). This is the ghost $G := \lambda x.(x - m)$. G on its turn waits to be given a number, same or different. Give it a number n and it will return the number $n - m$. Notice that $\lambda x.\lambda y.(x - y)$ also is a second order ghost, but a different one. Give it the number m and it will call the ghost $H := \lambda y.(m - y)$. Give H the number n and it will give you $m - n$, which is not the same as $n - m$. (For example, $3 - 5 = -2$, while $5 - 3 = 2$.)

Let us see about $\lambda x.x \cap g(T)$. Suppose that T is true. Then the function is

$\lambda x.x \cap 1$. x has two choices, 0 and 1:

$$(328) \quad (\lambda x.x \cap 1)(0) = 0 \cap 1 = 0, \quad (\lambda x.x \cap 1)(1) = 1 \cap 1 = 1$$

Suppose next that T is false. Then $\lambda x.x \cap g(T) = \lambda x.x \cap 0$.

$$(329) \quad (\lambda x.x \cap 0)(0) = 0 \cap 0 = 0, \quad (\lambda x.x \cap 0)(1) = 1 \cap 0 = 0$$

Now we are ready to write down the meaning of and. It is

$$(330) \quad \lambda y.\lambda x.x \cap y$$

How does this help us? We shall assume that if two expressions are joined to form a constituent, the meaning of the head is a function that is applied to the meaning of its sister. In the construction of S and T , the head is and in the first step, and and T in the second step. Suppose for example that Pete talks but John does not walk. Ignoring some detail (for example the morphology), the signs from which we start are as follows:

$$(331) \quad \begin{bmatrix} 1 \\ \text{CP} \\ \text{Pete talks} \end{bmatrix} \quad \begin{bmatrix} \lambda y.\lambda x.x \cap y \\ \text{C} \\ \text{and} \end{bmatrix} \quad \begin{bmatrix} 0 \\ \text{CP} \\ \text{John walks} \end{bmatrix}$$

Let's put them together.

$$(332) \quad \begin{bmatrix} 1 \\ \text{CP} \\ \text{Pete talks} \end{bmatrix} \bullet \left(\begin{bmatrix} \lambda y.\lambda x.x \cap y \\ \text{C} \\ \text{and} \end{bmatrix} \bullet \begin{bmatrix} 0 \\ \text{CP} \\ \text{John walks} \end{bmatrix} \right) \\ = \begin{bmatrix} 1 \\ \text{CP} \\ \text{Pete talks} \end{bmatrix} \bullet \begin{bmatrix} \lambda x.x \cap 0 \\ \text{C}' \\ \text{and John walks} \end{bmatrix} \\ = \begin{bmatrix} 0 \\ \text{CP} \\ \text{Pete talks and John walks} \end{bmatrix}$$

It is not hard to imagine that we can get different results if we put the signs together differently.

Types of Ghosts

We have said that there are ghosts, and that there are second order ghosts. They are higher up, because they can call ordinary ghosts to fulfill a task for them. To keep track of the hierarchy of ghosts, we shall use the following notation. Ghosts have *types*. Here are the types of the ghosts introduced so far.

$$\begin{aligned}
 & 0, 1 : t \\
 (333) \quad & \lambda x.x \cap 1 : t \rightarrow t \\
 & \lambda y.\lambda x.x \cap y : t \rightarrow (t \rightarrow t)
 \end{aligned}$$

The general rule is this:

Definition 12 *A truth value is of type t . A function from ghosts of type α into ghosts of type β is a ghost of type $\alpha \rightarrow \beta$. A ghost of type $\alpha \rightarrow \beta$ can only apply to a ghost of type α ; the result is a ghost of type β .*

Something of Type 1 is a truth value; by definition, it is either 0 or 1. So, it is known to us directly. Everything else needs some computation on our side. A ghost of type $\alpha \rightarrow \beta$ is asking for a ghost of type α to be given. It will then call a ghost of type β for an answer. Typically we had $\alpha = t$. A ghost of type $t \rightarrow \beta$ is waiting to be given a truth value and it will then call a ghost of type β .

Is there more to it than just truth values? There is! We shall assume, for example, that there is a type e that comprises all objects of the world. Every physical object belongs to e , including people, animals, stars, and even ideas, objects of the past like dinosaurs, objects of unknown existence like UFOs, they are all of type e . Ordinary names, like John and Pete can denote objects, and so their meaning is of type e , we say. A ghost of type $e \rightarrow t$ is a ghost that waits to be given an object and it will return ... a truth value. It says 0 or 1 (or 'yes' or 'no', however we like). These ghosts are called **(unary) predicates**. Intransitive predicates denote unary predicates. For example, run, talk and walk. We shall truly treat them as ghosts. We do not know how they work. This is part of the lexicon, or if you wish, this part of the knowledge of the language that we must acquire what it means that someone talks, or walks, or runs. In the absence of anything better we write run' for the ghost of type $e \rightarrow t$ that tells us if someone is running. And likewise we write walk' for the ghost that tells us if someone is

talking. And so on. Then we do however know that the meaning of

(334) John talks.

is $\text{talk}'(\text{john}')$. Big deal. But things get trickier very soon. There are also transitive predicates like *scold*. Since they form a category of intransitive verb together with a name, we conclude that they must be ghosts of second order: they are of type $e \rightarrow (e \rightarrow t)$. And so on. What this means is that semantics mirrors by way of typing ghosts the syntactic structure. If something is a transitive verb, not only does it need two syntactic arguments, also its meaning is a second order ghost, waiting to be given two objects before it will answer with a truth-value.

Semantics III: Basic Elements of Interpretation

On What There Is

Language is supposed to be able to talk about everything that exists. This is, of course, impossible. However, language comes very close in doing that. Such a task creates its own challenges. We notice, namely, that things we speak about are of different kinds: there are physical objects, people, places, times, properties, events and so on. Languages reflect this categorization of things in different ways. For example, many languages let the verb show which of the participants in a speech act is the subject: if it is the speaker, the hearer or someone else. We can perfectly well imagine a language that does not do that. But the fact that many languages do have such a system reflects the importance of the speaker and hearer in a speech act. To give another example: many languages classify things into various groups, called genders or classes. This classification proceeds along several lines. Many languages distinguish animate from inanimate, for example. We can certainly imagine that this contrast is relevant. For some languages it is so relevant that it enters the morphology.

To the extent that the classification of things is relevant to language, it is going to be reflected in the basic semantic *types*. We shall review a few basic types of things that are relevant to natural languages.

Number: Individuals and Groups

Look at the following contrast.

(335) John and Mary met.

(336) The senators met.

(337) *John met.

Evidently, in order to meet you need to meet with somebody. But apparently several people can meet in the sense that they meet with each other. One way to account for this difference is to say that the verb *meet* needs a group of people as its subject. Since lines and other things can also meet, we shall simply say that *meet* needs a group of things, without further qualification. There are basically two ways to form groups: you use the plural or you use *and*.

We say therefore that plural DPs and DPs coordinated with the help of *and* denote groups. Groups can evidently be subjects of verbs, and sometimes verbs specifically require groups as subjects. Other verbs take both.

(338) Paul is eating.

(339) Paul and the neighbouring cat are eating.

In the case of *eat* we may think of it as being an activity that is basically performed by each individual alone. If it is understood to be this way, the verb is said to be **distributive**. In means in the present case that the fact that Paul is eating and that the neighbouring cat is eating is enough to make the sentence (339) true. We can also say that the following inference is correct.

(340)
$$\frac{\begin{array}{l} \text{Paul is eating.} \\ \text{The neighbouring cat is eating.} \end{array}}{\therefore \text{Paul and the neighbouring cat are eating.}}$$

In general, if a verb *V* is distributive, then the following inferences go through:

(341)
$$\frac{\begin{array}{l} A \text{ Vs.} \\ B \text{ Vs.} \end{array}}{\therefore A \text{ and } B \text{ V.}} \quad \frac{A \text{ and } B \text{ V.}}{\therefore A \text{ Vs.}} \quad \frac{A \text{ and } B \text{ V.}}{\therefore B \text{ Vs.}}$$

Now, if there are groups of people, are there also groups of groups? There are! Here is an interesting pattern. You can get married as a group of two people by getting married to each other. This involves a one-time event that makes you married to each other. You can get married as a group by each marrying someone else. The sentence (342) can be read in these two ways. The reason why this is so lies in the fact that *get married* actually also is a verb that takes individuals.

(342) John and Sue got married.

(343) John got married.

Moreover, the above test of distributivity goes through in that case. But if we understand it in the non-distributive sense, the inference does not go through, of course. Now, let's form a group of groups:

(344) John and Alison and Bill and Sue got married.

There is a reading which says that John marries Alison, and Bill marries Sue. This reading exists in virtue of the following. *John and Alison* denotes a group, so

does Bill and Sue. The verb applies to such groups in the meaning ‘marry each other’. By our understanding of marriage, if several groups get married to each other, this means that all groups get married separately.

There are also verbs that encourage a group of groups reading.

(345) The ants were marching eight by eight.

Here, we think of a group of groups of ants, each group consisting of eight ants in a line.

Note that the same pattern can be observed with meet:

(346)	The senators meet.	
	The congressmen meet.	
	∴ The senators and the congressmen meet.	

However, reader beware: the conclusion has a reading where the subject is a single group which meets in a one-time event. This is not what can be said to follow from the premises. All that follows is that the senators met (with each other) and that the congressmen met (with each other).

Time

Properties can change over time. I can be tired in the evening, but maybe next morning I am not. Some properties have an inbuilt time dependency. For example, a cat is a kitten only through certain stages of its life. When it is old enough it ceases to be a kitten, even though it never ceases to be a cat. We picture time as consisting of a continuum of time points on a line. We write t, t', u for time points, and we write $t < t'$ to say that t is prior to t' , and $t > t'$ to say that it is after t' . For any two time points t and t' , either $t = t'$ (they are the same) or $t < t'$ or $t > t'$. This trichotomy translates into the three basic tenses of English. The present is used for something that happens now, the past is used for something that has happened before now, and the future is used for something that will happen later.

(347) John runs. (present)

(348) John ran. (past)

(349) John will run. (future)

We make reference to time through various ways. One is the words *now* and *yesterday*. The word *now* refers to the very moment of time where it is uttered; *yesterday* refers to any time point that is on the day before today. Today on the other hand is the day of now. Other words require some calculation.

(350) John realized on Monday that he had lost his wallet
the day before.

We do not know exactly when things happened. We know that John's realizing he had lost his wallet happened in the past (because we used past tense); and it happened on a Monday. His losing the wallet happened just the day before that Monday, so it was on a Sunday. Suppose we replaced the *day before* by *yesterday*:

(351) John realized on Monday that he had lost his wallet
yesterday.

Then John's realizing is in the past, and it is on a Monday. His losing his wallet is prior to his realizing (we infer that among other from the phrase *had lost* which is a tense called **pluperfect**). And it was yesterday. So, today is Monday and yesterday John lost his wallet, and today he realizes that. Or he realized yesterday that on that day he had lost his wallet. (Actually, the phrase *on Monday* hurts our ears here. We are not likely to say exactly what day of the week it is when it is today or yesterday or tomorrow. But that is not something that semantics concerns itself with. I can say: *let's go to the swimming pool on Thursday* even when today is Thursday. It is just odd to do so.)

That time is linear and transitive accounts for a few inference patterns that we note.

(352)	$\begin{array}{l} A \text{ Ved.} \\ B \text{ Ved.} \end{array}$	$\begin{array}{l} A \text{ Ved before B.} \\ B \text{ Ved before C.} \end{array}$
	$\therefore A \text{ Ved before B or A Ved after B} \\ \text{or A Ved at the same as B.}$	$\therefore A \text{ Ved before C.}$

Location

Space is as important in daily experience as time. We grow up thinking spatially; how to get furniture through a door, how to catch the neighbour's cat, how to not

get hit by a car, all these things require coordination of thinking in time and space. This is the reason why language is filled with phrases that one way or another refer to space. The most evident expressions are *here*, which functions the same way as *now*. It denotes the space that speaker is occupying at the moment of utterance. *In* is involved in the words *come* and *go*. If someone is approaching me right now, I can say

(353) He is coming.

But I would not say that he is going. That would imply he is moving away from me now. German uses verbal prefixes (*hin/her*) for a lot of verbs to indicate whether movement is directed towards speaker or not.

Space is not linear, it is organized differently, and language reflects that. Suppose I want to say where a particular building is located on campus. Typically we phrase this by giving an orientation and a distance (this is known as ‘polar coordinates’). For example,

(354) 200 m southwest of here

gives an orientation (southwest) and a distance (200 metres). The orientation is either given in absolute terms, or it can be relative to the way one is positioned, for example *to the right*. To understand the meaning of what I am saying when I say *go to the right* you have to know which way I am facing.

Worlds and Situations

We have started out by saying that sentences are either true or false. So, any given sentence such as the following is either true or false.

(355) Paul is chasing a squirrel.

(356) Napoleon lost the battle of Waterloo.

(357) Kitten are young cats.

We can imagine with respect (355) that it is true right now or that it is false. In fact, we do not even know. With (356) it is the same, although if we have learned a little history we will know that it is false. Still, we find ourselves thinking ‘what if Napoleon had actually won the battle of Waterloo ...’. Thus, we picture a

situation that is contrary to fact. The technical term is **world**. Worlds decide every sentence one way or another. There are worlds in which (355) is true and (356) is false, others in which (355) is false and (356) is true, others in which both are false, and again others in which both are false.

Seemingly then, any combination of saying this and that sentence is true or false is a world. But this is not quite true. (357) is different. It is true. But to suppose otherwise would be tantamount to violating the rules of language. If I were to say ‘suppose that kitten are not young cats but in fact old rats ...’ what I ask you is to change the way English is understood. I am not talking about a different world. Worlds have an independent existence from the language that is being used. We say then that (357) is **necessarily true**, just like $4+7=11$. If you do not believe either of them you are just not in the picture.

The denotation of a word like *cat* in this world is the set of all beings that are cats. They can change from world to world. We can imagine a world that has absolutely no cats. (If we go back in time, there was a time when this was actually true.) Or one that has no mice. But we do not suppose that just because there are different sets of cats in different worlds the meaning of the word changes — it does not. That’s why you cannot suppose that kitten are old rats. We say that the meaning of the word *cat* is a function cat' that for each and every world w gives some set, $\text{cat}'(w)$. We of course understand that $\text{cat}'(w)$ is the set of all cats in w . (Some people use the word **intension** for that function.) Likewise, the intension of the word *rat* gives for each world w the set of all rats in w , and likewise for the word *kitten*. It is a fact of English that

$$(358) \quad \text{kitten}'(w) \subseteq \text{cat}'(w), \quad \text{kitten}'(w) \cap \text{rat}'(w) = \emptyset$$

There are plenty of words that are sensitive not just to the denotation but to the meaning.

(359) John doubts that Homer has lived.

(360) Robin thinks that Napoleon actually won Waterloo.

Nobody actually knows whether or not Homer has existed. Still we think that the sentence ‘Homer has lived.’ has a definite answer (some ghost should tell us...). It is either true or not. Independently of the answer, we can hold beliefs that settle the question one way or the other, regardless of whether the sentence is factually true or not. Robin, for example, might be informed about the meaning of all English words, and yet is a little weak on history. So she holds that Napoleon won

Waterloo. John might believe the opposite, and Robin might believe that Homer has lived. Different people, different opinions. But to disagree on the fact that kitten are cats and not rats means not using English anymore.

Events

When I sit behind the computer typing on the keyboard, this is an activity. You can watch me do it and describe the activity in various ways. You can say

(361) He is typing on the keyboard.

(362) He is fixing the computer.

(363) He is writing a new book.

Both (361) and (362) can be manifested by the same physical activity: me sitting behind the computer and typing something in. Whether or not I am fixing the computer by doing so, time will tell. But in principle I could be fixing the computer just by typing on the keyboard. The same goes for writing a book. Thus, one and the same physical activity can be a manifestation of two different things. In order to capture this insight, one has proposed that verbs denote particular objects called **events**. There are events of typing, as there are events of fixing the computer, and events of writing a book. Events are distinct from the process that manifests them.

Events will figure in Lecture 18, so we shall not go into much detail here. There are a few things worth knowing about events. First, there are two kinds of events: **states** and **processes**. A state is an event where nothing changes.

(364) Lisa knows Spanish.

(365) Harry is 8 feet tall.

These are examples where the truth of something is asserted at a moment of time, but there is no indication that something changes. By contrast the following sentences talk about processes.

(366) Paul is running.

(367) The administrator is filling out the form.

(368) The artist is making a sculpture.

In the first example, Paul is changing place or posture. In the second the administrator is for example writing on some piece of paper which changes that very piece of paper. In the third example a statue comes into existence from some lump of material. Events have participants whose number and contribution can vary greatly. A process always involves someone or something that undergoes change. This is called the **theme**. In (366) the theme is Paul, in (367) the theme is the form, in (368) the theme is the sculpture. Events usually have a participant that makes the change happen; in (366) the actor is again Paul, in (367) it is the administrator, in (368) it is the artist. There need not be an actor, just as there need not be a theme; but mostly there is a theme. Some events have what is called an **experiencer**. In the next sentence, Jeremy is the experiencer of hate.

(369) Jeremy hates football.

Notice that experiencer predicates express states of emotion, so they fall into the category of verbs that express states rather than processes. Another class of participants are the **beneficiaries**; these are the ones for whose benefit an action is performed, like the wife of the violinist in the following example.

(370) The violinist composed a sonata for his wife.

The list of participant types is longer, but these ones are the most common ones.

Processes are sometimes meant to finish in some state and sometimes not. If you are said to be running no indication is made as to when you will stop doing so. If you are said to be running a mile then there is an inherent point that defines the end of the activity you engage in. Some verbs denote the latter kind of events: arrive, reach, pop, finish. The process they denote is finished when something specific happens.

(371) Mary arrived in London.

(372) The composer finished the oratorio.

In (371) the arriving of Mary happens at a more or less clearly defined time span, say when the train gets close to the station up to when it comes to a halt. Similarly for (372), where the finishing is the last stretch of the event of writing the oratorio. The latter is a preparatory action. So, you can write an oratorio for a long time, maybe years, but you can only finish it in a considerably shorter time, at the end of your writing it.

Semantics IV: Scope

Let us return to example (320), repeated here as (373).

(373) Pete talks and John talks or John walks.

We have said that under certain circumstances it may turn out to be both true and false, depending on how we read it. These interpretations are also called **readings**. In this lecture we shall be interested in understanding how different readings are also structurally different. The syntactic notion that is pivotal here is that of **scope**. Intuitively, the scope is that string part (constituent) that serves as an argument to some head. We shall fill this definition with life right away. In (373) we find two logical connectives, **and** and **or**. Each of them takes two CPs, one to the right and one to the left. This is all the requirements they make on the syntactic side. This means that syntactically the sentence can be given two different structures. They are shown in Figure 6. The structure of the individual CPs is not shown, to save space. Let us look at (a). We can tell from its structure what its meaning is. We shall work it out starting at the bottom. The complement of **or** is the CP **John walks**. So, **or John walks** is a constituent formed from **or** and **John walks**. This means that its meaning is derived by applying the meaning of **or** to the meaning of **John walks**. Notice that the latter is also the constituent that the node labelled **C** just above **or** c-commands (recall the definition of c-command from Lecture 12). The meaning of the lower **C'** node is therefore


(374) $(\lambda y. \lambda x. x \cup y)(g(\text{John walks}))$
 $= \lambda x. x \cup g(\text{John walks})$

This node takes the CP node above **John talks** as its sister, and therefore it c-commands it. The meaning of the CP that the two together form is

(375) $(\lambda x. x \cup g(\text{John walks}))(g(\text{John talks}))$
 $= g(\text{John talks}) \cup g(\text{John walks})$

This is the meaning of the lower CP. It is the complement of **and**. The two form a constituent, and its meaning is

(376) $(\lambda y. \lambda x. x \cap y)(g(\text{John talks}) \cup g(\text{John walks}))$
 $= \lambda x. x \cap (g(\text{John talks}) \cup g(\text{John walks}))$



(b)




Figure 6: Two Analyses of (373)

Finally, we combine this with the specifier CP:

$$(377) \quad (\lambda x.x \cap (g(\text{John talks}) \cup g(\text{John walks}))) (g(\text{Pete talks})) \\ = g(\text{Pete talks}) \cap (g(\text{John talks}) \cup g(\text{John walks}))$$

This is exactly the same as the interpretation (321).

Now we take the other structure. This time we start with the interpretation of the middle CP. It is now c-commanded by the C above the word *and*. This means that the two form a constituent, and its interpretation is

$$(378) \quad (\lambda y.\lambda x.x \cap y) (g(\text{John talks})) \\ = \lambda x.x \cap g(\text{John talks})$$

This constituent takes the specifier *Pete talks* into a constituent CP. Thus, it applies itself to the meaning of *Pete talks*:

$$(379) \quad (\lambda x.x \cap g(\text{John talks})) (g(\text{Pete talks})) \\ = g(\text{Pete talks}) \cap g(\text{John talks})$$

This constituent is now the specifier of a CP, which is formed by *Pete talks* and *John talks* and or *John walks*. The latter has the meaning $(\lambda x.x \cup g(\text{John walks}))$, which applies itself to the former:

$$(380) \quad (\lambda x.x \cup g(\text{John walks})) (g(\text{Pete talks}) \cap g(\text{John talks})) \\ = (g(\text{Pete talks}) \cap g(\text{John talks})) \cup g(\text{John walks})$$

And this is exactly (322).

Thus, the two different interpretations can be seen as resulting from different structures. This has motivated saying that our ghost *g* does not take sentences as inputs. Instead, it wants the entire syntactic tree. Only when given a syntactic tree the ghost can give a satisfactory answer. The way the meaning is computed is by working its way up. We assume that each node has exactly two daughters. Suppose we have a structure

$$(381) \quad \gamma = [\alpha \beta]$$

We assume that the semantics is arranged in such a way that if two constituents are merged into one, the interpretation of one of the nodes is a function that can

be applied to the meaning of its sister. The first node is then called the **semantic head**. If the meaning of α and β is known and equals $g(\alpha)$ and $g(\beta)$, respectively, then the meaning of γ is

$$(382) \quad g(\gamma) = \begin{cases} g(\alpha)(g(\beta)) & \text{if } \alpha \text{ is the semantic head} \\ g(\beta)(g(\alpha)) & \text{if } \beta \text{ is the semantic head} \end{cases}$$

The only thing we need to know is: is α the semantic head or is it β ? The general pattern is this: a zero level projection is always the semantic head, and likewise the first level projection. Adjuncts are semantic heads, but they are not heads (the latter of head is a syntactic notion and is different, as this case shows). Notice that by construction, the head ‘eats’ its sister: its meaning is a function that applies itself to the meaning of the sister. And the sister is the constituent that it c-commands. This is why c-commands has become such a fundamental notion in syntax: it basically mirrors the notion of scope, which is the one needed to know what the meaning of a given constituent is.

We shall discuss a few more cases where scope makes all the difference. Look at the difference between (383) and (384).

(383) This is the only car we have, which has recently been repaired.

(384) This is the only car we have which has recently been repaired.

The part **which has recently been repaired** is a clause that functions like an adjective; it is called a **relative clause**. (Recall that relative clauses are opened by relative pronouns in English.) Unlike adjectives, relative clauses follow the noun. Notice that **we have** also is a relative clause, though it is somewhat shortened (we could replace it **which we have**).

Suppose you go to a car dealer and he utters (383). Then he is basically saying that he has only one car. Moreover, this car has been under repair recently. Suppose however he says (384). Then he is only claiming that there is a single car that has been under repair recently, while he may still have tons of others. It is clear from which dealer you want to buy. To visualize the difference, we indicate

the scope of the operator *the only*:

- (385) This is the only [car we have], which has recently been repaired.
- (386) This is the only [car we have which has recently been repaired].

In the first sentence, *the only* takes scope over *car we have*. In the second sentence it has scope over *car we have which has recently been repaired*. It is not necessary to formalize the semantics of *the only*. We need only say that something is *the only P*, if and only if (a) it is a *P*, and (b) nothing else is a *P*. So, if *the only* takes scope only over *car we have*, then the car we are talking about is a car the dealer has, and there is no other that he has. So he has only one car. If the scope is *car we have which has recently been repaired*, then the car we are talking about has recently been repaired, it is one of the dealer's cars, and there is no other car like that. So, there may be other cars that the dealer has which have not been repaired, and there may be other cars that were not the dealer's but have been repaired; but none other that is both.

The difference in structure between the two is signaled by the comma. If the comma is added, the scope of *the only* ends there. The relative clause is then said to be **non-restrictive**; if the comma is not there, the relative clause is said to be **restrictive**. If we look at \bar{X} -syntax again we see that non-restrictive relative clauses must be at least D' -adjuncts, because they cannot be in the scope of *the*. In fact, one can see that they are DP-adjuncts. Let us see how this goes. First we notice that *the* can be replaced by a possessive phrase (which is in the genitive):

- (387) Simon's favourite bedtime story.
- (388) Paul's three recent attacks on squirrels.

The possessives are phrases, so they are in specifier of CP, quite unlike the determiner *the* itself. Notice that even though the possessive and the determiner cannot co-occur in English this is not due to the fact that they compete for the same position. In Hungarian they can co-occur:

- (389) Mari-nak a cipője
 Mary- the shoe- 3
Mary's shoe

Literally translated, this means *Mary's the her shoe*. The possessive *Marinak* (in the dative!) occurs before the actual determiner.

Now, take a DP which has non-restrictive DP.

(390) Simon's only bedtime story, which he listens to carefully

(391) Paul's only attacks on squirrels, which were successful

These DPs are perfect, but here we have *Simon's only* and *Paul's only*. We shall not go into the details of that construction and how it differs from *the only*. In the first DP, *Simon's only* takes *bedtime story* in its scope. It can only do so if the relative clause is an adjunct to DP.

Likewise one may wonder about the place of restrictive relative clauses. It is clear that they can be neither adjuncts to DP nor adjuncts to D' (because then the *only* cannot take scope over them). The restrictive relative clauses is therefore adjunct to either N' or NP. We shall not go into more detail.

So far we have seen that differences in interpretation manifest themselves in differences in structure. The next example is not of the same kind — at least at first sight. This example has to do with quantification. Suppose that the professors complain about office space and the administrator says

(392) Every professor has an office.

He might be uttering a true sentence even if there is a single office that is assigned to all professors. If this is to be understood as a remark about how generous the university is, then it is probably just short for

(393) Every professor has his own office.

in which case it would be clear that each professor has a different office. The first reading is semantically stronger than the second. For there is a single office and it is assigned to every professor then every professor has an office, albeit the same one. However, if every professor has an office, different or not, it need not be the same that there is just a single office. We can use 'stilted talk' to make the difference visual:

(394) There is an office such that it is assigned to every
professor.

(395) Every professor is such that an office is assigned to
him.

In the first sentence, *there is an office* takes scope over *every professor*. In the second sentence *every professor* takes scope over *an office*.

Returning to the original sentence (393), however, we have difficulties assigning different scope relations to the quantifiers: clearly, syntactically *every professor* takes *his own office* in its scope. This problem has occupied syntacticians for a long time. The conclusion they came up with is that the mechanism that gets the different readings is syntactic, but that the derivation at some point puts the object into c-commanding position over the subject.

There are other examples that are not easily accounted for by syntactic means. Let us give an example.

(396) John searches for the holy grale.

There are at least two ways to understand this sentence. Under one interpretation it means that there is something that John searches for, and he thinks it is the holy grale. Under another interpretation it means that John is looking for something that is the holy grale, but it may not even exist. This particular case is interesting because people are divided over the issue whether or not the holy grale existed. Additionally, it is not clear what it actually was. So, we might find it and not know that we have found it. We may paraphrase the readings as follows.

(397) There is the holy grale and John searches for it as that.

(398) There is something and John searches for it as the

(399) holy grale.

(400) John is searching for something as the holy grale.

Here, the meaning difference is brought out as a syntactic scope difference. The first is the strongest sentence: it implies that both speaker and John identify some object as the holy grale. The second is weaker: there is something of which only John believes that it is the holy grale. The third is the weakest: John believes that there is such a things as the holy grale, but it might not even exist.

Semantics V: Cross-Categorical Parallelism

We have said that nouns denote objects and that verbs denote events. It has been observed, however, that some categorizations that have been made in the domain of objects carry over to events and vice versa. They ultimately relate to an underlying structure that is similar in both of them. A particular instance is the distinction between mass and count nouns. A noun is said to be a **count noun** if what it refers to in the singular is a single object that cannot be conceived of consisting of parts that are also objects denoted by this noun; for example, *bus* is a count noun. In the singular it denotes a thing that cannot be conceived as consisting two or more busses. It has parts, for sure, such as a motor, several seats, windows and so on. But there is no part of it which is again a bus. We say that the bus is an **integrated whole** with respect to being a bus. Even though some parts of it may not be needed for it to be a bus, they do not by themselves constitute another bus. Ultimately, the notion of integrated whole is a way of looking at things: a single train, for example, may consist of two engines in front of several dozens of wagons. It may be that it has even been obtained by fusing together two trains. However, that train is not seen as two trains: it is seen as an integrated whole. That is why things are not as clear cut we might like them to be. Although we are mostly in agreement as to whether a particular is a train or not, or whether it is two trains, an abstract definition of an integrated whole is hard to give.

However, one thing is clear. The division into smaller and smaller units must stop. The train cannot consist of smaller and smaller trains. At some point, very soon, it stops to be a train. There is a difference with *water*. Although in science we learn that things are otherwise, in actual practice water is divisible to any degree we like. And this is how we conceive of it. We take an arbitrary quantity of water and divide it as we please — the parts are still water. Thus, *water* is not a count noun; it is a **mass noun**.

One problem remains, though. We have not talked about mass things, we have consistently talked about mass or count *nouns*. We said, for example, that *water* is a mass noun, not that the substance water itself is a mass substance. In this case, it is easy to confuse the two. But there are occasions where the two are different. The word *furniture* turns out to be a mass noun, even though what it denotes clearly cannot be indefinitely divided into parts that are also called furniture. But how do we know that something is a mass noun if we cannot ultimately rely on our intuitions about the world? There are a few tests that establish this. First, mass

nouns do not occur in the plural. We do not find *furnitures or *courage. On the surface, waters seems to be an exception. However, the denotation of waters is not the same as that of a plural of a count noun, which is a group. Waters is used, for example, with respect to clearly defined patches of water (like rivers or lakes). A better test is this one. Mass nouns freely occur with so-called **classifiers**, while count nouns do not.

- (401) a glass of water
- (402) a piece of furniture
- (403) a *glass/*piece of bus

Notice that one does not say *a glass of furniture, nor a *piece of water. The kind of classifier that goes with a mass noun depends on what it actually denotes. Some can be used across the board, like lot.

Notice that there is a distinction between a denotation that is divisible into like parts and one that is not, and a language that makes a distinction between the two. To the extent that the distinction between mass and count exists in the real world, it cannot be denied to exist by any other language. It can only be ignored. It is the same as gender distinctions: for some languages they are not morphologically relevant, but the division into various sexes, or other classes can usually be represented one way or another.

Let us now look at verbs. Verbs denote events, as we have said. Events are like films. We may picture them as a sequence of scenes, lined up like birds on a telephone cable. For example, scene 1 may have Paul 10 feet away from some squirrel; scene 2 sees Paul being 8 feet away from the squirrel; scene 3 sees Paul just 6 feet away from the squirrel; and so on, until he is finally right next to it, ready to eat it. Assume that the squirrel is not moving at all. This sequence can then be summarized by saying:

- (404) Paul is attacking the squirrel.

Similarly, in scene 1 someone is behind a blank paper. In scene 2, he has drawn a small line, in scene 3 a fragment of a circle. From scene to scene this fragment of a circle grows, until in the last scene you see a complete circle. You may say

- (405) He has drawn a circle.

While you are watching the film, you can say

- (406) He is drawing.

Or you can say

(407) He is spreading ink on the paper.

All these are legitimate ways of describing what is or was going on. Unfortunately, the director has decided to cut a few of the scenes at the end. So we now have a different film. Now we are not in a position to truthfully utter (405) on the basis of the film any more. This is because even though what that person began to draw looks like the beginning of a circle, that circle may actually never have been completed. Notice that the situation is quite like one where we stop watching the movie: we are witnessing part of the story and guess what the rest is like, but the circle may not get completed. However, (406) and (407) are still fine no matter what really happens. Even if the director cuts parts of the beginning, still (406) and (407) are fine. No matter how short the film is and no matter what really happen thereafter: the description seems adequate.

This is the same situation as before with the nouns. Certain descriptions can be applied also to subparts of the film, others cannot. Those events that can be divided are called **atelic**; the others being **telic**. This is the accepted view. One should note though that by definition, a telic event is one that ends in a certain state, without which the event would not be the same. In other words: if we cut out parts of the beginning, that would not hurt. But if we could out parts of the end, that would make a difference. An example is the following.

(408) John went to the station.

Here, it does not matter so much where John starts out from, as long it was somewhere away from the station. We can cut parts of the beginning of the film, still it is a film about John's going to the station. Telic events are directed towards a goal (that gave them their name; in Ancient Greek, 'telos' meant 'goal'). However, as it appears, the majority of nondivisible events are telic. A different one is

(409) John wrote a novel.

Here, cutting out parts of the film anywhere will result in making (409) false, because John did not write the novel in its entirety.

Now, how do we test for divisibility (= atelicity)? The standard test is to see whether *for an hour* is appropriate as opposed to *in an hour*. Divisible

events can occur with *for* an hour, but not with *in* an hour. With indivisible events it is the other way around.

- (410) John wrote a novel in an hour.
- (411) *John wrote a novel for an hour.
- (412) *John was writing in an hour.
- (413) John was writing for an hour.

So, divisible events can be distinguished from nondivisible events. However, let us see if we can make the parallel even closer. We have said that mass terms are divisible. But suppose also this: if you pour a little water into your glass, and then another little, as a result you still have water. You cannot say you have two water. You can only say this if, say, you have two glasses of water, so the bits of water are separated. (Actually, with water this still sounds odd, but rivers allow this use.) Also, it does not make sense to divide your portion of water in any way and say that you have two pieces of water in your glasses. Likewise, suppose that John is running from 1 to 2pm and from 2pm to 3pm and did not at all stop — we would not say that he ran twice. The process of running stretches along the longest interval of time as it possibly can. There is one process only, just as there is water in your glass, without any boundary. The glass defines the boundary of water; so if you put another glass of water next to it, there are now two glasses of water. And if John is running from 1 to 2pm and then from 3 to 4pm, he ran twice; there are now two processes of running, because he did stop in between. The existence of a boundary between things or events determines whether what we have is one or two or several of them.

Putting It All Together

Having taken a closer look at phonology, morphology, syntax and semantics, we shall revisit the pig picture of the first lecture. We said that there is one operation, called ‘merge’ and that it operates on all of these four levels at the same time. However, we had to make concessions to the way we construe the levels themselves. For example, we argued that the English past tense marker was [d], but that it gets modified in a predictable way to [t] or [əd]. Thus we were led to posit two levels: deep phonological and surface phonological level. Likewise we have posited a deep syntactic level and a surface level (after movement has taken place), and there is also a deep morphological level and a surface morphological

level. This throws us into a dilemma: we can apply the morphological rules only after we have the surface syntactical representation, because the latter reorders the lexical elements. Likewise, the deep phonological representation becomes apparent only after we have computed the surface morphological form. Thus, the parallel model gives way to a sequential model, which has been advocated for by Igor Mel'čuk (in his Meaning-to-Text theory). In this model, the levels are not parallel, they are ordered sequentially. We speak by organising first the semantic representation, then the words on the basis of that representation and the syntactic rules, then the morphological representation on the basis of the lexical, and the phonological on the basis of the morphological representation. Listening and understanding involves the converse sequence.

The paradigm of generative grammar is still different. Generative grammar assumes a generative process which is basically independent of all the levels. It runs by itself, but it interfaces with the phonology and the meaning at certain points. The transformations are not taken to be operations that are actually executed, but are ways to organize syntactic (and linguistic) knowledge. This makes the empirical assessment of this theory very difficult, because it is difficult to say what sort of evidence is evidence for or against this model.

There are also other models of syntax. These try to eliminate the distinction between deep and surface structure. For example, in GPSG the question words are generated directly in sentence initial position, there simply is no underlying structure that puts the object first right adjacent to the verb from which it is moved to the beginning of the sentence. It is put into sentence initial position right away. Other grammars insist on special alignment rules.

There are probably as many theories as there are linguists. But even though the discussion surrounding the architecture of linguistics has been much in fashion in the last decades, some problems still remain that do not square with most theories. We mention just one very irritating fact. We have seen that Malay uses reduplication for the plural. If that is so then first of all the plural sign has no substance: there is no actual string that signals the plural (like the English *s*). What signals the plural is the reduplication. This is a function ρ that sends a string x into that string concatenated with itself in the following way:

$$(414) \quad \rho(x) := x\hat{-}\hat{-}x$$

Thus, $\rho(\text{kerani}) = \text{kerani-kerani}$. This function does not care whether the input string is an actual word of Malay. It could be anything. But it is *this function*

which is the phonology of the plural sign. This means among other that the phonological representation of signs must be very complicated if the story of parallelism is to be upheld (recall the plural of *mouse* with a floating segment). We have to postulate signs whose phonology is not a string but a function on strings. Unfortunately, no other theory can do better here if that is what Malay is like. Thus, the door has to be opened: there is more to merge than concatenating strings. If that is so, we can try the same for syntax; there is more to syntax than concatenating constituents. It has been claimed that Chinese has a construction that duplicates entire constituents. Even if that story is not exactly true, the news is irritating. It means that entire constituents are there just to convey a single piece of meaning (here: that the sentence is a question).

But we need not go that far. Lots of languages in Europe have agreement of one sort or another. English still has number agreement, for example, between the demonstrative and the NP (*this flag* versus *these flags*), and with the verb. The agreement is completely formal. One says *these troops* and not *this troops*, even though one does say *this army*. However, the number that the demonstrative carries is semantically dependent on the noun: if the latter carries plural meaning, then the whole is plural, otherwise not. The semantic contribution of *these* is not plural irrespective of whether the noun actually specifies plural meaning. Take *guts*, whose meaning may be singular like *courage*. Unfortunately, it cannot really be combined with a demonstrative. (Otherwise we would expect *those guts*, and certainly not *that guts*.) A better example is Latin *litterae* ‘letter’ (which you write to a friend), a morphological plural derived from *littera* ‘the letter’ in the sense of ‘the letter A’. The letter you write is a single object, though it is composed from many letters. It controls plural agreement in any event. Now, if the plural morpheme appears many times in the sentence, but only once is it allowed to carry plural meaning — what are we to do with the rest of them? The puzzle has been noted occasionally, and again several solutions have been tried. Harris speaks of a ‘scattered morpheme’, he thought they are just one piece, distributed over many places. The same intuition seems to drive generative grammar, but the semantics is never clearly spelled out.

Language Families and History of Languages

Today, linguistics focuses on the mental state of speakers and how they come to learn language. To large parts, the investigation dismisses input that comes from an area of linguistics that was once dominant: historical linguistics. The latter is the study of the history and development of languages. The roots of historical linguistics go as far back as the late 17th century when it was observed that English, German, Dutch as well as other languages shared a lot of common features, and it was quickly observed that one could postulate something of a language that existed a long time ago, called **Germanic**, from which these languages developed. To support this, let us look at a few words in these languages:

	English	Dutch	German
(415)	bring	brenge[n] [brɛŋən]	bringen [brɪŋən]
	sleep	slape[n] [slapən]	schlafen [ʃlafən]
	ship	schip [sxɪp]	Schiff [ʃɪf]
	sister	zuster [zɪstər]	Schwester [ʃwestə]
	good	goed [xʊt]	gut [gʊt]

This list can be made longer. It turns out that the correspondences are to a large degree systematic. It can be observed that for example word initial [p] in Dutch and English corresponds to German [pf], that initial [s] becomes [ʃ] before [t] and [p]. And so on. This has led to two things: the postulation of a language (of which we have no record!), called **Germanic**, a set of words together with morphology and syntax for this language, and a set of rules which show how the language developed into its daughter languages. In the present case, the fact that Dutch [p] corresponds to German [pf] is explained by the fact that Germanic (not to be confused with German) had a sound *p (the star indicates reconstruction, not that the sound is illegitimate). This sound developed word initially into [p] in Dutch and into [pf] in German. This is called a **sound law**. We may write it in the same way as we did in phonology:

(416) *p → p / #_____ Dutch

(417) *p → pf / #_____ German

Often one simply writes *p > pf for the sound change. The similarity is not accidental; in the case of phonological rules they were taken to mean a sequential process, a development from a sound to another in an environment. Here a similar

interpretation is implied, only that the time span in which this is supposed to have taken place is much longer, approximately two thousand years!

The list of Germanic languages is long. Apart from the ones just listed also Danish, Swedish, Norwegian, Faroese, Icelandic, Frisian and Gothic belong there. Gothic is interesting because it is a language of which we only have written records, we do not know exactly how it was spoken.

As with the Germanic languages, similarities can be observed between French, Spanish, Italian, Rumanian and Portuguese. In fact, all these languages come from a language which we know very well: Latin. The development of Latin into these languages is well documented in comparison with others. This is important, since it allows to assert the existence of a parent language and changes with certainty, whereas in most cases the parent language has to be constructed from the daughter languages. This is so, for example, with **Celtic**, from which descended Gaelic, Irish, Welsh, and Breton. Cornish and Manx are also Celtic, but became extinct in the 19th century. A Celtic language, Gaulish, was spoken in the whole of France, but it was completely superseded by Latin. We have records of Gaulish only in names of people and places. For example, we know of the Gaulish king Vercingetorix through the writings of Caesar. The name is a Celtic name for sure.

Throughout the 19th century it became apparent that there are similarities not only between the languages just discussed, but also between Germanic, Latin, Greek, Celtic, Sanskrit, Old Persian, Armenian, Slavic, Lithuanian, and Tocharian. It was proposed that all these languages (and their daughters, of course) descend from a single language called **Indo-European**. When people made excavations in Anatolia in the 1920s and found remains of Hittite, it was recognised that also Hittite belongs to this group of languages. During the last 200 years a lot of effort has been spent in (re?)constructing the sound structure, morphology and syntax of Indo-European, to find out about the culture and belief and the ancient homeland of the Indo-Europeans.

The time frame is roughly this: the Indo-European language is believed to have been spoken up to the 3rd millennium BC. Some equate the Indo-Europeans with people that lived in the region of the Balkan and the Ukraine in the 5th millennium BC, some believe they originate further north in Russia, other equate them with the Kurgan culture, 4400 – 2900 BC, in the south of Russia (near the Caspian sea). From there they are believed to have spread into the Indian subcon-

continent, Persia and large parts of Europe. The first to arrive in central Europe and Britain were the Kelts who established a large empire only to be topped by the Romans and later by the Germans.

How the Language Looked Like

The sounds are believed to be these. Consonants are

		unaspirated		aspirated	
		voiceless	voiced	voiceless	voiced
(418)	velar	*k	*g	*kh	*gh
	palatal	*k̂	*ĝ	*k̂h	*ĝh
	apico-dental	*t	*d	*th	*dh
	labial	*p	*b	*ph	*bh

Other people assume instead of the palatals a series of labiovelars (*k^u, *g^u, *k^{uh}, *g^{uh}). The difference is from an abstract point of view irrelevant (we do not know anyway how they were exactly pronounced...) but it makes certain sound changes more likely. Another set is *y, *w, *r, *l, *m and *n, which could be either syllabic or non-syllabic. Syllabic *y was roughly [i], nonsyllabic *y was [j]. Likewise, syllabic *w was [u], nonsyllabic *w was [w]. The nonsyllabic *r was perhaps trilled, nonsyllabic *m and *n were like [m] and [n]. Syllabic *l was written ʎ, similarly m̥ and n̥. The vowels were *i (= syllabic *y), *e, *a, *o and *u (= syllabic *w).

Here are some examples of roots and their correspondences in various I–E languages:

*w_lk^uos ‘wolf’. In Latin we find *lupus*, in Greek *lykos*, in Sanskrit *vṛkaḥ*, Lithuanian *viľkas*, in Germanic *wulfaz, from which English *wolf*, and German *Wolf* [wɔlf].

*dek_m ‘ten’. In Sanskrit *daśa*, Latin *decem*, pronounced [dekem] or even [dekē], with nasalized vowel, in Greek *deka*, Germanic *tehun, from which Gothic *taihun*, German *zehn* [tse:n], and English *ten*.

Here is an example of verbal conjugation. Indo–European is believed to have had not only singular and plural but also a dual (for two). The dual was lost in

Latin, but retained in Greek and Sanskrit. The root is *bher ‘to carry’.

		Sanskrit	Greek	Latin
(419)	Sg	1 bhar-ā-mi	pher-ō	fer-ō
		2 bhar-ā-si	pher-eis	fer-s
		3 bhar-ā-ti	pher-ei	fer-t
	Du	1 bhar-ā-vaḥ	–	–
		2 bhar-ā-thaḥ	pher-e-ton	–
		3 bhar-ā-taḥ	pher-e-ton	–
	Pl	1 bhar-ā-maḥ	pher-o-mes	fer-i-mus
		2 bhar-ā-tha	pher-e-te	fer-tis
		3 bhar-ā-nti	pher-o-nti	fer-u-nt

To be exact, although Latin *fero* has the same meaning, it is considered to belong to another inflectional paradigm, because it does not have the vowel ‘i’. In Attic and Doric Greek, the 1st plural was *pheromen*. Thus, there has been a variation in the endings.

The verb *bher is also found in English in the verb *bring* (often, root vowels become weak, giving rise in the case of *e to a so-called ‘zero’-grade *bhr̥).

How Do We Know?

The reconstruction of a language when it is no longer there is a difficult task. One distinguishes two methods: comparison between languages, and the other internal reconstruction. The latter is applied in absence of comparative evidence. One observes certain irregularities in the language and proposes a solution in terms of a possible development of the language. It is observed, for example, that irregular inflection is older than regular inflection. For example, in English there are plurals in *en* (*oxen*, *vixen*) and plurals in vowel change (*women*, *mice*). These are predicted by internal reconstruction to reflect an earlier state of the language where plural was formed by addition of *en* and vowel change, and that the plural *s* was a later development. This seems to be the case. Likewise, this method predicts that the comparative in English was once formed using *er* and *est*, but at some point got replaced by forms involving *more* and *most*. In both cases, German reflects the earlier stage of English. Notice that the reasoning is applied to English as it presents itself to us now. The change is projected from present day English. But how can we ascertain that we are right?

First and foremost, there are written documents. We have translations of the bible into numerous languages (including medieval Georgian (a Caucasian language)), and we have an Old English bible (King Alfred's bible), and a Gothic bible, for example. Latin and Greek literature has been preserved to this day thanks to the effort of thousands of monks in the monasteries (copying was a very honorable and time consuming task in those days). Also other languages have been preserved, among which Avestan, Sanskrit and Hittite (written mostly in cuneiform). The other languages got written down from the early middle ages onwards, mostly in the form of biblical texts and legal documents. Now, this provides us with the written language, but it does not tell us how they were spoken. In the case of Hittite this is very obvious: the writing system was very different from ours and it had to be discovered how one had to read it. For Sanskrit we know from the writings of the linguists of those days, among which Paṇini (500 BC) is probably one of the latest, how the language was spoken. This is because we have explicit descriptions from them of how the sounds were produced. For Latin and Greek matters are less easy. The Greeks, for example, did not realize the distinction between voiced and voiceless consonants (they knew they were different but couldn't say what the difference was). In the middle ages all learned people spoke Latin, but the Latin they spoke was very different from classical Latin, both in vocabulary and pronunciation. By that time, people did not know how things were pronounced in the classical times (= first century BC). So how come *we* know?

One answer is: through mistakes people make when writing Latin. Inscriptions in Pompeii and other sites give telling examples. One specific example is the fact that *m* after vowels was either completely lost or just nasalized the preceding vowel. One infers this from the fact that there are inscriptions where one finds *ponte* in place of *pontem*. Also, in verse dating from that time the endings in vowel plus *m* counted as nonexistent for the metre. (This in turn we know for sure because we know what the metre was.) This is strong evidence that already in classical times the final *m* was not pronounced. The next method is through looking at borrowings into other languages. The name Caesar (and the title that derived from it) was borrowed into many languages, and appears in the form of *Kaiser* in German, taken from Gothic *kaisar*, and *césar* [sɛza:r] in French. So, at the time the Goths borrowed the *c* it was pronounced [k]. And since French descends from Latin we must conclude that the Gothic borrowing is older. Moreover, there was a diphthong. The diphthong was the first to disappear, becoming plain long [e:], and then [k] changed into [s] in French and [tʃ] in Italian and Rumanian. A third source is the alphabet itself. The Romans did not distinguish *v*

from u. They wrote v regardless. This shows that the two were not felt to be distinct. It is unlikely that v was pronounced [v] (as in English vase). Rather, it was originally a bilabial approximant (the nonsyllabic *w mentioned above), which became a labiodental fricative only later.

Historical explanations are usually based on a lot of knowledge. Languages consist of tens of thousands of words, but most of them are not indigenous words. Many words that we use in the scientific context, for example, come from Latin and/or Greek. Moreover, they have been borrowed from these languages at any moment in time. The linguistic terminology (**phoneme**, **lexeme**) is a telling example. These words have been artificially created from Greek source words. Learned words have to be discarded. Another problem is that words change their meaning in addition to their form. An example is German *schlimm* ‘bad’, which originally meant ‘inclined’. Or the word *vergammeln* ‘to rot’, which is from Scandinavian *gamall* ‘old’. English *but* derives from a spatial preposition, which is still found in Dutch *buiten* ‘outside’. From there it took more abstract meanings, until the spatial meaning was completely lost. If that is so, we have to be very cautious. If meanings would be constant, we could easily track words back in time; we just had to look at words in the related languages that had the same meaning. But if also the meaning can change — what are we to look out for? Linguists have done a lot of effort in determining in which ways the meanings of words can go and which meanings are more stable than others.

Two Examples Among Many

As an example of the beauty and danger of historical linguistics, we give the history of two words.

The first example shows that once we know of the rules, the resemblances become very striking. The English word *choose* has relatives in Dutch and German. If we look at the verbs that these languages normally use, we might be disappointed: the Dutch word is *kiezen*, and the German is *wählen*. The relation between Dutch and the English are easier seen. First notice that the Dutch verb has a PPP *gekozen* ‘chosen’, which has the o in place of the i.e. The change from [e] to [o], called **Ablaut**, is widely attested in the Indo-European languages and is believed to be very old. Now, [k] often becomes [tʃ] before either [e] or [i] (like Latin [ke] became [tʃ] in Italian, often with loss of e in pronunciation).

However, this change occurred in English only, not in Dutch. However, we still have to see why it could occur in English. The Old English word in fact was *ceōsan*. (When no star appears that means that we have written records.) The pronunciation of *c* changed and incorporated the *e*, and the infinitive ending got lost like with other verbs. The German case seems hopeless. In fact, *wählen* does *not* come from a related root. However, in German we do find a verb *kiesen* in similar meaning, although it is now no longer in use. Strangely enough, there PPP of the verb *auserkiesen* (two prefixes added, *aus* and *er*) is still in use: *auserkoren* ‘chosen’. Notice that the ablaut is the same as is Dutch, which incidentally also uses the circumfix *ge- -en* as does German. Finally, in German the PPP has *r* in place of *s* (which would be pronounced [z]). The change from *s* to *r* in between vowels (called ‘rhotacism’) is a popular change. Latin has plenty of examples of this.

Now, the root from which all this derives is believed to be Germanic **keusa-* ‘to try out, choose’. Once we have progressed this far, other words come into sight: Latin *gustare* ‘to taste’ (from which via French English got *dis-gusting*), Greek *geuomai* ‘I taste’, Sanskrit *juṣāti* ‘he likes’, Old Irish *do-goā* ‘to choose’. From this the root **geus* ‘to choose’ has been proposed. The Latin word presents the zero-grade **gus*. In West Germanic we have *kuz-i* ‘vote, choice’. But beware: French *choisir* does *not* come from Latin — there is no way to explain this with known sound laws. For example, it cannot be derived from *gustare*, known sound laws predict *goûter*, and this what we find. Instead it was taken from — Gothic! Indeed, French has taken loanwords from Germanic, being occupied/inhabited in large parts by Germanic tribes. The name France itself derives from the name of a tribe: Frankon.

With respect to Dutch and German the reconstruction is actually easy, since the languages split around the 17th century. Certain dialects of North Germany still have [p] where others have [pf]. It often happens that a word is not attested in all languages. For example, English *horse* corresponds to Dutch *paard* and German *Pferd*, with same meaning. The sound laws allow to assert that *paard* and *Pferd* descend from the same word, but for English *horse* this is highly implausible. There is a German word, *Ross* ‘horse’, and a Dutch *ros*, which sound quite similar, but they are far less frequent. What we have to explain is why the words are different. Now, we are lucky to have source confirming that the word was sometimes spelt *hros* sometimes *ros* in old German. In Icelandic, it is still *hross*. The loss of *h* before *r* is attested also in other cases, like the word

ring. But the change happened in English, too, and the only reason why the h was preserved in *horse* is that the r changed places with o.

Finally, where did the Dutch and German get their words from? Both words come from the same source, but it is not Germanic, it is Medieval Latin *paraveredus*, which in turn is *para* + *veredus*. *para* comes from Greek (!) *para* ‘aside’, and *veredus* is Celtic. In Kymric there is a word *gorwydd* ‘horse’. So, *paraveredus* had a more special meaning: it was the horse that was running on the side (to replace the one that one is riding on when it gets tired). The Indo–European root that has been constructed is *ek^wos. Latin *equus*, Greek *hippos* and Sanskrit *aśvaḥ*. Had we not known that the Latin word is *equus*, we would have had to guess from French *cheval* and Spanish *caballo*.

Thus, roots do not survive everywhere. Words get borrowed, changed, returned and so on. There are not too many roots that are attested in all languages, mostly kinship terms, personal pronouns and numbers.

Other Language Families

In addition to Indo–European, there is another language family in Europe: the **Uralic** language family. The languages that are said to belong to that family are Finnish, Estonian, Lappish, Hungarian and a number of lesser known languages spoken in the north of Russia. The affiliation of Hungarian is nowadays not disputed, but in the 19th century it was believed to be related to Turkish. Unfortunately, the written records of these languages are at most 1000 years old, and the similarities are not always that great. Finnish, Estonian and Lappish can be seen to be related, but Hungarian is very much different. This may have to do with the fact that it was under heavy influence from Slavic, Turkish (the Turks occupied Hungary for a long time) and Germanic (not the least through the Habsburg monarchy).

The affiliation of Basque is unknown.

Other recognized language families are: **Semitic** (including Hebrew, Ethiopic, Amharic, Aramaic and Arabic), **Altaic** (Turkish, Tungus and Mongolian), **Dra-vidian** (spoken in the south of India: Tamil, Telugu, Kannada and Malayalam), **Austronesian** (Malay, Indonesian, Malagassy, languages spoken on Macronesia, Micronesia and Polynesia), **Eskimo–Aleut** (Inuit (= Eskimo), indigenous lan-

guages spoken in Canada, Greenland, Western Siberia and the Aleut Islands). The list is not complete. Sometimes languages are grouped together because it is believed that they are related, but relationships are actually hard to come by. This is the case with the **Caucasian** languages (Georgian and many others). It is believed that the people who live there have not moved for several millenia. This has given rise to a dazzling number of quite distinct languages in a relatively small area in and around the southern part of the Caucasian mountains.

Probing Deeper in Time

It has been tried to probe deeper into the history of mankind. One way to do this is to classify people along genetic relations (a large project with this aim has been led by Cavalli-Sforza), another has been to establish larger groupings among the languages. Although genetic relationships need not coincide with language relationships, the two are to a large degree identical. Two hypotheses are being investigated starting from Indo-European. Joseph Greenberg proposed a macrofamily called **Eurasiatic**, which includes Indo-European, Uralic, Altaic, Eskimo-Aleut, Korean, Japanese and Ainu (spoken in the north of Japan) and Chukchi-Kamchatkan. It has been suggested on the other hand by Russian linguists that there is an even larger macrofamily called **Nostratic**, which originally was believed to include all Eurasiatic families, Afro-Asiatic (spanning north Africa including Semitic), Dravidian and Kartvelian (= Caucasian). Greenberg does not reject the existence of Nostratic but wants to put it even farther back in history, as a language which developed among other into Eurasiatic. Moreover, the position of modern Nostraticists has come closer to that of Greenberg's views. At a larger scale there are believed to be twelve such macrofamilies in this world all ultimately coming from a single language ... Examples of words that are believed to have been passed to us by the single ancestor language are *aq'wa* 'water', *tik* 'finger', and *pa1* 'two'. This is highly speculative, but the evidence for Eurasiatic is not as poor as one might think. Common elements in Eurasiatic are for example: first person in *m*, second person in *t/n*. Greenberg has collected a list of some 250 words or elements that are believed to be of Eurasiatic origin.