

The Shape of Information

The Shape of Information

How Language Gets Hold of the World

Net Koene

Stichting Neerlandistiek VU Amsterdam
Nodus Publikationen Münster
2007

Bibliographic information published by Die Deutsche Bibliothek

Die Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliographie;
detailed bibliographic data is available in the internet at <http://dnb.ddb.de>

Cahiers voor Taalkunde 27

Stichting Neerlandistiek VU
De Boelelaan 1105
NL-1081 HV Amsterdam

Nodus Publikationen
Klaus D. Dutz — Wissenschaftlicher Verlag
Postfach 5725
D-48031 Münster

© copyright 2007 Net Koene, Amsterdam
No reproduction without permission
All rights reserved

ISBN: 978-3-89323-533-9 (Nodus Publikationen)
ISBN/EAN: 978-90-8880-005-4 (Stichting Neerlandistiek VU)

For Bram, Sarah and Heleen

Contents

Introduction: The Aim of this Book	1
Chapter 1: One Structure, a Multiplication of Meaning	5
1.1 Noun Phrase Plus Pitch: Referring of Not Referring	5
1.2 The Shape of Information: Mental Images	11
1.3 Emphasis: Double Images	18
1.4 Noun Phrase Plus Rhythm: Single or Repetitive Images	25
Chapter 2: Language and World	29
2.1 The Human World	30
2.2 How Forms of Language Exist	39
2.3 Forms of Language and the Human World	41
Chapter 3: The Skeleton of Information	48
3.1 Understanding Information	50
3.2 Four Types of Information	54
3.3 Recognizing the Real Thing	57
Chapter 4: The Shape of Negative Information	72
4.1 Negative Predicates	72
4.2 Four Types of Negative Information	73
4.3 Negative Subjects	78
Chapter 5: The Shape of Emphasis	82
5.1 Pieces of Information	82
5.2 A Few Combinations	83
5.3 More Combinations	91
Chapter 6: More Noun Phrases	94
6.1 Definite Noun phrases	94
6.2 Russell 1905	103
6.3 Proper Names	104
6.4 Material Noun Phrases	106
Chapter 7: Some Philosophical Issues	108
7.1 Awareness and Self	108
7.2 Subjectivity	108
7.3 Verbal Information and the Notion of Truth	109
7.4 Meaning, Representation, Reference	110
7.5 Empirical Semantics	112
Tribute	117
References	
Index	

Introduction: the Aim of this Book

Language gives people grip on their world, and apparently, they are able to share this grip. This is rather amazing: between speaker and hearer, there is nothing but sound waves, but to the hearer, these sound waves are the speaker's words, they are information, and he can recognize the reality talked about if it comes in sight.

Now, when we try to explain this, there is the complicating fact that already the simplest constructions can be used to call up different meanings; for example, an indefinite noun phrase, like *a parrot*, may in very different ways be connected to the world talked about: referring to a specific parrot, hovering above a number of parrots all equally coming into consideration (compare 'specific versus non-specific' use), or concerning each and every parrot (compare 'generic' use). And a simple combination of the noun phrase with a negative predicate seems to cause enormous complications: *not* may be understood to deny just the predicate, or the whole utterance, or it may seem to zoom in on an arbitrary small part of it. And definite subjects, proper names, material subjects, etc., seem to be equally problematic, each with their own ambiguities and nuances.

Remarkably, in actual communication these facts are not problematic at all; speakers and hearers are not aware of any such ambiguity, and the hearer just gets the meaning intended. The question we should ask is: exactly what goes on in the head of the hearer? In fact there are two questions: in the first place, what happens in the mind of the hearer when he turns the sound forms into understanding and when he recognizes the real thing, and in the second place, what is the kind of existence of whatever it is that goes on in the head of the hearer. These are the questions this book is meant to address. Its second chapter develops a theory about the brain and its mind in order to get a theoretical framework in which to explain the data. One might choose to read this second chapter first.

The first chapter aims to get a better grip on the data, focusing on simple sentences with an indefinite subject: what forms does the hearer recognize in the speaker's sequence of sounds and what piece of information does each of these forms contribute to his understanding. As a hearer myself, I have tried as closely as possible to stay with the hearer, each step of the way. To illustrate my approach I shall give one example. Let's take *A parrot is not for sale*, which may, among others, be understood as 'some particular parrot is not for sale', and as 'no parrot is for sale'. Usually, this difference is described in terms of the 'scope' of *not*, as 'predicate negation' versus 'sentence negation', and formal semantics tries to capture their different connections with the world in terms of truth conditions: when there are two parrots in a pet shop, one with a price tag on its cage, and the other with a sign 'Benny the shop's pet', given the first understanding the sentence is true, and given the second understanding it is false. In my opinion however, neither the notion of 'scope' nor the notion of 'truth conditions' has a place in the explanation of these facts, because these notions do not capture what goes on in the head of the hearer. In the process of understanding there is no intermediate stage where the hearer reverses the order between the elements, and the process of connecting the information to reality has nothing to do with truth conditions. Rather than arguing against these ideas, in this book I offer an alternative and a better explanation.

The basic idea is very simple and close to common sense: the speaker's sounds call up in the hearer's mind an idea of what the thing talked about would look like, feel like, sound like, smell like, etc., if it were in view, but only insofar as the information goes; the idea remains

fragmentary. We might just as well call this idea a mental image; problems one might have with this notion can be dealt with. Each form heard in the speaker's sequence of sounds systematically contributes its fragment to the image taking shape. Listen to a speaker's isolated *a parrot* and notice what it evokes: obviously it does not include a claim of existence of some parrot; what it does call up is just a dim parrot-shaped outline.

Listening to the sounds of language we can hear pitch and rhythm making a difference. A small number of basic informative elements is conjectured, among which a minimal pitch pair. Listen again to *a parrot*: if in its final syllable its intonation meets the low base level, it makes the hearer's parrot-shaped outline grow into a complete parrot; although its details remain unspecified, the hearer understands them to be the factual details of some actual parrot in the world talked about. However, if *a parrot*'s final syllable remains above base level, the outline remains an open sketch that still allows for any combination of complementary features. When a predicate is added, it calls up another feature-image and adds this fragment to the image-under-construction. If a positive predicate feature is added to a complete parrot, it specifies one of its hitherto unspecified details, and if added to an open outline it fills up a gap, thereby narrowing down the range of combinations of features still allowed for. If the predicate is negative, its feature is added as an example of what is to be excluded from the world talked about. If it is added to a complete parrot, it is excluded from its unspecified but determined other features ('predicate negation'). If it is added to an open outline, this feature is excluded from the range of combinations still allowed for ('sentence negation').

The image called up is the hearer's idea of what the real situation would be like if it came in view. As long as the situation talked about is out of sight, to the hearer the image stands-in for it. If the situation talked about indeed comes in view, in case the information is positive the hearer expects his image to fuse with the real thing. In case the information is negative the hearer expects his image of what is absent not to clash with what is in fact there. Let's suppose that the hearer enters the pet shop; just given the noun phrase *a parrot*, he sees two parrots that equally come into consideration, if he hears *a parrot* (meeting base level) *is not for sale*, the image of a complete parrot fuses with Benny, as the unnamed one is no option. But if he hears *a parrot* (above base level) *is not for sale*, the open outline with the excluded predicate feature clashes with the unnamed parrot. This example illustrates my approach.

The isolated minimal pitch pair determines only one place in the actual intonation contour; the hypothesis predicts that all possible intonation contours meeting the base line of pitch at the relevant point will result in a 'predicate negation' understanding, and all others will result in a 'sentence negation' understanding.

As said above, in chapter 1, I isolate a small (not exhaustive) number of basic informative elements in pitch and rhythm that come along with the speaker's sounds. The subject noun phrase goes with two minimal pairs, both located in its final syllable: next to the pitch pair sketched above a rhythmic pair that distinguishes single and generic meanings; two minimal pairs result in four possible combinations. Apart from that, there may be an optional extra high peak in pitch, that may single out any small or larger part of the utterance. With each location, the number of possible combinations of prosodic form doubles. All these prosodic forms are independent of, and freely combinable with, any distribution and realization of pitch accent. Each isolated form systematically makes its own contribution. Together, they predict a range of possible messages that is as large as the number of possible combinations of basic prosodic forms. This range includes not only the cases sketched, but also the case where *not* seems to zoom in on a small part of the utterance. It includes the 'specific versus non-specific' difference in positive

utterances. It predicts a subtle but unmistakable nuance in utterances with a negative subject. It includes ‘generics’, and it correctly predicts two very different kinds of generic statements. Also, it predicts a range of other nuances, some examples of which are given.

The notion of mental images may seem problematic. The often mentioned ‘homunculus problem’, the little man in the head supposedly necessary for looking at mental pictures, is a pseudo problem, but it is a fact that people’s images are fleeting and that they cannot be observed by anyone else, and therefore, they may seem to form a weak link between solid forms of language and a solid world. This is a misunderstanding. Chapter 2 aims to get hold of not only these images but also of the forms of language heard and the world in the hearer’s sight. The hardest problem is that of the world: how does the hearer know at all of the world around him, the pet shop, the parrots? No doubt he needs a brain, but not just a brain enabling him to interact with his surroundings, but a brain that makes him know what it is that he is interacting with. In other words, he needs a brain that makes him aware of it all. Although hard to understand, it seems reasonable to accept as a fact that the parrots show up in his awareness as unifications of phenomenal features precisely during some special coordinated activity in his brain in networks of neurons in different specialized areas, triggered by the input on his senses originating from the parrots.

The same goes for the sound forms of language: the hearer is aware of the word *parrot* precisely during this special activity in the network of neurons in areas specialized in language triggered by the speaker’s sequence of sounds.

Both the forms of language and the phenomenal world exist to him, and exclusively to him, as long as this special activity goes on in the corresponding networks of neurons in his brain. None of it survives fading from his awareness when this activity in these networks subsides. His brain does not make a copy to be kept in store, nor something that, without causing serious misunderstanding, could be called a representation. What happens is this: the mechanisms in the brain keep together the particular combination of neurons used in making it show up by strengthening their connections. This particular combination is ready to flash into action again upon a small input. In effect, these strengthened complex networks record how to call the phenomenal parrots, as well as the word *parrot*, into existence again any moment. Someone’s life of experiences leaves a growing complex of such records, overlapping and criss-crossing and changing with new experience, connections getting strengthened and weakened. Thanks to these records people experience both their world and the speaker’s language as familiar.

In the absence of external signals these ready-to-use networks of neurons may be activated from inside the brain as well, and if someone’s record of a parrot becomes active, automatically the parrot shows up to him, as-if-there: an image of the parrot. Homunculus exit.

However, there is a real difficulty we have to overcome when we use mental images in a theory of communication. The problem is not that images are subjective; everything that shows up in someone’s awareness is subjective, language and the phenomenal world just as well as images. However, when their appearance is triggered by an external input, people take what shows up to themselves as exemplary for what, given the same input, would show up to someone else. This is how people experience both their forms of language and their phenomenal world as something they share with others, their world synchronized by signals from nature, and words and prosody synchronized by sounds they can produce themselves. The problem with images is that they are not synchronized by an external trigger, and therefore, people cannot take their own images as exemplary.

It is here that language comes in. Words by themselves are arbitrary sound forms. In someone's head, they get tagged on aspects of the world when they simultaneously show up in this person's awareness. In time, a more or less stable record of the word *parrot* will get linked to an ever more branching complex of records of experiences of parrots. And now, a speaker's sequence of sounds that calls *parrot* to the hearer's awareness accesses this branching complex of networks linked to it. From this point of departure the action may flash back and forth at random, making a growing and shrinking image take shape, morphing from one appearance into another, one shape into another, one colour into another, one parrot into another, as long as the starting point in the speaker's word is kept in mind. Other than the mental pictures of his own imagination, the hearer may, up to a point, take these word-triggered images as exemplary for what, given the same input, would show up to someone else: it is at least close enough for them to recognize the same real things as fitting if they came in view.

So, although a speaker cannot actually squeeze meaning into his sounds and send them along to the hearer, his sounds can take remote control over the already existing mechanisms in the hearer's brain that have been evolved to make the world take shape around him. This is the magic of language.

The question addressed in this book is: what happens in the head of the hearer when he understands the speaker's information and recognizes the real thing talked about. Elaborating on chapter 2, in chapter 3 thru 5, a notation is developed to represent each step of the hearer's process: the sound forms, the fragmentary images under construction and the world in his view as each of them shows up in the hearer's awareness during this special activity in the corresponding neural networks. The notation developed to represent images and world is analogous: it shows the fragmentary images called up to the hearer insofar as they are authorized by the speaker, it shows what room is left for the hearer's imagination and it represents the hearer's world as and insofar as it comes in his view and insofar as it has his attention. It shows where image and world may fuse and where they clash, and it shows partial and complete recognition. These representations are not supposed to be present as such in the hearer's head, but to linguists they hypothetically represent what in fact takes place in the hearer's head.

In chapter 6, I briefly show that other kinds of noun phrases, like definite noun phrases, proper names and material noun phrases allow for the same range of combinations of prosodic form, resulting in the same range of possible meanings, the whole range accounted for by exchanging the information contributed by an indefinite noun phrase for the information contributed by a definite noun phrase, a proper name or a material noun phrase. This explains some well-known facts and correctly predicts a range of further nuances and ambiguities.

The book concludes with some remarks about awareness, truth, subjectivity, meaning and representation, and a discussion on linguistics as an empirical investigation.

Chapter 1

One Structure, a Multiplication of Meaning

A speaker has nothing but sounds to get his message across, and it works. How? Take the perspective of a hearer: what forms does he recognize as informative in the sound waves that enter his ears, what piece of information does each of these forms contribute, how does he put them together, how does he recognize the real thing talked about if it comes in sight? The aim of this book is to come to understand what actually happens here.

There is an interesting complication: a speaker can use even the simplest subject-predicate construction to send a whole range of different messages. It works: the speaker gets his information across. My investigation starts here: how does it work? In chapter 1 a first exploration.

1.1 Noun Phrase plus Pitch: Referring or Not Referring

1.1.1 ‘Predicate Negation’ versus ‘Sentence Negation’

Take the following example:

(1) *A sister of John’s is not interested.*

One may understand (1) as a statement about some sister: she is not interested, or one may understand that no sister is interested. With context:

(1/1) *A sister of John’s is not interested,
(but she will have a look anyway)
(another sister however is very keen)*

(1/2) *A sister of John’s is not interested,
(except for his only brother, his siblings do not care)*

In (1/1), *a sister of John’s* seems to be not included in the denial of *not* (‘predicate negation’), and in (1/2), it seems to be included (‘sentence negation’), so the difference might seem to be a matter of the ‘scope’ of *not*. However, the question is: what is actually happening when a hearer turns the forms he recognizes in a speaker’s sequence of sounds into understanding, what exactly are the separate pieces of information, and how do they add up to ‘predicate negation’ in case (1/1), and to ‘sentence negation’ in case (1/2). Now, if I am correct, the difference comes with the use of the subject noun phrase; in case (1/1), one understands *a sister of John’s* to refer to a particular (albeit not uniquely identified) sister, and as a consequence, one understands the rest of the sentence to be a (negative) statement about that sister; in case (1/2), one does not understand *a sister of John’s* to refer to a particular sister, and as a consequence, *not interested* is not about such a sister. The difference does not depend on the presence of *not*, or its ‘scope’, but if *not* is present, it reinforces a difference already there.

One of the first questions now coming to mind is: if a speaker can use (1) to get two different messages across, what is there between speaker and hearer to make the difference? There is no audible difference in words or syntax. If I am correct, this is how it works: if the last syllable of the noun phrase is spoken in a low pitch, it is used to refer, and if it is kept somewhere up in the air, it is not used to refer. More precisely, in ‘Dutch School’ terminology (see Tribute for references):

HYPOTHESIS

Between the top of the final syllable of a noun phrase and its end, the pitch contour either meets the ‘low declination line’ at least at one point (possibly followed by a very late rise to mark the phrase boundary), or it remains above it. The hearer hears a difference.

NOTATION

↓ pitch contour meeting the low declination line
 ↑ pitch contour not meeting the low declination line

THE FIRST IDEA

↓ makes the noun phrase refer to something that is there
 ↑ makes the noun phrase not refer to something that is there

if pitch makes a subject noun phrase refer (↓), it will not be understood as included in the denial of a negative predicate;

if pitch makes a subject noun phrase not refer (↑), it will be understood as included in the denial of a negative predicate.

From here, in the notation of forms of language, I shall drop punctuation marks and initial capitals of orthography, as their relationship to what is audible is not always clear. Thus:

(1/1) *a sister of John's[↓] is not interested* about some sister, ‘predicate negation’
 (1/2) *a sister of John's[↑] is not interested* not about some sister, ‘sentence negation’

This still leaves many questions unanswered, but already these conjectures can be put to a preliminary test. Both ↓ and ↑ can be realized in a large number of very different pitch contours, they are independent of both the distribution and the realization of pitch accent and can be combined with any of them. The conjectures predict that (1), with any pitch contour that goes with ↓, will be understood to be about some sister (‘predicate negation’), and with any pitch contour that goes with ↑, will be understood not to be about some sister (‘sentence negation’).

Before I can illustrate these predictions, a few words about pitch accent. ‘Dutch School’ investigations have demonstrated that what one hears as accent is realized by a steep pitch movement properly located at the top of the syllable; this can be a steep rise, a steep fall, both of which one can hear as an accent that comes out on a different pitch, or it can be a steep rise immediately followed by a steep fall; in this case one does not hear any change in pitch, but just an accent. Gradual pitch movements as well as steep movements on other locations are not heard as accents. Pitch accent as well as its absence at suitable locations is informative in its own right; apparently it has something to do with first time or repeated mentioning. I shall return to that.

NOTATION

| as in pitch accent, realized by a steep pitch movement properly located on top of the syllable
s|*ister*


I shall represent pitch accent only if it is relevant in the discussion.

Now take an arbitrary distribution of pitch accent, for both cases the same, for instance, add one accent in the subject, take s|*ister*, and one accent in the predicate, take |*interested*:

(1/1) a s|*ister* of John's[↓] is not |*interested* about some sister, 'predicate negation'
(1/2) a s|*ister* of John's[↑] is not |*interested* not about some sister, 'sentence negation'


Both cases still can be realized in different ways; at this point, I want to switch to Dutch intonation; as Dutch is my native language, my representation of actual contours can be more precise; Dutch intonation seems to have a smaller pitch range than (British) English intonation, but I believe the informative prosodic forms are the same for both languages. The notation here is based on 'Dutch School' notation; I have drawn what I hear in (1) translated into Dutch: *een z|uster van Jan^{↓/↑} is niet geïnteress|eerd*, but the predictions carry over to English intonation contours as well. Anyhow, given this in itself arbitrary distribution of pitch accent, the prediction for contours (a) thru (f) is as follows:

(a)



een z|uster van Jan[↓] is niet geïnteress|eerd 'about' ; 'predicate negation'

(b)




een z|uster van Jan[↑] is niet geïnteress|eerd 'not about'; 'sentence negation'

Contour (a) very clearly makes (1) a statement about some sister: start low, give *zuster* (*sister*) a rising pitch accent, gradually fall to base level before or at *Jan* (*John's*), which makes it low at the crucial point, and then at the very last moment give *Jan* (*John's*) a late rise, which indicates a phrase boundary, immediately resume pitch at base level, and give *geïnteress|eerd* (*interested*) the rising-falling pitch accent that sounds like just an accent coming out on the same pitch. And contour (b) very clearly makes (1) be understood as 'sentence negation': start low, give *zuster* (*sister*) a rising pitch accent again, but now keep the contour high across the subject-predicate boundary, and then come down with a falling pitch accent on *geïnteress|eerd* (*interested*) again.

With the same distribution of pitch, two other interesting contours to compare are


(c)



een z|uster van Jan[↓] is niet geïnteress|eerd

‘about’; ‘predicate negation’

(d)




een z|uster van Jan[↑] is niet geïnteress|eerd

‘not about’; ‘sentence negation’

both start with the same rising-falling pitch accent, and both have a rise that is not an accent between the two pitch accents, still (c) goes with ‘predicate negation’, and (d) goes with ‘sentence negation’. (c) can do without this rise:

(e)




een z|uster van Jan[↓] is niet geïnteress|eerd

‘about’; ‘predicate negation’

The late rise in (a) and (c) marks the noun phrase boundary and goes together with [↓]; together with [↑], the phrase boundary can be marked as follows:

(f)




een z|uster van Jan[↑] is niet geïnteress|eerd

‘not about’; ‘sentence negation’

I believe in all these cases the predictions are correct. Changing anything about the accents, adding accent, relocating them, dropping one or even both accents, placing an accent on *not*, does not make any difference to the predictions; given the hypothesis, the only factor that matters is pitch at the final syllable of the noun phrase.

Given the illustrations, it is clear that it is not the distribution or realization of pitch accent that makes the difference between the two understandings. Also, the overall contour only is important insofar as it determines pitch at the relevant point; for instance, take the very frequent ‘hat’ contour of (b), shift its accents to any location, for instance its rise from *zuster* (*sister*) to *Jan* (*John*) to *niet* (*not*), and its fall from *geïnteresseerd* (*interested*) to *niet* (*not*) to *Jan* (*John*), in any combination that still makes a ‘hat’, and then it is exactly when this whole ‘hat’ is pushed to either side of the noun phrase boundary that ‘sentence negation’ reverses to ‘predicate negation’:


(g)



een zuster van J|an[↑] is niet geïnteress|eerd

‘not about’; ‘sentence negation’


(h)



een z|uster van Jan[↑] is n|iet geïnteresseerd

‘not about’; ‘sentence negation’


(i)



een zuster van J|an[↑] is n|iet geïnteresseerd

‘not about’; ‘sentence negation’


(j)



een z|uster van J|an[↓] is niet geïnteresseerd

‘about’; ‘predicate negation’

(k)



een zuster van Jan[↓] is n|iet geïnteress|eerd

‘about’; ‘predicate negation’

Besides pitch accent, there are other informative prosodic forms enriching form and information, notably peaking pitch, and rhythm (see below). They are independent of the pitch pair discussed here, and do not influence the predictions about its information.

1.1.2 Speech and Writing

Calling (1) ambiguous would be missing the point. The fact is that written representations of language incompletely represent its informative form. In the following I would therefore like to make a few remarks on speech and writing.

It may seem that written text very well presents its information, and that pitch and rhythm just add some emotion, but I believe they have a systematic function not only in understanding speech but just as well in understanding writing.

If all goes well, a human child naturally learns to listen and speak; writing and reading does not come with human nature but with human culture; writing systems were developed over the centuries to preserve the information communicated by speech. A writing system is as good as its success in representing what is relevant in speech, that is, as good as it enables a skilled reader to reconstruct the original forms of language; writing systems can be seen as implicit theories of which forms are relevant in language; our alphabetic system is reasonably good in representing the verbal part, but obviously it is far from perfect in representing prosodic form.

Reading, both aloud and silent, is the skill of reconstructing the informative forms from the graphs. Now, if one reads (1) aloud, one needs to add pitch and rhythm to make sense at all, and this constructed form is not ambiguous. And to make sense of (1) in silent reading, one needs the information that comes with the intonation anyhow, so, where it is not in the representation written down, the reader at the spot has to decide at random whether or not to understand the noun phrase to refer to something. The same goes for information contributed by other prosodic forms.

Readers, just as well as hearers, end up with one meaning; they do not choose between alternative finished products of understanding, but they make their choices in the course of composing the one understanding they arrive at. So in reading there is a chance of misunderstanding that in listening is not there; on the other hand, by its permanence, writing has a compensation, albeit an unsystematic one: from writer to reader subtle nuances just get lost, but if a reader happens to end up with an understanding that turns out to be inconsistent with context or word meaning, he can start over and try again. It is only in these cases, which do happen once in a while, that an everyday reader may become aware of the fact that there is more than one possible reading for cases like (1) on paper.

To say that it is the context that disambiguates would be missing the point. The fact is that context unsystematically may rule out incorrect reconstructions of those systematic forms of language that are neglected in writing.

Where in this book I refer to the hearer, this must be taken to include the reader listening to his own reconstruction of the sound form.

1.1.3 Parallel Data

If a noun phrase systematically can be used either to refer or not to refer to something that is there, there is no particular reason why that phenomenon would be limited to sentences with a negative predicate, like (1). Given the conjectures, parallel data can be expected for positive sentences, and indeed they can be found. Compare (1) with its positive counterpart (2): the so-called ‘specific versus non-specific’ distinction.

- (1) *a sister of John's is not interested*
 (2) *a sister of John's is interested*

a sister of John's[↓] makes both (1) and (2) into a statement about some sister of John's:

- (1/1) *a sister of John's*[↓] *is not interested*
 (but she will have a look anyway)
 (another sister however is very keen)
 (2/1) *a sister of John's*[↓] *is interested*
 (but she cannot come)
 (another sister however could not care less)

And given *a sister of John's*[↑], both the negative and the positive case are not about a particular sister; (1/2) describes a negative fact, and (2/2) describes a positive fact:

- (1/2) *a sister of John's*[↑] *is not interested*
(except for his only brother, his siblings do not care)
 (2/2) *a sister of John's*[↑] *is interested*
(but I do not know which one)

Parallel data can also be found in case the subject noun phrase is negative. Compare:

- (3) *not a leaf*[↓] *stirred*
 (4) *not a blade of grass*[↑] *grew there*

In a way, (3) is about all the leaves, they all did not stir, and according to (4), there is no grass. Note that this cannot possibly be a matter of reversing the order between *not* and the noun phrase ('scope'). Or compare

- (5) *not a word*[↓] *was true*
 (6) *not a word*[↑] *was spoken*

(5) is about all that had been said, according to (6), there was silence. Just given the incomplete representation of form in writing, it may seem that it is the predicate that makes the distinction, but what the predicate in fact does is rule out an incorrect reconstruction of form. (7) is an example that on paper may seem to be ambiguous, but the sound is unambiguous:

- (7) *not an emotion showed on his face*

(7/1) *not an emotion*[↓] *showed on his face* refers to someone's emotions, the information is that he succeeds in hiding them; (7/2) *not an emotion*[↑] *showed on his face* says that there were no emotions to be seen on his face, perhaps because he did not have any feelings.

If a noun phrase systematically can be used either to refer or not to refer, there is no particular reason why that phenomenon would be limited to indefinite noun phrases. Parallel data can be expected for all kinds of noun phrases, like definite noun phrases (*the winner*), proper names (*Sophie*), and material noun phrases (*water*), both in positive and in negative sentences. Indeed they can be found (see chapter 6).

1.2 The Shape of Information: Mental Images

1.2.1 Next Questions

This is the conjecture: a hearer recognizes a combination of forms in the speaker's sequence of sounds. Each form contributes its own precise bit of information; together, they result in the instant end-product of understanding. Thus, in (1/1), the words *sister*, *John*, *interested*, the noun phrase that is built from *sister* and *John*, its indefiniteness, [↓], *not*, each make their own precise contribution, and together, in combination with the contribution of the predicate, the instant end-

product of understanding is ‘predicate negation’. And in (1/2), exactly the same combination except for \uparrow instead of \downarrow results in ‘sentence negation’. The inevitable next question is: what exactly are those separate pieces of information.

It is clear that there is no one-to-one connection between the phrases of language and the things talked about; evidently, we cannot take noun phrases as direct representations of real things. The combination *a sister of John's* \downarrow may seem to come close, it may seem to stand-in for the sister referred to, but the indefiniteness of the noun phrase immediately complicates matters: although the speaker uses it to refer to a specific sister, he might just as well use it to refer to another sister, and the hearer, on his side of the sound waves, cannot identify the particular one referred to when the sisters line up before him. The information that comes across is small but precise: it does not include identification of the specific sister referred to; in combination with more than one sister in view it is inconclusive. The question is: what exactly is the contribution of each of the forms, so that they add up to this partly specific partly inconclusive result.

The combination *a sister of John's* \uparrow is even more elusive; it is easier to say what \uparrow does not contribute than to say what it does contribute; *a sister of John's* \uparrow does not pick out someone fitting the description; listen to the positive (2/2) *a sister of John's* \uparrow *is interested*, it seems as if it just narrows down the hearer's attention to what might come into consideration. A variety of even more puzzling examples is easy to find:

- (8) *an ambulance* \uparrow *is on its way*
- (9) *one of them* \uparrow *is redundant*
- (10) *one of us* \uparrow *is terrible*
- (11) *a star* \uparrow *is born*
- (12) *a problem* \uparrow *arises*
- (13) *an answer* \uparrow *remains forthcoming*

(8) *an ambulance* \uparrow *is on its way*, on the scene of an accident, obviously is not about a particular ambulance; as far as the information goes, any ambulance will do. Given (9) *one of them* \uparrow *is redundant*, *one of them* fits each of the persons, but it does not settle down on one of them. The next example seems to be similar (this is an anecdote I heard years ago on BBC radio, but I do not remember the names of the people involved): one person does a rather unflattering imitation of someone else, not in time aware of that person's presence. When, next, they look each other in the eye, (10) *one of us* \uparrow *is terrible* is what the other says. In (11) *a star* \uparrow *is born* and (12) *a problem* \uparrow *arises*, *a star* does not refer to some star, and *a problem* does not refer to some problem; they only seem to activate the hearer's ability to recognize what might come into consideration. (13) *an answer* \uparrow *remains forthcoming* is even more puzzling; what does *an answer* say? Obviously not that there is something that fits the noun phrase, but still enough for the hearer to recognize it when it finally comes.

Whatever else we can say, there is the unmistakable fact that forms of language, words, phrases, sentences, do not mirror the logic of the real situations they seem to represent. Obviously, phrases do not directly stand-in for real things; what would *a sister of John's* \downarrow stand-in for? An unidentified sister? She is not exactly a logical beauty, as there are no such things in reality as sisters who may be different persons. Or *an ambulance* \uparrow , *one of them* \uparrow , *one of us* \uparrow , *a problem* \uparrow , *an answer* \uparrow , what would they stand-in for, an indiscriminate ambulance, an undecided one of them, one of us, a not yet quite existing problem, a non-existing answer? They are not

exactly logical beauties either. Now one thing has to be clear: as soon as we try to improve the forms of language so that they do mirror the logic of the situation, we lose track of what we want to explain. The problem is not the forms of language as they are; as they are, they work perfectly well; the most intriguing fact about all these cases is that, without any trouble, any ordinary hearer just gets the information. The questions to be asked are: what happens, how does he turn the forms of language as he picks them up into this instant understanding, and what exactly does each form contribute? At this point there are more questions than answers, but it is essential to be clear about what questions we want to address.

1.2.2 Understanding Information

These are not easy questions, for one reason because it is not at all clear what kind of thing verbal information is, and what it is to understand information. Now, whatever it is, information is not just a form; it only is information if it is about something, a person, a situation in the world, a vague idea, whatever, and to understand information is to grasp that thing. Then, to explain this understanding we do not primarily need to get hold of the things concerned as they physically and logically are; what we really need to get hold of is the hearer's grasp of those things. There are some clear facts to be accounted for:

- If the information regards a real situation, the hearer does not need its presence to get that grasp.
- A speaker's sequence of sounds can make a hearer grasp absent things as-if-there, in three dimensions, shape, colour, texture, sound, smell, etc.
- The hearer's grasp of the speaker's real thing most of the time remains fragmentary.
- Verbal information may go without any claim of existence of something that fits the information; whatever it is about may well be unrealized.
- However, if the information concerns some non-fictional and realized situation, if the real thing happens to come in sight and if the hearer gets a good enough look at it, he will recognize it as fitting the information. It is in this way that forms of language are connected to the real thing.

So we need a concept of information, understanding, and recognition that accounts for the hearer's grasp of the (realized, unrealized, real, fictional) things talked about both in their absence and in their possible presence. Now in the absence of the things talked about, hearers swap the forms they hear for an idea of what those things would look like, feel like, sound like, smell like, taste like, how they would handle them, etc., if they were factually there. Call it a mental image: whatever it is that actually shows up in the hearer's mind. In the presence of the real things, during experience with them, all this gets recorded, and in their absence, it can be called up again. Mental images have all we are looking for: other than words and phrases, images are thing-like, they can be made to act for what is not there as-if it were there, shape, colour, texture, everything, but other than real things they may stand-in for, which are bound to the laws of nature and logic, images have both the limitations and the freedom of the human mind: they are fragmentary, they can be tentative, they can morph into something different, and they do not depend on anything existing. And then again, they may seamlessly fuse with the real solid things talked about as soon as they happen to come in view, the mind's eye and the real eye, so to speak, brought into visual line.

I am talking about full-blown images: those private, phenomenal things as-if-there that show up in one's mind, including, if you like, 'qualia' of all sorts. This is not some naive idea; I know that images seem to be highly elusive, but also I am convinced that they are the only bridge there is between the forms of language, which do not have any intrinsic value, and the world. Indeed images are inescapably private, they are thing-like to the one to whom they show up, but they are not thing-like enough to make them perceptible to anyone else. The problem is methodological: when we propose hypotheses in which mental images play a major part, there may be a problem in distinguishing explanation from speculation. But if we want to explain the process of understanding verbal information, we cannot ignore what the understanding consists of; I believe what we need to do is to try to get hold of them in a methodologically satisfactory way despite their subjectivity, incorporating them in a theory of verbal communication, finding ways to distinguish between conjectures that cover and conjectures that fail to cover what actually happens in the process. I am convinced this can be done; in this chapter, assume that it can.

1.2.3 Positive Information

Take a person who looks at something in front of his eyes, say a cat. His perceptual system constructs it as a unification of features, shape, colour, texture, etc. Then take away the real cat but let this person keep in mind this unification of features, which is now an image of the cat, the cat as-if-there. Let the details fade one by one. At each stage, if the real cat is supplied again, the image of the cat and the cat in view will flow together again, but with each feature faded, the remaining image less and less determines what real cat it is an image of, and other things than the original one may show up for real that may fuse with it.

Also, in the absence of the real thing, a person can re-create by himself any feature of an earlier recorded experience, on its own, in old combinations, or in new combinations. These images shrink and grow between a tiny fragment on its own and a detailed picture, up to a complete unification of features, something virtually as tangible as the real thing, like in a dream or a hallucination — anything short of its substance. Imaging is how infants first conceive of what is not present, and it is only when people can conceive of what is not present, that they can go on and make arbitrary symbolic forms take its place.

The general idea of this book is this:

WORKING HYPOTHESIS

Verbal information calls up in the hearer's mind, piece by piece, an image of what the information is about.

The hearer gets the picture as far as the information goes; his image remains fragmentary. It is private; given the same information, different hearers will each form their own image, however, if the real thing comes in view, they will all recognize it because it can fuse with the ghost version thereof existing in their own minds.

Obviously, a speaker cannot squeeze these images into sound and send them along to the hearer. Before any verbal communication can take place, the members of a speech community must build up a basic vocabulary: each of them must link the words to aspects of the world and during shared experiences they must make sure to link the same words to the same things. With a growing vocabulary, real things gradually get covered by more and more words that are linked to

their different features and aspect; each person individually stores these words and features together, but different people will store the same words linked to roughly the same things. Thus accomplished, a word can call up a feature-image.

WORKING HYPOTHESIS

a word calls up (or keep from fading) in the hearer's mind the feature-image connected to it.

A speaker cannot literally convey a feature-image, but he can call it up and thus authorize it. The feature-image called up by a word is not a fixed picture; it may remain dim, it may be detailed and it may change its appearance, as long as it fits the speaker's word.

Words like *cat*, *sister*, etc., only authorize unconnected fragments. To connect fragments, the speaker has forms such as the noun phrase and the predicate. Suppose the speaker has Tina in view, who is a sister of John's, then the word *sister* tags only an aspect of Tina, but a noun phrase, like *a sister* covers the whole of her, implicitly including all her other features and aspects. The predicate on the other hand tags an aspect. Switch to the hearer's perspective: the idea is that the speaker's noun phrase starts to build a unification of features, and that the predicate includes its feature within this unification under construction; the prosodic information now fits in very well:

WORKING HYPOTHESES

a noun phrase	calls up the outline of a unification of features, filling in the features called up by its word(s), and leaving room for complementary features. The resulting image will be fragmentary. The hearer may tentatively add all kinds of unauthorized complementary detail. The range of what the hearer can think of is largely determined by the features that in the hearer's past experience have gone together with the authorized features.
↓ combined with a noun phrase	fills up the room in the outline of the unification of features called up by the noun phrase with a set of unspecified features.
↑ combined with a noun phrase	keeps open the room in the outline of the unification of features called up by the noun phrase.
a predicate	includes the feature-image it calls up in the outline of the unification of features called up by the subject noun phrase.

To simplify the discussion, abbreviate example (2) as (14):

(14/1) *a sister*[↓] *is interested*

<i>sister</i>	calls up a <i>sister</i> -feature,
<i>a sister</i>	calls up the outline of a unification of features, filling in the <i>sister</i> -feature and leaving room for complementary features,

↓	
<i>is interested</i>	fills up the room with a set of unspecified complementary features, includes the <i>interested</i> -feature in the unification of features.

The resulting picture is a picture of a complete person, a unification of features including a *sister*-feature, complementary features, and an *interested*-feature; as long as the real sister is out of view, this mental picture stands-in for her to the hearer. As it is largely unspecified, the hearer may guess what the real person in fact looks like; he may tentatively add a choice of complementary details for himself, erase them, fill in something else, etc., but in fact the person represented is what she is.

If the real person comes in view, mental stand-in and real thing should fuse together, unite. However, there still is the matter of the indefiniteness of the noun phrase; given this combination, the picture stands-in for a specific but non-identifiable person; the hearer may be left with inconclusive information. I shall return to that in chapter 3.

(14/2) *a sister*[↑] *is interested*

<i>sister</i>	calls up a <i>sister</i> -feature,
<i>a sister</i>	calls up the outline of a unification of features, filling in the <i>sister</i> -feature and leaving room for complementary features,
↑	keeps open the room for complementary features,
<i>is interested</i>	includes the <i>interested</i> -feature in the unification of features.

The resulting image is a fragmentary sketch: the outline of a unification of features including a *sister*-feature, room for complementary features left open, but an *interested*-feature included. It is up to the hearer whether or not, and how, tentatively to fill in unauthorized details of his own; the range of what he can think of is any arbitrary choice of features; as far as the information goes, each as good as any other. If any interested sister happens to come in view, the sketch can fuse with it, which still does not make (14/2) into a piece of information about that sister.

For a discussion of cases (8) thru (13), see chapter 3.

1.2.4 Negative Information

This is the idea: information about a real situation piece by piece builds up a simulated perception of the real thing, as if it is there. But what if the information is negative? Negative information states absence, but the world does not contain non-existent facts and absent features for the hearer to see for real; so how does he get this absence in his mental picture of the world? In constructing the picture to look like what is factually there, the hearer cannot simply leave out what is not there, since what is absent from the world does not exist, but what is absent from the hearer's image of the world is what he does not know about the world. The image is the hearer's fragmentary grasp of the world, with lots of unspecified space. Negative information limits the hearer's possibilities for speculation about what is there.

So what shape does negative information take? I believe people understand negative information in much the same way as they understand positive information: words do just what they are meant to do, they call up the feature-images linked to them, and also the constructions, as well as the prosodic forms, do exactly what they are meant to do, only the image becomes an

example of what is excluded from the world. Non-existent facts and absent features are not there in the world, but they are there in one's grasp of the world.

WORKING HYPOTHESIS

not combined with marks the part of the image called up by this predicate or this
a predicate or a noun phrase as 'absent'. An image marked 'absent' is
noun phrase irreconcilable with the original image.

Abbreviate (1) to (15) *a sister is not interested*, and compare (15/1) and (15/2):

(15/1) *a sister*[↓] *is not interested*

<i>sister</i>	calls up a <i>sister</i> -feature,
<i>a sister</i>	calls up the outline of a unification of features, filling in the <i>sister</i> - feature and leaving room for complementary features,
↓	fills up the room with a set of unspecified complementary features
<i>is not interested</i>	includes the <i>interested</i> -feature marked 'absent' in the unification of features.

The resulting picture is a unification including a *sister*-feature, complementary features, and the *interested*-feature excluded. It represents a complete person; to the hearer, it stands-in for someone there. As it is largely unspecified, the hearer may guess what the real person may be like, as long as he does not include an *interested*-feature, which is irreconcilable with the information. If the real situation comes in view, the picture must unite with a person that has the *sister*-feature but lacks an *interested*-feature. If there is someone in view that has the *sister*-feature and also the *interested*-feature, it does not come into consideration. If on the other hand there is more than one option, the hearer has to keep his choice on hold.

(15/2) *a sister*[↑] *is not interested*

<i>sister</i>	calls up a <i>sister</i> -feature,
<i>a sister</i>	calls up the outline of a unification of features, filling in the <i>sister</i> - feature and leaving room for complementary features,
↑	keeps open the room for complementary features,
<i>is not interested</i>	includes the <i>interested</i> -feature marked 'absent' in the unification of features.

The resulting image is a fragmentary sketch: the outline of a unification of features including a *sister*-feature, room for complementary features left open, an *interested*-feature excluded. It is up to the hearer whether or not, and how, to fill in unauthorized details of his own, as long as he does not include the *interested*-feature. If the real situation comes in view, the hearer will expect that he will not find the two features combined. Should anyone happen to have both the features this would clash with the information.

Compare again the two negative cases (15/1) and (15/2): assume a situation where Rose is in fact interested, and Tina is not. Assume that the hearer on his own authority has somehow gathered this information. Now if the speaker says (15/1) *a sister[↓] is not interested*, the hearer can fuse the picture called up with Tina, and he does not bother about Rose; however, if the speaker says (15/2) *a sister[↑] is not interested*, Rose clashes with the sketch called up.

Next compare (5) and (6), where the noun phrase is negative:

(5) *not a word[↓] was true*

This information marks as ‘absent’ a unification of features including a *word*-feature, unspecified complementary features, and a *true*-feature.

(6) *not a word[↑] was spoken*

This information marks as ‘absent’ a unification of features including a *word*-feature, room for complementary features left open, but including a *spoken*-feature.

1.3 Emphasis: Double Images

1.3.1 Not Zooming In

In (1/1), *not* seems to limit its influence to the predicate (‘predicate negation’); in (1/2), it seems to cover the whole sentence (‘sentence negation’). But *not* may also seem to zoom in on a small part of the sentence; in fact, this can be any part of the sentence; for instance, in (1), it can seem to zoom in on *sister*:

(1/3) *a sister of John’s is not interested*
(implication: an (other-than-sister) of John’s is interested)

(1) Only has this implication if it is spoken with a rather specific ‘fall-rise’ intonation; a first approximation: intonation singles out *sister*, not by pitch accent but by extra height of the pitch contour, possibly in combination with a pitch accent; this peaking pitch is followed by a gradual fall, and a rise without accent again at the end. This sound form is unambiguous. So there are questions about the form and questions about the message.

Take a closer look at the information. Most striking is the positive implication; it is almost as if *not* just narrows down on *sister*: *a (not-sister) of John’s is interested*. However, this positive implication comes along with a regular ‘sentence negation’: (1/3) says that no sister of John’s is interested, and the intonation implies that someone else related to John is interested. Shift the peaking pitch over to *John*:

(1/4) *a sister of John’s is not interested*
(implication: a sister of (other-than-John)’s is interested)

(1/4) still says that no sister of John's is interested, but now the intonation implies that someone else's sister is interested.

Next take

(16) *she is not John's elder sister*
(implication: she is John's younger sister)

Now the implication comes along with a regular 'predicate negation': (16) says about this *she*-person that she is not John's elder sister, and the intonation implies that she is another sister.

In fact, any case where *not* seems to zoom in on a narrow spot of the sentence is a special case either of 'predicate negation' or of 'sentence negation'. The point is that in all these cases, as in all the cases discussed earlier, the subject noun phrase is used either to refer or not to refer to something there. The implication is extra.

Then take a closer look at this 'fall-rise' form. In the first place, note that it completely fits in with my conjectures: in cases (1/3) and (1/4), the gradual fall of the intonation contour after its peak has the final syllable of the noun phrase well above base level, leaving the image called up an open sketch. In (16), the pitch contour on *she* is still at base level, well before reaching its peak at *elder*, and the statement, as well as the extra implication, is about this *she*-person.

And then, apparently, there is something extra in the form, responsible for the extra implication. I believe we should take the 'fall-rise' apart; what catches the ear are the following characteristics:

HYPOTHESIS

In the 'fall-rise' intonation, the hearer recognizes as informative, next to ↓ or ↑ in the final syllable of the noun phrase, the following basic forms:

- pitch rising above the average peaks of intonation, that singles out a narrow spot of the sentence,
- an overall absence of pitch accent, with a possible exception on the spot singled out,
- a final rise.

NOTATION


—	as in	—	pitch rising above the average peaks.
		<i>sister</i>	
—	as in	s— <i>ister</i>	absence of pitch accent on otherwise suitable locations. That is, any pitch movement as long as it is not both steep and properly located in the syllable.

I shall not discuss 'suitable locations' for pitch accent; I shall represent absence of pitch accent only where relevant in the discussion.

/	final rise, steep but too late in the syllable to lend an accent, or gradual.
---	---

Given these notations, we can unambiguously represent the forms of the cases discussed; again, the complete contour given is roughly based on ‘Dutch School’ notation; I have drawn what I hear in the corresponding Dutch utterance, and the conjectured skeleton is a hypothesis for English ‘fall-rise’ intonation as well:


(l)



een zuster van Jan is niet geïnteresseerd

(1/3) *a s—ister of J—ohn's[↑] is n—ot —interested/*


(m)



een zuster van Jan is niet geïnteresseerd

(1/4) *a s—ister of J—ohn's[↑] is n—ot —interested/*

(n)



ze is niet Jans oudere zuster

(16) *sh—e[↓]'s n—ot J—ohn's —elder s—ister/*

If I am correct, each of the forms: peaking pitch, absence of pitch accent, and final rise, adds its own piece of information. The final rise contributes something like ‘to be continued’. Words without pitch accent repeat what just has been said, and the overall absence of pitch accent in these cases makes them suitable for their typical use, namely, to correct another speaker’s preceding information by repeating it, just adding *not*, as in:

(other speaker:

a sister of John's[↑] is interested)

(1/3) *a s—ister of J—ohn's[↑] is n—ot —interested/*

Peaking pitch is essential to get the effect; we must try to get hold of precisely the piece of information it contributes on its own. Take another variant of (1), (1/5); like (1/3), it is suitable to correct a preceding statement concentrating on *sister*:

(other speaker:

a brother of John's[↑] is not interested)

—
(1/5) *a s|ister of J—ohn's[↑] is n—ot —interested*

Compare the form of (1/3) and (1/5): in both, peaking pitch singles out *sister*, in both, the intonation contour does not come down to base level until well behind the subject-predicate boundary, and in both cases, there is an absence of pitch accent at every spot not singled out. There are two differences: (1/5) does, and (1/3) does not have a pitch accent combined with peaking pitch on *sister*; in (1/3), *sister* is just high, in (1/5), it has a strong and steep rising pitch accent. The reason for this difference is that in (1/5) it does not, and in (1/3) it does repeat the other speaker's *sister*. And the final rise is missing in (1/5). So here we may have a good view of the information contributed by peaking pitch isolated from the information of the final rise.

Compare the information given by (1/3) and by (1/5), both in the context in which their information fits. At first, they strike us as very different; in (1/3), *not* seems to erase just *sister* in the preceding statement, leaving a positive statement with a gap still to be filled; the presence of *not* seems absolutely essential for the effect; in contrast, in (1/5), the emphasized *sister* seems to overwrite the parallel word *brother* in the preceding statement; here, the presence of *not* seems coincidental; both corrected and correcting statement are negative, and *not* does not seem to play any part in the correction. Now listen again to what (1/5) says: its regular information is that no sister of John's is interested, but simultaneously, as an extra, it contradicts what the other speaker said, it contradicts that no brother of John's is interested, and so, what this implies is that there is an interested brother. And so, the extra information given by (1/3) with its 'fall-rise', and (1/5) with an additional accent and without a final rise, is not so different after all; both imply that someone related to John other than as a sister is interested. Only, in the contexts in which they fit, (1/3) still leaves a gap, and (1/5) already has its gap filled by the preceding *brother*.

So it is reasonable to assume that in (1/3) and (1/5) the same is going on: peaking pitch on its own somehow, but not by making *not* zoom in, results in this extra implication. The question is what peaking pitch by itself contributes; in (1/5), this seems obvious. As the presence of *not* is coincidental, it can be left out:

(other speaker:

a brother of John's[↑] is interested)

—
(2/3) *a s|ister of J—ohn's[↑] is —interested*

We can hear what peaking pitch contributes by just lowering its peak to average; if we do, we add an interested sister to the other speaker's interested brother, but if we keep the peak, it makes the interested sister replace that brother, contradicting what the other speaker just said. A speaker who peaks his pitch speaks with emphasis; he reinforces his regular message by, simultaneously, contradicting its opposite, zooming in on the spot singled out by pitch peak. For the regular information, the conjecture is that it takes shape as a unification of features under construction,

forming the hearer's expectation of what the world is like. So for the extra information, the conjecture is that it likewise takes shape as a unification of features under construction, but now reversing the hearer's expectation.

WORKING HYPOTHESIS

given a sentence that calls up its regular image

—

doubles the regular image,
erases the spot singled out by the pitch peak form, the gap to be filled by an
opposite to be found in context or situation,
reverses this image, that is, from an image of presence into an image of absence, or
vice versa.

The working hypothesis applied to the positive (2/3) abbreviated to (14/3):

(14/3) $\overline{a\ s|ister^\uparrow\ is\ \text{---}interested}$

The regular message is given by

(14/2) $a\ sister^\uparrow\ is\ interested$

The resulting image is the outline of a unification of features including a *sister*-feature and an *interested*-feature, room for complementary features left open.

The extra implication is given by:

—

sister calls up an extra image: the outline of a unification of features including an
opposite-of-*sister*-feature and an *interested*-feature, room for
complementary features left open.

This extra image is reversed from an image of presence into an image of
absence,
so what the hearer expects is the absence of the combination of those
features in one unification.

s|ister new,
---interested repeated.

A suitable preceding context would be *a brother[↑] is interested*; it would
make the hearer expect the absence of the combination of a *brother*-feature
and an *interested*-feature in one unification.

The working hypothesis applied to the corresponding negative (1/5) abbreviated to (15/5):

(15/5) $\overline{a\ s|ister^\uparrow\ is\ n\text{---}ot\ \text{---}interested}$

The regular message is given by

(15/2) $a\ s\text{---}i\text{---}s\text{---}t\text{---}e\text{---}r^{\uparrow}\ is\ not\ interested$

The resulting image is the outline of a unification of features including a *sister*-feature, room for complementary features left open, but an *interested*-feature marked ‘absent’ included, excluding it from whatever details the hearer may fill in for himself, what the hearer expects is the absence of the *interested*-feature combined with the *sister*-feature in one unification.

The extra implication is given by:

—
sister calls up the extra image: the outline of a unification of features including an opposite-of-*sister*-feature, room for complementary features left open, but an *interested*-feature marked ‘absent’ included, this extra image is reversed: what the hearer expects is the presence of the combination of the two features.
s|ister new,
 —*interested* repeated.
 A suitable preceding context would be $a\ brother^{\uparrow}\ is\ not\ interested$; it would make the hearer expect the combination of a *brother*-feature and an *interested*-feature in one unification.

And finally the working hypothesis applied to (1/3), the case where *not* seems to zoom in, abbreviated to (15/3):

(15/3) $a\ s\text{---}i\text{---}s\text{---}t\text{---}e\text{---}r^{\uparrow}\ is\ n\text{---}o\text{---}t\ \text{---}i\text{---}n\text{---}t\text{---}e\text{---}r\text{---}e\text{---}d\ /$

Both the regular message and the extra implication are as above. The difference is this:

$s\text{---}i\text{---}s\text{---}t\text{---}e\text{---}r$ repeated,
 —*interested* repeated.

A suitable preceding context would be $a\ sister^{\uparrow}\ is\ interested$, which does not supply the required contextual opposite, so the hearer is left with the expectation of the combination of a feature still to be found and an *interested*-feature in one unification.

/ to be continued.

So if I am correct, the positive implication of negative sentences with ‘fall-rise’ results from the denial of *not* and the reversing pitch peak cancelling each other out in the extra image called up by peaking pitch.

1.3.2 A Multiplication of Combinations

I have hypothetically isolated a rather small (not exhaustive) number of basic informative forms in the pitch contour. Combined with a simple subject-predicate structure, they are: the obligatory choice of \downarrow versus \uparrow , that is: the noun phrase either or not meeting the low declination line at its end; optionally a peaking pitch; any distribution of pitch accent, that is: a steep pitch movement properly located in the syllable, and final rise. There is a multiplication of possible combinations of these forms, each combination the (possibly incomplete) skeleton of one or more possible intonation contours. Hypothetically, these basic forms each contribute their own precise piece of information; together with the information of words and syntactic form, they unite into the instant end-product of understanding. In the cases discussed above, on a more or less intuitive basis I have shown the connection between the particular combination of forms chosen and the end-product. Of course this needs more precision.

There are other combinations than the ones discussed; in these cases the connection between form and understanding should come out correctly as well. Two examples are worth mentioning here.

Take (1/5) but replace \uparrow by \downarrow ; if the contour is down to base level at *John’s*, this very clearly changes the meaning from ‘sentence negation plus an implication’ to ‘predicate negation plus an implication’, even if we added a final rise to the sentence (which would make a ‘fall-plus-rise’ intonation) for the speaker to continue with *but...*:

(1/6) $\overline{\quad}$
a s|ister of J—ohn’s \downarrow is n—ot —interested (/)

The regular picture is a unification of features including a *sister*-feature, complementary features, but an *interested*-feature excluded. The hearer expects someone specific with a *sister*-feature but without an *interested*-feature; if there is another sister that does have the *interested*-feature, she is of no concern. The extra picture is a unification including an opposite-of-*sister*-feature, complementary features, and an *interested*-feature excluded. This extra picture is reversed; the hearer expects the absence of such a person.

Another interesting variant of (1) is (1/7):

(1/7) $\overline{\quad}$
a s—ister of J—ohn’s \uparrow is n|ot —interested

Here, peaking pitch does not cancel out the denial of *not*, but just the opposite, it reinforces the denial. Intuitively, it works as follows:

The regular image is an open sketch: the outline of a unification of features including a *sister*-feature and an *interested*-feature marked ‘absent’. The hearer expects the absence of the *interested*-feature combined with the *sister*-feature in one unification. The extra image is the outline of a unification of features including a *sister*-feature, and an *interested*-feature marked with the opposite of ‘absent’. This extra image is reversed, so also the extra implication makes the hearer expect the absence of the combination of a *sister*-feature and an *interested*-feature.

One remark about the connection between the skeleton of informative forms and the actual intonation contour. It may happen that two or three basic forms have to be realized on the same syllable, for instance if the final syllable of the noun phrase happens to be a syllable with word accent (not discussed in this book), for instance,

—
a c|at↓

The actual pitch here is determined by these three combined basic forms: the accent needs to be realized by a steep pitch movement on top of the syllable, it needs to have an extra high component, and pitch needs to touch the low declination line between the top of its syllable and its end; the actual pitch will be a steep fall from higher than average height, either or not immediately preceded by a steep rise to this extra high level. If start and finish of the pitch movement are at the same level, what one hears is a strong accent rather than a change in pitch. Anyhow, the hypothesis is that a hearer at this point recognizes these three forms in the intonation, and that they each contribute their own piece of information.

1.4 Noun Phrase plus Rhythm: Single or Repetitive Images

1.4.1 Repetitive Stand-ins and Repetitive Sketches

Next listen to

- (17) *a tortoiseshell cat is a female cat*
 (18) *a vixen is a female fox*

These are general statements. Again, it is the use of the noun phrase that makes the difference, again, this general use has a prosodic form of its own, and again, this form is located in the final syllable of the noun phrase: it is lengthened. In all examples discussed above, the noun phrases were not lengthened.

HYPOTHESIS

Hearers recognize in a speaker's sequence of sounds whether or not the final syllable of the noun phrase is lengthened.

NOTATION

- | | | |
|---|-------------------------|---|
| ↔ | as in <i>a vixen</i> ↔ | the final syllable of the noun phrase is lengthened |
| * | as in <i>a sister</i> * | the final syllable of the noun phrase is not lengthened |

This rhythmic form is independent of, and freely combinable with the pitch pair in the same syllable, so ↓↔ or ↑↔ next to ↓* or ↑*. As can be expected, parallel to the two types of non-general statements (both positive and negative) discussed above, either or not referring to something there, there are two types of general statements: ↓↔ makes the statement refer to all things concerned, and ↑↔ keeps open sketches of all things concerned. Compare, on the one hand:

- (17) *a tortoiseshell cat*[↓]_↔ *is a female cat*
 (19) *a red tabby*[↓]_↔ *is a male cat*
 (20) *a promise*[↓]_↔ *is a promise*
 (21) *a good neighbour*[↓]_↔ *is worth more than a far friend*

These are statements about tortoiseshell cats, red tabbies, promises, good neighbours. They all say that if there is something (anything) that is a tortoiseshell cat, a red tabby, etc., on a regular basis it has the predicate feature. Take (17): one recognizes a cat as a tortoiseshell because of its colour, and one is informed that in that case, it is female as well. This kind of general statement rather easily allows for an occasional exception, for instance (19) is far from an iron truth, but still, as a rule, it goes. Like (20) and (21), many proverbs have this form.

And then, on the other hand, take cases like

- (18) *a vixen*[↑]_↔ *is a female fox*
 (22) *a quillback*[↑]_↔ *is a freshwater fish*
 (23) *a planet*[↑]_↔ *is a heavenly body moving around the sun* (and whatever further requirements astronomers agree upon)
 (24) *a perpetual motion machine*[↑]_↔ *is a device that operates indefinitely by creating its own energy*

Obviously, they do not inform the hearer about vixen, quillbacks, planets, and perpetual motion machines. Rather, the predicate seems to give defining or otherwise necessary characteristics for fitting the noun phrase; the whole information seems to be that the noun phrase feature always is linked to, does not go without, the predicate feature. For instance (18) seems to say what it takes to be called a vixen. Other than in cases like (17) and (19) thru (21), such cases do not allow for exceptions: in order to find a vixen we have to look for a female fox, and if we found a male fox instead, we just would not call it a vixen. And other than the cases (17) and (19) thru (21), these cases do not claim to describe existing things, as is clear in (24), which informs the hearer what perpetual motion machines would be like if they existed.

So a speaker systematically has four different ways to use the same noun phrase, and in the final syllable of that noun phrase, he efficiently gets his intentions across by this choice of prosodic forms. Compare (25) thru (28):

- (25) *a tomcat*[↓]* *is an especial beauty*
 (26) *a tomcat*[↑]* *is a good choice*
 (27) *a tomcat*[↓]_↔ *is a good mouser*
 (28) *a tomcat*[↑]_↔ *is a male cat*

A little story to give these examples some context: a woman wants to buy a kitten; she prefers a pedigree lady cat; an acquaintance knows a good address: *They love animals and have many pets.* (25) *a tomcat*[↓]* *is an especial beauty.* *He will be the father of this new litter they have.* The woman goes there, bringing her little son, to find that there are only two male catkins left to choose from. The owner says (26) *a tomcat*[↑]* *is a good choice.* *These here will look just like their father. Look at him!* When the mother hesitates, he adds (27) *a tomcat*[↓]_↔ *is a good mouser.* And

now the little boy asks *what is a tomcat?* The answer is (28) *a tomcat*[↑]↔ *is a male cat*. In orthographic writing, these four sentences are largely identical, however, their complete forms, if correctly reconstructed, are not; given its prosodic form, each very accurately says exactly what it says.

As can be expected, parallel to the two types of positive general statements, there are two types of negative general statements. Compare:

- (29) *a big boy*[↓]↔ *doesn't cry*
 (30) *a vegetarian*[↑]↔ *does not eat meat*

(29) is about all big boys, although it has its exceptions, and (30) says what it takes to be called a vegetarian.

So the question is what piece of information this rhythmic form ↔ contributes, which by now has turned into the question: what shape does this information take.

WORKING HYPOTHESIS

- ↔ combined with turns the unification of features under construction into a self-
 a noun phrase replicating image

Intuitively, (27) pictures the unification of features including a *tomcat*-feature, a set of unspecified complementary features, and a *good mouser*-feature, and with each set of complementary detail that the hearer thinks of, the picture self-replicates, forming a row of tomcats that are also good mousers. (28) sketches the outline of a unification of features including a *tomcat*-feature, room for complementary features left open, but a *male cat*-feature included, and with each possible detail the hearer happens to think up, the sketch self-replicates into a row of such open sketches that all at least combine the *tomcat*-feature and the *male cat*-feature, whatever other features the hearer may think of.

1.4.2 The Multiplication of Combinations

As illustrated above, given a simple syntactic form, a few basic informative forms in pitch can be put together in a multiplication of different combinations. Now, including rhythmic form, the range of combinations is doubled. Each change in prosodic form changes the message, subtly or drastically. In the context of the general use of noun phrases, I'd like to add another illustration: take this still rather small boy who has fallen off his bike and is crying; now his sister, hardly more than a year his senior, might say

- (29/1) *a b|ig b—o y*[↓]↔ *doesn't cr|y*

To her brother, she pictures a model of a big boy, a picture he wants to fit, so, he probably will grow into this model and dry his tears. Alternatively, his sister might say

- (29/2) *a b|ig b—o y*[↑]↔ *doesn't cr|y*

Now she sketches the combination of a big boy-feature and a cry-feature as irreconcilable by definition; no chance of him fitting. He will feel belittled. Probably he will cry even harder. Alternatively, his sister might say

(29/3) $a\ b|ig\ b\text{---}o\ y^{\uparrow}\leftrightarrow\ \text{doesn't}\ cr|y$

Now it gets nasty. Besides sketching the permanent combination of a *big boy*-feature with the absence of a *cry*-feature, in which sketch by definition he does not fit, she sketches an opposite-of-*big*-which-is-*small*-boy-feature, with a *cry*-feature permanently linked to it. This one fits. Which hurts more than a grazed knee. Probably he will kick his sister as hard as he can to give her something to cry about, the little baby herself.

I want to add a final example that combines almost all that seems to be problematic. The context given is to ensure the correct reconstruction of prosodic forms; it is these forms and not the context that constitute the message. Suppose two people are trying to think of names for male, female, and young animals. One of them proposes $a\ drake^{\uparrow}\leftrightarrow\ \text{is}\ a\ \text{male}\ swan$, and the other corrects

(31) $a\ dr\text{---}ake^{\uparrow}\leftrightarrow\ \text{is}\ n\text{---}ot\ a\ m\text{---}ale\ sw\text{---}an/$
(*a cob is a male swan*)

Listen to this: understanding is immediate. Try to paraphrase the understanding and it turns out to be virtually undoable. If a speaker wants to give this information, apparently (31) is the form to use. It works, it says exactly what it says. Again, the question of this book is: how does it work?

In this first chapter, I have proposed some hypotheses about a number of basic prosodic forms that a hearer picks out as informative from a speaker's sequence of sounds, and their respective contributions to the information given. The hypotheses about prosodic form are precise; I believe they are correct. The conjectures about their information are working hypotheses; I believe they have some credibility because they suggest the existence of a very simple and elegant system behind a range of seemingly different and problematic phenomena. However, in order for them to really count as an explanation, these working hypotheses need to be made more precise; for each different combination of forms, their pieces of information put together should precisely account for both the logic and the nuances of the end-product of understanding.

Chapter 2

Language and World

This book is about understanding people's understanding of language. To avoid confusion, I shall refer to the speaker, the hearer, the ones to be explained, in the third person: 'he', 'she', 'people', and to us who do the explanation, in the first person: 'I', 'we'.

If verbal information is about something that is there, what then is the connection between the forms of language and the real thing? The idea of this book is that a speaker's words in the hearer's mind call up a fragmentary ghost version of the real thing, to stand-in for it when it is out of sight, and to fuse with it when it is there. Now images may seem to be rather elusive, so it may not be self-evident that we can use these ghosts in the private mind to explain the connection between language and world; how to get hold of them? But I believe there are more urgent problems: how to get hold of those forms of language, and how to get hold of that world; when we talk about language and world, what exactly is it that we are talking about?

Take the forms of language: they may have an air of objectivity because people seem to share them, but what are they, what is in fact their way of existence? Not an easy question at all. Take reality, and the questions become even harder; it seems reasonable enough to suppose that reality exists, and it seems clear enough that forms of language may refer to what is there, but, to start with, there is the problem that speakers and hearers do not talk about Objective Nature but about some kind of human version of it, a phenomenal world of colour, texture, sound, smell, beauty. What is this phenomenal world, what kind of existence does it have, what does it have to do with Nature, how do people share it? And the really tough problem is the question behind these questions: how do people know at all that there is a reality around them? How is it that they can break out of the boundaries of their own body? There is an explanatory gap: we, or at least specialized scientists, may come to understand how organisms more or less successfully interact with what is there around them, but it is something else to understand how organisms come to know what it is they are interacting with; the question is, what makes people have a world around them. This is the most basic question to be asked in this book, and the toughest one; in the discussion about language and world it cannot be skipped: before organisms are capable of using arbitrary sequences of sounds to stand-in for what is there, they must have a notion of what is there. So the question is, again: what exactly is it that we are talking about?

No doubt the brain has to be made part of the answer: given the neuroscientific evidence, it is reasonable to suppose that everything a person knows completely depends on this person's own brain. People each have this neural machine of their own, that opens up the world around them, space and things in space, that enables them to picture what is out of sight, and that enables them to use and understand forms of language to stand-in for what is there. However, when we are talking about human language and the world it refers to, we are not talking about neural structures and processes. Indeed, whatever theory we might come up with to explain verbal communication, it should agree with neuroscience's results, but still, my questions cannot be rephrased in terms of the brain without getting lost: where in speakers' and hearers' individual brains to look for the words shared between them, where to look for the real things they know that are there in the human world, where to find the virtual worlds of fiction? It does not help to call these neural structures and processes 'representations' of the real thing; in order for them to

be representations of something else there should be someone to whom they stand-in for the real things.

In order to be able to find answers to my questions about language, I shall have to make some basic assumptions about brain and mind, in agreement with what I understand to be known about the brain. I am not a neuroscientist, and everything I know about the brain I have read in publication of those who are (see references in Tribute). However, my assumptions shift the focus from the mechanism of the brain to the brain's awareness. In agreement with these assumptions, I shall propose a general theoretical framework to get hold of both the human world and the forms of language in terms of what the individual neural machine makes emerge in the private mind. Then within this general framework, the working hypotheses given in chapter 1 can be made more precise. The notation to be used is mine; it is adapted to function in a theory of verbal communication.

2.1 The Human World

2.1.1 Living in a Phenomenal World

Imagine an omniscient perspective Ω on everything; we cannot take that perspective, but we can conjecture what we would see if we were able to take it. Part of everything is some particular human organism; represent this organism as H. Focus on this H, starting out on his way to master the universe. First there is, say, some nearby cat; represent this cat as C. One of the questions eventually to be answered is how H can use and understand phrases like *a cat*, *the cat*, *Kitty*, to stand-in for this C, but H still has a long way to go.

The first question is how it is that H knows C to exist. Of course, H has senses to pick up signals originating from C, and a neural system to put the information to use; however, the question is not how, from perspective Ω , H's brain makes its organism interact with C, but the question is how this activity inside H's neural machine creates its own perspective, and separates C from H's own organism into something in its own right, doing its own thing in the space around. This is the basic question, and it is the toughest question; the more we think about it, the more it seems to escape understanding. Fortunately, we can be sure that it happens: H's brain, set in motion by input on a few sensitive spots at the boundaries of its body, calls the source of this input into an existence separated from its own person. To be precise, we can be sure that it happens during the time H's brain succeeds in making H aware of this cat. I believe this is as far as it goes; for all we know, it ceases to happen when the brain ceases to make H aware of the cat. So, the problem of the world around people (how is a neural machine able to create a notion of space around them, and things there in their own right) cannot be separated from the problem of awareness (how is a neural machine able to make awareness emerge in an organism): how on earth is a neural machine able to make its own person aware of other things around him. But anyhow, from perspective Ω we can see that the cat is there to H precisely as long as his brain turns the signals picked up by his senses into awareness of the cat.

When we talk about H's cat, that is, the cat that H eventually will be able to talk about, we are not talking about the physical organism C as it is, nor about H's brain's processing C's signals, but about C as and as far as H's brain makes it show up to him. Now, even while we do not understand how a neural machine can create a perspective on what is there, knowledge of the organisation of H's brain will give a better idea of H's world. It is known that the perceptual regions of H's brain are organized in parallel, hierarchical, feed-back, and criss-cross connected areas, each specialized in responding to its own kind of input; each area contains an enormous

amount of neurons, connected by a multiplication of synapses. Take a brief moment of H's awareness of the cat, and trace the action triggered by C's input on H's senses: in each area relevant to the perception at this particular moment, a certain choice of all the neurons available is actually active, together, a certain complex of sub networks of neurons all over the cortex, that is, a specific combination of neurons out of the uncountable number of possible combinations. At any given moment of H's awareness of the cat, this is a specific complex, but all the time neurons and networks of neurons will drop out of action and others will join in. We need a representation of this complex of networks of neurons; as it is impossible to find a notation that captures its complexity and dynamics, I shall use a simple notation:

${}^H(N \text{---} \text{---} N)_C$ the complex of parallel, hierarchical, criss-cross and feedback connected networks of neurons active in H's brain while it makes H aware of C.

Whenever relevant in the discussion, parts of this complex can be specified:

${}^H(N \text{---} N_i \text{---} N)_C$ the same, including the network N_i .

So H is aware of the cat while ${}^H(N \text{---} \text{---} N)_C$ is active in his brain. We might say that H's brain figures out the source of the signals; we might want to think of this process as some sort of complex calculation, but if indeed it is a calculation that takes place here, it is one of a remarkable kind: we, observing H and his brain, may be able to see the problem, that is the cat that is there and its signals hitting H's senses, and perhaps some time in the future we, or at least specialized scientists, may be able to follow in detail the calculation as well, but H's brain will present its solution, that is the cat taking shape, exclusively to a 'self' simultaneously emerging in the brain's organism; and this 'self' in its turn does not know of any problem, let alone of his brain's process to solve it: the cat just shows up. We can imagine ourselves from perspective Ω to see it all, problem and solution, but as long as we are human, we would need H's brain in action in our own head to see what solution it comes up with.

If I correctly understand the literature, what happens is something like this: this whole complex ${}^H(N \text{---} \text{---} N)_C$ bursts into some as yet not identified kind of coordinated activity, and it is precisely during this special activity that awareness emerges, not as a product to be delivered at the end of a production line, but more like the other side of the coin: each side does not exist without the other: there is this organism, there is this neural machine within the organism, this neural machine is capable of some very special coordinated kind of activity, and precisely during this activity, the organism is a self with a perspective on the world around. I shall represent this special activity, whatever it may turn out to be, as A :

A the special coordinated activity that goes on in H's brain while it makes H aware of something

The basic question was: how is this organism H able to know that things exist, which was rephrased as the question: how does H's neural machine create its person's awareness of things that are there. Narrow down this tough question to the question what are the necessary and sufficient conditions for this to happen, and it has become a question that we, or at least specialized scientists, might be able to answer in the future. In the meantime, I shall just assume

that H's brain achieves what it achieves. If I am correct, for H's brain at a particular moment to decloak the cat to its self, it takes two things:

1. this complex of networks of neurons ${}^H(N \text{---} \text{---} N)_C$
2. this special coordinated activity A in this complex.

Represent them together as

${}^H A (N \text{---} \text{---} N)_C$ the special coordinated activity A going on in the complex of networks of neurons ${}^H(N \text{---} \text{---} N)_C$ that is necessary and sufficient for H to be aware of C.

So it is during, precisely during, and only during this ${}^H A (N \text{---} \text{---} N)_C$ that C is there in H's world. In this ${}^H A (N \text{---} \text{---} N)_C$, it is the special activity A that is the neural correlate of H being aware, having a perspective, and it is the specific complex network of neurons ${}^H(N \text{---} \text{---} N)_C$ bursting into that action that determines the actual content of this awareness.

For the purpose of our theory of verbal information about the world, we want to get hold of the cat in H's world, that is not so much C as it objectively is, as well C as his brain makes it show up to H. Take colour; C itself reflects light on a certain wavelength that enters H's eyes and activates a certain combination of neurons in the area specialized in colour in H's brain; say, ${}^H N_{\text{ginger}}$ is triggered into ${}^H A N_{\text{ginger}}$, which to H is the sensation of ginger; this sensation is how H's brain succeeds in making its 'self' aware of the physical quality of C to reflect light of that particular wavelength. Other features emerge from their own networks activated by signals from other physical aspects of C, and in coordinated action A , H's brain gathers together what it makes of all the signals that its senses can register, and so this cat C shows up to H as a unification of shape, colour, texture, sound, smell, etc.


If I correctly understand the literature, each of these features is contributed directly to H's awareness by the correlated sub network of neurons joining action A . Thus, ${}^H A N_{\text{ginger}}$ taking part in ${}^H A (N \text{---} N_{\text{ginger}} \text{---} N)_C$ directly contributes colour to the cat that takes shape to H.

So C takes on shape, colour, texture, sound, smell, etc. to H: a phenomenal cat. For the purpose of our theory of verbal information about the world, we need to get hold of C as it shows up to H. As it is impossible to find a notation that captures the complexity and dynamics of such a unification of features, I shall use a simple notation:


${}^H \backslash \underline{\bullet \bullet \bullet} /_c$ the unification of features c that shows up to H while his brain makes him aware of C.
Whenever relevant in the discussion, features can be specified:

${}^H \backslash \underline{\bullet \bullet_{\text{ginger}} \bullet} /_c$ the same c including the colour ginger.

Given these notations, we can represent the equation between what goes on in H's brain and what shows up in his mind from perspective Ω as follows:

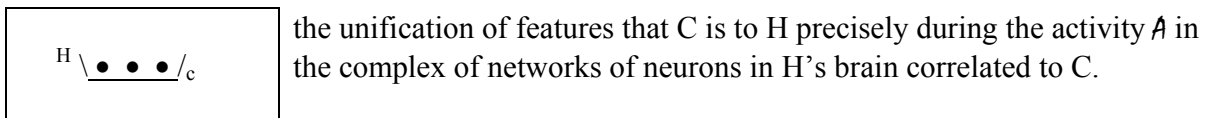


or, when specification is relevant in the discussion:



$H \bullet \bullet \bullet / c$ does not show up in H's awareness without ${}^H A (N \text{---} \text{---} N)_C$ going on in his brain, and ${}^H A (N \text{---} \text{---} N)_C$ does not go on in his brain without $H \bullet \bullet \bullet / c$ showing up to him. In the notation, H's perspective is put in a box: what shows up in H's awareness (except in the text, where this would take up too much space).

What we need for the purpose of our theory is a representation of H's cat, and therefore we can narrow down the notation to H's perspective:



There are major differences between C and $H \bullet \bullet \bullet / c$, all because of the fact that C is a physical organism, and $H \bullet \bullet \bullet / c$ is a product of the brain:

- C is what it is, $H \bullet \bullet \bullet / c$ is C as and as far as H's brain manages to make it take shape,
- C and the changes in C are determined by its physical and biological nature; $H \bullet \bullet \bullet / c$ is as plastic as H's brain is: from split second to split second, neurons in that complex coalition of millions of neurons, or even complete sub networks, may drop out of action, and others may join in; these changes may or may not be caused by changes in C and its signals,
- two cats in the world do not overlap. To H, both cats will be similar, as his brain will use partly the same complex of networks of neurons to make each of them take shape,
- C has the quality that makes it reflect light on a certain wavelength, whether or not there is light. But to H, the cat will lose its colour when ${}^H N_{\text{ginger}}$ drops out of action,
- for C, life and death is a matter of its biology. $H \bullet \bullet \bullet / c$ does not have existence beyond H's awareness; as soon as A dies out in ${}^H (N \text{---} \text{---} N)_C$, $H \bullet \bullet \bullet / c$ is just gone.
- C is objectively there; its signals are there, and can be registered and turned into the appearance of a cat by any organism with a suitable neural machine. $H \bullet \bullet \bullet / c$ is the result of such a process, and can only be private; as long as it is impossible for two persons H and M to share a brain, there is no way for $H \bullet \bullet \bullet / c$ to become $M \bullet \bullet \bullet / c$ or vice versa.

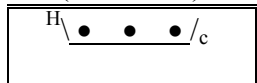
The main points:

The BASIC QUESTION is how H knows that C exists.

The idea is that C is there to H precisely as long as his brain makes him aware of C; the question is rephrased as the question how H's brain makes him aware of C; this is narrowed down to the question what the necessary and sufficient conditions are for C to show up in H's awareness.

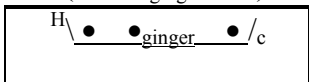
The **BASIC ASSUMPTION** is that C is there to H precisely as long as some special coordinated activity *A* goes on in a complex of networks of neurons ${}^H(N \text{---} \text{---} N)_C$ spread over the cortex, triggered into action by signals originating from C picked up by H's senses, that makes C show up to H as a unification of features, each feature contributed by the activity in its own correlated sub network in its own specialized area of H's brain.

NOTATION

$${}^H A (N \text{---} \text{---} N)_C \ll C$$



the equation between on the one hand the special coordinated activity *A* going on in the complex of parallel, hierarchical, criss-cross, and feed-back connected networks of neurons ${}^H(N \text{---} \text{---} N)_C$ triggered by input from C, and on the other hand the unification of features ${}^H \underline{\bullet \bullet \bullet} / c$ that shows up in H's awareness.

Whenever relevant in the discussion, parts of the equation can be specified:

$${}^H A (N \text{---} N_{\text{ginger}} \text{---} N)_C \ll C$$



the same, including the network ${}^H N_{\text{ginger}}$ contributing the colour ginger.

What we need for the purpose of our theory is a representation of what is there to H:



the unification of features that C is to H precisely during the activity *A* in the complex of networks of neurons in H's brain correlated to C,

or, partly specified,



the same, including the colour ginger.

2.1.2 Recording the World

Thus, if I am correct, what people experience around them is private and fades into non-existence when it disappears from awareness. This may seem counterintuitive; people do not experience the world as occasional, they self-evidently take it to go on to exist behind their backs; and they do not experience the world as private either, they take their world of experience as the world that is there, and that is there to others as well. If ${}^H \underline{\bullet \bullet \bullet} / c$ fades into non-existence the moment ${}^H A (N \text{---} \text{---} N)_C$ quits the action, how does H know that the cat does not vanish along with its appearance to him? It is far too early to bring in language as an explanation; H has to be able to conceive of the cat's ongoing existence before he can make an arbitrary sequence of sounds stand-in for it. The question splits up into two sub questions: firstly, how does H recognize the cat as the familiar cat when it shows up again, and secondly, how does H know of its existence in between. The second question will be addressed in 2.1.4.

Suppose H sees C for the first time; his brain manages to figure out this cat from whatever his senses register. But this input varies enormously with the circumstances, the distance, the angle of perception, available light, movement of the cat, etc., which makes the task of figuring out the same source on each occasion enormously complex. So H's brain speeds up the process by making a record: while it makes $^H \bullet \bullet \bullet /_c$ show up, it also strengthens together the neurons used to get this job done; it lowers synaptic thresholds, and grows new connections. Thus, what remains in H's brain after $^H \bullet \bullet \bullet /_c$ has faded into non-existence is this complex of neurons $^H(N- -N)_C$, which before was just an arbitrary combination of neurons out of an unimaginable number of possible combinations, but is now ready to flash together in action again any time; so on renewed encounter with C, this time-consuming analysis of what hits the senses needs not completely be done over again.

This record $^H(N- -N)_C$ kept in H's brain is not a copy of the real thing, not a stand-in off duty, not something that without causing serious misunderstanding could be called a representation. What the brain stores is the ready-to-use combination of neurons that when in action *A* is the neural side of the cat taking shape to H again. It is easy to trigger as a whole; a glimpse of the real thing may suffice for a flash of instant recognition.

These strengthened networks keep changing; the brain may keep the complexes together, permanently lowering the synaptic thresholds, and growing new connections, but the links may also decay, rapidly or over time. But then again, with each new use, the connections are re-strengthened, and new connections are made; the brain's mechanisms strengthen together the neural networks of everything it happens to bring together in awareness; each new experience may add its own complexes of neurons to the old ones, forming larger coalitions; sound, smell, texture, etc. become tied to shape and colour. And then, a small input in any of those regions may flash the whole complex network into action; H will recognize his own cat when he hears it mew in the distance, or when it rubs against his leg in pitch dark; he will not just recognize the sound, or the feel, it will be the entire cat that is there.

With H's growing experience of reality, more and more of such records of whatever it was that he experienced will be left in his brain, criss-crossing and overlapping each other. This experience is not a detached analysis of the things as they are; H's brain will record phenomenal details together with his experience of how to deal with them, together with the emotions that go with the experience, connecting networks of perception with networks in the motor areas with networks in the areas of emotion. Apparently, the brain indiscriminately records together whatever simultaneously happens to show up in awareness, whether related in nature or coincidental; it makes idiosyncratic associations along with links that mirror connections in the real thing. It is this mechanism of the brain of linking together whatever shows up, that will make it possible to link arbitrary sound forms to content, but we are not yet ready to address language.

The world surrounding him is familiar to H thanks to his brain that holds its neural networks together; these neural networks keep adapting to what happens around him, and thus, in an ongoing process, H's world is constructed and reconstructed as his brain's succession of hypotheses of what is there, entangled with his capacity to deal with it, mixed in with his emotions.

The main points:

The QUESTION is how H knows of C's ongoing existence, especially, how he recognizes C as the familiar thing when it shows up again. Neuroscience's answer is that between the moments C shows up as a unification of features, H's brain keeps a record of what combinations of neurons

to activate in order to repeat the experience, by strengthening the connections between those neurons.

NOTATION

${}^H(N \text{---} \text{---} N)_C$ the complex of networks of neurons, connected in parallel, hierarchical, criss-cross, and feedback, strengthened together in H's brain, that is a ready-to-use record of how to make H aware of C again, taking the shape ${}^H \backslash \bullet \bullet \bullet /_C$.

As above, it is not possible to capture the plasticity of these records in the notation.

One remark before I go on. These strengthened neural networks may form ready-to-use pathways for other neural activity than the time consuming A , and so, earlier experience may influence H's non-conscious reactions. However, there is a big difference between conscious and unconscious perception: it is only while H is aware of something that it becomes something in its own right to him, separated from his own self; it is only then that they become real objects of knowledge from H's perspective rather than things he can be seen seemingly intelligently interacting with from perspective Ω .

2.1.3 Sharing the Human World

Take two human organisms, H and M, both looking at C. From perspective Ω we can see that both H and M have a set of senses tuned to the same kind of signals, and that both pick up signals from C. To H, C shows up in the shape of ${}^H \backslash \bullet \bullet \bullet /_C$, and its record ${}^H(N \text{---} \text{---} N)_C$ only is there in H's brain; to M, C shows up in the shape of ${}^M \backslash \bullet \bullet \bullet /_C$, and its record ${}^M(N \text{---} \text{---} N)_C$ only is there in M's brain. From perspective Ω we can compare them; what we see is both similarities and differences; for both of them, the action results from the impact of C on a brain of a general human design, but in detail, the brains differ from the start, and a different history of experience enlarges the differences. H and M themselves however cannot compare their cats; as long as they cannot share a brain, there is no way for them to see the other's cat, and usually, they do not feel the need to compare: to H, ${}^H \backslash \bullet \bullet \bullet /_C$ equals the cat that is there, so the cat that is there to M as well, and to M, ${}^M \backslash \bullet \bullet \bullet /_C$ equals the cat that is there, so the cat that is there to H as well. And from perspective Ω , the best H and M can do, confined as they both are to their own brain, is what they do already: generalize the world of their own experience.

I believe people's worlds are inescapably private: what their own brain is able to make of what their own senses are able to register of what in fact is there. And I believe for them to ignore its subjectivity, either because they do not know better, or because they do not care, or because they do not really have an alternative, is as close as they can come to sharing their world with others. This does not mean that we, in trying to understand what happens, can ignore its subjectivity as well. But we can generalize the notation ${}^H \backslash \bullet \bullet \bullet /_C$, that is the cat showing up to H, to $\backslash \bullet \bullet \bullet /_C$, that is the cat showing up to H, taken by him as exemplary for what, given the same input, would show up to someone else. This does not represent an abstraction from H's mind; separated from the individual mind unifications of features do not have existence.


The main points:

The **QUESTION** is how two people H and M share the world.

The **ASSUMPTION** is that they take the things that show up before their senses as exemplary.

What we need for the purpose of our theory is a representation of what is there to H taken as exemplary.

NOTATION



the unification of features that C is to H precisely during the activity A in the correlated complex of networks of neurons in H's brain, taken by H as exemplary for what would show up to someone else in the same circumstances.

2.1.4 Having a World Not Depending on Immediate Experience

So H recognizes C the moment it shows up again. But how does H know of the cat's existence when it is out of sight? A brain, triggered by perceptual input originating from the real thing, takes electrical impulses, from neuron to neuron, away from the senses, to the primary ('phenomenal') areas, already taking them apart for specialized analysis, and from there to more and more abstract areas deeper down the brain. The pathways for the electrical signals are unidirectional, but all the way, there are loopholes, neural connections in opposite direction, taking impulses back up the road to keep up and coordinate the action.

When C is in sight, ${}^H\backslash \bullet \bullet \bullet /_c$ takes shape to H; when it fades from H's awareness, the complex ${}^H(N- -N)_C$ remains behind, ready-to-use, and easy to trigger into action as a whole, making the cat take shape again as the same unification of features. The trigger may be any glimpse, or sound, or touch of the real cat. But also, the action can be set off from inside; via the loopholes in the system, the brain itself may trigger those networks of neurons as if they were glimpses from outside. That is, from inside, the brain can take control over the neural machinery that has been evolved to make the world take shape, it can trigger the record ${}^H(N- -N)_C$ into ${}^HA(N- -N)_C$, and this ${}^HA(N- -N)_C$ in H's brain does not go without ${}^H\backslash \bullet \bullet \bullet /_c$ showing up in H's awareness: a simulated perception, the cat as-if-there. It is in this way that H conceives of the cat in its absence; it is not there all the time, but it is there, familiar and well, the split second H needs it, as if it has never been away. Where I speak of mental images, I refer to whatever it is that actually shows up in H's awareness when the neural networks that record earlier experience are re-activated from within.

Adapt the notation for the record to the fact that its activity does not depend on the presence of C in the world:

NOTATION

${}^H(N- -N)_C$	as above: the network of neurons that is the ready-to-use record of how to make H aware of C again, taking the shape ${}^H\backslash \bullet \bullet \bullet /_c$.
${}^H(N- -N)_c$	the network of neurons that is the ready-to-use record of how to call ${}^H\backslash \bullet \bullet \bullet /_c$ to H's awareness.

As above, the notation does not capture the plasticity of these records.

Now in a lot of ways the imaged cat may differ from the perceived cat; the first is fleeting and may seem elusive, the second seems solid enough. But the difference is emphatically not that the perceived cat is objective and scientifically unproblematic, and the imaged cat subjective and obscure. It seems clear to me that private imaging is not the toughest problem for a theory of

language and world. The hard problem is the world: how is it possible that a brain makes its organism aware of the surrounding reality. If we accept the fact that it is able to do that, imaging is not a whole new mystery, as it just takes internal control over a neural machinery already there to do the really tough job.

From perspective Ω , both perception and imaging are inescapably private, as they both cannot be but what H's own brain brings to his own mind, and what M's brain brings to M's mind. And when we look at H and M, it is not as if their percepts are similar and their images are incomparable: if the more or less similar $^H\bullet\bullet\bullet/c$ and $^M\bullet\bullet\bullet/c$ leave, respectively, the records $^H(N\text{---}N)_C$ and $^M(N\text{---}N)_C$, both can be triggered into a more or less similar simulated perception again.

No doubt people with a similar neural apparatus, adapted during experience with the same reality, will be able to create similar images. Still, people self-evidently generalize what they see, but they equally self-evidently take their images to be private. The reason for this is not so much that the real thing of perception is missing; the point is that there is no external trigger to synchronize what shows up in their minds. Their brain can re-create an earlier experience of reality more or less as if it is there again, but it cannot re-create the matching external signals that triggered the first experience; the as-if reality of imaging does not produce the signals it would if it were real, for other people's senses to pick up, and for their brains to get synchronized.

So at this point, H and M share their human world around them; also, both are able to conceive of things out of sight by making it take shape as-if-there, but they are not —yet— able to share their as-if-there.

In comparing images with percepts, perhaps most notable is their fragmentary and changeable character; most of the time, an image is only a partial re-enactment of an earlier experience. This difference is not so strange: perception is bound to what is there, it starts at the senses with a bombardment of signals, then hits the primary areas at the back of the brain, making phenomenal features show up, and goes from there to secondary areas and from there further down the brain, each network in each area that joins the co-ordinated action adding its own detail. In contrast, apparently, the inside trigger may start anywhere, reactivating from there the sub networks in a choice order, plugging in at any accessible point in the perceptual pathway, going with the perceptual direction from there, or going upstream using the feedback pathways, simulating the perception of a choice combination of features, possibly but not necessarily all the way down to the primary areas. If it is only one network that is activated, it is only one feature that shows up; if more than one network is co-ordinately active, it is their unification that shows up; if action in one of those networks is shifted to action in another combination of neurons in the same area, the corresponding feature morphs into another feature; for instance, an isolated sensation of red may grow into a red circle, red may change into orange, a circle may grow into a ball, etc.

The main points:

The **QUESTION** is how H knows of C's ongoing existence, especially when it is out of sight.

The **BASIC ASSUMPTION** is that H's brain takes internal control over its neural machinery to make the world take shape, by activating from within bits and pieces of the ready-to-use records of earlier experience, making less or more detailed fragments of real things take shape as-if-there.

$^H A_{\text{ginger}}$ does not go on in H's brain without the colour ginger showing up in H's awareness, and

${}^H A (N \text{---} \text{---} N)_c$ does not go on in H's brain without ${}^H \underline{\bullet \bullet \bullet} /_c$ taking shape to him. This is how H first conceives of what is not present.

Features and things imaged might be represented in the same way as features and things perceived: ${}^H \bullet_{\text{ginger}}$ and ${}^H \underline{\bullet \bullet \bullet} /_c$, but H does not take the images as exemplary and therefore the notation cannot be generalized. Before I can use them in the theory of verbal communication, this notation requires some adaptation.

2.2 How Forms of Language Exist

2.2.1 Hearing and Recording the Words

The next question is: what is language? How does it exist? Focus again on H, assume that somehow on his way to master the universe he has acquired forms of language. Postpone the question how he got there, and the question what he has them for, and focus on the question what is their way of existence.

It is almost as if forms of language, like the real things in the world, somehow have an existence of their own, independent of the human mind. When we take perspective Ω , and compare H who perceives the world, and H who hears forms of language, we can see both differences and similarities. As for the differences: when H sees the cat, there is this C that betrays its presence by signals that H's senses can register; given this input, H's brain conjectures the origin of the signals, and presents its results in H's awareness. We can distinguish between what shows up to H, this phenomenal cat ${}^H \underline{\bullet \bullet \bullet} /_c$, and the object of observation, C. But when H hears a speaker's *cat*, obviously there is no Verbal Presence that betrays its presence by the sequence of sounds that triggered H's hearing, and the word *cat* that his brain comes up with is not a guess what caused the sound waves, but a guess what the speaker said.

As for the similarities: the word *cat* that H hears is not so different in its way of existence from the phenomenal cat: the areas of H's brain specialized in figuring out forms of language from sound waves are not fundamentally different from the areas specialized in figuring out reality; just like the cat itself, the word comes to H's awareness precisely during this special activity A in a certain complex of networks of neurons, in specialized areas of his brain. And like the cat he sees, the word he hears results from a public trigger, for everyone to pick up. And also this network, by its very use, gets strengthened together to become a ready-to-use record, so that H will instantly recognize the sound form in all the noise that enters his ears. And also, H can re-activate those records from inside, thinking up the words for himself.

To represent forms of language, their neural correlate and their records, we can use notations that mirror the notations used above: the complex networks of neurons involved can be represented as ${}^H (N \text{---} \text{---} N)_{cat}$, the action in this complex triggered by the speaker's sequence of sounds as ${}^H A (N \text{---} \text{---} N)_{cat}$, and the sound form it makes show up as ${}^H cat$. When this complex of neurons drops out of action A , the word ${}^H cat$ slips from H's mind, until reactivated. In between, ${}^H (N \text{---} \text{---} N)_{cat}$ holds the record.

The main points:

The QUESTION is: what is the way of existence of forms of language.

The **BASIC ASSUMPTION** is that a word is there to H precisely as long as activity A goes on in the correlated complex of networks of neurons in the language areas of H's brain, triggered into action by the speaker's sequence of sounds entering H's ears, or triggered into action from inside H's own brain, which calls up the sound form to H.

NOTATION

$$\begin{array}{c} {}^H A (N \text{---} \text{---} N)_{cat} \\ \hline {}^H cat \end{array} \quad (< < (S))$$

the equation between on the one hand activity A going on in the complex of networks of neurons ${}^H(N \text{---} \text{---} N)_{cat}$ triggered by a speaker's sequence of sounds or from inside, and on the other hand the sound form ${}^H cat$ in H's awareness.

$${}^H(N \text{---} \text{---} N)_{cat}$$

the complex of networks of neurons strengthened together in H's brain that is the ready-to-use record of how to hear ${}^H cat$ again.

As above, the notation does not capture the plasticity of these records, however, words and their records are comparatively stable.

Forms of language do not have existence independent of the private human mind. They come into existence whenever needed; in between, the strengthened complex of neurons keeps a record.

2.2.2 Sharing the Words

How do people share their words if words cannot come out of the private mind? Compare words and images: both can be triggered from inside H's brain, both do not have existence outside the human mind, but words seem to be there to share, while images seem to be inescapably private. This is what makes the difference: H's brain that simulates perception cannot reproduce external signals that would have originated from the real thing had it been there, but H's brain that makes a word come to his mind can produce corresponding vocal sounds, which are precisely the sounds that would have triggered the appearance of this same word in his own mind had it not already been there, and that now may synchronize other people's minds. The networks of neurons required for making those vocal sounds are linked to the networks of neurons required for hearing the sound form; they are to be included in the notation ${}^H(N \text{---} \text{---} N)_{cat}$.

Like people's worlds, people's words are inescapably private: what their own brain brings to their own mind. But people ignore this; they self-evidently take their own words as the words of others as well. This is as close as people can come to sharing their words: any time, people of the same speech community can simulate the existence of their words to each other by uttering vocal sounds that trigger some pre-strengthened complex of neurons within the language areas of the other's neural machine; it is the words that emerge from their own brain that hearers take to be the speaker's words. It is always the same record in H's own brain that is flashed into action by rather different inputs from different speakers, so each time the same sound form emerges. Ironically, this reinforces the impression that the word has an existence of its own.

In the notation, we can switch from ${}^H cat$ to the generalized cat .

The main points:

The **QUESTION** is: how do people share their words if words cannot come out of the private mind. It is a fact that people can produce vocal sounds corresponding to the words they think up,

synchronizing the sound forms in other people's minds. The ASSUMPTION is that people share their words by taking what shows up to themselves as exemplary for what would show up to others. What we need for the purpose of our theory is a representation of the words that come to H's awareness, which he takes as exemplary.

NOTATION

<i>cat</i>

the sound form that comes to H's mind precisely during activity A in the complex of networks of neurons ${}^H(N- -N)_{cat}$, triggered by a speaker's sequence of sounds or triggered from inside H's brain, taken by H as exemplary for what would show up to someone else given the corresponding input.

2.3 Forms of Language and the Human World

2.3.1 Having One's World Covered with Words


In 2.1 above, we followed the organism H on his way to master the universe up to the point where his brain has managed to provide him with a familiar world that goes on to exist when out of sight, and then in 2.2 we jumped to the point where he has mastered the sound forms of language. Now go back to the earlier stage: H is a pre-verbal infant within a community that already has speech. Apparently, he learns to recognize those sequences of sounds that seem to be important to the others. But what is the gain? In themselves, these combinations of sounds are completely arbitrary; they only become useful when connected to something that is relevant in itself, that is, certainly for the first words, something concrete, something to see, to grasp, to squeeze (or, in case of the cat, something that scratches back). And so we approach one of the central questions of this book: what is the connection between verbal information and the world. We have to start with the words: what is the connection between words like *cat* and the real thing. When we look at the real thing C as it is and at the word *cat* we see two very unlikely partners: not only are they incongruous, but they have a completely different way of existence, C objectively there in its own right, and *cat* that does not escape the human mind. There is only one place where those incomparable things can meet, and that is the private mind emerging from the individual brain; it is only here that they have a way of existence in common, C not as it is but as it shows up to H precisely during ${}^HA(N- -N)_C$, and *cat* as Hcat during ${}^HA(N- -N)_{cat}$. Although they can meet now, the unification of features ${}^H\bullet\bullet\bullet/c$ and the sound form Hcat still are incongruous; however, in awareness, things that have nothing to do with each other can become connected: if only they simultaneously show up—the people around H will have to arrange that—H's brain will find and strengthen pathways linking the two responsible networks together, and record the connection ready-to-use again. H's brain will have to do some strengthening, re-strengthening of what returns again and again and pruning what accidentally shows up together, in order to fine-tune his word-world connections with those of others, tagging Hcat to the cat-features, Hginger to the ginger-feature, etc., but eventually, the record of the word, ${}^H(N- -N)_{cat}$, grows a direct link to the records of the cat-features, which is a sub network (in fact, a plastic complex of networks) of the records of this particular ginger cat, of this other grey tabby, etc., and that may take part in a vast amount of other unifications of features. Thus, the comparatively stable network that records the word *cat* becomes connected to an ever more

branching complex of complexes of networks that record all kinds of cats and experiences with cats, changing with each new experience.

Now we can represent the link between the records, the action in this complex, and what shows up in H's awareness precisely during this action as follows; respectively:


${}^H N_{\text{cat}} | (N- -N) \text{cat}$ the records of the cat-feature and the record of the word *cat* linked together.

${}^H A (N_{\text{cat}} | (N- -N) \text{cat})$ these records active together.



the feature and the word tagged to it as they show up to H during this activity.

In this latter notation, the status of 'cat' and '*cat*' differs; 'cat' is part of the notation, it is there for our benefit, an index to identify the feature that is represented; as such, as a word, as an index, it does not show up in H's awareness, while *cat* represents the sound form that shows up to H. Now in this notation, the identifying index has become redundant, as the word that H tags to the feature also identifies it, so:




Note that this complex feature, like the whole world around, needs to be there already to H before the word can be attached to it. It is the feature that gives sense to the word, not the other way around. Note also that we can only switch from the notation ${}^H \bullet_{\text{cat}}$ to ${}^H \bullet | \text{cat}$ if it represents what is known to H himself. This switch is not just a switch to a more elegant notation for the same thing; it represents something more, a change in H's brain, a tag connected with this phenomenal feature. Also, the notation changes from a representation of a phenomenal feature into a representation of a phenomenal feature that in principle can be talked about, not only by us, but by H himself.

In an ongoing process, H picks up from others more and more sound forms to cover the world of his experience. His first vocabulary will have to be learned simultaneously with his exploration of the things and features they are to cover; if not, the words will have no link to reality. These first words will cover concrete things, phenomenal features, down-to-earth aspects of reality; for these concrete first words, it is relatively easy to synchronize experience between him and members of his speech community-to-be. There is no rigid body of first vocabulary that is the same to everyone, but every new member of the speech community will have to learn a fair amount of words directly from others and in direct and conscious connection with the real thing, otherwise there will be no bridge between language and reality. Once a firm bridge has been established, new words can build upon what has been achieved.

Thus H's brain will grow a forest of precisely placed verbal tags. They change H's originally largely phenomenal world: not only does it become a world that can be talked about, but the words change what they label. H's idea of the world starts with perception and experience; with first experience, his brain analyzes the real thing as a unification of building

stones, turning signals picked up into phenomenal shape; it cuts up what is there into features that are not so much building stones of the objective thing, but of his grasp of those objective things; it categorizes those building stones, recording them, arranged by type, in specialized areas, and it holds unifications together by preserving links. With each new experience, H's brain enlarges and revises this complex of parallel, hierarchical, and criss-cross networks of neurons that record H's world. With the addition of verbal tags, this process goes on, the first vocabulary of words rather haphazardly pinned here and there on the reality that shows up to the eye, and then in increasing number also connected to each other; then, words may make connections between things in the world that do not systematically have phenomenal features in common, continuing the process of analyzing the world. Running ahead of H's development: later on, new words may abstract away from the experiential features, tagging aspects hidden from the eye, supposedly but not visibly there, up to the point where phenomenal shape may seem the least important of characteristics, or even a distortion of the thing as it is. However this may be, phenomenal experience is and remains the first and the final link between people's words and what is there. In order not to complicate the discussion, let us return to H at the point where he is, learning to pin words on what is in sight. Anyhow, this is an ongoing process of getting hold of the world. Each new experience makes its own small or major change in the criss-cross of records, and each word added makes its small or major change, and each change in the whole reshuffles everything connected to it.

Let us assume that there is something in H's view. What we need for the purpose of our theory is a notation to represent it. To H, it is a unification of features, that is, his brain's state-of-the-art conception of what is there. As illustrated above, we can represent individual features via the verbal tags that H himself would put on them. Now suppose we would ask him to think up all suitable words recorded in his brain, in order to capture the whole thing that is there to him in the words he has available; he would start the job, come up with more and more words, construct pyramids of them, connecting them in parallel, and on top of each other, but trying to sort out their connections, he would not finish the job: too complex, too many loose ends, too many hesitations. It will not do for us to use H's imaginable pyramids of words to represent what is in front of him; from perspective Ω , we can see the reason why: how to mirror the immense complexity of the brain's state-of-the-art parallel, hierarchical, criss-cross, and feed-back connections of unequal strength, how to keep track of its ongoing changes? Fortunately, for our purpose we do not need such a full-blown representation; we can restrict it to what is relevant; what we need is a representation of the real thing as it comes to H's mind, and out of all the words that H has available, recorded in his brain, he will have only a few simultaneously active. We can assume that, roughly, the words active in his mind correspond to those aspects of the real thing that have H's attention. For instance:




The main points:

A **CENTRAL QUESTION** is: what is the connection between words and real things; they do not have the same kind of existence, and they are incongruous.

As a start to the answer, the **BASIC ASSUMPTION** is that it is only as they show up in the private mind that they have a way of existence in common. Also, the brain connects unrelated things that simultaneously show up in awareness by linking the records ready-to-use together, tagging the words to their features.

What we need for the purpose of our theory is a representation of the connection as it shows up in H's awareness:


NOTATION



the (complex) feature ${}^H cat$ is linked to as they both show up to H precisely during activity A in their respective correlated networks linked together.

As for unifications of features, what we need for the purpose of our theory is a representation of such unifications enriched with the words tagging features insofar as those words are active:

NOTATION



the unification of features c , including the feature tagged by *cat* and the feature tagged by *ginger*, and other features not in focus, that shows up to H precisely during activity A in the correlated complex of networks of neurons, including ${}^H N_{cat}$ linked to the records of the cat-features, and including ${}^H N_{ginger}$ linked to the records of the experience of ginger.


2.3.2 Sharing a Word-Covered World

Before any verbal communication between H and others will have any chance of success, H will have to learn a basic vocabulary of words and adjust it to the vocabulary of the others, that is, H's brain will keep sorting out combinations of words and things, one or a few at a time, until in new experiences he applies the same words to the same things as others do. From perspective Ω , the comparison between the pyramids of words as recorded in H's brain, and the pyramids of words as recorded in M's brain will show similarities and differences. However, once they both have this vocabulary, given the same world in view, and the same selection of words from that vocabulary, H and M will roughly pin the same words on the same things. If this is the case, there is a basis for verbal communication.

People take the world that shows up to them as the world that is there; we can say that without thinking they take that world as exemplary for what would show up to others in the same circumstances. When it comes to covering the world with words, people do the same thing: when given a choice of words they take their own distribution of those words over what is there as exemplary for how someone else would distribute these words over this reality.

What we need for the purpose of our theory is a representation of unifications of features, tagged with the words that are active, and taken as exemplary:

NOTATION




the unification of features the same as above, taken by H as exemplary, that is, assuming that someone else, given the same words, would distribute them in the same way over what is there.

2.3.3 Sharing a World Not Depending on Immediate Experience

We have arrived at the central question of this book: how is it possible that sounds can get information from a speaker to a hearer.

As said already, experience with a cat leaves a ready-to-use neural network in H's brain that records how to make it show up. A glimpse from this real cat can trigger this record into action, making the cat show up in H's sight; the same goes, *mutatis mutandis*, for M. H will take this cat as exemplary. Now in the absence of the cat, H's brain can trigger this record into action from within, making the cat show up in H's mind. Again, the same goes, *mutatis mutandis*, for M, but this time, H will not take his image as exemplary. Why? As argued above, it is not because H's percept would be objective and his image subjective; both are what his own brain brings to his own mind. Neither is it because H's image of the cat would be totally different from M's image of the cat. The reason why H generalizes what he sees and not what he images, is that the first is, and the second is not triggered by a public source that synchronizes different brains and minds into a hopefully similar show. The problem with other people's images is their lack of synchronization. Both H and M can call up the cat as-if-there for themselves, but they cannot make it leave their private space, they cannot re-create the external signals that would originate from the cat if it were there, to activate the other person's record of the cat, to make the cat show up there.

It is here that language comes in: the cat and the neural networks that record its appearance are linked ready-to-use to the sound of *cat* and the neural networks that record it, and what people can do is produce the sound, which calls up the word, which flashes the records of the cat into action as well, triggering them in this way; this means that where they cannot directly project their stand-in cat in the other person's space, they can give each other something that is almost as good. This is the magic of language: someone else's words conjure up the cat in its absence; forms of language give shape to what is out of sight. The speaker's sequence of sounds almost replaces a glimpse of the real thing: while it is not there, it shows up as-if-there. This is the basic assumption about verbal information in this book: with his sequence of sounds, the speaker takes remote control over the hearer's brain's machinery to make the world take shape around him. What might happen is this:




However, the real cat is complex, and only some of its features are covered by *cat*. The other way around, the word *cat*, or rather the neural network that records it, is linked to a branching complex of overlapping and diverging networks each recording its own cat, of which the record of this particular cat is only one. Had the action started with *cat* taken another pathway, there would be another cat in H's mind. And if S utters his synchronizing *cat*, there is no guarantee whatsoever that the cat called up in H's mind will, apart from being cat-shaped, be similar to M's cat. And so the synchronization that words can achieve remains fragmentary. Especially when the words tag aspects that are not concrete, the actual image that takes shape to H will be far less similar to the image that takes shape to M than a synchronized perception would be.

Nevertheless, there is a basis for communication: given both a choice of words and the world in view, H and M will roughly distribute those words in the same way over that world. It is not the in itself arbitrary sound form, incongruous as sound form and real thing are, that gives him the

key to this distribution, but it is for both H and M their private, fleeting, shape shifting image called up by the sound form that in one shape or another can fuse with what is there.

H does not generalize his self-triggered images, but up to a point, he does generalize his images called up by S, which he takes to be similar enough to the images of someone else to recognize the same reality. I shall use the term ‘word-triggered image’ to refer to whatever it is that actually shows up in H’s awareness when corresponding neural networks are accessed and activated via the words heard in S’s sequence of sounds. This word-triggered image is not a fixed picture; when the action halts close to its point of departure, the image remains dim, when it flashes through one or another of the pathways accessed, the image becomes more detailed, when it retreats and takes another path, the image changes its appearance. It is the part of the image that remains close to its point of departure that is authorized by S; it is this part that H takes as exemplary for what would show up to other people. But it is the whole realm of possible appearances that forms the meaning of the words to him. What we need for the theory of verbal information is a representation of these word-triggered images insofar as generalized. I shall use the following notation:



I believe these word-triggered images are the sole basis of verbal communication. Different though in detail the image that shows up to H may be from the image that shows up to M, H’s image is the shape that this information takes to H, it is what it means to him, and it is what connects the speaker’s words with the real thing, and the same, *mutatis mutandis*, goes for M: M’s image is the shape that this information takes to M, it is what it means to him, and it is what connects the speaker’s words with the real thing. This real thing is the same for H and M.

The main points:

The **QUESTION** now becomes: how can words be used to get information from S to H.

The **BASIC ASSUMPTION** about verbal information is that a speaker with his sequence of sounds takes remote control over the hearer’s brain’s machinery to make the world take shape around him. A speaker’s words call up images to the hearer.

Up to a point, people generalize these word-triggered images. People conceive what is absent from view by imaging it as-if-there, triggering the neural networks that record earlier experience; words synchronize images by triggering the records of the features they are linked to.

What we need for the purpose of our theory of verbal information is a representation of word-triggered images insofar as these images are taken as exemplary:

NOTATION

\boxed{cat} > $\boxed{o|^{cat}}$

the feature-image that shows up to H triggered by the word *cat* he hears in the speaker's sequence of sounds, precisely during activity *A* in the records of both *cat* and the feature-image, taken by him as similar enough to the image called up to someone else to enable fusion with the same real things.

Let us return to the organism H on his way to master the universe; he has created a world around him, he is able to conceive what is absent from his view by imaging it, he has tags on what is there; these tags are roughly the same as those used by other members of his speech community; H and others are able to activate corresponding tags in each other's mind, and in doing so, call up corresponding fragmentary images in each other's minds. What H needs now is the art of piecing the fragments together.

Chapter 3

The Skeleton of Information

Recapitulating: between speaker S and hearer H, there is nothing but sound waves, but H grasps the thing talked about whether or not it is in view, and if the real thing is in view, he recognizes it. The general question of this book is: how does it work. The general idea of this book is this: H hears a complex of simple forms in S's sequence of sounds, these forms call up an image in H's mind; each form systematically makes its own small but precise contribution to the picture; the separate contributions unite into a fragmentary shape, which to H stands-in for what is out of view, and fuses with the real thing when it is within H's view.

The special problem is that S can use the same subject-predicate construction to get a multiplication of messages across. The special idea of this book is that in addition to syntactic form there is a systematic role for prosodic forms; a small number of simple choices in pitch and rhythm enable a multiplication of possible combinations; along with the other forms, each of these prosodic forms makes its own small but precise contribution to the emerging picture. In chapter 1 I have formulated some hypotheses about what precisely these prosodic forms are, and some working hypotheses about what they each contribute.

To give these ideas more precision we must get hold of the forms of language that H picks out from S's sequence of sounds, we must get hold of the images that they, fragment by fragment, call up in H's mind, and we must get hold of the real thing itself as it comes within H's view. In order to achieve that, in chapter 2 the basic questions were asked: what is their way of existence.

The hard problem is how H knows at all that there is a world around him. Certainly, he needs a brain; it turns signals picked up by his senses into an idea of their origin. This takes some special coordinated activity in complexes of networks of neurons dividing labour over the expanses of the cortex. This special activity in H's brain is necessary and sufficient for H's awareness to emerge; the combination of neurons involved determines what shows up in awareness: a unification of features, each feature contributed by the action in its own sub network in its own area of H's brain. The existence of a biological activity that simultaneously is its organism's awareness is hard to understand but seems to be a fact. The basic assumption of this book is that things around him exist to H separated from himself precisely as long as this action goes on in the corresponding networks of neurons, that is, precisely as long as his brain makes him aware of them. In between, the networks strengthened together keep ready-to-use records of how to call them into existence again. Then, images take shape when, in the absence of the thing itself, internal triggers take over from signals from nature; they may activate small networks to make a single feature show up in his mind, and complexes of networks to make unifications of features show up as-if-there. The whole criss-cross of networks of neurons strengthened over the years by experience with what is there forms H's ability to make a world around him take shape whether or not it is physically there before his eyes.

The sound forms of language are made in specialized areas of the brain. Each form exists to H precisely as long as his brain makes him aware of it, that is, precisely as long as the special action goes on in its corresponding networks of neurons. In between, strengthened networks keep ready-to-use records. These records can be activated by sound waves produced by S, so S can

trigger forms of language into existence to H; also, H can activate the records from within, and in his turn, he can produce corresponding sound waves. Sound forms have no use unless systematically linked to people's ability to conceive the world, sounds linked to features, networks linked to networks. When linked, words triggered into existence in their turn trigger the corresponding piece of world into shape.

Suppose S wants to get something across to H. Assume that he has some cat in view. It shows up to him as a unification of features; this takes an enormous amount of co-ordinated action in his brain. H would share this experience if a similar unification of features would show up to him; this would take a similar complex activity in H's brain, which would be triggered if H's senses would pick up the same signals as S's senses. This however is not the case. So S takes over with the means he has available, putting words into sequences of sounds that on H's side become words again, that in their turn trigger the corresponding piece of world into shape. Replacing the input from nature by his sequences of sounds, S takes direct control over H's brain's already existing mechanism to make the world around him existing to him. The Human H of chapter 2 has become a Hearer. And in case H does get the cat in view after all, input from nature and S's words join forces in making it show up.

By now, we can return to the working hypotheses of chapter 1, and give them more precision. The ultimate goal is a formal empirical theory that captures the actual process that goes on when H makes sense of S's sequence of sounds. What we need in the first place is an accurate representation of the forms recognized by H in S's sequence of sounds; in the second place an accurate representation of the fragments of shape each of these forms contribute as well as of their combination; and in the third place an accurate representation of the real thing when it comes in sight, if and insofar as it comes in sight. Taken together, it should work. For lack of time, I can only make a start.

In the first place, forms of language will, as above, be represented in a more or less orthographical form; it is far from perfect in representing the phonemes that build up words, but as a representation of informative form it has the advantage that we already are skilled in reconstructing those sound forms from writing. Punctuation marks will be left out, as they do not systematically represent what comes with the sounds; I shall add representations of pitch and rhythm as introduced in chapter 1; hypothetically, they represent what H's ears actually pick up from the sounds. This representation is precise but incomplete: there is more to prosodic form than can be discussed here. As for structure, at the moment my idea of what is there on either side of the sound waves is not clear enough to attempt a representation of syntactic form. For the time being, I shall limit the discussion to simple third person singular subject-predicate constructions, and just assume that H recognizes the noun phrase, the predicate, and their connection. To simplify in the discussion the comparison between the different types of information, I shall each time use *a cat (not) predF*; *a cat* as a standard example of a noun phrase, and *predF* to represent a predicate in finite form, which can be replaced by any real sentence. This representation of form calls for improvement, but this does not affect the gist of the argument below. The gain of chapter 2 is that it precisely defines what kind of thing it is that these graphs, however imperfect and provisional, represent: what H's brain brings to H's mind precisely as long as the sequence of sounds produced by S and picked up by H's ears flashes the corresponding networks of neurons into this special action that is H's awareness.

In the second place, the idea is that the forms of language call up an image. The working hypotheses of chapter 1 described the contribution of each of the forms to this image, so, they replaced the forms that call up the image by a description of the image called up. The gain of chapter 2 is that it starts to replace the description of the image by a direct representation of this shape, congruent with what is represented. Also, chapter 2 precisely defines what kind of thing it is that is represented: what H's brain brings to H's awareness precisely as long as the networks of neurons in the areas of experience in H's brain, triggered into action via the words called up by S's sequence of sounds, are active.

In the third place, in chapter 2 I made a start representing what is there in H's sight, defining what exactly it is: what shows up to him precisely as long as signals originating from the real thing, hitting his senses, keep corresponding networks of neurons in action: the real thing to H.


First I shall focus on the process of understanding in the absence of the real thing (see 3.1 and 3.2); next I shall focus on the process of recognition when the real thing is in sight (3.3).

3.1 Understanding Information

The word-triggered feature

So S has this cat in view. On his side of the sound waves, he might think up a complex of state-of-the-art pyramids of words to cover each of its discussible aspects, but in fact, only a few of them will be active; the rest, via the links to their records, is available but not active. Each of these words, whatever its place in the hierarchy of the imaginary pyramids, when turned into vocal sounds, triggers the appearance of a corresponding simple or complex feature in H's awareness. Take *cat*:

NOTATION



S has the whole cat in view, but he has to make a choice from all the words available; with his choice, he can authorize only so much of the picture; by choosing *cat*, he calls up a cat-shape, but leaves out colour, smell, sound, etc. However, with this single feature, S does plug in at a particular point in H's recorded ability to conceive the world: within the whole criss-cross of networks of neurons strengthened over the years by H's experience with what is there, S plugs in at the records of H's experiences of cats, where they, roughly, overlap: $o|^{cat}$ shows up. Now from this point of departure, any of those more complex records of cats may flash into action as well, colouring the cat-shape, giving it detail and identity. And each point arrived at may be a point of departure itself, shaping a chain of associations. As for the records of cats: with $o|^{cat}$, S's *cat* accesses the criss-cross of overlapping and diverging records of phenomenal features, characteristics, actions, etc., which in H's past experience have gone together with the cat-shape. From perspective Ω , we can see them as a stock of potential complementary features, features excluding each other, located in the same area of the brain, like 'fat or skinny', and features that go together, like 'fat and ginger'. To H, these features only come into existence when their

records are flashed into action. From perspective Ω , we can see that the stock of potential complementary features determines the range of unifications of features that H may assemble in his mind: if $o|^{cat}$ would access just ‘fat or skinny’, and ‘black or ginger’, H could think up only a fat black cat, a skinny black cat, a fat ginger cat, or a skinny ginger cat, but with each recorded feature, the range of thinkable combinations multiplies, soon to become virtually infinite. This stock of potential complementary features in H’s brain is not an objective analysis of the reality of cats, but it is H’s brain’s state-of-the-art record of his interaction with cats. The stock is of enormous complexity, with its parallel, hierarchical, and criss-cross links of unequal strength, links strengthening and changing with each new experience, and gradually deteriorating with lack of experience.

So this whole complex of knowledge already is there in H’s brain before S tries to get his information across; H needs it to give flesh and blood to S’s information; S cannot actually send meaning along with the sound waves, but he can authorize what shows up in H’s mind by calling it up. He can, however, only call up what is already there. It is clear that any effort on our part to chart H’s stock of complementary features of $o|^{cat}$ could by far not do justice to the complexity and dynamics of H’s grasp of the reality of cats, and would be outdated before it could halfway be finished. But we can get hold of the information that does come across: S very efficiently plugs into H’s complex and plastic ability to conceive what is there with precisely measured bits of information.

S and H have a shared vocabulary of words; each of S’s words that H picks up calls up the corresponding feature. By themselves, word-triggered features are unconnected. The basics of connection are simple: the working hypothesis of chapter 1 is that the noun phrase calls up the outline of a unification of features, filling in the feature(s) called up by its word(s), and leaving room for complementary features. Also, it accesses a stock of records of such complementary features, which determines the range of unifications that H might think of. Prosodic forms either fill up this room or leave it open, and either make a single image or replicate it again and again. The predicate includes another feature-image in the unification of features under construction.

The noun phrase skeleton

So S has this cat in view. Out of all the words in principle available, he chooses *cat*. This *cat* only covers part of the cat’s features, excluding for instance colour, and to H, it calls up just the word-triggered feature-image $o|^{cat}$. To cover the whole cat, including colour and all other characteristics, S can make a noun phrase construction, like *a cat*, or *the cat*. On H’s side of the sound waves, this *a cat*, or *the cat*, still does not contribute colour or any other characteristic, but more than just *cat*, it calls up a cat-shaped outline, a skeleton of a unification of features consisting of the feature $o|^{cat}$ and room for features that would complement it into a full blown cat. I shall represent the image this far as follows:

NOTATION

forms recognized:	image called up	
<i>a cat</i>	\o ^{cat} _____ / ^{a cat}	>
		the skeleton of a unification of features fitting <i>a cat</i> , consisting of o ^{cat} and room for complementary features.

Different noun phrases like *a cat*, *the cat*, *Sophie*, *water*, basically work the same way: they all call up a fragmentary skeleton; they all cannot authorize more than what they explicitly specify; for a discussion of their differences, see chapter 6. For now, I shall focus on the indefinite noun phrase.

Prosodic information

The information of the prosodic forms ↓ or ↑ and * or ↔ in the final syllable of the noun phrase builds upon the noun phrase skeleton. As for the pitch pair, in chapter 1 I stated as working hypotheses: ↓ fills up the room in the outline of the unification of features called up by the noun phrase with a set of unspecified complementary features; ↑ keeps open this room. I shall represent this information as follows:

NOTATION

o o	unspecified features
...	room kept open


As for the rhythmic pair, in chapter 1 I stated as working hypotheses: ↔ turns the unification under construction into a self-replicating image, * makes a single image. To keep the notation simple, I shall keep the single image as it is; omitting other information, I shall represent a replicating image as follows:

NOTATION

{ } { \ _____ / } is self-replicating \ _____ /

The possible combinations of noun phrase skeleton and prosodic information results in four images:

forms recognized	image called up	
<i>a cat</i> ^{↓*}	\o ^{cat} o o / ^{a cat}	>
		a single closed picture consisting of the skeleton of a unification of features fitting <i>a cat</i> , containing o ^{cat} and unspecified complementary features.
forms recognized	image called up	
<i>a cat</i> ^{↑*}	\o ^{cat} ... / ^{a cat}	>
		a single open sketch consisting of the skeleton of a unification of features fitting <i>a cat</i> , containing o ^{cat} and room for complementary features left open.



Unauthorized details

These notations represent the image that shapes precisely the information that S authorizes. H may stick to it. Or he may tentatively fill in details for himself: $o|^{cat}$ accesses a stock of recorded complementary features that determines the range of cats that H may assemble in his mind. This range is virtually infinite, but H will take into consideration only a fraction out of this range.

If the image includes $o\ o$, any detail that comes to H's mind is a guess of what S actually is talking about. In case H takes into consideration more than one complement, his mind jumps from one guess of what is there to another, from this fat ginger cat to this skinny ginger cat to this skinny black cat, etc.

If the image includes \dots , that is, has its room kept open, H may try out any complementary detail, morphing detail into detail, shrinking a fat ginger cat into a skinny one, taking out its colour, thinking it black again, etc.


As long as H only has a single image, there only is room for one set of complementary features at the most, so in case H takes more than one complement into consideration, he has to clean up that room first; each new option goes at the expense of the earlier one, skinny replacing fat, the skinny black cat replacing the fat ginger cat.

If the image includes $\{ \backslash ______ / \}$, that is, if it self-replicates, H adds each new option he thinks of to the earlier ones.

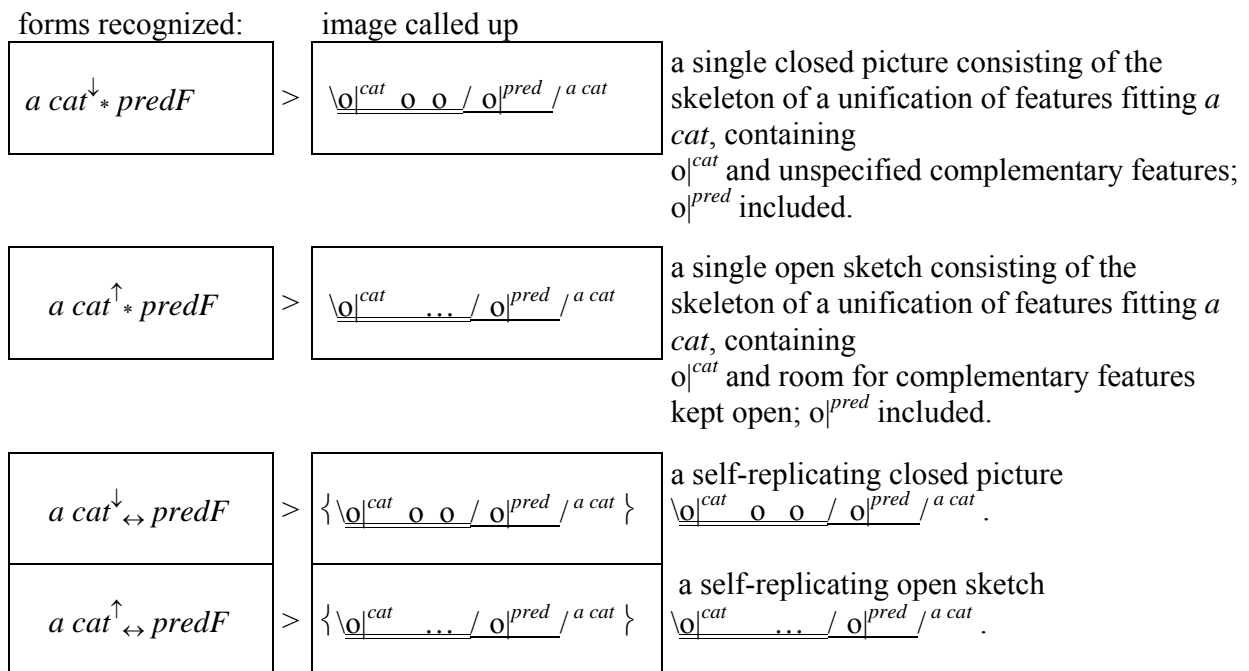
The subject-predicate skeleton

For lack of time, I can only make a start. Represent predicates as $predF$, where $pred$ stands for the symbolic part of the form, and F represents its constructive part. The symbolic form may be a verb, and the constructive form its finite form: *sleeps*. Or it may be a passive: the past participle of a transitive verb plus a finite form of *to be*: *is invited*. Or the predicate may be a noun phrase or an adjective plus a finite form of *to be*, *to remain*, etc.: *is a cat*, *is beautiful*, etc. This notation leaves a lot to be wished for; it ignores the fact that predicates may have a complex inner structure, it ignores tense, and it does not represent the necessary agreement between grammatical subject and predicate, in person (first, second, third) and number (singular, plural). Anyhow, the hypothesis is that all these different predicates, all provisionally represented as $predF$, in the same way make their contribution to the image under construction: the working hypothesis of chapter 1 is that a predicate includes the feature-image(s) it calls up in the outline of the unification of features called up by the subject noun phrase. Omitting prosodic information, I shall represent this combination as follows, keeping the subject-predicate structure visible:

NOTATION

**3.2 Four types of Information**

The possible combinations of noun phrase skeleton, prosodic information and predicate, result in four images:



Remember the *tomcat* examples (25) thru (28) given in chapter 1. For lack of time, I cannot address their composed predicates here, but they very well illustrate the four types of information:

- (25) $a \text{ tomcat}^{\downarrow} * \text{is an especial beauty}$
 (26) $a \text{ tomcat}^{\uparrow} * \text{is a good choice}$
 (27) $a \text{ tomcat}^{\downarrow \leftrightarrow} \text{is a good mouser}$
 (28) $a \text{ tomcat}^{\uparrow \leftrightarrow} \text{is a male cat}$

In all four cases the image called up insofar as authorized remains fragmentary. In all cases it is up to H whether or not he takes unauthorized detail into consideration. Look at the four cases in more detail:

TYPE $a \text{ cat}^{\downarrow} * \text{ predF} > \underline{\backslash \text{O}^{\text{cat}} \text{ O } \text{O} / \text{O}^{\text{pred}} / a \text{ cat}}$: a stand-in for the real thing.

What S authorizes is the picture of a complete cat, many of its features unspecified; one further feature is specified. The whole picture gives H a ghost version of the cat S is talking about to stand-in for it while it is out of view.

In this type of information, the noun phrase seems to say ‘what the information is about’; in fact however, this ‘subject’ information is composed of smaller parts: the noun phrase itself calls up just a skeleton with gaps, \downarrow compensates as far as it goes for the missing features, and * makes sure that there is just the single picture. It is as if S, with the use of \downarrow *, puts a single set of complementary features in a closed box, sending it over to H together with O^{cat} , but of course, H cannot open that box at the other side of the sound waves to reveal its content. Take for example (25) (*they love animals and have many pets*) *a tomcat* \downarrow * *is an especial beauty*. Given the fact that the picture as far as authorized by S is partly unspecified, H can still assemble different cats in his mind; he may take a Persian into consideration, replace it by an Angora, and so on, but each option is very possibly a wrong guess of what S’s one and only actual cat looks like. The cat may in fact be a Siamese.

Note that the picture called up in H’s mind represents to him the real thing as long as it is out of his view, and the notation $\underline{\backslash \text{O}^{\text{cat}} \text{ O } \text{O} / \text{O}^{\text{pred}} / a \text{ cat}}$ is a theoretical representation to us of the mental picture that actually takes shape to H.

TYPE $a \text{ cat}^{\uparrow} * \text{ predF} > \underline{\backslash \text{O}^{\text{cat}} \dots / \text{O}^{\text{pred}} / a \text{ cat}}$: a fragmentary sketch

What S calls up is just an unfinished sketch of a single cat, one extra feature included. It is obvious that things in the world are not fragments of what they are, but because of the way H conceives of the world, it is easy enough for him to image fragmentary things. They do not stand-in for something there; possibly, there is not even something there that fits the noun phrase; compare (13) *an answer* \uparrow * *remains forthcoming*; possibly there is something and possibly there is more than one thing fitting the noun phrase; if so, the sketch fits them all in the same manner: (26) *a tomcat* \uparrow * *is a good choice*, (9) *one of them* \uparrow * *is redundant*.

S may use this type of information to sketch a design for the future, or an expectation, a wish, a demand, a prediction, or an otherwise imaginable but not, or not yet, or not completely realized situation. I shall, for briefness’ sake, in the following use ‘future’ to take all these unrealized situations together.

Alternatively, S may use this type of information to give a partial description of a completely realized situation; the detail he leaves out may be unknown or irrelevant to S; compare (2/2) *a sister of John’s* \uparrow * *is interested*, and (8) *an ambulance* \uparrow * *is on its way*.

Again, for this second type of information, given the fragmentary character of the sketch as far as authorized by S, H may try out details of his own to fill the gaps, morphing features into others, shifting back and forth between alternative ways to add flesh and blood to the skeleton. Only in this case, as long as the situation is not realized, each result is as good as another.

TYPE $a \text{ cat}^{\downarrow \leftrightarrow} \text{ predF} > \{ \underline{\backslash \text{O}^{\text{cat}} \text{ O } \text{O} / \text{O}^{\text{pred}} / a \text{ cat}} \}$: a repetitive stand-in.

What S calls up is this self-replicating series of complete but partly unspecified cats in H’s mind; repetitively, the predicate feature is included.

Obviously, S may use this type of information as a general statement. Many proverbs are of this type: (20) *a promise*^{↓↔} *is a promise*; (21) *a good neighbour*^{↓↔} *is worth more than a far friend*, etc., but also generalizations made for the occasion like (27) *a tomcat*^{↓↔} *is a good mouser*. Also, S may use this type of information to talk about a specific thing, an exemplary case of its kind; for instance about this cat that just brought in this pathetic little mouse:

(32) *a cat*^{↓↔} *is a predator*.

For this third type of information again, given the fact that the repetitive picture is partly unspecified, H may assemble a range of cats. Only now, they do not replace each other, but crowd H's mind.

TYPE *a cat*^{↑↔} *predF* > { o^{cat} ... / o^{pred} / *a cat* } : a repetitive fragmentary sketch.

What S calls up is this self-replicating series of unfinished sketches of cats; repetitively, the predicate feature is included. This very simple image shapes very abstract information. This type of information has three obvious uses: to define to H the meaning of a word new to him, that is, to add this word to his vocabulary, to define the meaning of some new word (or redefine an old one) within the whole speech community, and to refine the meaning of a word H already knows.

An in itself arbitrary sound form becomes a word if it is systematically linked to something, some content. In order to be useful in communication, it must be made sure that different people link the same word to the same things. The basic way to achieve that is to share a direct experience. Let's assume that H does not yet know the meaning of *cat*; to remedy this, catch his attention, say *cat*, and point at a cat. When H has recorded the link, he has a shape that comes with the word *cat* also in the physical absence of the cat. Now when H's vocabulary grows, this shared direct experience may be replaced by a shared language-triggered experience. Let's assume that H has an image that comes with the words *male cat*, but that he does not yet know the meaning of *tomcat*. Now, to remedy this, say (28) *a tomcat*^{↑↔} *is a male cat*; what happens is this: *a tomcat* calls up a virtually empty outline in H's mind, the prosodic forms make it into a self-replicating open sketch, and the predicate each time adds its image, permanently linking it to *tomcat* as well. The information does not work the other way around: although a male cat in fact happens to be a tomcat, this is not what this type of information says; compare (22) *a quillback*^{↑↔} *is a freshwater fish*: the self-replicating open sketch only has room for complementary features of o^{quillback}, but o^{freshwater fish} may be a feature of all kinds of other fish.

In the same way, S may introduce completely new words into the vocabulary of his speech community, or redefine words already in use, as in (23) *a planet*^{↑↔} *is a heavenly body moving around the sun* (and whatever further requirements astronomers agree upon). So the question whether Pluto is a planet does not depend on what Pluto is, but on what people agree to call a planet.

And if H already knows the meaning of a word, like *whale*, that is, if he already has an image that comes with it, this type of information can be used to refine that image:

(33) *a whale*^{↑↔} *is a mammal*

in the self-replicating open sketch of a whale, again and again $o|^{mammal}$ is included, so it becomes part of the image that *whale* by itself calls up.

For this type of information again, given the fragmentary character of the sketch as far as authorized by S, insofar as it accesses complementary features, H may assemble any combination in his mind, morphing features, combinations of features, and entire things into each other; now, H does not shift back and forth between them, but he keeps them all.

These fragmentary sketches do not stand-in for things in the world. Possibly, there is not even something there that fits the noun phrase; compare (24) *a perpetual motion machine*^{↑↔} *is a device that operates indefinitely by creating its own energy.*

3.3 Recognizing the Real Thing

Assume that S did give some piece of information, and that H did get the picture. Now, if H gets a good look at the real thing, he will recognize it as fitting the information. This recognition has to be explained, but we should be very clear about what exactly is in need of an explanation.

First we need to be clear about what sort of process we are talking about. There is a relatively simple story and there is a complicated story. What happens is this: S's sequence of sounds flashes networks of neurons in H's brain into special action, making H hear forms of language; they access complex and branching networks in the areas of experience. Insofar as the action flashes through these networks, it makes the corresponding shape, colour, etc., show up. Its start is authorized by S. Simultaneously, signals from the world picked up by H's senses trigger networks of neurons in the areas of experience, making the real thing show up. From perspective H, the authorized part of the image fuses with the real thing in sight; it disappears in it. From perspective Ω , the neural networks active in making the authorized part of the image show up are included in the neural networks active in making the real thing show up. Possibly, H's language-triggered image contained unauthorized details; in that case he may have to erase them in order to be able to fuse image and real thing. This is from perspective Ω : away from its point of departure in the words, the action flashed through distantly connected networks; now when the real thing comes in sight there is a mismatch: although the networks hit by sight are included in the branching complex of networks accessed by the words, it is another branch that happened to be active; in these networks, the action should subside. Anyhow, from perspective H, the image called up by verbal information fuses with the real thing in view; the mind's eye and the real eye are, so to speak, brought into visual line, like two eyes that are brought into visual line.

This is as close as the relationship between verbal information and reality gets: in H's brain, signals from outside and sound waves from S join forces in calling the real thing to H's awareness. Apart from these repeated brief moments of contact in a person's mind, there is no way for a connection between language and reality to exist. What we need to explain is the connection whenever it exists. This is the relatively simple story.

There is, however, a more complicated story. H all the time connects information from different sources; he may very well be acquainted with the things that S talks about also if he does not have them directly in view. These other cases of fusing information should be explained as well.

Suppose H does not get the real thing in view because it is beyond the rather limited range of his senses, too far away, or in the past, or it is something that the human senses are not tuned to. In that case, H's brain can simulate its presence by activating and recombining ready-to-use records of bits and pieces of experience. This as-if presence may be triggered by different sources, including re-called memory, independent verbal information, and preceding context. We may have a case of recognition any time there are two shapes in H's awareness, originating from two sources: the fragmentary image called up by S on the one hand and the real thing, present or as-if-present on H's own authority, on the other.

The basic case of recognition happens when the thing is present, and the information concerns concrete, phenomenal aspects. I believe all other cases of recognition are built upon this basic case, both when the information is more abstract, and when the real thing is not literally in view. For brevity's sake, I shall keep talking about the real thing in view; read this to include all other cases where H has the real thing in mind based on other than S's most recent piece of information.

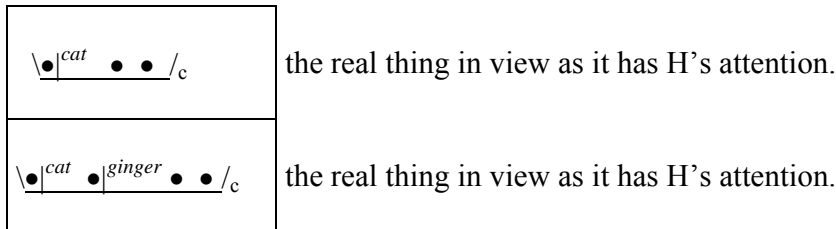
Note that the facts are different when a fragmentary image called up by S sketches some kind of future. As the real situation does not exist, H cannot possibly take it in view, and he does not take it in mind as something there in its own right either. So there is no recognition for us to explain. The future is open; anything may still happen. However, as soon as the future happens to become reality before H's eyes, H will recognize what is now present as fitting the earlier sketch for the future, some earlier possibility become real, others expired. This postponed recognition does indeed need an explanation.

Also we have to explain fiction: S makes up a world that does not exist, so H cannot possibly take it in view, ever. Still, to H, it gradually takes shape like a real world that is just out of sight. As the story unfolds, H all the time recognizes characters and things talked about: on the one hand, there is the fragmentary image called up by S's most recent words, and on the other, there is the world already there, pictured as-if-there in its own right, called up by what S said earlier. S has taken control over H's brain's mechanism to make reality take shape around him, bound to the laws of its nature, with an irrevocable past, an unfolding present, and an open future. Of course it is a simulation, but a fictional world can become almost as tangible as the real world, S calling up the crucial fragments, and H supplying all further details from his own past experience. It is tempting to go into detail about what S can do with this world he has completely under control, including its physical laws, but I have to stick to the issue: recognition. The point is this: when we want to explain verbal information, we do not just need to explain its connection with the really real thing, but we also need to explain what happens when H makes sense of fiction. I believe, if we can explain H's recognition of real things just out of view, we are on our way to explain the gradual build-up of fictional worlds as well.

The process of recognition is still more complicated. H does recognize the real thing if it is in sight, but only insofar as it shows up to him. And what shows up to him not only depends on what in fact is there, but just as much on H's brain's capacity to figure it out: on sight, phenomenal features will readily show, but abstract aspects will not; if the real thing is as-if-there rather than physically present, phenomenal features may be dim as well. On the other hand, H may very well have other than perceptual knowledge of the real thing. The point is that more often than not, there is partial rather than complete recognition. Both will have to be accounted for.

What we need for the purpose of our theory of verbal information about a real or fictional world is a representation of the world as and insofar as it shows up in H's awareness, triggered by perception or other sources of information, independent of that new verbal information, taken by H as exemplary for what would show up to someone else; this representation needs to be specified insofar as the features represented have H's attention.

In chapter 2, I represented the real thing C as and insofar as it takes shape to H, which he takes as exemplary, as $\backslash \bullet \bullet \bullet /_c$; what has H's attention can be represented by adding (H's) verbal tags to the features concerned. Given the world in view, what has H's attention is what S, talking about this world, mentions. So in order to get a precisely tailored representation of the real thing in view insofar as relevant in the theory, cut loose the symbolic words from S's text as they come, let H distribute them on his own authority, that is, on view, or based on independent information, over the real things as best as he can, and enrich the representation accordingly. Thus, $\backslash \bullet \bullet \bullet /_c$ may become, the moment S mentions *cat*, the moment S talks about colour:



In chapter 2, this notation represented something physically present. As defined, it takes shape to H precisely during activity *A* in a corresponding complex of networks of neurons, triggered into action by signals originating from the real thing. As the strengthened complex of neurons keeps a ready-to-use record that can be re-activated by other triggers as well, the real thing can take the same shape when not physically present. So we can use the same notation for the real thing out of view but as-if-there in H's mind. And as this as-if presence does not require the physical presence of the thing that takes shape, it does not necessarily require its physical existence; new things as-if-there can be composed from bits and pieces of earlier experience. So we can in principle use the same notation to represent fictional things. However, H will be a bit more cautious in generalizing his world that is out of sight, especially if it is a fictional world, and his ongoing process of piecing together new information and what he otherwise has in mind may be a very individual process, but on occasion, and whenever necessary, everything tagged by words can be taken as exemplary within the same speech community. To simplify the discussion, suppose generalization.


So I shall use this same notation to represent the world talked about, whether it is in view or just taken in mind, whether or not fictional, as long as it is, or has been, realized or imagined as realized. To explain H's recognition of what fits the information, I do not need a representation of future or otherwise unrealized situations until they become real and are no future any more.

So H has the image called up on the one hand, and the world in view on the other; what happens is that he scans the world for what matches with the noun phrase skeleton, which gives him an inventory of what comes into consideration, then he includes the predicate feature, and expects that the world is such that he can fuse his image with it; this is the case if there is a match there

for the whole image. To get hold of this process, represent both sides as discussed above, represent H's focus on what is relevant in the world by means of his distribution on his own authority of S's words as they come, visualize the fragmentary image as transparent, and visualize the world in view as solid.

First, given the noun phrase skeleton, H scans the world, searching for a match, which is there if the elements of the unifications as far as specified are the same. In reality, there may be a vast range of things that H would recognize as fitting if they came in view, but in a given situation only one, two, or a few of them will actually come into consideration. A match will directly show in the representation. In order to keep both the image and the real thing visible, I shall in the notation put the image directly above the real thing to represent the image superimposed on the real thing. Omitting prosodic information:


NOTATION




Thus scanning the world, it may be that there is only one thing that comes into consideration, but placing and replacing the image, it may be that H spots more than one such thing. At this point, the kind of noun phrase used (definite, indefinite, etc.) is important; I shall postpone its discussion (see chapter 6), and for now stick to indefinite noun phrases, which, apart from the fact that they only fit 'countable' things, do not give further identifying information.

For instance, let's assume that H has two cats in view, now on the one hand, H has the noun phrase skeleton called up by S, and on the other, after first inspection of the world, he has an inventory of what comes into consideration:

S's noun phrase skeleton



cat distributed
on H's own authority




Given this combination, H can superimpose the skeleton on the first cat, and he can superimpose the skeleton on the second cat.

Now include the information given by prosodic form: if H has a single picture or sketch, with each replacement from one thing that fits to another, H has to remove it from its earlier place. If he has a self-replicating picture or sketch, with each replacement from one thing to another, he copies it. If he has a closed picture, it completely covers the real thing; if he has an open sketch, it covers only part of the real thing.


$$\text{TYPE } a \text{ cat}^{\downarrow} * > \backslash \underline{\text{O}^{\text{cat}} \text{ O O}} / ^a \text{ cat}$$

H has this single stand-in; with the real thing in sight, they should unite, therefore, when S intends to call up a single stand-in, and the world is in view, he will try to make this stand-in precise enough in order for H to be able to identify the real thing. So for this type of information, the most natural illustration of recognition would assume that only one thing would come into consideration. However, to get a clear view of the differences between the four types of information, and to pin down exactly the information contributed by each of the basic forms, assume that there are precisely two of those things; also assume that H spots them; now on the one hand, H has the picture called up by S, and on the other an inventory of what comes into consideration:


S's picture



cat distributed
on H's own authority




That is, on the one hand, H has this single closed picture of a cat, and on the other hand, he has these two cats. In combining them, he can unite the picture with the first cat, separate them, unite the picture with the second cat, and vice versa. Now, as he only has this single picture, if he unites it with the first cat he loses the second option, and if he unites it with the second cat he loses the first option, and because he does not know which of the two S has in mind, he has to keep the choice on hold. H just understands S's information; this is not the problem; the information is not imprecise, but this very precise piece of information does not include identifying clues, so, the combination of this information and this world in view is inconclusive, and H may get confused. As said above, S tends to try to avoid such confusions, but they do occur. I shall represent the process by forking into two options to superimpose the single picture on what is there, covering it completely; only one path to be followed, but the choice on hold:




$$\text{TYPE } a \text{ cat}^{\uparrow} * > \backslash \underline{\text{O}^{\text{cat}} \dots} / ^a \text{ cat}$$

On the one hand, H has the sketch called up by S, and on the other he has an inventory of what comes into consideration:

S's sketch




cat distributed
on H's own authority



In combining them, he can fuse the sketch with the first cat, separate them, fuse the sketch with the second cat, and vice versa. Due to the fact that the sketch only represents a fragment, it does not stand-in for and therefore it does not unite with a solid cat in the world; as far as the information goes, there is no choice to be made between the two cats in view. Still, H only has a single sketch available to fuse with what is there. So what happens is that he shifts this single sketch back and forth between them, narrowing down what comes into consideration without splitting options.


This dynamic connection is easy to visualize but not so easy to freeze on paper. I shall use the following notation:




TYPE $a \text{ cat} \downarrow \leftrightarrow > \{ \backslash \text{O} | \text{cat} \text{ o o } / \text{a cat} \}$

On the one hand, H has the repetitive picture called up by S, and on the other he has an inventory of what comes into consideration:


S's sketch



cat distributed
on H's own authority




That is, on the one hand, H has this self-replicating closed picture, and on the other, he has two cats. In making combinations, he can fuse the picture with the first cat, and he can fuse the picture with the second cat; now, in the process, the picture duplicates; H gets just as many pictures as he has cats in view. Because the pictures are complete stand-ins, in each case he unites picture and real thing. In notation:




TYPE $a \text{ cat} \uparrow \leftrightarrow \{ \backslash \underline{\text{O}}^{\text{cat}} \dots / \text{a cat} \}$

On the one hand, H has the repetitive sketch called up by S, and on the other he has an inventory of what comes into consideration:


S's sketch



cat distributed
on H's own authority




That is, on the one hand, H has this repetitive open sketch of a cat, and on the other he has these two cats. Combining, he can fuse the sketch with the first cat, and he can fuse the sketch with the second cat; now, in the process, the sketch duplicates. Because these sketches only represent fragments, they do not unite with the real thing, but they do give a match. In notation:




The predicate

Suppose H has made the connection between the noun phrase and the world. Now, take the predicate: on S's authority, H includes the predicate feature within this connection. This gives him an expectation of what the thing(s) concerned will be like: he expects them to be such that his image can fuse with them. Whether he will in fact be able to fuse image and real thing depends on two factors: what is in fact the case, and what he knows to be the case. Let us for now assume that H is by himself able to distribute the predicate over the inventory of what comes into consideration. Complete recognition will show in the representation: the transparent image superimposed on the solid real thing will disappear against this background. Given, as above, two cats in view, there are three possibilities: one of the cats, say, the first one, has \bullet^{pred} , both cats have it, or neither of them has it:


cat distributed
pred distributed
 on H's own authority
 situation 1



situation 2



situation 3




Now compare the four types of information again:

$$\text{TYPE } a \text{ cat}^\downarrow * \text{ pred}F > \underline{\underline{\bullet|^{cat} \bullet \bullet / \bullet|^{pred} /}} a \text{ cat}$$

Example (2/1) and the abbreviated (14/1) in chapter 1 are of this type: (2/1) *a sister of John's*[↓] *is interested*.

Let's assume again that H has two cats in view. For the start of the process, and his first inspection of the world, see above. Now, including the predicate, on the one hand, H has:


cat distributed
 on S's authority, the fork



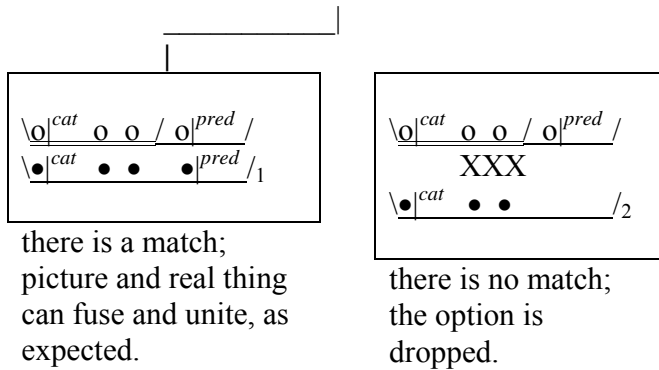
H expects that the world is such that S's picture of what is there can fuse with it when he alternatively tries out the options. This is the case if there is at least one match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected:

cat distributed
pred distributed
 on H's own authority
 situation 1




In combination:



Of course the situation in sight may be different; if it is situation 2, H's combination of information remains inconclusive, and he still has to keep the choice between options on hold.

Variations in recognition

H can only recognize what he has in view, and only insofar as he has it in view. In terms of the theory: H's recognition of what fits the information depends on his ability to distribute by himself the words used over the world in view. It may be that H is unable to distribute *cat*; in that case, there is no recognition at all. It may be that H can distribute *cat* but is not sure whether his inventory is complete; in that case, recognition of what it is that S talks about remains tentative. It may be that H is able to distribute *cat* but unable to distribute *pred*, or that he is uncertain about its distribution. As long as he cannot distribute *pred* himself, the superimposed image can on S's authority project into his world's unspecified space; the process of recognition ends halfway, as in:



It may also be that H cannot distribute *pred* because S gives it as a future feature. In that case, its recognition will not be an issue until that future becomes reality.

So far, I did not question the correctness of S's information. In the process of understanding and recognition, correctness of the information is not the first issue; when H has the world in view, he will try to somewhere fit in the picture called up; it may disappear against the background of the solid world, or partly or completely project into its unspecified space. Unless H distrusts S from the start, it is only when he cannot find a way to reconcile picture and world that he will question the information. For instance, in the illustration above, if it is situation 3 that he gets in view, he will reject the information, because he cannot fuse S's picture of what is there with what is in fact there.

Every once in a while it does happen that H knows for sure that there is nothing that comes into consideration; H does get the picture, but he does not even need to hear the predicate to know that it will not fuse with what is there. For example, when S says

(34) *(John could not make it because) a sister of his[↓]* is very ill*


while H knows very well that John is an only child.

TYPE $a\ cat^{\uparrow} * predF > \underline{\backslash O^{cat} \dots / O^{pred} /} a\ cat$

Example (2/2) and its abbreviation (14/2) in chapter 1 are of this type: (2/2) *a sister of John's[↑]* is interested.*

Let's assume again that H has two cats in view. For the start of the process, and his first inspection of the world, see above: H shifts his single sketch back and forth between the two cats, narrowing down what comes into consideration without splitting options. Now, including the predicate, on the one hand, H has:


cat distributed
on S's authority




H expects that the world, if realized, is such that S's sketch of what is or may be there can fuse with it when he shifts it back and forth between the cats. This is the case if there is at least one match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected:


cat distributed
pred distributed
 on H's own authority
 situation 1



Now, when trying out the combination:




There is a match: the sketch can fuse with the first cat, so:



Variations in recognition

For this second type of information, all variations of recognition described for the first type of information may happen as well. One case: take example (2/2) abbreviated, (14/2) *a sister[↑] is interested*; assume that H has the two people in sight that come into consideration; it is one of them that is in fact interested, but H is not, by himself, able to distribute *interested*. Now this is as far as he gets:

sister distributed;
 on S's authority



Let's now have a look at (9) *one of them[↑] is redundant*. Again, H cannot distribute the predicate, but in this case not because he does not know the facts, but because there is no fact to be known as yet: *redundant* still hovers above the people concerned, for the future to settle on one or another.

And take (10) *one of us[↑]* is terrible*, said by someone who overheard an unflattering imitation of himself. At first sight, it is much like the case above: it seems to be undecided where *terrible* will settle. But the point here is that the first and the second person apparently disagree about its distribution: they each will see the other as fitting the sketch.

As already said, S's sketches of what is or may be there do not claim simultaneous existence of anything fitting the noun phrase; if S says for instance (13) *an answer[↑]* remains forthcoming*, H perfectly well understands the information, but he will not expect the answer to be already there; because of the fragmentary character of the sketch, it may still take a variety of shapes, but the sketch is specific enough for postponed recognition, when the answer has taken its concrete form.


It is clear that correctness of the information is not an issue that can be decided upon if S sketches an unrealized situation.

TYPE $a\ cat \downarrow \leftrightarrow predF > \{ \ \underline{\backslash o|^{cat} \ o \ o \ / \ o|^{pred} \ / a\ cat} \}$

Example (27) in chapter 1 is of this type: *a tomcat[↓]↔ is a good mouser*

Let's assume again that H has two cats in view. For the start of the process, and his first inspection of the world, see above: H unites his stand-in picture with the first cat, replicates it, and unites the copy with the second cat. Now, including the predicate, on the one hand, H has:


cat distributed
on S's authority




H expects that the world is such that S's picture of what is there fuses with it when he superimposes it on one thing that comes into consideration after the other, replicating the picture. This is the case if each time there is a match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected:

cat distributed
pred distributed
on H's own authority
situation 2:



Now, when trying out the combination:



in both cases, the picture can fuse and unite with what is there, as expected.

Variations in recognition

For this type of information, all variations of recognition described for $a\ cat^{\downarrow} * predF$ may happen as well.


It may be that there is only one cat in view; if so, the information makes it into an exemplary cat, representing all others. Alternatively, it may be that only a few cats are directly in sight, but any number of them might still turn up, getting their own copy of S's picture. Also, it may happen that among the majority of cats that do indeed have this predicate feature, there are one or two without it. As long as they remain out of sight, H will not mind, and if one comes in sight amidst a majority of cats that do fit the information, H will not so much recognize this one as a counterexample to refute the general rule, but rather as an exception: one cat that cannot unite with the picture among the majority of cats that can.

TYPE $a\ cat^{\uparrow} \leftrightarrow predF > \{ \underline{\underline{\O|^{cat} \dots / O|^{pred} / a\ cat}} \}$

Example (28) in chapter 1 is of this type: $a\ tomcat^{\uparrow} \leftrightarrow is\ a\ male\ cat$.

This is abstract information, it is not primarily meant to describe some situation that is there. However, H may have cats in view, and if so, he will expect them to fit the information. Let's, once again, assume that H has two cats in view. For the start of the process, and his first inspection of the world, see above: the reduplicated sketch fits both cats, covering a fragment of them. Now, including the predicate, on the one hand, H has:


cat distributed
on S's authority




H expects the world, if realized, to be such that S's sketch fuses with it each time he superimposes it on one cat after the other, replicating the sketch.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected:

cat distributed
pred distributed
 on H's own authority
 situation 2



Now, when trying out the combination:



in both cases, sketch and image can fuse as expected; they will not unite, as the image is only fragmentary.

Variations in recognition

The connection between this abstract information and the world is rather remote; however, if H gets a good enough view of something fitting, he will recognize it as fitting. Compare the following cases:

- (33) $a\ whale \overset{\uparrow}{\leftrightarrow} is\ a\ mammal,$
 (28) $a\ tomcat \overset{\uparrow}{\leftrightarrow} is\ a\ male\ cat$
 (22) $a\ quillback \overset{\uparrow}{\leftrightarrow} is\ a\ freshwater\ fish$
 (23) $a\ planet \overset{\uparrow}{\leftrightarrow} is\ a\ heavenly\ body\ moving\ around\ the\ sun$ (and whatever further requirements astronomers agree upon)

Assume that *whale* is in L's vocabulary and that he has two of them in sight, but that he does not know on his own authority that they are mammals; now S's information enriches what he sees.

Assume that both *tomcat* and *quillback* are not yet in H's vocabulary; so H does not recognize tomcats or quillbacks right in front of his nose. However if he is able to distribute the predicate, and make an inventory on that basis, he gets a clue, although he has no way of knowing whether it includes too much, as happens to be the case in example (22).

Finally, assume Pluto in sight, along with other heavenly bodies; there may be uncertainty about the precise distribution of *planet* over what is there; now (23) gives a criterion, and depending on what astronomers agree to include in the definition, Pluto is or is not to be tagged *planet*.

Again, H will usually just accept S's information. Occasionally, however, he may find it incorrect. With this type of information, two kinds of incorrectness are to be distinguished: if H is able on his own authority to distribute the subject independent of the predicate, he can find S's information false, that is, false in relation to the facts. For example, if the information is

(35) *a whale*[↑]_↔ *is a fish* (incorrect)

However, if S uses this type of information to make the distribution of the subject noun phrase dependent upon the presence of the predicate feature, the information may also be incorrect, but only in the sense of ‘this is not what the speech community agreed upon’, which, for example, is the case if the information is

(36) *a tomcat*[↑]_↔ *is a young cat* (incorrect).

Chapter 4

The Shape of Negative Information

So far, the idea has been that verbal information takes phenomenal shape, and that this shape on occasion fuses with the real thing in view. Given this idea, negative information may seem problematic: it states what is absent, but how could absence take shape, and how would the shape of absence fuse with what is not there? The working hypothesis of chapter 1 is that in H's mind, absence takes shape just like presence, but *not* marks its part of the image as 'absent', that is, it is irreconcilable with the original image of presence; the shape of absence works as an example of what to exclude, and with the world in view, it works as a filter to sieve out what would clash with the information.

4.1 Negative Predicates

Distinguish negative noun phrases, *not a cat*, and negative predicates, *not predF*. First look at the negative predicate, for instance, *is not black*. To H, the word *black* does what it is programmed to do: it calls up $o|^{black}$. *Not* turns this feature into an example of what is excluded; building upon the representation of H's images of presence, we can now directly represent his images of absence as well:

NOTATION

$$\langle o|^{black} \rangle$$

$$\langle o|^{black} \rangle \text{XXX} o|^{black}$$

$o|^{black}$ turned into an example of what is excluded,
 $\langle o|^{black} \rangle$ irreconcilable with $o|^{black}$

$$\langle o|^{black} \rangle$$

$$\text{XXX}$$

$$\bullet|^{black}$$

$\langle o|^{black} \rangle$ clashing with $\bullet|^{black}$


Included in the noun phrase skeleton just like a positive predicate feature, as in $\langle o|^{cat} \rangle / \langle o|^{black} \rangle / a \text{ cat}$, $\langle o|^{black} \rangle$ acts as a doorkeeper to prevent the addition of $o|^{black}$ to the picture.

As for recognition: H cannot directly take in sight what would fit $\langle o|^{black} \rangle$, as it is not there; if there is, for instance, a ginger cat, its $o|^{ginger}$ cannot fuse with this $\langle o|^{black} \rangle$. The only thing that would give a match is a black cat, but this is exactly the thing that S did exclude. So H cannot fuse the negative part of this image with what is there, but he can directly recognize what would clash with it. What he expects to find is the absence of a clash.

Take perspective Ω : *a cat* calls up a noun phrase skeleton in H's awareness, and accesses a stock of potential complementary features recorded in a vast criss-cross of strengthened networks of neurons, determining a virtually infinite range of cats that H may assemble in his mind; most of these records will not be activated, so their features will not reach H's awareness. A predicate, activating the corresponding networks, calls a feature to H's awareness, probably one of the potential complementary features. If the predicate is positive, like *is ginger*, it includes this feature in the skeleton. If the predicate is negative, like *is not black*, it includes this feature in

the skeleton as a doorkeeper and filter. Then a non-black cat in view, for instance a ginger one, does not hit the same networks of neurons as *not black* does; there is no match. However, an (unexpected) black cat in view does hit the same networks of neurons: there is a match, but a clashing one. In the representation superimposing the image of absence on the real thing, the clash will directly show:

NOTATION



Not reinforces the differences between the four types of information discussed. Compare $a\ cat^{\downarrow*}$ and $a\ cat^{\uparrow*}$: when the predicate is positive, o^{pred} included in $\underline{o}^{cat}\ o\ o /^{a\ cat}$ specifies some feature of a closed unification, leaving it up to H to find out what those other features in fact are, and included in $\underline{o}^{cat}\ \dots /^{a\ cat}$, it fills one of the gaps in the room left open, leaving the range of cats that H still may assemble virtually infinite. The difference is unmistakable, but often enough rather subtle, especially when both types describe a realized situation, as both make single pictures that are to fuse with what is there. But when the predicate is negative, $\langle o^{pred} \rangle$ included in $\underline{o}^{cat}\ o\ o /^{a\ cat}$ excludes o^{pred} from the set of unspecified other features; the picture stands-in for some cat that lacks the predicate feature ('predicate negation'), and $\langle o^{pred} \rangle$ included in $\underline{o}^{cat}\ \dots /^{a\ cat}$ excludes o^{pred} from the stock of features available to H to detail the sketch. That is, in effect it excludes the subsection of cats with the feature o^{pred} from the whole range of cats that H may still assemble ('sentence negation').

4.2 Four Types of Negative Information

The types of information in more detail:


TYPE $a\ cat^{\downarrow*}\ not\ predF > \underline{o}^{cat}\ o\ o / \langle o^{pred} \rangle /^{a\ cat}$

Example (1/1) and the abbreviated (15/1) in chapter 1 are of this type: (1/1) *a sister of John's^{↓*} is not interested.*

If the world talked about is not in view, to H the picture stands-in for the real thing. Given the fact that the picture as far as authorized by S is partly unspecified, H can still tentatively assemble a range of different cats in his mind, each new one replacing the earlier one, as long as he does not include o^{pred} , each cat he thinks up may very possibly be a wrong guess of what S's one and only cat looks like.

Let's assume again that H has two cats in view. For the start of the process, and his first inspection of the world, see chapter 3: H has two alternative options for fusing his single picture with what is there, the first or the second cat. Now, including the predicate, on the one hand, H has:


cat distributed
on S's authority, the fork




H expects that the world is such that S's picture of what is there can fuse with it when he alternatively tries out the options. This is the case when within the inventory of what comes into consideration, there is at least one that does not match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected:

cat distributed
pred distributed
on H's own authority
situation 1:



When trying out the combination:



there is a match;
picture and real thing
cannot fuse; the
option is dropped.

there is no clash;
picture and real thing
can fuse and unite.

TYPE $a \text{ cat}^{\uparrow} * \text{ not pred}F > \underline{\text{ o } | \text{ cat } \dots | \langle \text{ o } | \text{ pred } \rangle} / a \text{ cat}$


Example (1/2) and the abbreviated (15/2) in chapter 1 are of this type: (1/2) *a sister of John's[↑] is not interested.*

If the world talked about is not in view, the image blocks the inclusion of $\text{ o } | \text{ pred}$ in the open sketch of a cat. Given the fragmentary character of the sketch as far as authorised by S, H may try

out any detail of his own to fill the gaps, morphing features into other features, cats into other cats, as long as he does not include $o|^{pred}$.

The information does not include a claim of existence of anything fitting the noun phrase, but let's assume that H has two cats in view. For the start of the process, and the world on first inspection, see chapter 3: H shifts his single sketch back and forth between the two cats, narrowing down what comes into consideration without splitting options. Now, including the predicate, on the one hand, H has:


cat distributed
on S's authority



The image excludes $o|^{pred}$ from occurring together with $o|^{cat}$ and $\bullet|^{cat}$. H expects that the world is such that S's image does not clash with it when he shifts it back and forth between the cats. This is the case if there is no match.


On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected:

cat distributed
pred distributed
on H's own authority
situation 3:




Trying out the combination, there is no problem. However, on second inspection the world might be different, such as for instance

cat distributed
pred distributed
on H's own authority
situation 1:



Now, when trying out the combination:



there is a match: the sketch clashes with the first cat.

Compare again the closed picture called up by $a\ cat^{\downarrow} * not\ predF$ and the open sketch called up by $a\ cat^{\uparrow} * not\ predF$: if situation 1 is in view, the closed picture unites with the second cat, and the first cat is irrelevant, but the open sketch clashes with the first cat (see also section 7.3 below: Verbal information and the notion of truth).


$$\text{TYPE } a\ cat^{\downarrow} \leftrightarrow not\ predF > \{ \underline{o}^{cat} \ o \ o \ / \langle o \rangle^{pred} \ / \ a\ cat \}$$

An example of this type given in chapter 1 is (29) $a\ big\ boy^{\downarrow} \leftrightarrow doesn't\ cry$.

If the world talked about is not in view, to H the picture is a repetitive stand-in for real things there. Given the fact that the pictures as far as authorized by S are partly unspecified, H may assemble a crowd of cats in his mind, as long as he does not include o^{pred} .

Let's assume again that H has two cats in view. For the start of the process, and his first inspection of the world, see chapter 3: H fuses his stand-in picture with the first cat, replicates it, and fuses the copy with the second cat. Now, including the predicate, on the one hand, H has:


cat distributed
on S's authority




H expects that the world is such that S's repetitive picture of what is there fuses with it when he superimposes the copies on one thing that comes into consideration after the other. The world is as expected if in each case there is no match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, *cat* distributed, *pred* distributed, is, on H's own authority, as expected, it is situation 3. When trying out the combinations, there will be no problem. However, on second inspection the world may be different, such as for instance

cat distributed
pred distributed
 on H's own authority
 situation 1:



Now, when trying out the combination:



there is a match;
 picture and real thing
 cannot fuse, which is
 not as expected.

there is no match;
 picture and real
 thing can fuse and
 unite, as expected.


TYPE $a \text{ cat}^{\uparrow} \leftrightarrow \text{not pred}^F > \{ \underline{\text{o}^{\text{cat}} \dots / \langle \text{o}^{\text{pred}} \rangle} / a \text{ cat} \}$

An example of this type given in chapter 1 is (30) *a vegetarian*[↑] *does not eat meat*.

Given the fragmentary character of this repetitive sketch, H may assemble a crowd of cats in his mind, thinking up any combination of complementary features; features, combinations of features, and whole cats morphing into each other, as long as he does not include o^{pred} .

This is abstract information, it is not primarily meant to describe some situation that is there. However, H may have cats in view, and if so, he will expect them to fit the information. Let's assume that H has two cats in view. For the start of the process, and his first inspection of the world, see chapter 3: the duplicated sketch fits both cats, covering a fragment of them. Now, including the predicate, on the one hand, H has:

cat distributed
 on S's authority



The repetitive image again and again excludes o^{pred} from occurring together with o^{cat} . H expects that the world is such that S's image does not clash with it when he superimposes the copies on one thing that comes into consideration after the other.

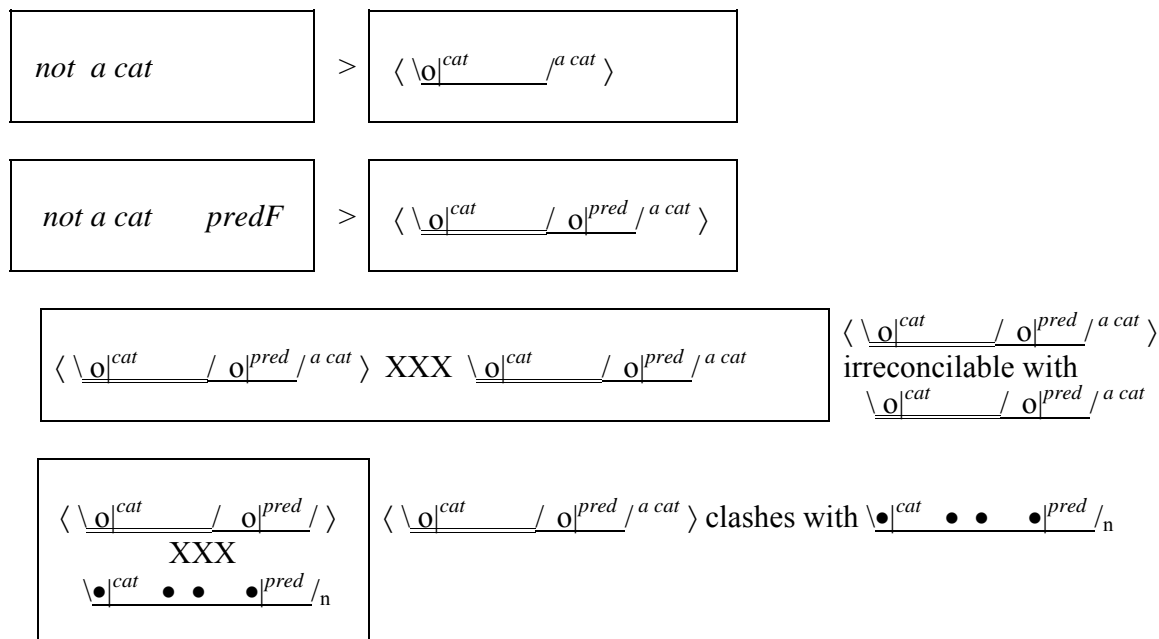
On the other hand, H has the world in view. Assuming that the world, on second inspection, *cat* distributed, *pred* distributed, is, on H's own authority, as expected, it is situation 3. When trying out the combinations, there will be no problem. However, on second inspection the world may be different; when we assume situation 1, it is easy to see that the first cat gives a match, and, as a result, clashes with the information.

4.3 Negative Subjects

S's subject-predicate constructions take shape to H as noun phrase skeletons that include the predicate feature. If the information is positive, it shapes (possible) presence, which is to fuse with the real situation if it comes in sight. If the predicate is negative, the predicate feature shapes absence. If the subject noun phrase is negative, the whole image shapes absence. The shape of absence does not fuse with what is there; what H expects is the absence of a clash; if a match happens to be in sight, H recognizes it as irreconcilable with the information.

In *not a cat* the range of the form *not* is the whole noun phrase; *not* turns the noun phrase skeleton as a whole into a doorkeeper in order to block what is irreconcilable with it, and a filter to sieve out what would clash with the information. Omitting prosodic information, in notation:

NOTATION



Including prosodic information, two types of information in more detail:


TYPE *not a cat*[↓]* *predF* > $\langle \underline{\backslash o}^{cat} \quad \underline{o} \quad \underline{o} \quad / \underline{o}^{pred} / a \text{ cat} \rangle$

Examples of this type given in chapter 1 are: (3) *not a leaf*[↓]* *stirred*, and (5) *not a word*[↓]* *was true*.

When H does not have the world talked about in view, this picture of absence blocks the addition of a (any) cat that has o^{pred} among its features to his idea of the world.

Let's assume again that H has two cats in view. Now, after a first inspection of the world, and including the predicate, on the one hand, H has:


cat distributed
on S's authority



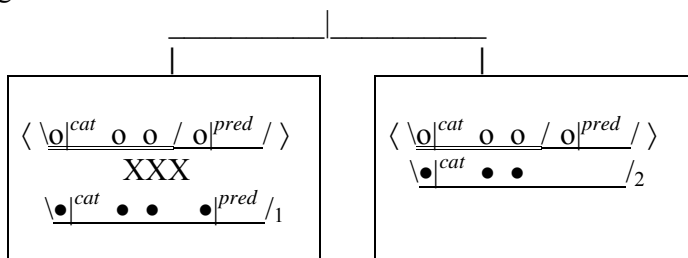
H expects that the world is such that S's picture of absence does not clash with it if he tries out both options. There is no clash if in both options there is no match.

On the other hand, H has the world in view. Assuming that the world on second inspection, *cat* distributed, *pred* distributed, is, on H's own authority, as expected, it is situation 3. When trying out the combinations, there will be no problem. However, on second inspection the world may be different, like, for instance, in situation 1:

cat distributed
pred distributed
on H's own authority
situation 1:



Now, when trying out the combination:



there is a match; so
the first cat clashes
with S's picture


TYPE $not\ a\ cat^{\uparrow} * predF > \langle \underline{O}^{cat} \dots / \underline{O}^{pred} / a\ cat \rangle$

Examples of this type given in chapter 1 are: (4) $not\ a\ blade\ of\ grass^{\uparrow} * grew\ there$, and (6) $not\ a\ word^{\uparrow} * was\ spoken$.

When the world is not in view, the sketch blocks the combination of the cat-feature and the predicate feature in H's idea of what is or might be there.

This type of information does not claim presence or existence of anything fitting the noun phrase. Still, there may be something in view that comes into consideration. Let's assume again that H has two cats in view. Now, after a first inspection of the world, and including the predicate, H has:


cat distributed
on S's authority




H expects that the world is such that it does not clash with S's sketch of absence when he shifts it back and forth between the first and the second cat. This is the case if there is no match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, *cat* distributed, *pred* distributed, is, on H's own authority, as expected, it is situation 3. When trying out the combinations, there will be no problem. However, the world may be different; like, for instance, in situation 1:

cat distributed
pred distributed
on H's own authority
situation 1:



Now, when trying out the combination:



the first cat gives a match, so
the sketch clashes with it.

If H has in fact some cats in view, the difference between $not\ a\ cat^{\downarrow} * predF$ and $not\ a\ cat^{\uparrow} * predF$ becomes very subtle; given the first type of information, he will successively cover the cats with the picture, and given the second type of information, he will shift the sketch back and forth between them, but both shape absence, so neither of them is to fuse with what is there, and in both cases a single match would render the information incorrect.

As repeatedly discussed, there is a systematic difference between single pictures and repetitive pictures; however, in pictures of absence, this difference is all but neutralized. Let's assume again that H, on first inspection, has two cats in view. Now compare, for instance,

$$\begin{aligned} not\ a\ cat^{\downarrow} * predF &> \langle \underline{\underline{0}}^{cat} \ 0 \ 0 \ / \ 0^{pred} \ / \ a\ cat \rangle, \text{ as illustrated above, and} \\ not\ a\ cat^{\downarrow} \leftrightarrow predF &> \{ \langle \underline{\underline{0}}^{cat} \ 0 \ 0 \ / \ 0^{pred} \ / \ a\ cat \rangle \}. \end{aligned}$$

The latter would give a repetitive shape of absence to cover each cat in view with its own copy of the picture, rather than a single shape placed and replaced from one to the other. However, in both cases, H does not expect any match to be there at all, so in both cases, just a single match already clashes with S's information. I believe there still may be a subtle difference between single and repetitive pictures of absence, but I shall not go into that.

Chapter 5

The Shape of Emphasis

Fall-rise intonation seems to make *not* zoom in on a small spot of the sentence. The idea is that the fall-rise form can be broken into basic forms that each make their own contribution. In chapter 3, the example was (1/3), abbreviated to (15/3):

(1/3) $\overline{a\ s\text{---}i\ s\ t\ e\ r\ o\ f\ J\text{---}o\ h\ n\ 's\ ^\uparrow\ * \ i\ s\ n\text{---}o\ t\ \text{---}i\ n\ t\ e\ r\ e\ s\ t\ e\ d\ /}$

For the discussion here, this is the type

$\overline{a\ c\text{---}a\ t\ ^\uparrow\ * \ n\text{---}o\ t\ p\ r\text{---}e\ d\ F\ /}$

In the preceding chapters, the contribution made by the noun phrase, the predicate, *not*, and the prosodic forms \uparrow and $*$ has been addressed. What is left to discuss is the final rise, pitch accent, especially its distribution and absence, and peaking pitch.

5.1 Pieces of information

Peaking pitch

The most important contribution is that of peaking pitch. What happens is that S, speaking with emphasis, not only calls up the regular image of what he asserts or excludes, but simultaneously an extra image of what he rejects, or rejects to exclude. In chapter 1, the working hypothesis was that peaking pitch doubles the regular image, erases, in this extra image, the spot singled out by the pitch peak, the gap to be filled by an opposite to be found in context or situation, and reverses the result. By now, I can give a notation that directly represents rather than describes this extra image:

NOTATION

$o|^{(-/cat)}$ as in $\backslash o|^{(-/cat)} / o|^{pred} / a^{(-/cat)}$

the feature-image linked to *cat* to be replaced by an opposite feature-image, to be found in the context or situation.

$\cup \cup$ as in $\cup \backslash o|^{(-/cat)} / o|^{pred} / a^{(-/cat)} \cup$

reversed expectation of what the world in question is or is to be like.

Pitch accent and its absence

Pitch accent and its absence differ in function from the other prosodic forms discussed; they do not contribute to the build-up of the image under construction, but their 'old-new' distinction apparently helps to knot ongoing information together. For lack of time, this book is mostly limited to isolated sentences, so mostly ignores pitch accent. Also this chapter focuses on emphasis rather than pitch accent. Emphasis is used to make a stronger case by simultaneously

(A) without context:

H simultaneously gets two pictures. The regular one is a single closed picture of a cat with the predicate-feature, a stand-in for some real cat that is there. Given the partly unspecified character of the picture insofar as authorized, H can still alternatively assemble a range of cats in his mind, each one replacing the earlier one, but each is a possibly wrong guess of what S's actual cat looks like. The extra picture is a single closed picture of something-other-than-cat with the predicate-feature. In the absence of any clues, H will tentatively choose the first thing opposing cats that comes to his mind, not improbably, dogs. With respect to this extra picture, H's expectation of what is there is reversed, so it acts as a doorkeeper: where the regular picture projects some cat with the predicate feature into his idea of the world talked about, this reversed extra picture prevents H from even tentatively adding any dog with the predicate feature to his idea of that world.

(B) Given the pitch accent on *cat* and the absence of a pitch accent on *pred*, a suitable context would be

(other speaker P:
 $a \text{ dog}^{\downarrow} * \text{pred}F > \underbrace{\backslash \text{dog} \text{ o o } / \text{o}^{\text{pred}} / a \text{ dog}} \quad)$
 $a \text{ c}|\text{at}^{\downarrow} * \text{pr}—\text{ed}F > \text{the regular and the extra} \quad \underbrace{\backslash \text{(-/cat)} \text{ o o } / \text{o}^{\text{pred}} / a \text{ (-/cat)}} \quad \underbrace{\quad}$

First, H gets P's information; on this basis he expects some unspecified dog to be there with the predicate-feature. Next, he gets S's information. Now S's extra picture can be superimposed on P's picture, filling up its gap $\text{o}^{(-/cat)}$ to become o^{dog} , and reversing H's expectation that there will be such a dog there, into the expectation of the absence of such a dog:


$$\underbrace{\underbrace{\backslash \text{(-/cat)} \text{ o o } / \text{o}^{\text{pred}} / a \text{ (-/cat)}}}_{\underbrace{\backslash \text{dog} \text{ o o } / \text{o}^{\text{pred}} / a \text{ dog}}}$$

This removes on S's authority P's single dog with the predicate feature from H's idea of what is there, and blocks the addition of any other such dog.

(C) H has the world in sight:


For the regular picture, see chapter 3. Now given the emphasized *cat* rather than the ordinary *cat*, H, on first inspection, along with distributing *cat*, makes an inventory of the $(-/cat)$ category closest by, depending on the context and the situation in view; assuming that there are two dogs next to the two cats:

cat distributed
inventory of ($\text{---}/\text{cat}$)
on H's own authority



Now, after a first inspection, including the predicate, on the one hand, H has:


given the ($\text{---}/\text{cat}$) inventory
on S's authority, the fork



The extra picture called up by S reverses H's expectation that the world is such that the picture can fuse with it somewhere when he alternatively tries out the options, into the expectation that the world is such that there is no way to fuse the picture with it. This is the case if in both options there is no match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected, given both the regular and the extra picture:

cat distributed
inventory of ($\text{---}/\text{cat}$)
pred distributed
on H's own authority



When trying out the combinations, there will be no problem. It is easy to see that a dog that would have the predicate-feature would clash with the information, and therefore, the illustration can be skipped.

TYPE $a\ c|at^{\uparrow} * pr-edF$ > the regular $\backslash o|^{cat} \dots / o|^{pred} / a\ cat$
 > the extra $\cup \backslash o|^{(-/cat)} \dots / o|^{pred} / a\ (-/cat) \cup$

Example (2/3) and its abbreviation (14/3) in chapter 1 are of this type:

(2/3) $a\ s|ister\ of\ J-ohn's^{\uparrow} * is\ -interested$

(A) without context:

H simultaneously gets two single open sketches. The regular one includes the predicate feature in a cat-skeleton, and leaves the room for complementary features open. The extra sketch includes the predicate-feature in a something-other-than-cat skeleton, leaving the room for complementary features open, and reversing H's expectation of what is or may be there, so it acts as a doorkeeper: it prevents H from even tentatively combining the predicate feature with this other-than-cat feature in his idea of what is or may be there.

(B) Given the pitch accent on *cat* and the absence of a pitch accent on *pred*, a suitable context would be

(other speaker P:
 $a\ dog^{\uparrow} * predF$ > $\backslash o|^{dog} \dots / o|^{pred} / a\ dog$)
 $a\ c|at^{\uparrow} * pr-edF$ > the regular and the extra $\cup \backslash o|^{(-/cat)} \dots / o|^{pred} / a\ (-/cat) \cup$

First, H gets P's sketch; on this basis he includes the predicate-feature into the open skeleton of a dog: there is or may be such a dog. Next, he gets S's information. Now S's extra picture can be superimposed on P's sketch, filling up its gap $o|^{(-/cat)}$ to become $o|^{dog}$, and reversing H's expectation that such a dog might be there, into the expectation of the absence of such a dog:

$$\cup \backslash o|^{(-/cat)} \dots / o|^{pred} / a\ (-/cat) \cup$$

$$\backslash o|^{dog} \dots / o|^{pred} / a\ dog$$


This, on S's authority, removes from H's idea of what is or may be there any dog with the predicate feature that he might have thought up, and it blocks the addition of any other such dog.

(C) H has the world in sight:

Assume that S's double sketch is a partial description of a realized world. For the regular sketch, see chapter 3. But given the emphasized *cat* rather than the ordinary *cat*, along with distributing

cat, on first inspection, H makes an inventory of the ($\text{---}/cat$) category closest by, depending on the context and the situation in view; let's assume, as above, that H has two cats and two dogs in view. Now, including the predicate, on the one hand, H has:

given the ($\text{---}/cat$) inventory
on S's authority



The extra sketch called up by S reverses H's expectation that the world is such that the sketch can fuse with it when he shifts it back and forth between the two dogs, into the expectation that the world is such that they cannot fuse. This is the case if there is no match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected, it is the same world as above. When trying out the combinations, there will be no problem. It is easy to see that a dog, having the predicate-feature, would clash with the information, and therefore the illustration can be skipped.

—

TYPE $a \ c|at^{\downarrow} \ * \ not \ pr\text{---}edF$ > the regular $\backslash \underline{0}^{cat} \ 0 \ 0 \ / \langle \underline{0}^{pred} \rangle / \ a \ cat$

> the extra $\cup \backslash \underline{0}^{(-/cat)} \ 0 \ 0 \ / \langle \underline{0}^{pred} \rangle / \ a \ (-/cat) \cup$

Example (1/6) and the abbreviated (15/6) in chapter 1 are of this type:

—

(1/6) $a \ s|ister \ of \ J\text{---}ohn's^{\downarrow} \ * \ is \ not \ \text{---}interested \ (/)$

(A) without context:

H simultaneously gets two pictures. The regular one is a single closed picture of a cat that lacks the predicate-feature, a stand-in for some unspecified real cat, expected to be there. The extra picture is a single closed picture of something-other-than-cat that lacks the predicate-feature, let's say a dog; for this extra picture, H's expectation of what is there is reversed, so it acts as a doorkeeper: it prevents H from even tentatively adding any dog without the predicate feature to his idea of the world talked about.


(B) Given the pitch accent on *cat* and the absence of a pitch accent on *pred*, a suitable context would be

(other speaker P:
 $a \ dog^{\downarrow} \ * \ not \ predF$ > $\backslash \underline{0}^{dog} \ 0 \ 0 \ / \langle \underline{0}^{pred} \rangle / \ a \ dog$)

—

$a \ c|at^{\downarrow} \ * \ not \ pr\text{---}edF$ > the regular and the extra $\cup \backslash \underline{0}^{(-/cat)} \ 0 \ 0 \ / \langle \underline{0}^{pred} \rangle / \ a \ (-/cat) \cup$

First, H gets P's information; on this basis he expects some unspecified dog that lacks the predicate-feature to be there. Next, he gets S's information. Now S's extra picture can be superimposed on P's picture, filling up its gap $o|^{(-/cat)}$ to become $o|^{dog}$, and reversing H's expectation that there will be such a dog there, into the expectation of the absence of such a dog:




This, on S's authority, removes P's dog that lacks the predicate feature from H's idea of what is there, and blocks the addition of any other such dog.

(C) H has the world in sight:

For the regular picture, see chapter 4. Again, along with distributing *cat*, H makes an inventory of the $(-/cat)$ category closest by; let's assume, as above, two dogs and two cats in H's view. Now, including the predicate, on the one hand, H has:

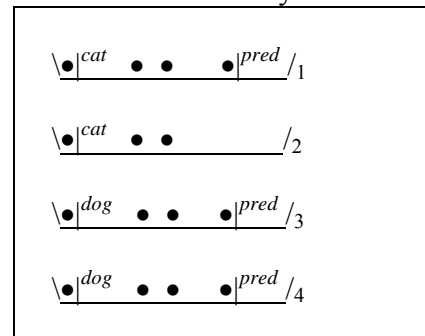
given the $(-/cat)$ inventory
on S's authority, the fork



The extra picture called up by S reverses H's expectation that the world is such that the picture can fuse with it when he alternatively tries out the options, into the expectation that the world is such that he cannot do this. This is the case if in both options there is no match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected, given both the regular and the extra picture:

cat distributed
inventory of $(-/cat)$
pred distributed
on H's own authority



When trying out the combinations, there will be no problem.

TYPE $a\ c|at^{\uparrow} * not\ pr\text{---}edF$ > the regular $\backslash o|^{cat} \dots / \langle o|^{pred} \rangle / a\ cat$
 > the extra $\cup \backslash o|^{(-/cat)} \dots / \langle o|^{pred} \rangle / a^{(-/cat)} \cup$

Example (1/5) and the abbreviated (15/5) in chapter 1 are of this type:

(1/5) $a\ s|ister\ of\ J\text{---}ohn's^{\uparrow} * is\ not\ \text{---}interested$

(A) without context:

H simultaneously gets two single open sketches. The regular one excludes the predicate feature from the cat-skeleton; this excludes a cat with the predicate-feature from H's idea of what is or may be there. The extra sketch excludes the predicate-feature from a something-other-than-cat skeleton, but reversing H's expectation; so a something-other-than-cat with the predicate-feature takes shape in H's idea of what is or may be there.

(B) Given the pitch accent on *cat* and the absence of a pitch accent on *pred*, a suitable context would be

(other speakerP:
 $a\ dog^{\uparrow} * not\ predF$ > $\backslash o|^{dog} \dots / \langle o|^{pred} \rangle / a\ dog$)
 $a\ c|at^{\uparrow} * not\ pr\text{---}edF$ > the regular and the extra $\cup \backslash o|^{(-/cat)} \dots / \langle o|^{pred} \rangle / a^{(-/cat)} \cup$

First, H gets P's sketch of some realized or unrealized situation; on this basis he excludes the predicate-feature from the otherwise open skeleton of a dog; H expects the absence of a match, he expects that he will not find the predicate feature combined with the dog-feature. Next, he gets S's information. Now S's extra sketch can be superimposed on P's sketch, filling up its gap $o|^{(-/cat)}$ to become $o|^{dog}$, and reversing H's expectation of the absence of the combination of the predicate feature with the dog-feature, into the expectation of their combination in what is or may be there:


$$\cup \backslash o|^{(-/cat)} \dots / \langle o|^{pred} \rangle / a^{(-/cat)} \cup$$

$$\backslash o|^{dog} \dots / \langle o|^{pred} \rangle / a\ dog$$

(C) H has the world in sight:

For the regular sketch, see chapter 4. Let's assume, as above, that H has two dogs and two cats in view. Now, including the predicate, on the one hand, H has:

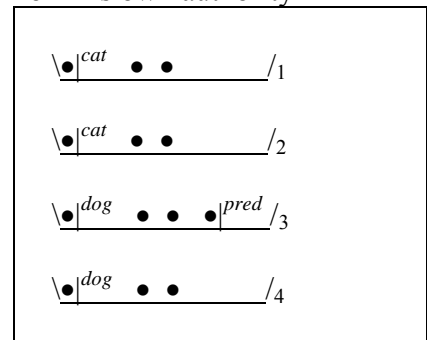
given the ($\text{---}/cat$) inventory
on S's authority



The extra image called up by S reverses H's expectation that the world is such that the sketch does not clash with it when he shifts it back and forth between the two dogs, into the expectation that the world is such that it does clash with it. This is the case if for at least one of the dogs there is a match.

On the other hand, H has the world in view. Assuming that the world, on second inspection, is as expected:

cat distributed
inventory of ($\text{---}/cat$)
pred distributed
on H's own authority



TYPE $a\ c\ \text{---}at^{\uparrow} * \text{not } pr\ \text{---}edF /$

This is at last the fall-rise that seems to make *not* zoom in, resulting in a positive implication, as in chapter 1, (1/3), and its abbreviation (15/3):

(1/3) $a\ s\ \text{---}ister\ of\ J\ \text{---}ohn's^{\uparrow} * \text{is } \text{not } \text{---}interested /$

(A) without context:

Both the regular and the extra image are as in the preceding case. However, the accented $c|at$ of that case, and the unaccented $c\text{---}at$ of this case have a different impact in combination with the gap $o|^{(-/cat)}$: $c|at$ is 'new' in an 'old' surrounding; H will tend to assume that $c|at$ replaces a preceding opposite; and $c\text{---}at$ is 'old' in an 'old' surrounding, and therefore suggests that the opposite has not already been given; the final rise then adds to the impression that the image still has a gap to be filled.

(B) As said before, in this type both the regular and the extra image are the same as in the preceding case. But because of their difference in pitch accent, they fit in a different context. Compare them again. Here, the complete absence of pitch accent makes it fit a context like

$$\begin{array}{l}
 \text{(other speaker P:} \\
 a \text{ cat}^{\uparrow} * \text{ pred}F \quad > \quad \underbrace{\backslash \text{O} |^{cat} \dots / \text{O} |^{pred} /}_{cat} \quad) \\
 \text{---} \\
 a \text{ c---at}^{\uparrow} * \text{ not pr---ed}F / \quad > \text{ the regular} \quad \underbrace{\backslash \text{O} |^{cat} \dots / \langle \text{O} |^{pred} \rangle /}_{a \text{ cat}} \\
 > \text{ the extra} \quad \cup \underbrace{\backslash \text{O} |^{(-/cat)} \dots / \langle \text{O} |^{pred} \rangle /}_{a \text{ (-/cat)}} \cup
 \end{array}$$

In an ongoing text, this makes a lot of difference, as it knots the information together in a different way: first, H gets P's sketch for some realized or unrealized situation, including the predicate-feature into the skeleton of a cat. Next, he gets S's information, and this time, it is S's regular sketch rather than the extra one of the case above that overrules what P said, superimposing it on P's sketch:

$$\boxed{
 \begin{array}{c}
 \underbrace{\backslash \text{O} |^{cat} \dots / \langle \text{O} |^{pred} \rangle /}_{a \text{ cat}} \\
 \underbrace{\backslash \text{O} |^{cat} \dots / \text{O} |^{pred} /}_{a \text{ cat}}
 \end{array}
 }$$

and it is S's extra image rather than the regular one that implies something else instead.

(C) H has the world in sight:

As in the preceding case, H expects the absence of cats with the predicate-feature, and the presence of something-other-than-cat with the predicate-feature. However, it is far less likely that he already is able to make an $\text{O} |^{(-/cat)}$ inventory.

5.3 More Combinations

Emphasized *not*

A few words about another variant of (1):

$$(1/7) \quad a \text{ s---ister of J---ohn's}^{\uparrow} * \text{ is n|ot ---interested}$$

Simplified, this is the type

$$a \text{ c---at}^{\uparrow} * \text{ n|ot pr---ed}F$$

The regular image is $\underbrace{\backslash \text{O} |^{cat} \dots / \langle \text{O} |^{pred} \rangle /}_{a \text{ cat}}$. Suppose H has two cats in view, then he expects that the world is such that the regular image does not clash with it when he shifts it back and forth over these two cats. This is the case if there is no match. The hypothesis is that peaking pitch gives H an extra image for which the expectation is reversed, by copying this one, in the

copy erasing the spot singled out, to replace it by a contextual opposite. Now here the spot singled out is *not*. Intuitively, what happens also in this case is that the contribution of *not* to this extra sketch gets erased, to be replaced by the contribution of a contextual opposite; the contribution of *not* is to mark the image as ‘absent’; if we for now assume that its opposite is to mark the image as ‘present’, the result is $\cup \underline{\underline{O|^{cat} \dots / O|^{pred} / a \ cat}} \cup$. This extra sketch reverses H’s expectation that the world is such that $\underline{\underline{O|^{cat} \dots / O|^{pred} / a \ cat}}$ fuses with it when it is shifted back and forth over the two cats in view, into the expectation that the world is such that they cannot fuse in this way. This is the case if there is no match.

Thus, S, simultaneously, in two ways brings the same message: the sketch of absence $\underline{\underline{O|^{cat} \dots / \langle O|^{pred} \rangle / a \ cat}}$ does not clash, and the sketch of presence $\underline{\underline{O|^{cat} \dots / O|^{pred} / a \ cat}}$ does not fuse with what is there. With context:

$$\begin{array}{lcl}
 \text{(other speaker P:} & & \\
 a \ cat^{\uparrow} * \ pred F & > & \underline{\underline{O|^{cat} \dots / O|^{pred} / cat}} \) \\
 \text{---} & & \\
 a \ c\text{---}at^{\uparrow} * \ n|ot \ pr\text{---}ed F & > \text{ the regular} & \underline{\underline{O|^{cat} \dots / \langle O|^{pred} \rangle / a \ cat}} \\
 & > \text{ the extra} & \cup \underline{\underline{O|^{cat} \dots / O|^{pred} / a \ cat}} \cup
 \end{array}$$

So S in one statement twice overrules P’s statement: by contradicting it, and by reversing it.

This, however, requires some further thinking. Between a negative statement containing *not*, and a positive statement not containing *not*, there is a whole range of modal adverbs qualifying the probability of finding the world fitting the image. I believe they could be smoothly included in the present theory, but in this book I shall not elaborate on that.

Refining the picture

Within the types of information discussed, there is a special category that may seem more difficult to explain; take for example

$$\begin{array}{l}
 (39) \quad \underline{\underline{she's^{\downarrow} * \ not \ pr\text{---}etty /}} \\
 \quad \quad \quad \underline{\underline{she's \ gorgeous}}
 \end{array}$$

If verbal information would be logic, this certainly would be a contradiction.

To get an example where *not* zooms in on an indefinite subject look again at

$$(1/3) \quad \underline{\underline{a \ s\text{---}ister \ of \ J\text{---}ohn's^{\uparrow} * \ is \ not \ \text{---}interested /}}$$

As said before, H immediately understands its double message; its implication is clear: someone-other-than-sister is interested, and a continuation like

$$(a) \quad a \ brother^{\uparrow} * \ is \ interested$$

would not come as a surprise. Now, given the very same (1/3), having understood what it says, H, subsequently, is very well able to make sense of the continuation

(b) *the whole family*[↑] * *is interested*

Although this does not come as expected, it comes as a surprise and not as a contradiction. The question, therefore, is: what happens here?

Note that understanding the information is getting the picture; S's verbal information does not say to H what the world is like if his statement is true (see also chapter 7), but it calls up an image of what the world can be expected to look like. And in these cases, the effect is a rhetorical trick that S plays on H's expectation: in the reversed sketch, he erases the sister-detail, leaving a gap there, and makes H expect something-other-than-sister filling the gap, but then S, instead of the expected opposite-of-sister, draws in the whole family, which includes the sister again: a more accurate version of what he just erased. Switch to H's perspective: let's assume that, before S said anything, he already had this interested sister in mind. Now, hearing S say (1/3), on S's authority, he removes this interested sister from his idea of the situation, only to re-enter her, along with her whole family, when S continues with (b). This process has nothing to do with truth.

A complex combination of basic forms

In this book, I hypothetically isolated a small set of basic informative forms, each form making its own precise contribution to the fragmentary image that takes shape in the hearer's mind. They unite into a multiplication of possible complex forms, each calling up its own image.

There are more possible combination than illustrated here; for instance, I did not illustrate a combination of peaking pitch, calling up a double image, and a lengthening of the final syllable of the noun phrase, calling up repetitive images. However, if the hypotheses about the separate forms and their contributions are correct, all combinations should come out correctly. For instance, in chapter 1, I mentioned

(31) *a dr—ake*^{↑↔} *is not a m—ale sw—an* /

Its meaning is abstract and complex; nevertheless, understanding is immediate. According to the theory, it calls up the double sketch:

the regular	{ \o drake ... /< \o male swan > / a drake }
the extra	↻ { \o (-/drake) ... /< \o male swan > / a (-/drake) } ↻

which nails down the understanding.

Chapter 6

More Noun Phrases

The question about verbal information is: how does it work. For lack of time, in this book I discussed only a very small fragment of language: the simplest structures, built from words linked to concrete, down-to-earth aspects of what is there. There is a lot that has not been covered. However, if we want to get hold of the processes of understanding and recognition, we should start with the basics; we will not be able to understand how speakers give information by using complex structures if we do not understand how they give information by using simple structures. If we understand how simple structures work, we have a start. And we will not be able to understand how abstract words connect with what is there if we do not understand how hearers recognize concrete real things as fitting concrete words; however abstract some information may be, in order to make a connection between this information and the real world, one ultimately has to rely on one's own senses, one's own concrete experience.

As illustrated, small though the fragment may be, there is already a whole range of facts to be accounted for; this is where my investigation started. Throughout this book, the idea is that there is a simple system behind these facts: a few basic forms that each systematically contribute their own piece of information, multiplied into a range of possible combinations. As illustrated, just exchanging one form by another will change the message, subtly or drastically, but always systematically. This far, in the illustration of possible combinations, I have each time included an indefinite subject noun phrase. It is to be expected that an exchange of the indefinite for another kind of noun phrase will allow for the same range of possible combinations, and the exchange of this one piece of information will account for their range of ambiguities and nuances. In the following, I shall illustrate definite noun phrases, proper names, and material noun phrases.

6.1 Definite Noun Phrases

Before illustrating the range of possible combinations, the question we should ask ourselves is: precisely what piece of information does S, by choosing a definite noun phrase, get across to H. Take S's perspective. Suppose he has some cat in view that he wishes to talk about. He tags it with *cat*; however, this tag does not cover features like colour and other characteristics. To cover the entire thing including its other features, S can construct a noun phrase; for instance, if a feature in a separate unification of features can be tagged *cat*, S can label the whole unification *a cat*, and if a feature in a unification of features can be tagged *cat*, S can exclusively label the whole unification *the cat*.

Obviously, if there is something there that uniquely has the feature $o|^{cat}$, and S wants to say something about it, he can use the noun phrase *the cat*. However, this does not mean that *the cat* informs H that it refers to the one and only thing there with the feature $o|^{cat}$. Switch perspective to H; note two facts, one about existence, and one about exclusivity. In the first place, a claim of existence of something fitting the definite noun phrase is not included in the information contributed by this noun phrase itself. In this respect, the definite noun phrase does not differ from the indefinite noun phrase; compare for instance:

- (40) *the logic of it escapes me*
 (41) *the answer remains forthcoming*

Note that there is no negation or other element in these examples that can be held responsible for undoing a claim of existence; from the start, this information is not there. *Existence*-examples fit in here; compare

- (42) *the Loch Ness Monster exists* (correct or not)

If a claim of existence were included in *the Loch Ness Monster*, (41) would have been tautological, which clearly it is not. Negative examples are even easier to find:

- (43) *the connection between these two events is not clear to me*
 (37) *the Bogeyman doesn't exist*

In the second place, it is not so that a phrase like *the cat* informs H of the uniqueness of the tag *cat*, the feature $o|^{cat}$; compare the following example:

- (38) *the neighbour called the man wanted to borrow a cup of sugar*


(38) certainly does not say that the neighbour is the only man there. Read any text, and you will find plenty of similar examples. What is the case here is that the label *the man* as a whole, in this situation, and at this moment, exclusively applies to one person.

So it seems reasonable to suppose that a definite noun phrase as a whole, by itself, without prosodic forms, contributes the information of exclusive applicability at the moment of its utterance. To incorporate this information in the theory, compare *a cat* and *the cat*. They have a lot in common. In the first place, *cat* built into the noun phrase in both cases calls up $o|^{cat}$, and plugs in at the same point in H's brain, accessing the criss-cross of overlapping and diverging records of phenomenal features, characteristics, actions, etc. which in H's past experience have gone together with the cat-shape. In the second place, in both cases, the noun phrase calls up a cat-shaped outline; in notation:

<i>a cat</i>	>	$\backslash o ^{cat} /^{a cat}$	the skeleton of a unification of features fitting <i>a cat</i> , consisting of $o ^{cat}$ and room for complementary features.
<i>the cat</i>	>	$\backslash o ^{cat} /^{the cat}$	the skeleton of a unification of features fitting <i>the cat</i> , consisting of $o ^{cat}$ and room for complementary features.

Both are equally fragmentary in what they authorize. In both cases, the stock of recorded complementary features accessed determines the range of cats that H still tentatively can assemble in his mind; in both cases, as far as H is informed, the skeleton may still grow into any cat. The difference is this: given just the unification of features under construction $\backslash o|^{cat} /^{a cat}$, each cat it can grow into is equally good, but the unification of features under

construction $\backslash \underline{0}^{cat} /_{the\ cat}$ is exclusive, whatever it is or whatever it may grow into, whatever H thinks up or finds out, it excludes other things, also other things that may have the feature $o|^{cat}$.




may grow into any discrete unification of features.


may, in the situation concerned, at the moment of speech, grow into an exclusive unification of features.

The difference is especially clear when the world concerned comes in view, when H on first inspection is to find out what comes into consideration. Compare: given *a cat*, H has

S's noun phrase skeleton




cat distributed
on H's own authority




That is, H's inventory of what comes into consideration may contain any number of cats. In contrast, given *the cat*, his search stops with the first thing tagged *cat* that comes in view; it may be the only thing tagged *cat* in fact there, or the most prominent one, the one closest by, the one in focus, but anyhow, this one with the exclusion of anything else. So

S's noun phrase skeleton



cat distributed
on H's own authority



So, if the situation is realized, and in view, and if H is able to distribute *cat*, he is able to identify the one and only thing that comes into consideration in that particular situation, and at that moment of speech.

When S uses a noun phrase, he composes his information of at least three different forms: the noun phrase construction itself, the pitch form \downarrow or \uparrow , and the rhythmic form $*$ or \leftrightarrow ; when the noun phrase is definite, this does not make a difference. Also, the process of recognition is the same; but as can be expected, the exclusivity of the label up to a point neutralizes the differences between the four types of information that are so clear in case of an indefinite noun phrase. Still the differences are unmistakably there. Remember the four *a tomcat* examples:

- (25) *a tomcat*[↓]* *is an especial beauty*
 (26) *a tomcat*[↑]* *is a good choice*
 (27) *a tomcat*^{↓↔} *is a good mouser*
 (28) *a tomcat*^{↑↔} *is a male cat*

Now listen to the following examples and look at the types in more detail:

- (44) *the winner*[↓]* *is ecstatic*
 (45) (as a partial description of a realized situation; at the ceremony:)
 (*and*) *the winner*[↑]* *is ...* (follows a name)
 (46) (as a design for a still unrealized situation, before the final decision:)
the winner[↑]* *is on the short list*
 (47) (in the rules of the game:)
the winner^{↓↔} *gets the pool*
 (48) (in the rules of the game:)
the winner^{↑↔} *is the first to arrive*

TYPE *the cat*[↓]* *predF* > \o|^{cat} o o / o|^{pred} / *the cat*


Called up in H's mind, there is a single closed picture consisting of the skeleton of an exclusive unification of features, containing o|^{cat} and unspecified complementary features; o|^{pred} included. Examples of this type are

- (44) *the winner*[↓]* *is ecstatic*
 (49) *the murderer*[↓]* *is insane*

In this combination, the information does claim the existence of something specific fitting the label at this moment. If H does not have the real situation in view, to him the picture stands-in for the real thing that is supposedly there. As far as H is informed, he may still tentatively assemble a range of cats in his mind, each replacing the earlier one, each one very possibly being a wrong guess of what the actual thing looks like.

Let's assume that H has the real situation in view, and that first inspection makes him identify the one and only thing that comes into consideration, say, \●|^{cat} ● ● /₁, then, including the predicate, H has


on S's authority



and H will expect the world to be such that S's picture can fuse with it, uniting picture and real thing, so he expects this specific cat to have the predicate feature. I shall omit further illustration.

Compare the definite $the\ cat^{\downarrow} * predF$ and the indefinite $a\ cat^{\downarrow} * predF$. Both call up a single stand-in for something supposedly there, in both cases the stand-in is to fuse and unite with the real thing if it comes in view. The difference comes when H is to determine what comes into consideration: given $the\ cat$, there should be only one thing; $the\ cat$ makes him stop at the first thing found, and S's information is clear, as illustrated above. But $a\ cat$ leaves H as many options as there are things tagged cat , which may constitute a problem as he may have to keep the choice on hold; with two cats in view, H has

on S's authority, the fork:



Of course also given a definite subject noun phrase there are all kinds of variations in recognition. Often enough, H will be able to identify the thing meant but will be unable to distribute $pred$ on his own authority, so he gets new information about the thing identified, to be added in the world's unspecified space. Also, it may happen that he cannot distribute $pred$ because it mentions some future characteristic.

Note that there is no guarantee that H will be able to identify the thing meant, even if the information claims its exclusive existence; for example

(50) $the\ murderer^{\downarrow} * whoever\ he\ may\ be\ is\ a\ monster$

Usually, H will not question the information until he finds that he cannot reconcile it with what he knows otherwise. Of course this may happen; the thing identified may turn out not to have the predicate-feature. And every once in a while it will happen that H knows for sure that there is nothing that comes into consideration; he does get the picture, but he does not even need to hear the predicate to know that it will not fuse with what is there. For example, if he hears

(49) $the\ murderer^{\downarrow} * is\ insane$

when he is absolutely sure the death was an accident.

TYPE $the\ cat^{\downarrow} * not\ predF > \underline{o|cat\ o\ o} / \langle o|pred \rangle / the\ cat$

Called up in H's mind, there is a single closed picture consisting of the skeleton of an exclusive unification of features, containing $o|^{cat}$ and unspecified complementary features, but $o|^{pred}$ pictured as absent, excluded. An example of this type is:

(51) $the\ murderer^{\downarrow} * is\ not\ legally\ accountable$

Further elaboration is unnecessary.

TYPE $the\ cat^{\uparrow} * predF > \underline{\backslash o|^{cat} \dots / o|^{pred} /}^{the\ cat}$


Called up in H's mind, there is a single open sketch consisting of the skeleton of an exclusive unification of features, containing $o|^{cat}$ and room for complementary features kept open; $o|^{pred}$ included. Given such a fragmentary sketch, the situation may be realized, unrealized, or unfolding. Examples of this type are

- (52) $the\ murderer^{\uparrow} * is\ one\ of\ them$
- (45) $(and)\ the\ winner^{\uparrow} * is \dots$ (follows a name)
- (46) $the\ winner^{\uparrow} * is\ on\ the\ short\ list$
- (53) $the\ tide^{\uparrow} * is\ coming\ in$
- (41) $the\ Loch\ Ness\ Monster^{\uparrow} * exists$
- (42) $the\ logic\ of\ it^{\uparrow} * escapes\ me$

Given this fragmentary information, what H can still assemble in his mind may go in all directions, but any which way it goes, it goes with the exclusion of all other possibilities; as long as the situation is not realized, each possibility is as good as any other, but as soon as the situation is realized, all but one expire. Still, H may or may not be able to identify that one exclusive thing.

Let's assume that the situation is realized, that H has it in view, and that first inspection makes him identify that thing, let's take again $\backslash \bullet |^{cat} \bullet \bullet /_1$, then, including the predicate, H has


on S's authority



H will expect the world to be such that S's sketch can fuse with it, so he expects this specific cat to have the predicate feature. Further illustration is unnecessary.

Compare the definite $the\ cat^{\uparrow} * predF$ and the indefinite $a\ cat^{\uparrow} * predF$. The difference is especially clear when H is to determine what comes into consideration: with the real situation in view, $the\ cat$ makes him stop at the first thing found, and S's information is as illustrated above, but in case of $a\ cat$, H may find more than one suitable thing. In the event that there are two of them, H will continue to shift S's sketch back and forth between them:

on S's authority



Compare the two definite types *the cat*[↓]* *predF* and *the cat*[↑]* *predF*. The difference between them can be very subtle, especially if the second is a partial description of a realized situation, and if H can in both cases identify the one and only thing that comes into consideration. Take for example

- (54/1) (Question: *what about his siblings?* Answer:) *the sister*[↓]* *is interested*
 (54/2) (Question: *who would buy such things?* Answer:) *the sister*[↑]* *is interested*

Still, there is a difference: (54/1) is about the sister, and (54/2) states a fact. The difference is most notable in examples showing that a claim of existence is not included in *the cat*[↑]*, like

- (41) *the Loch Ness Monster*[↑]* *exists*
 (40) *the logic of it*[↑]* *escapes me*

TYPE *the cat*[↑]* *not predF* > $\backslash \underline{\text{O}}^{\text{cat}} \dots / \langle \text{O}^{\text{pred}} \rangle / \text{the cat}$


Called up in H's mind, there is a single open sketch consisting of the skeleton of an exclusive unification of features, containing O^{pred} and room for complementary features left open, but O^{pred} pictured as absent, excluding it.

Because of the exclusivity of the label, the difference between the two definite types *the cat*[↓]* *not predF* and *the cat*[↑]* *not predF* is much subtler than in the parallel indefinite case illustrated in chapter 4, especially if H can identify the one and only thing that comes into consideration. Again, the difference is most notable in examples showing that a claim of existence is not included in *the cat*[↑]*, like

- (42) *the connection between these two events*[↑]* *is not clear to me*
 (43) *the Loch Ness Monster*[↑]* *does not exist*


As can be expected, because of the exclusivity of the definite label, there is a big difference with the parallel indefinite case. Let's assume that H has the world in view, that there are in fact two cats, but that H, given *the cat*[↑]* *not predF*, has identified the one closest by as the one thing that comes into consideration. Now he has

on S's authority



So it is only this one specific cat that is expected to lack the predicate feature. But given *a cat*[↑]* *not predF*, both cats come into consideration, and H has

on S's authority



So it is both cats that are expected to lack the predicate feature.

TYPE *the cat* $\downarrow \leftrightarrow predF$ > { $\underline{\underline{\o|^{cat} \ o \ o \ / \ o|^{pred} / the \ cat}}$ }

Called up in H's mind, there is a self-replicating closed picture, each of these copies consisting of the skeleton of an exclusive unification of features, containing $\o|^{cat}$ and unspecified complementary features, $\o|^{pred}$ included. Examples of this type are:

- (47) *the winner* $\downarrow \leftrightarrow$ *gets the pool*
- (55) *the receiver* $\downarrow \leftrightarrow$ *is as bad as the thief*


The combination is interesting: the picture called up is both self-replicating and exclusive in its situation. H may assemble a range of cats (winners, receivers) in his mind, but each exclusive in its situation, so each cat (winner, receiver) he imagines needs its own situation; for each, the predicate feature is included.

Both the definite and the indefinite noun phrase can be used to make general statements about whatever fits the description; compare


- (27) *a tomcat* $\downarrow \leftrightarrow$ *is a good mouser*
- (47) *the winner* $\downarrow \leftrightarrow$ *gets the pool*

The difference is that H pictures his tomcats each as one amongst many, and his winners each on their own. The situation may come in view. If 'in view' is taken literally, this will be only one situation at a time, but H may very well have more than one situation in mind, based on independent information. In the event of two situations, each with its own thing that exclusively comes into consideration, we might, a bit ad hoc perhaps, represent the combination of information as follows: H has

on S's authority



In chapter 3, the parallel indefinite case was represented as



TYPE *the cat*[↓] ↔ *not predF* > { $\backslash \underline{o|^{cat} \quad o \quad o} / \langle o|^{pred} \rangle / \textit{the cat}$ }

Called up in H's mind, there is a self-replicating closed picture, each copy consisting of the skeleton of an exclusive unification of features, containing $o|^{cat}$ and unspecified complementary features, $o|^{pred}$ pictured as absent, excluded. An example of this type is:

(56) *the leopard*[↓] ↔ *cannot change his spots*

Further illustration is unnecessary.

TYPE *the cat*[↑] ↔ *predF* > { $\backslash \underline{o|^{cat} \quad \dots} / \langle o|^{pred} \rangle / \textit{the cat}$ }

Called up in H's mind, there is a self-replicating open sketch, each copy consisting of the skeleton of an exclusive unification of features, containing $o|^{cat}$ and room for complementary features kept open, but $o|^{pred}$ included. This sketch shapes rather abstract information, as it repetitively links one feature-image to another in very fragmentary sketches. Examples of this type are:

(48) *the winner*[↑] ↔ *is the first to arrive*

(57) (according to the constitution) *the president*[↑] ↔ *is elected*

(58) *the joker*[↑] ↔ *is an extra playing card in certain card-games*

(59) *the tortoise*[↑] ↔ *is a reptile*

As is the case for the parallel indefinite noun phrase, the information defines or refines meaning. Further illustration is unnecessary.

TYPE *the cat*[↑] ↔ *not predF* > { $\backslash \underline{o|^{cat} \quad \dots} / \langle o|^{pred} \rangle / \textit{the cat}$ }

An example is

(60) *the whale*[↑] ↔ *is not a fish*

Further illustration is unnecessary.

As extensively discussed for the indefinite noun phrase, also for all illustrated types of the definite noun phrase, further variations can be made by adding a pitch peak to any part of the sentence. And also here, varying distributions of pitch accent bind information in ongoing text together in different ways. I shall not illustrate all this; it should come out correctly.

6.2 Russell 1905

In his famous 1905 paper, Russell listed some ‘logical puzzles’ that arise when one supposes that definite descriptions stand for external objects. Russell put the blame on natural language; his solution was to improve language, to design a logical form in which those problems do not arise. My approach, obviously, is different; I want to understand how language, as it is, works as it does, and we cannot come to understand something by changing it into something else. I believe when we carefully isolate precisely the piece of information contributed by a definite description on its own, it does not include a claim of existence of something there fitting it, so, when we get hold of what happens between speaker and hearer, none of Russell’s puzzles comes up, and there is no need to improve the forms of language.

Still, I think Russell’s analysis is very perceptive. He sees two readings of

(61) *the king of France is not bald*

one that claims existence of a referent for *the king of France*, and one that does not claim existence for such a person. I believe he is right. Ignore for a moment the fact that Russell’s analysis is a logical form, and my notation represents a mental shape, and compare them: Russell’s ‘primary occurrence’ of *the king of France*, and my representation for (61/1) *the king of France*[↓]* *is not bald*

$\exists x (\textit{king of France} (x) \wedge \sim \textit{bald} (x) \text{ (plus uniqueness criterion))}$

$\underbrace{\langle \textit{king of France} \text{ } \text{ } \text{ } \rangle}_{\text{ } \text{ } \text{ } } / \langle \textit{bald} \rangle / \textit{the king of France}$

Both represent a whole person, both composed from three parts: a claim of existence of something, its *king of France*-aspect, and its definiteness: Russell’s ‘*king of France* (x)’ is mirrored in my $\langle \textit{king of France} \text{ } \text{ } \text{ } \rangle$, Russell’s existential quantifier ‘ $\exists x$ ’ is more or less mirrored in the set of unspecified complementary features ‘o o’ in the noun phrase skeleton. Then, in Russell’s analysis, the definiteness of the noun phrase is captured as the uniqueness criterion ‘ $\forall y (\textit{king of France} (y) \rightarrow x = y)$ ’, which is the logical form of the idea that there is only one person who fits *king of France*; in my representation, the definiteness gets shape in the fact that the label *the king of France* as a whole, when used, exclusively fits one person. In both analyses, the (negative) predicate information is added as an extra aspect: ‘ $\sim \textit{bald} (x)$ ’ and ‘ $\langle \textit{bald} \rangle$ ’.

A crucial difference is that in Russell’s analysis, the whole person represented is given in advance: ‘ $\exists x$ ’, while in my representation, he has to be put together from different fragments. It is in Russell’s ‘secondary occurrence’ that this difference becomes obvious: to undo the now unwanted claim of existence that comes with the existential quantifier, Russell needs *not*, turned

into the logical operator ‘~’, and relocated to the front; compare my representation for (61/2) *the king of France*[↑]* *is not bald* that does not include a claim of existence from the start:

~ ∃ x (*king of France* (x) ∧ *bald* (x) (plus uniqueness criterion))

\o|king of France ... / \o|bald \ / the king of France

There is a very good reason for this difference: an analysis meant to mirror the logic of the reality described, to capture truth and falsehood of the description, has no place for fragmentary, unfinished things. A notation meant to represent how information takes shape in people’s minds has just as much place for fragments as for whole things. Of course we cannot really ignore the difference between a logical analysis and a mental shape, which makes my analysis, in the end, incomparable with Russell’s.


6.3 Proper Names

The question to ask is: what exactly does S’s use of a proper name say to H. A proper name is a word, a symbol, and it can form a noun phrase by itself, so what does it contribute as a symbol (as compared to *cat*), and what does it contribute as a noun phrase (as compared to *a cat, the cat*). As a word, when linked to someone, a proper name is applied as an exclusive tag. When H knows the person concerned, S’s use of his or her proper name calls up in H’s mind an image of that person, and plugs in on H’s records of his experiences with this person, ready to be reactivated. However, proper names, other than more descriptive words, often are used without H’s prior acquaintance with the one carrying that name; take *Sophie*, if H does not know Sophie, it just calls up the feature-to-carry-that-name, o|^{Sophie}, which besides being probably female, reveals nothing of the real Sophie’s characteristics.

Then, as a noun phrase form, a proper name calls up a skeleton of a unification of features, including the feature-to-carry-that-name; take *Sophie* again, it calls up

\o|Sophie / Sophie

If H knows her already, in his mind, this skeleton may flash into the whole Sophie, but proper name noun phrases are not different from the other noun phrases discussed: the use of a proper name noun phrase by itself does not include a claim of existence of someone (the one and only) carrying that name. What the proper name does contribute is the uniqueness of the unification under construction throughout all situations: when reality is not in view, whatever thing the skeleton grows into, it must be the same whatever the situation, and with reality in view, it is to be one and the same thing in all situations that is the one and only thing that comes into consideration:



may grow into a unique unification of features.

As illustrated in the chapters above, when S uses a noun phrase, he composes his information by adding his choice between ↓ and ↑, and between * and ↔. When the noun phrase is a proper name,

that does not make a difference. But uniqueness of the unification under construction neutralizes much of the differences between the four main types of information, even more so than the exclusivity of the definite case. Still, the nuances are there, and they are unmistakable. Below, I shall give some examples, without further illustration. Note the subtle differences between (62/1) and (62/2), and (63/1) and (63/2):

- (62/1) (Question: *what about his siblings?* Answer:) *Sophie*[↓]* *is interested*
 (62/2) (Question: *who would buy such things?* Answer:) *Sophie*[↑]* *is interested*
 (63/1) (Question: *what about his siblings?* Answer:) *John*[↓]* *is not interested*
 (63/2) (Question: *who would buy such things?* Answer:) *John*[↑]* *is not interested*

and to distinguish the twins Jack and Tom, (64) obviously is not about Jack:

- (64) (Question: *how can I tell them apart?* Answer:) *Jack*[↑]* *is the one with the scar*

A good example of the fact that a proper name noun phrase may call up just a fragmentary image, a unique unification under construction containing nothing more than an almost empty feature-to-carry-that-name, is the following: having read in one publication about the Morningstar which is the Eveningstar, and in another publication about Hesperus which is Phosphorus, all of them apparently Venus, you may still not know which is Hesperus and which is Phosphorus. If so, for your information:

- (65) *Hesperus*[↑]* *is the Morningstar* > \o|*Hesperus* ... / o|*the Morningstar* / *Hesperus*

It is a fact that proper names can be used without claiming existence of something fitting. This is not a flaw of illogical natural language, but just part of its elegant and efficient system. This never causes a problem in actual communication, people just use the form, and understand the information. Take:

- (66) *Santa Claus*[↑]* *exists* (correct or incorrect)

Although the whole (66) claims existence for Santa Claus, this claim obviously is not contained in the proper name noun phrase. And of course there are the negative examples:

- (67) *Santa Claus*[↑]* *does not bring the presents*
 (68) *Old Bogey*[↑]* *does not exist*

A proper name noun phrase is unique, and still it can be used to make general statements, calling up repetitive pictures:

- (69) (*whatever you say*) *Jennifer*^{↓↔} *is in the opposition*
 (70) *Martin*^{↓↔} *speaks the truth (never even a white lie)*
 (71) *Martin*^{↓↔} *does not cheat*


They call up a self-replicating stand-in, each time for one and the same person, each time attributing (or excluding) the predicate-feature.

Finally, listen to

(72) $Jekyll^{\uparrow \leftrightarrow}$ is a respectable doctor and $Hyde^{\uparrow \leftrightarrow}$ is an evil monster
 (73) $Suzy^{\uparrow \leftrightarrow}$ is her little daughter and $Suzan^{\uparrow \leftrightarrow}$ is her big girl

6.4 Material Noun Phrases

Noun phrases like *water* are not that different from other, especially indefinite, noun phrases. *Water* is a word that can form a noun phrase by itself; the word *water* tags some material, the noun phrase *water* fits arbitrary quantities of the material tagged *water*, possibly overlapping. The main difference with noun phrases like *a cat* is that *a cat* fits arbitrary separate things tagged *cat*.



may grow into any, possibly overlapping, unification of features, any quantity of the matter.

On inspection of the world, the skeleton may fuse with any smaller or larger quantity of water; anything from half a drop to an ocean may come into consideration.

Again, when he uses noun phrases like that, S composes his information by adding his choice between \downarrow and \uparrow , and between $*$ and \leftrightarrow . Examples of the four basic types are easy to find:

(74) $water^{\downarrow *}$ dripped down

(75) $snow^{\downarrow *}$ covered the fields

(76/1) $rain^{\downarrow *}$ did not stop him (wet to the bone, he kept on trying)

(74), (75), and (76/1) say something about a certain quantity of water, snow and rain. (77) and (76/2) are not about a certain quantity of soap and rain:

(77) $soap^{\uparrow *}$ is in stock

(76/2) $rain^{\uparrow *}$ did not stop him (it was dry all afternoon, something else must have happened)

The feature-images $o|^{soap}$, and $o|^{rain}$, respectively, are part of a fragmentary sketch, where the noun phrase skeleton leaves the existence of something fitting undecided.

Next:

- (78) *arsenic*[↓] ↔ *is poisonous*
 (79) *water*[↓] ↔ *freezes at zero degrees*
 (80) *Adam's ale*[↓] ↔ *is the best brew*
 (81) *money*[↓] ↔ *begets money*
- (82) *good wine*[↓] ↔ *needs no bush*
 (83) *money*[↓] ↔ *does not smell*

Having the world in view, they are about any quantity of arsenic, water, money, etc., that comes into focus. As these quantities may overlap, in H's mind they may tend to flow together into one big mass, so that the noun phrase used this way may seem to refer to the whole of the matter concerned.

Finally, some examples of the abstract information of repetitive fragmentary sketches:

- (84) *spelter*[↑] ↔ *is a zinc-based alloy*
 (85) *Adam's ale*[↑] ↔ *is water*
 (86) *water*[↑] ↔ *is H₂O.*

Chapter 7

Some Philosophical Issues

7.1 Awareness and self

A person's awareness completely depends on its organism's brain: it is there precisely during a special kind of very complex coordinated activity in this brain. It is hard to understand how a neural mechanism can create awareness, but it seems to be a fact that it does; in the future, scientists may be able to find out in detail what it takes for a brain to make awareness emerge.

Awareness is not an epiphenomenon of the brain in action. It is only by creating awareness that the brain manages to separate other things from a self simultaneously emerging in the brain's organism. This may have been the evolutionary advantage of this complex and relatively time-consuming brain process to develop: to give its organism a perspective on other things in their own right.

What we are used to call our 'self' emerges precisely during the times that our brain makes us aware. During and between those moments, this 'self' is held together by the vast criss-cross of more or less easily accessible records of earlier experience, actions, and thoughts which determine the organism's 'self', that is, what this organism more or less coherently is during its existence in awareness, familiar with its environment, taking responsibility for both its conscious and its automatic actions, and picking up the thread of earlier thinking.

7.2 Subjectivity

Forms of language have an air of objectivity, and mental images are clearly subjective, but both are the private products of the individual brain. Such products can be taken as exemplary for what shows up to someone else if they can be synchronized by an external trigger, which at first sight is obviously the case for forms of language, and not the case for images. However, an image can be language-triggered, and if so, also the image has an external trigger to synchronize what shows up.

These image-triggering forms can be divided into words with a symbolic function and constructive forms. Suppose a speaker S's sequence of sounds synchronizes the symbolic sound form *cat*, which in its turn calls up an image. So both the word and the image are synchronized, but still there is a difference, still it seems to be logical to generalize the word, and far more questionable to take the image as exemplary for what shows up to someone else.

This has a reason: the sound form is just a tag without intrinsic value, a public entrance that plugs into the private mind, the sole function of which is to access what does have intrinsic value. So the less this form changes over time, and the more it remains constant within the speech community, the better it functions. In contrast, the word-triggered image is fleeting, morphing, expanding and shrinking as it gives shape to each individual person's state-of-the-art grip on a

reality which, although it is what it is, is far too complex and dynamic to be cut up in neat and lasting categories. So the more flexible the image is, the better it functions.

Insofar as we can speak of degree of subjectivity, word forms are less subjective, and their word-triggered images are more subjective. But it is only the images that get hold of the things that are there, and different though they may be from person to person, it is the same real things that they get hold of.

As for constructive forms of language, the story is different. Their function is to build up complex images from the fragments called up by the symbolic words. So both the information of symbolic words and of constructive forms of language take shape in images, but if we speak of degree of subjectivity, this construction of complex images is less subjective than the images triggered by words: the choices that can be made are limited, and determined by the speaker. For instance, the image that a speaker's *cat* calls up to one person may be very different from the image it calls up to someone else, the sound connected to *mew* in the mind's ear may be different, but their combination given *the cat mews* will be the same. Also for prosodic information the choices are limited, and they come across. Without a second thought, people generalize the build-up of their language-triggered images.

7.3 Verbal Information and the Notion of Truth

So far, the notions of truth and falsehood have hardly come up for discussion. H does not need them to be able to understand verbal information, and to recognize the real thing. It is the other way around: first H needs to be able to understand verbal information, that is, to give it shape in his mind, and to recognize what fits, that is, to fuse the mental shape and the real thing in view, and it is only then that H can begin to develop a notion of truth. To get this notion, H needs to be able to keep the mental shape and the real thing in view separated long enough to compare them, to become aware of the fact that they indeed can fuse together. Perhaps it is easier to get the notion of false information, because its shape anyhow remains separated from the real thing in view. And once H has learned to compare image called up and real thing, it may be easier to get this notion of true information as well.

So the notion of truth does not play a role in H's process of understanding information. Usually, he will just expect the real situation to fit the information, and, if he has more than one source of information, just try to blend partly overlapping pieces of shape together. Still, it is important for H to be aware of the fact that the real thing may not be as expected, that the information may be false. However, to be able to check the information, H, on his own authority, needs to have complete knowledge of the real situation, either from direct perception or otherwise. This requires the situation at least to be realized.

If S's sequence of sounds calls up a shape of presence, H will expect it to fuse with the real thing. If the real thing is completely in view and the shape does indeed fuse with it, H will consider the information true. And if there is no way to fuse the shape called up with the real situation, H will consider the information false. In contrast, if S's sequence of sounds calls up a shape of absence, H will expect that it will not clash with anything in view. If indeed it does not clash, he will consider the information true, but if it does clash, he will consider the information false.

Take Russell's famous (1905) description of the non-existing King of France ('primary occurrence'), in this book (61/1) *the king of France*[↓]* *is not bald*. In a theory of verbal information the question whether the absence of a referent for the *king of France* makes it false or lack a truth value is of secondary importance; the important question is how the information takes shape to H and how he recognizes what fits; now in this case, obviously, already on first inspection of the world, looking for what comes into consideration, H can establish that there is no way to fuse the shape called up with what is there (therefore, the information is blatantly false). The same goes for example (49) with its positive predicate: *the murderer*[↓]* *is insane* when H is absolutely sure the death was an accident (see section 6.1) and for the indefinite (34) (*John could not make it because*) *a sister of his*[↓]* *is very ill* while H knows very well that John is an only child (see section 3.3).

If the information relates to some kind of future, the notion of truth, as far as I am concerned, is irrelevant. Only when this future becomes reality and therefore is no future anymore can it be established, in hindsight, whether it has come out true or false. This, however, does not make it true or false at the time the future was still open.

Formal semantics tries to capture the connection between language and world in terms of truth conditions. Now, in daily life, the notion of truth is important to H; the notion of truth conditions is not. The facts of language phrased in terms of truth conditions are real enough, but we cannot use that notion in a theory of verbal communication because it does not play a role in the head of the hearer. Instead, the theory should explain which situations H would recognize as fitting S's information. This is, for instance, how the 'truth conditional' difference between the examples (1/1) and (1/2) can be accounted for; see 4.2 above. Then, if H's notions of truth and falsehood were added to the theory, it should correctly predict the range of situations in which H would call the information true, and the range of situations in which H would call the information false. Beyond that, there is no need for a theory of verbal communication to cover the notion of truth.

7.4 Meaning, Representation, Reference

If I am correct, in a theory about the connection between language and world the notion of representation is at least as important as the notions of meaning and reference supposedly are, but we should be careful in using each of them. It may seem that words, phrases and sentences have meaning, that they represent something in the world and that they refer to real things. However, by themselves forms of language do not contain meaning, do not act on behalf of real things and do not refer to something that is there. It is only to people that forms of language have meaning, it is only when used that they can replace what is out of view, or refer to what is in view.

For instance, a representation is something used by someone to stand-in for something else, therefore it only makes sense to call something a representation if we can specify to whom this representation takes the place of the real thing. Not only in semantics the word representation has been sloppily used; for instance, it is incorrect to say that the neural network that is active during a person's perception of something is a representation of it. In terms of the notation used in this book: the neural network ${}^H(N \text{---} \text{---} N)_C$ in a person H's brain is not a representation of C, nor is this network in action, ${}^H A(N \text{---} \text{---} N)_C$: to whom would they stand-in for the real thing? As said, a representation requires a subject; now, during ${}^H A(N \text{---} \text{---} N)_C$ triggered by signals from C,

H becomes aware of the real thing C in the shape of ${}^H\langle \bullet \bullet \bullet \rangle_c$, it exists to him. Still, ${}^H\langle \bullet \bullet \bullet \rangle_c$ is not a representation of C, that is: to H, C's phenomenal appearance is not a representation of C as it is; it does not stand-in for the real objective thing to him but as far as he knows it IS the real thing in sight.

However, the notion of representation does make sense when the real thing is out of H's sight but still in his mind: the shape this takes in H's awareness stands-in to him for the real thing during its absence: ${}^H\langle \bullet \bullet \bullet \rangle_c$ in more or less detail showing up as-if-there to him when ${}^H A(N - -N)_c$, or part of it, is triggered from within. Note that H, just like he does not need a homunculus to see the real thing does not need one to see the image either, as the image results from the same activity in (part of) the same neural networks that made the percept.

Let's go back to language; as argued, a speaker S's utterance calls up an image to H. Note that this image to H does not represent the meaning of S's utterance but it IS its meaning. In my opinion, the whole idea that H, understanding S's utterance, would make a mental representation of its meaning, in whatever form or picture, is misguided and creates more problems than it solves, as it does not explain what meaning itself is, let alone how it is supposed to represent that meaning or what its role would be in H's process of getting the meaning. If I am correct, H's understanding consists in the information taking shape in his awareness, shaping fragments of whatever it is S talks about as it would look like, feel like, sound like, taste like, etc., if it were in view. And insofar as this image is a representation to H, it is a representation of the real thing: as argued, this fragmentary unification of features under construction may or may not be or grow into a complete stand-in for something supposedly there in the world but out of sight.

What S's utterance means to H is what shows up in his awareness when he understands that utterance; in a narrow sense, this meaning consists of the feature-images called up by the words as far as authorized, assembled according to the prosodic and structural information in a more or less fragmentary unification of features. This thin image is the meaning that H takes as exemplary for what would show up to someone else. Although it is necessarily private, H will tend to take it as THE meaning because other people will recognize the same things as fitting.

This narrow-sense meaning accesses a branching complex of networks of neurons recording part of H's personal history. To H, the meaning of S's utterance in a broad sense is formed by the whole range of unauthorized possible appearances that this thin image may flash into, including personal memories and idiosyncratic associations, which all may come to H's awareness.

Then, if the situation talked about happens to be in H's view, the image called up in H's mind melts together with what he recognizes as fitting at that moment, narrowing down the range of possible appearances to the one actually meant by S at that moment. We might call the thing recognized, if you like, the reference of S's utterance at that point.

The connection between language and world builds upon the basic words, but by themselves, words do not contain meaning. It is not a word's meaning that determines whether it applies to something, but it is someone's repeated experience of applying the word to things that shapes and re-shapes its meaning to the person concerned; this meaning is state-of-the-art, growing and adapting, and shared between people only insofar as they already have made sure to distribute the word in the same way over what is there. And then it is what a word means to a person at a given time that enables him to recognize what fits at that moment. Therefore, as linguists, in order to get hold of the relationship between language and world, we should not try to define an abstract

notion of meaning determining reference in the world, but we should try to get hold of that relationship when it happens. And when we succeed we can understand what it is that we are used to call meaning.

Using these notions of meaning, representation and reference, we should clearly distinguish between the level that needs to be explained, that is the level of speaker and hearer, and the level of the linguist trying to do the explanation. In the theory of verbal information presented here, the notation aims to get hold of what actually happens, therefore it includes, for instance, a notation to represent the image called up to H by S's utterance, however, this representation is not supposed to be a representation to H, in H's head, of the meaning of S's utterance, but it is meant to be a theoretical representation to us linguists of S's information as it actually takes shape in H's head.

7.5 Empirical Semantics

Question: is the relationship between language and world suitable material for empirical investigation? In my definition, empirical research is the investigation of what is the case; it tries to get hold of the facts, it asks questions that have an answer in reality, and it searches for the answers in reality. Everything that is the case in principle is suitable material. So the question is: is the relationship between language and world something that is the case? Is it a fact? Does it exist?

If we think of it as a direct connection between forms of language and reality, the answer has to be negative. However, it is a fact that hearers understand a speaker's forms of language, and are able to recognize the real thing talked about as fitting the information. It is a process, it works; people make the connection; the question 'how does it work' can be guaranteed to have an answer in reality. There can be no doubt that this is a legitimate empirical question.

Still it remains to be seen whether we can get hold of this reality. Efforts to formalize some abstract language-world relationship before addressing the questions of people's understanding put the cart before the horse; to investigate something that does not exist certainly does not help to get answers from reality. We need to be clear about what kind of reality we are dealing with, and we need to find a way that is up to scientific standards to investigate that reality.

Empirical investigation is a process of trial and error, an alternation of and interaction between, efforts made in order to better get hold of the facts and efforts made in order to better explain them. What is essential is a commitment to find out what is the case, with all means available separating solid facts from obscure phenomena, and explanation from speculation.

Natural science has become the textbook case of good empirical research, an example of how to get results. Here, experiment and observation play a central role in drawing the dividing line. This may lead to the wrong conclusion that what is not observable is obscure by definition; applied to semantics, especially private images would end up on the wrong side of the line. But I believe empirical semantics needs to develop its own criteria.

A brief comparison. Scientific observation is sophisticated and controlled perception. Everyday perception happens when the senses of some person, let's say H, pick up signals from something there, let's say C, and H's neural system figures out the source of the signals; what shows up in H's awareness is his brain's best guess of what is there. Scientific observation may enhance the range of the senses by instruments, and theory may improve the interpretation of

what is picked up. Suppose H sees a unification of features $H \setminus \bullet \bullet \bullet / c$. In daily life, he takes this to be the real thing, in science, he tries to look through this appearance to the real C he assumes to be behind it. If such a C is indeed there, some other person M can pick up its signals as well; what M sees is $M \setminus \bullet \bullet \bullet / c$. In daily life, M takes this to be the real thing, in science, he tries to look through it. Anyhow, H and M agree that there is something there. If however M in the same circumstances as H does not see $M \setminus \bullet \bullet \bullet / c$, the conclusion is that H's $H \setminus \bullet \bullet \bullet / c$ was a figment of his imagination, that what he thought to be there does not exist. Drawing the dividing line of science, this would be: if M can see what H sees, it is there, so it is suitable material for empirical investigation; if M cannot see what H sees, it does not exist, and it has no place in science.

Now take images. Obviously, M cannot see H's images. Does this mean that they are illusions, nothing more than ghosts we should ignore because of their non-existence? Wrong conclusion. Switch back to perception. When H observes C, what he sees is the unification of features $H \setminus \bullet \bullet \bullet / c$; when we say that M can see what H sees, we do not mean that M sees $H \setminus \bullet \bullet \bullet / c$; for this, M would need H's brain to make sense of the signals picked up by H's senses. What we mean is that there is some C there that triggers its appearance to both H and M. This does not mean that C exists and $H \setminus \bullet \bullet \bullet / c$ is non-existent: in those brief moments it shows up to H, it is there in H's mind. To deny the existence of the shape that the observation takes to H, would be to deny that the observation takes place. The same goes for images. When H imagines something that M does not see, when we call this an illusion, we mean that there is no C there to trigger this appearance in H's mind. This does not mean that this $H \setminus \bullet \bullet \bullet / c$, in the brief moments it shows up to H, is non-existent; of course it is there, the illusion is there. The fact, which of course is a fact, that M cannot see it in H's mind would be a very wrong reason to deny its existence, or to deem it unsuitable for empirical investigation. There is nothing obscure about ghosts as long as they stay where they belong, which is inside people's minds, during those brief moments they show up in their awareness, and as long as they are not reported to have shown up in reality in their own right, whining and blowing out candles.

Interestingly, methods to detect in H's brain the occurrences of his phenomenal experiences are rapidly improving. Of course such observation does not take anything like a shape that echoes H's images, so it is not of immediate help to us now.

Natural science by definition investigates the external reality; the data are to be gathered by experiment and observation, by scientists who agree on what they see, and try to look behind its appearance. Our own material is a person's process of turning the sound forms picked up into understanding of what is meant, and recognizing the real thing if and as it comes in view. This is material of a different kind: not only mental images, also the phenomenal world is inescapably private, and it is misunderstanding the nature of language to suppose that its forms are observable. On the other hand, it is a mistake to suppose that, because all this is unobservable, it is not empirical material. It is not the existence of the material that is the problem. Verbal communication happens. It works. And because it happens, the question 'how does it actually work' has an answer in reality.

The problem is how to get hold of this kind of material. Forget about observation. We will have to find our own methods to ascertain the facts, and our own criteria to exclude speculation. If indeed this material only exists in those brief moments one is aware of it, the only access is introspection. What remains essential is a commitment to find out what actually is the case, and an open mind to the fact that our conjectures may be wrong. Agreement will be part of the criteria. Natural scientists have to agree on what is there, thinking away its appearance; in empirical semantics we need to agree about the appearances: we need to establish to what extend

we can take what shows up to ourselves as exemplary for what in the same circumstances would show up to someone else. As we cannot directly compare what shows up, the test is whether it works in communication, and ultimately, whether we would recognize the same things there as fitting what shows up to each of us privately.

As compared to physical nature, this introspective material has a big advantage. There is no guarantee whatsoever that the scientist's brain is up to the mysteries of nature and the universe, largely hidden from human perception; theories becoming ever more abstract ask for efforts becoming ever more ingenuous to get answers from reality. In contrast, introspective accessibility of verbal information comes with the material: what we want to understand is how we understand whatever it is that we do understand, we want to investigate what can be guaranteed to come to our awareness, and we do not need sophisticated experiments to know what just came to our mind. And to a large extent, agreement comes with the material as well, as we investigate how we understand each other. When we discuss words, we take it for granted that we share them. And when we discuss meaning, most of the time we take agreement for granted as well; when linguists paraphrase sentences, when philosophers of language wonder about how the forms of natural language only roughly seem to cover their meaning, how phrases sometimes are and sometimes are not stand-ins for something real, when logicians point out that what seems to be the same negative sentence goes with different truth conditions, all the time this discussion is about meaning already instantly and effortlessly understood from the original form actually given, and without a second thought agreed upon. In fact, there is a large amount of seemingly diverse and problematic but nevertheless firm data, as illustrated in chapter 1 and 6.

Summarizing the argument thus far: when we investigate the process of understanding verbal information, existence of the material is not the problem, its accessibility is not the problem, and agreement on the data, the forms heard, and the instant end-products of understanding, is not the problem either. The problem is how to explain the facts.

This is my approach: given some verbal information, the forms of language and the instant end-products of its understanding show up in my awareness. These I can try to get hold of. As I only have access to my own mind, I split up myself into the one who does the investigation, which is me, and on the other hand the people investigated, the speaker S, the human and hearer H, these other people M and P. I take what shows up in my mind as exemplary for what would show up to actual speakers and hearers, put differently, I assume agreement both on the complex sound form and on the instant end-product of its understanding. Given this start and this finish, the idea of what I do next is simple enough: I try to break down the complex form into its basic forms, and figure out what each of these forms contributes to the instant end-product. Where understanding is immediate, its breaking up into pieces is not so easy; it is a process of trial and error, all the time checking intuition, resulting in hypotheses about what actually goes on.

It is essential that the theory starts where H starts: the forms of language recognized as informative in S's sequence of sounds. As prosodic form makes a difference, we cannot ignore it, so I tried to pin down what it is in pitch and rhythm that makes the differences.

From there, the theory has to follow H's process as closely as possible; he turns this complex of forms into an understanding that can stand-in for whatever it is that S talks about as long as it is out of view, and that makes H recognize the real thing if it is in view. As an investigator taking H's position, I see a mental picture that has enough resemblance with the real thing to fuse with it if it is actually there. For H, this works very well. But for an empirical theory, there is a problem. This problem is not the subjective nature of imaging, as the whole

process investigated takes place inside the private mind. But how to deal with the undeniable fact that images are fleeting? How to get hold of something that seems to be elusive by its very nature? This is the real problem of mental images. But we need to get hold of them if we are ever going to understand how people understand verbal information, because they are, as far as I can see, the only bridge there is between language and world. I believe in this book I have come up with a characterization that may work. And so, for each of the sound forms hypothetically recognized as informative in S's sequence of sounds, I have tried to isolate what exactly it contributes to the image under construction.

So in this book, I presented an empirical theory about the process of understanding verbal information. I started with hypotheses about especially prosodic form: what forms exactly does a hearer pick out as informative from a speaker's sequence of sounds, and I phrased some working hypotheses about what they each contribute to the understanding. Then I phrased some basic assumptions about the way of existence of the material investigated, and against this background, the working hypotheses could be made more precise.

Given the subjectivity of all the material, the question again arises how to draw the line between explanation and speculation. Here are some thoughts:

As already said, from the start I have assumed that within a speech community we agree on the basic facts of verbal communication: when hearing a real life speaker uttering a sequence of sounds, you and I shall take both the sound forms we hear and the understanding we arrive at as exemplary for what shows up to the other. In case of doubt, this agreement might be checked, but most of the time, this will not have any priority; most of the material will be clear enough. It is when we try to represent this start and finish of the process of understanding that things get hypothetical: do we indeed get hold of what we do hear and what we do understand? In a number of ways, these hypotheses can be tested:

As for the hypotheses of prosodic form, I believe this is what people actually pick up, but the hypotheses may be wrong. They can be tested independently from the rest of the theory; the larger part of the argument has already been given in chapter 1. The limited set of precisely defined basic prosodic forms enable a multiplication of possible complex prosodic forms; for any combination, the end-product of understanding can be predicted and tested. Possible meanings can be provisionally paraphrased closely enough for a reconstruction of the intended meaning. See for example the contours (a) thru (n) in chapter 1. In Koene (1994), I discussed in more detail prosodic forms and their information as reported by other linguists; it includes for instance a discussion on the difference between 'fall-rise' and 'fall-plus-rise', see also the difference between (1/3) and (1/6) as accounted for in section 1.3.2 above.

As for the basic pieces of information, the working hypotheses have some initial credibility because they suggest a simple system behind a large amount of seemingly diverse and problematic data. (See Koene (1994) for a discussion on other views on these data). To me, these working hypotheses feel intuitively right, but to really count as an explanation of the facts, they need more precision. This is what I tried to do next: to find a precisely defined notation that each step of the way exactly represents what goes on. The test is whether it works, that is: do the pieces of shape hypothetically called up by the basic forms hypothetically picked out from the

sequence of sounds, put together, indeed account for both the logic and the nuances of H's instant end-products of understanding? And do they correctly predict what real situations H would recognize as fitting the information if they came in view? Ideally, the theory should completely work by itself. I made a start, but more often than I would have liked I had to compromise on formalization for lack of time. Still, the theory seems in principle able to cover a whole range of seemingly diverse and problematic phenomena.

Then the whole theory is embedded in basic assumptions about brain and mind; everything conjectured about the contents of the mind simultaneously is a conjecture about the organisation of and the activity in the brain. So in principle, on an indisputably scientific level, the theory is open to refutation. I am not a neuroscientist, and everything I think I know about the brain I learned from those who are; I distilled my ideas from literature to suit my own purposes. The ideas may very well need serious revision. Or I may be completely wrong in my basic assumptions. But in this discussion on method, the point is: these are testable hypotheses.

My commitment in this whole investigation has been to find out what actually happens; this is, if you like, a commitment to the truth. In the investigation of language, it is here that truth matters: we have to aim for a theory that shows what is the case.

Tribute

This book has a long history. It took years for the questions addressed in it to crystallize, and years again to search for answers; in the process, the questions sharpened, and called up new questions. Looking back, in the course of the years a number of publications in different areas of research have had serious impact; I would like to pay tribute here to their authors.

The first on the list definitely has to be Karl Popper; having come from an environment where thoughts were not free, the reading of *The Logic of Scientific Discovery* (1959), *Conjectures and Refutations* (1963) and *Objective Knowledge* (1972) changed my world. What especially struck me was this idea that all knowledge of what is there, including knowledge from observation, is theory-impregnated, and that empirical science tries to come nearer and nearer to the truth by critically going back and forth between conjectures and data. When I read other literature on philosophy of science, this reinforced the idea that nothing should be taken for granted, especially not the obvious.

However, I did not have the faintest idea how to apply all this to language, which was my subject. So when I read Chomsky's *Syntactic Structures* (1957), it came as an eye-opener that language has a system, that we can conjecture what this system is like, and check the conjecture against the data of language. Still, for myself I had no sense of direction, and soon enough I was rather discouraged by the fact that Deep Structure fails to represent meaning.

So, despite passing all the exams, and graduating in Dutch language and literature, I had no idea how to be a linguist. Highly frustrated with my own inability even to find a question to work on, I continued to study philosophy to learn more about philosophy of science, but when I started my studies at Amsterdam University, I plunged right into the middle of discussions on philosophy of language and truth-conditional semantics, discussions on Frege (1892), the Russell (1905)-Strawson (1950) debate, modal logic, Montague seminars, etc. I thought I knew one or two things about language, but this completely different way of looking at it quite confused me, and it took some time for things to reshuffle in my mind. With hindsight, I can distinguish three major influences:

The first thing I picked up was that semantics should explain the relationship between language and world, where the world could be represented in a model. This idea was new to me. Before I got any further knowledge of formal model theoretic semantics, I self-evidently took this to be an empirical question about what actually is the case, and it seemed only reasonable to me to look for the answer where language and world in fact meet, that is, at the point where a hearer makes sense of a speaker's words, and recognizes the real thing. The fact is that as a rule a hearer just instantly understands a speaker's verbal information, and, self-evidently, the question for me was: how does it work. This question has kept me fascinated ever since. Intuitively, I rephrased this question into: how does a hearer turn the speaker's forms of language into a mental model of what the speaker talks about. Soon enough, I found out that this was hardly the aim, and hardly the kind of model of formal semantics, but still I think that without model theoretic semantics the empirical question would never have occurred to me. Anyhow, for a very brief period of time, I had the naïve idea that for me as a linguist it would not be that difficult to explain what as a hearer I just instantly understood.

But if I had hoped that going from phrases of language to the hearer's mental picture of the world to the world itself would be easy, philosophy of language and formal semantics soon enough opened my eyes: not only do the noun phrases of language seem to have very different

logical properties, depending on whether they are proper names, indefinite descriptions, definite descriptions, mass terms, etc., but also often enough the same sentence seems to have different readings; there is ‘scope ambiguity’ of negative sentences; also, there are the ‘generic’ readings next to the ‘non generic’ ones. And it was phrasing the connections between these forms of language and the world in terms of truth conditions that made it absolutely clear that there is no way of reasoning these facts away. So if I wanted to answer my empirical question, I had to deal with them. The fact is that a speaker can use what seems to be the same form to send quite different messages, and still the hearer just instantly gets the meaning intended. It works. The question remained: how does it work. And the answer still had to be found in what actually happens.

Then the third influence has been a negative one: fruitful as logical analyses are in pinpointing the problems, they cannot possibly be part of a solution. To the ear of the hearer, a logical analysis, or any formal or informal re-phrasing, is a never completely successful effort to capture the content of the already instantly and effortlessly understood original forms of language. They are Procrustes’ beds, and useless in an explanation of the original’s understanding. Worse, discussing the problems in terms of quantifiers, logical operators, scope, truth, truth conditions, etc., blocks one’s view of the question that matters: how do forms as they are work as they do.

To answer this question, it is essential to start where the hearer starts: the sound waves that enter his ears. Now as soon as I focused on what I heard, it was clear that it would be impossible to ignore intonation: it is there, and it makes all the difference. To this day, it strikes me as irrational to on the one hand ignore informative forms that are audibly there, and invent on the other, forms, logical forms, or whatever other forms, that are nowhere to be seen or heard.

Hearers very early in life learn to recognize what is relevant in all the vocal sounds that enter their ears, including prosodic forms. So I should get hold of all those forms. Now, alphabetic writing gives linguists an enormous head start: over the centuries writing systems succeeded increasingly better in representing those aspects of sound that are relevant; writing as it is now can be considered a rather successful early theory of language. But obviously, it largely misses out on prosodic form, and it has been far from easy for me as a linguist to pinpoint what it is that as a hearer I just picked out. Fortunately, at the time there were the ‘Dutch School’ investigations: Cohen & ’t Hart (1967), ’t Hart & Cohen (1973), ’t Hart & Collier (1975), Collier & ’t Hart (1981), etc., with their elegant notational system. Following a well-thought-through method, the authors had found out what exactly it is in the (Dutch) intonation contour that is perceptually relevant to the hearer (leaving aside the possible questions of intonational meaning), that is, which changes in pitch make the contour sound different to the ear, which characteristics cause a pitch movement to be heard as a pitch accent, etc. This, in my opinion, is a textbook case of good, solid, useful research. Although my hypotheses on basic prosodic forms and their information are independent of the particular notational system used to write down complete contours, I do not know whether I would have found my way in intonation without the accurate visualization of pitch this research made possible.

Focusing on what, ignoring prosody, seemed to be ambiguous forms of language, including prosody, I found an elegant system: a speaker can use a noun phrase, whether definite, indefinite, or material, whether the sentence is positive or negative, in four different ways, differentiating the form with a minimal pair in pitch and a minimal pair in rhythm, both located in the final syllable of the noun phrase. Intuitively, this simple system covered a whole range of seemingly problematic and seemingly different phenomena, including ‘scope ambiguity’ and

‘generic’ readings. I wrote my dissertation *De eigen semantische systematiek van de natuurlijke taal* (1984), *Form and Interpretation, a One to Four Relation* (1987) and *Ambiguity: Syntactic and Prosodic Form in Empirical Semantics* (1989).

Given these conjectured informative prosodic forms, the inevitable next question was exactly what they each contribute to the hearer’s mental picture under construction of the world talked about. That is, assuming that my intuition was right that understanding is drawing a mental picture of the world talked about. Unfortunately, at the time and place where I was, this was not done, either because it should not be done for methodological or philosophical reasons, or because it could not be done for practical reasons. However, as I did not see what else would bridge the gap between the real thing and the incongruous forms of language, I did not see a way around mental pictures, elusive or not. So I had my question, I had my problems, I had some hypotheses, and I had a sense of direction, but I was on my own.

I read what I could find that had any relevance to my questions. I wrote a book that did not get published: *Understanding negative sentences* (1994). It contains some rather extensive discussions and a matching bibliography, among which Ladd’s *The Structure of Intonational Meaning* (1980), and Horn’s *A Natural History of Negation* (1989).

Still, my theory, from a strong start in the forms of language, kept ending in the middle of nowhere. I needed to get hold of those mental pictures in a methodologically satisfactory way. I found Johnson-Laird’s *Mental Models* (1983) very inspiring, I studied Johnson-Laird’s *The Computer and the Mind* (1988), I studied Kosslyn’s *Image and Brain* (1994), but still I was at a loss about how to bring things together.

Then, in the special issue on imagery of Consciousness and Cognition (1996), Hubbard’s article *The Importance of a Consideration of Qualia to Imagery and Cognition* took me by surprise. He claims the importance of ‘qualia’, both in perception and imaging: qualities of reality as it shows up in our awareness, the colour, feel, texture, sound, smell, taste of the world. He argues that when we know the world by experience, there is something really important that escapes phrasing in words. Hubbard’s claim struck me as more than right: indeed we cannot ignore how the world takes phenomenal shape, that is, we cannot ignore how we are aware of what is there (Searle is right that we cannot separate ‘qualia’ from consciousness, from our awareness of something; *The Mystery of Consciousness*, 1997, p. 9). The important point to me is not so much that we cannot communicate phenomenal knowledge—we cannot—but that it is conscious experience that gives us first knowledge: if physical reality would not appear to us in colour, texture, sound, smell, etc., we would not have a world at all for us to describe in words. And these words of natural language do not label Objective Nature, but its human version, nature as it shows up to each of us; for instance the word *red* does not label a thing’s quality to reflect light on a certain wavelength, it labels, for each of us privately, the colour that shows up in our awareness of that thing, the ‘quale red’ if you like. Now how would I be able to explain the connection between (speaker’s) words and (hearer’s) reality if I ignored what the basic words of language refer to?

This forced me to rethink the whole question of what happens in the process of understanding language and recognizing the real thing. Not only mental pictures are difficult to get hold of, the world itself is at least as problematic. And then, we cannot take the forms of language for granted either: what on earth is their way of existence? Reading Hubbard made me realize that all of it depends on what we privately are able to be aware of. If I really wanted to explain this process, I could not ignore what we are aware of. And what we privately are able to be aware of, the world we experience around us, the words we hear, the pictures in our mind, it

all depends on the machinery and activity of the individual brain. I learned a lot from Scientific American's Special Issue on the brain (September 1992); it has leading scientists in the field explain their well-established results; it gives non-specialists enough food for thought, and it is as relevant today as it was at the time of its first publication.

What stuck with me when I was reading Crick & Koch: *The Problem of Consciousness*, and subsequently Crick: *The Astonishing Hypothesis* (1994), is the fact that a person's consciousness depends on the presence of a special although as yet not identified kind of activity in this person's brain. This does not solve the hard brain-mind problem how neural firing can cause subjective experience (I agree with Searle 1997 again, for instance p.28), but still, it seems to be a pretty hard fact that it does, and this fact is the basis for the mind-brain equation I defined in chapter 2.

What stuck with me when I was reading Zeki: *The Visual Image in Mind and Brain* is the fact that what we see is an invention of the brain: it results from a complex division of labour over specialized areas of the brain, cooperating in presenting a unified picture of the world. This fact is the basis for my characterization of what is there in people's world as unifications of features. Kandel & Hawkins' paper *The Biological Basis of Learning and Individuality* is about the neuronal mechanisms of learning; it seems that short-term memory consists of the change of strength of existing synaptic connections, and long-term memory consists of the growing of new neural connections at the same site. Subsequently, I read and re-read Fuster: *Memory in the Cerebral Cortex* (1995): the cortical networks of perception and memory coincide (for instance p. 114), perception and memory are inseparable.

Damasio & Damasio's *Brain and Language* is about the brain's external objects, its language, and the mechanisms that connect the two (Damasio et al. 1996 in fact trace some connections). It is fascinating. As they say, the brain uses the same machinery for language as it does for other entities. It is their notion of 'record' that put me on the right track to get hold of mental images: the brain does not hold permanent 'pictorial' representations of objects, etc., but it holds a record of the neural activity during interaction with that object, in the form of a pattern of synaptic connections (similar to Fuster's networks of neurons). I believe this notion of 'record' replaces the generally used but questionable and confusing notion of 'representation' with a notion that works: an experience of something does not store its representation, not something that can stand-in for the real thing in its absence, but it records the combination of neurons to activate to have it (re)appear, and it is this reappearance that, in the physical absence of the real thing, can represent it, stand-in for it (as for imaging, see for instance Damasio 1994, p. 101).

In all publications, neuropathological evidence plays a large part. There is overwhelming evidence that damage in one of the specialized areas of the brain may literally make the corresponding part of the person's world vanish: it is the very mechanism evolved to call it into existence that is damaged.

At this point, at last I had my questions together from start to finish, and I could really set to work. Which took another decade.

The references given below therefore reflect a very personal journey. As a bibliography, the list is seriously incomplete in two ways. Firstly, there are the many publications I read over the past decades but did not mention here; they must have had their influence, and I may have

borrowed ideas that I cannot trace back to their origin anymore; I can only hope that I did not forget anyone. Secondly, there is the virtual and virtually infinite list of publications that I have not read but should have read had I had another lifetime.

Next to the people mentioned above, I am indebted to Theo Janssen. We are I think on common ground where we believe that it is only in the use of language that its meaning comes into existence (compare Janssen 2007). I much benefitted from his encouraging and accurate comments on an almost finished manuscript, which among others made me refine my ideas on mental images and explicitly phrase my ideas on correct and incorrect use of the notion of representation.

I am very grateful to Andrea Koster for taking the time and trouble to check my language, weed out phrases in Dutch-English and point out otherwise unintelligible prose; this has been a lot of work. Of course remaining faults are my own responsibility, especially as I kept rewriting the text.

In my personal life, Bram de Waard was there with his undemanding support. He is my favourite Maecenas. He never tired of helping me with my numerous hardware and software problems; where I despaired, for him the fun started, performing surgery on my computer, rattling its keyboard. For better for worse, he never ever put pressure on me either to finish or to quit an enterprise that for so many years took so much of my time with so little to show for it; this is what created room for me to think. I dedicate this book to him and our two daughters.

REFERENCES

- Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.
- Cohen, A. & Hart, J. 't (1967). On the Anatomy of Intonation. *Lingua* **19**, 177-92.
- Collier, R. & Hart, J. 't (1981). *Cursus Nederlandse Intonatie*. Leuven: Acco.
- Crick, F. (1994). *The Astonishing Hypothesis: The Scientific Search for the Soul*. London: Simon & Schuster Ltd.
- Crick, F. & Koch, C. (1992). The Problem of Consciousness. *Mind and Brain. Scientific American* **267**, no.3, 111-117.
- Damasio, A.R., & Damasio, H. (1992). Brain and Language. *Mind and Brain. Scientific American* **267**, no.3, 63-71.
- Damasio, H., Grabowski, T.J., Tranel, D., Hichwa, R.D. & Damasio, A.R. (1996). A Neural Basis for Lexical Retrieval. *Nature* **380**, 499-505.
- Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* **100**, 25-50. Translated as: On Sense and Reference. P.T. Geach & M. Black (Eds.). *Philosophical Writings of Gottlob Frege*. Oxford: Basil Blackwell, 1952, 56-78.
- Fuster, J.M. (1995). *Memory in the Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Primate*. Cambridge, Mass: MIT Press.
- Hart, J. 't & Cohen, A. (1973). Intonation by Rule: a Perceptual Quest. *Journal of Phonetics* **1**, 309-27.
- Hart, J. 't & Collier, R. (1975). Integrating Different Levels of Intonation Analysis. *Journal of Phonetics* **3**, 235-55.
- Horn, L.R. (1989). *A Natural History of Negation*. Chicago: University of Chicago Press.
- Hubbard, T.L. (1996). The Importance of a Consideration of Qualia to Imagery and Cognition. *Consciousness and Cognition* **5**, 327-385.
- Janssen, Th. A.J.M. (2007). A Speaker/Hearer-Based Grammar: The Case of Possessives and Compounds. Mike Hannay & Gerard Steen (Eds.). *Structural-functional studies in English grammar*, 353-87. Amsterdam: Benjamins.
- Johnson-Laird, P.N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge: Cambridge University Press.

Johnson-Laird, P.N. (1988). *The Computer and the Mind: An Introduction to Cognitive Science*. Cambridge, Mass: Harvard University Press.

Kandel, E.R. & Hawkins, R.D. (1992). The Biological Basis of Learning and Individuality. *Mind and Brain. Scientific American* **267**, no.3, 53-60.

Koene, N. (1984). *De Eigen Semantische Systematiek van de Natuurlijke Taal*. Diss. University of Amsterdam.

Koene, N. (1987). Form and Interpretation, a One to Four Relation. *Theoretical Linguistics* **14**, 1-33.

Koene, N. (1989). Ambiguity: Syntactic and Prosodic Form in Empirical Semantics. R. Bartsch, J. van Benthem, & P. van Emde Boas (Eds.). *Semantics and Contextual Expression*, 57-73. Dordrecht: Foris.

Koene, N. (1994). *Understanding Negative Sentences*. www.netkoene.com.

Kosslyn, S.M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, Mass: MIT Press.

Ladd, D.R. (1980). *The Structure of Intonational Meaning: Evidence from English*. Bloomington: Indiana University Press.

Popper, K.R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson; New York: Basic Books.

Popper, K.R. (1963). *Conjectures and Refutations: The Growth of Scientific Knowledge*. London: Routledge & Kegan; New York: Basic Books.

Popper, K.R. (1972). *Objective Knowledge: An Evolutionary Approach*. London: Oxford University Press.

Russell, B. (1905). On Denoting. *Mind* **14**, 497-93.

Strawson, P.F. (1950). On Referring. *Mind* **59**, 320-44.

Searle, J.R. (1997). *The Mystery of Consciousness*. London: Granta.

Zeki, S. (1992). The Visual Image in Mind and Brain. *Mind and Brain. Scientific American* **267**, no.3, 43-50.

INDEX

- absent claim of existence, 13, 26, 55, 57, 68, 75, 80, 94, 95, 100, 103, 104, 105
- awareness, 3, 4, 30–34, 37, 38, 39, 40, 41, 42, 44, 46, 48, 49, 50, 57, 58, 59, 72, 83, 108, 111, 113, 114, 119
- conscious and unconscious perception, 36
- hard problem, 29, 48
- narrowed down to necessary and sufficient conditions, 32, 33
- basic assumptions, 34, 38, 40, 44, 45, 46, 48, 115, 116
- brain-mind equation, 33, 34, 40, 42, 120
- context, 10, 21, 22, 23, 20–23, 28, 58, 82, 83, 84, 86, 87, 89, 91, 92
- correctness of information, 66, 68, 70, 98
- empirical investigation, 4
- empirical question, 117, 118
- empirical theory, 49
- semantics as empirical investigation, 110–16
- method of investigation, 114–15
- fiction, 13, 29, 58, 59
- future, 55, 58, 59, 65, 67, 98, 110
- intonation
- Dutch contours, 7–9, 20
- Dutch School, 6, 7, 20, 118
- 'fall-plus-rise', 24, 115
- 'fall-rise', 18, 19, 21, 24, 18–24, 90, 115
- broken into basic forms, 19, 20, 82
- final rise, 19, 20, 21, 24, 82, 83, 90
- minimal pitch pair, 2, 6–9, 118
- peaking pitch, 9, 18, 20, 21, 24, 19–24, 24, 82, 91, 93
- pitch accent, 6, 7
- presence and absence, 6, 19, 20, 84, 86, 87, 89, 90, 91
- language
- connection language and world, 29, 41–44
- inventory of what comes into consideration, 59, 60, *ff.*, 84, *ff.*, 96
- recognition, 57–71, 84–86, 86–87, 89–90, 91, 94, 97–103
- connection negative information and world, 72–81, 88–89
- kind of existence, 39–40
- magic of language, 4, 45
- neural record, 40
- shared, 40–41
- meaning, 46, 56, 102, 110–12, 114
- meaning and reference, 110, 111, 112
- mental image, 13–18, 48
- homunculus problem, 3, 111
- images and percepts, 38
- kind of existence, 37–39
- methodological problem, 14, 119
- neural record, 37
- self-triggered, 46
- unification of features, 14, 37, *passim*
- word-triggered feature-image, 15, 47, 50–51
- word-triggered image, 44–47
- authorized, 4, 46, 54, 55, 57, 73, 76, 84, 111
- complementary features, 50, 51, 52, 53, 72, *passim*
- unauthorized details, 15, 16, 17, 53, 54, 57, 111
- negation
- negative subject, 11, 18, 78–81
- not* zooming in, 18–24, 82, 90
- surprise, 92–93
- 'predicate negation' versus 'sentence negation', 1, 2, 5–9, 12, 17–18, 18, 19, 24, 73–76
- 'scope' of negation, 1, 5, 11, 118, 119
- noun phrase
- definite noun phrase, 11, 52, 60, 94–104
- compared with indefinite noun phrase, 98, 99, 100, 101
- difference from indefinite noun phrase, 95–96
- information, 94–96
- parallel indefinite noun phrase, 96–97
- Russell's analysis, 103–4
- general or not general, 25–27
- indefinite noun phrase, *passim*
- 'generic', 1, 2, 3, 118, 119
- 'specific' versus 'non-specific', 1, 3, 10
- material noun phrase, 1, 4, 11, 52, 94, 106–7, 118
- information, 106
- parallel indefinite noun phrase, 106–7
- proper name, 1, 4, 11, 52, 94, 104–6
- information, 104
- parallel indefinite noun phrase, 104–6
- referring or not referring, 5–11
- perspective Ω , 30, 31, 32, 36, 38, 39, 43, 44, 50, 51, 57, 72
- phenomenal world
- covered with words, 41–44
- shared, 44
- kind of existence, 30–34
- neural record, 34–36, 37
- shared, 36–37
- unification of features, 14, 32
- theoretical representation, 32, 33, 34, 37
- unification of features covered with words
- theoretical representation, 44
- world in sight, 50
- as it has the hearer's attention, 59, *passim*
- pitch accent, 2, 6, 7, 8, 9, 18, 19, 20, 21, 24, 82, 83, 91, 103, 118
- qualia, 14, 119

- representation
 - incorrect use of the notion, 3, 29, 35, 110–11, 111
 - mental picture as stand-in for the real thing, 16, 55, 61, 111
 - theoretical representation, 4, 112
 - of mental stand-in, 55
- rhythm
 - minimal pair, 25–27
 - information, 52
- self, 31, 32, 36, 108
- subjectivity, 14, 36, 108–9, 115
 - degree of subjectivity, 109
- truth, 76, 104, 109, 110
 - truth value, 110
- truth conditions, 110
- verbal information
 - from speaker to hearer, 44–47, 48–49, 51
 - working hypothesis, 14
 - information final rise, 83
 - information minimal pitch pair
 - working hypothesis, 15
 - theoretical representation, 52
 - information minimal rhythmic pair
 - working hypothesis, 27
 - theoretical representation, 52
 - information *not*
 - working hypothesis, 17
 - theoretical representation, 72
 - information peaking pitch
 - working hypothesis, 22
 - theoretical representation, 82
 - information pitch accent
 - presence and absense, 82–83
 - multiplication of possible combinations, 24, 26, 27–28, 91–93, 94, 103
 - negative information, 16–18, 72–81
 - noun phrase, 51–52
 - working hypothesis, 15
 - noun phrase skeleton, 52
 - old and new, 82, 83, 90
 - predicate
 - working hypothesis, 15
 - subject-predicate skeleton, 54
 - types of information, 52–53, 54–57, 63, 64–71, 73–78, 78–81, 83–91, 97–103
 - words
 - working hypothesis, 15
 - word-triggered feature-image, 47, 50–51
- vocabulary, 14, 42, 43, 44, 51, 56, 70
- writing, 49, 118
 - incomplete representation of form, 11, 27
 - speech and writing, 9–10