Consciousness at the Cutting Edge between World and Brain

The Interaction between Perception, Imagination, Thought, and Language

Renate Bartsch

Preface

This book consists of three parts, which explore consciousness from different angles: First, from the philosophical point of view, showing the central tasks of consciousness for representation, intentionality, and for designing and following rules, as they have to be performed in order for language, understanding, and especially denotational semantics to be possible; second, from the perspective of connectionist and neuro-cognitive, especially neuro-dynamical approaches, which model at least partially how, under a large-scale view, the brain might work in achieving the central tasks of consciousness; and third, from a view which combines the philosophical with the neuro-dynamical approach in order to find and determine a standpoint about some disputed questions in consciousness research.

The first part consists of chapter 1, which deals with the function consciousness has, mainly with respect to language, thinking, intentionality, and the possibility to devise rules and norms and to learn and to follow these. The central task performed by consciousness is entertaining representations such that we can evaluate these under different points of view. An important strain of argumentation in this chapter is the difference between causal semantics and denotational semantics, and what role consciousness plays there in order to make intentionality of semantics possible. A further point of discussion is the nature of rules and norms and how they function in language and in the linguistic reconstruction of language. Furthermore two standpoints will be argued for, namely that consciousness is possible without language, and that free will is possible because of the special contribution of consciousness, which consists in providing representations of situations and actions such that comparison and judgement is possible.

The second part consists of chapter 2 and 3, which explore the contributions dynamic conceptual systems and connectionist models can make to explaining the nature and function of representations. Chapter 2 elaborates the possibilities of dynamic conceptual semantics and the possibilities and the shortcomings of plain connectionist models consisting just of input, output, and one hidden map of units. The shortcomings are evident with respect to modelling contiguity relationships, classification, and the understanding of basic sentences. Chapter 3 is the central chapter of the book, which shows how an architecture of connectionist maps with circuits of activation in principle works as a model for perception, imagination, and for understanding situations and basic sentences. The representations are percepts or imaginations of situations and of basic sentence inscriptions, i.e. of sentence utterances or written instances. Central is the notion of an episodic map on which activation circuits involving conceptual maps and sensorial fields get expressed in the form of the representations which constitute consciousness. The material basis of episodic maps are the primary sensorial and pre-motoric fields, emotional fields, and also proprioceptic fields, which express feelings concerning our own body. Sensorial fields of the different modalities, or sensorial maps, receive input from the respective sensors. Pre-motoric maps form gestures for realising the motoric output. These maps have connections to higher order maps, conceptual maps, in which groups of neurones indicate when the system has classified and ordered input with respect to previous input under similarity or contiguity relationships. From the higher maps activation is sent back to the primary fields. Such circuits of activation, namely those that hit the primary sensorial, emotional, and pre-motoric fields, receive an expression in consciousness. They have some short term stability in that the firing of the neurones involved is co-ordinated in a certain oscillation. By hitting the primary fields in this way the phenomenal qualities and forms are brought about which constitute representations of situations, objects, and linguistic utterances. These episodes are the conscious expressions of the primary fields in their interaction with the conceptual maps. The role of episodic and conceptual maps in understanding situations and sentences is explicated in the architecture of maps, in which smaller and larger circuits form the constituent structure of our perception and imagination of situations and utterances of basic sentences. Special attention is given to the possibility of integrating an episodic memory into the architecture, such that current remembering consists in creating episodes on the episodic maps, organised from out the episodic memory map in interaction with conceptual maps. In chapter 4 the ability to evaluate representations, and herewith the relationship between understanding and interpretation (in the modeltheoretic sense), is elaborated on as a special capacity of consciousness. Thinking is seen as manipulating representations on episodic maps, with constraints given by the control through evaluation. Thinking and judging is discussed with respect to the question of whether all thinking and judging requires consciousness, and whether all thinking has to be in language.

The third part of the book consists of 4 discussions of controversial questions in consciousness research. The issues are whether consciousness is an internal monitoring device of brain states, or rather a monitoring of the external, whether all conscious states involve thought or judgement, whether there are different kinds o consciousness, and whether there is a one-one correspondence between (a certain kind of) brain states and conscious states. The different arguments are discussed and a standpoint is taken, derived from the evaluation of the arguments and the position on consciousness developed in this book, namely that consciousness is a product of the episodic maps, i.e. of primary sensorial fields in their interaction with conceptual maps, and that therefore all consciousness consists of episodes in the form of representations, with or without their evaluations.

Table of contents

Chapter 1

Consciousness and intentionality in perception, semantics, representations, and rules 1

1. Introduction: The case of blindsight and the episodic character of consciousness 1

2. Consciousness is necessary for denotational semantics and following rules. 7

3. Consciousness is necessary for having general knowledge and structural semantics 19

4. Consciousness is necessary for an intentional creation and improvement of language 24

5. Linguistic ability is not a generative system of rules 27

6. Linguistic theory formation is rational reconstruction 31

7. Linguistic structuring is flexible and is not fixed by a single set of rules 42

8. Learning from rules is possible via the construction of examples 46

9. Consciousness and intentionality are prerequisites for cognition and knowledge 52

10. There is conscious thought without language 57

11. Free will has to be learned on the level of consciousness 60

Chapter 2

Formal Constraints on the Relationship between Understanding and Connectionist Neural Network Activity 65

2.1 The non-classical cognitive model 72

2.2 Projecting Structures of Dynamic Conceptual Semantics onto Structures on Activation Patterns in Connectionism 75
2.2.1 The structure of Dynamic Conceptual Semantics 78
2.2.2 The desired structural properties of neuronal network processing. 92
2.3 The problem of systematicity and binding in simple connectionist networks 101

3. Conclusions with respect to a one-map architecture 108

Chapter 3

An architecture of episodic and conceptual maps for understanding basic sentences 111

3.1 Episodic and conceptual maps 111

3.2 Understanding situations and linguistic expressions 122
3.2.1 Understanding a situation 127
3.2.2 Understanding a sentence 129

3.3 The Binding-operation on different levels of description 134

3.4 Episodic memory in an architecture of neural maps 136

Chapter 4

Evaluation, Thought, Imagery, and Language 143

4.1 The relationship between understanding and interpretation 143

4.2 Ways of thinking. Thought, Imagery, and Language 150

Chapter 5

Controversial Issues in consciousness research 161

5.1 Is consciousness an internal monitoring device of brain states, or rather a monitoring of the external? 161

5.2 Do all conscious states involve thought/judgement? 172

5.3 Are there different kinds of consciousness? 179

5.4 Is there a 1-1 correspondence between brain states and conscious states? 188

Conclusion 201

References 205

Index 213

Chapter 1

Consciousness and intentionality in perception, semantics, representations, and rules

This chapter discusses the role of consciousness in the distinctions between reception and perception, between a purely causal and a referential or denotational semantics, and between linguistic ability and linguistic analysis, including representations and rules. The first two topics are treated by designing several thought experiments based on the phenomenon of blindsight. It is argued that reception, causal semantics, and linguistic ability do not require consciousness, while a denotational semantics, a notion of truth and reality, linguistic analysis, forming representations and rules, and following these require consciousness and imagination, like any design activity does. They presuppose a linguistic or a picturing medium in which they are formed. The medium is interpreted, i.e. connected to the world, via a neural network background established in training as our linguistic ability, which does not contain symbolic or picturing representations. Rather it functions as a system of dispositions in the ability to interpret pictorial and linguistic representations and rules. Dispositions towards the production of certain neural activation patterns are implicit in the connectivity between neurones. They are not language-like or picture-like representations. It is further pointed out that rules and representations can only indirectly function in changing or forming linguistic ability by serving in consciously constructing series of examples, which in learning processes can be a basis, a training set, on which linguistic abilities then are formed or reformed.

1. Introduction: The case of blindsight and the episodic character of consciousness

It is generally agreed upon that experiences have an underlying neuro-physiological basis, though it remains a problem what kind of relationship can possibly exist between two ontologically so different realms as, on the one hand, neuro-physiological states and, on the other, phenomena of consciousness. Physiological states and processes in the brain are necessary for experience to occur, but it is not the case that every such state or process has a conscious expression, since many such states are not accompanied by experiential awareness in consciousness. What kinds of features of such states or processes or what relationships between these result in consciousness is not yet known. With respect to states and processes that usually go together with awareness, there can be, due to an accident, a missing link in the causal relationship between reception of data submitted via the senses and the brain states or activities that express themselves in perceptual experience, such that we are not consciously aware of the fact that our body correctly measures, for example, the colour differences which are present before us. Without awareness we still might react bodily to colour and visual movement and form registrations in our brain. Dennett (1991) and Zeki (1992) refer to research about blindsight by Weiskrantz (1986), where the V1-region of the visual cortex does not function properly due to a lesion. Still, patients with this kind of lesions, who report seeing nothing in the blindsight field, registrate differences of light frequencies and motion in this field, though they do not see the colours or motion. Their ability to correctly assent or dissent to judgements of others about what colours, forms, and movements are presented before them is not perfect, but they make significantly better than random "guesses" about these, when these guesses are elicited by presenting to them in their blind field a choice, say between two forms X and O. Zeki (1992: 48) characterises their situation as "Blindsight patients are people who see, but do not understand. Because they are unaware about what they have seen, they have not acquired any knowledge." He suggests that the V1-area might play a role in pre-processing incoming visual data, transferring them through V2, which probably preserves the contiguity order of the input and sends the form, colour, and motion aspects through to the specialised categorising areas V3, V4, V5, and receives them back in an analysed form, recurrently after having been processed in these specialised areas of the visual cortex, which deal with form, colour, and motion, whereby V4 in majority contains cells involved in processing light wave length but also cells that respond to aspects of form.

The phenomenon of blindsight, radicalised in extent and form in philosophical thought experiments, has been used in philosophical argumentation about the nature of consciousness. Vision (1998) criticises Dennett's (1991) claim that blindsight shows that phenomenal consciousness is not central to seeing and that by training and accepting the available reactions as information about the surroundings that fall within the blindsight area, consciousness about the colour, form and motion properties of these surroundings could be attested, although the phenomenal content of these is not available. Without phenomenal content the person might thus have access to visual information and be conscious of this information. This position, which Vision (1998) calls 'doxasticism', denies that vision need be phenomenal. Of course, this position implies the claim that a robot with uptake of visual or other information about its surroundings can be claimed to be conscious of its surroundings. The criticism of the claim that the blindsight person sees or has visual experience about the blindsight area stresses the fact that this person has consciousness merely about experiencing certain non-visual reactions about which others explain to him that they must be caused by what can be seen in the area that falls within his blind field. From the information given by others he can infer by induction that these kind of reactions are indicative of certain colours, forms, or motions. Having this information involves consciousness about other kinds of perceptual input and about certain steps of inference via linguistically supplied information, but it is not seeing that there is a certain object in the blindsight field. Vision (1998) also criticises McGinn (1991), who takes blindsight to indicate that there must be deeper layers in consciousness that are not accessible to introspection. He argues that such a conclusion does not follow from the blindsight phenomenon. Also animals with a blind field can evade obstacles that are located within the blind field. There is no intervening unconscious belief in the blindsight person that he has seen that there is, for example, a car in front of him. The motoric reactions do not require such a belief but can directly be linked to certain input without forming such a belief. The blindsight person has no disposition to say that there is a car in front of him. We can add that from Weiskrantz (1986) case studies we might expect that probably he will be prompted to answer 'Yes' when asked whether there is a car in front of him, but he does not know why he gives that answer. To take this as evidence of a deeper layer of consciousness, which includes hidden beliefs, goes too far, because also here stabilised connections between brain states from visual input processing to motoric dispositions and behaviour cause the right linguistic reaction, although the blindsight person does not know why he reacts in this way. No consciousness about the reason for answering 'Yes' to the question is involved. If we call this a phenomenon of deeper or hidden consciousness we must take all conditioned reflexes between input into the body and reactions to them as evidence of deeper consciousness. This would inflate the notion of consciousness such that it becomes totally useless for discriminating mental phenomena from purely bodily phenomena.

An auditory case of unconscious awareness is that we do not hear sounds below the normal bass-range though we react towards these. We feel insecure and anxious without knowing that this is due to our body experiencing low frequency sounds above certain decibels. Based on these facts of unconscious awareness, let us engage in the following thought experiments, which have the purpose of getting some clearer idea about the role of consciousness in linguistic usage and thinking. We cannot say what consciousness is in reductionist terms. Rather we experience it as the basic fact of life, and we can try to understand it by way of clarifying its role in language use and thought.

It seems plausible that consciousness about data requires some basic categorisation and ordering of the data under several aspects, i.e. linking them to other similar data by partial identity of basic features and relationships under certain perspectives, and connecting them in contiguity relationships. Awareness of the data in perception, then, can be understood as a selection from the categorised previous input under relevance and attention, and hereby getting directed towards certain aspects. Episodic memory may consist in integrating the variously categorised data, percepts, into world- or situation-maps under space-time, causal, means-ends, and other relationships, which relate objects to each other or to actions and events, or to interests of the perceiving subject. These are kinds of contiguity relationships. Over these contiguity-maps a second layer of categorisations, namely generalisations by feature similarity and relational similarity could be organised by mapping episodes analysed on the first level onto maps in which generalisations take place. Processes of this kind, involving association by contiguity and categorisation by similarity, can be modelled in connectionist models. These second level maps organised by similarity under certain perspectives, furtheron serve for categorisations of incoming data and are thus used for filtering and transforming the incoming data such that they, in the forms so processed, can be sent back onto the pairs of contiguity preserving and primary sensoric maps, which hereby become episodic maps. From the episodic maps episodes emerge in consciousness. These recurrent ways in the whole construction of the neuro-physiological background of our conscious data require circuits of fibres, such that both kinds of maps, the contiguity preserving maps and the similarity preserving maps, i.e. the categorising and ontologically generalising maps, are linked sufficiently to each other and to the primary sensoric maps, which are hit first by the effects of the sensory data.

The architecture of visual episodes

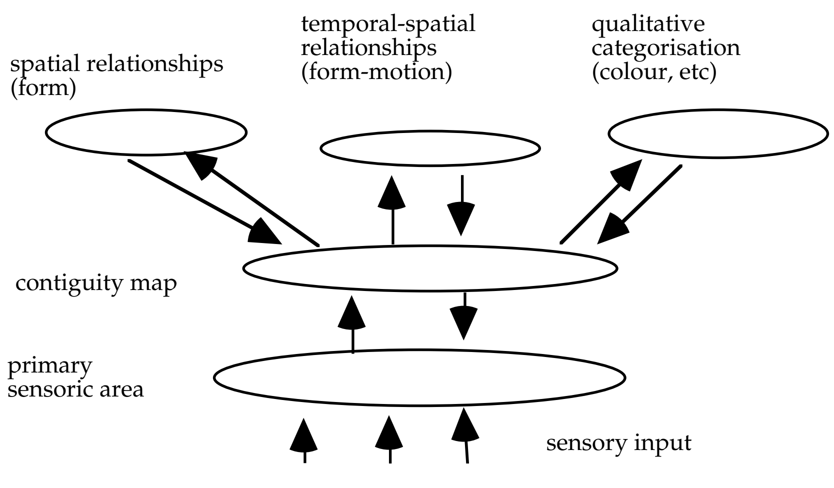

The episodic maps seem to be strongly related to consciousness, while the generalising ones, the concept representing ones, are related to consciousness merely indirectly via the episodic maps. We can, for example, not imagine a concept. Rather we only can perceive or imagine series of situations that exemplify the concept, which may originate from a stimulus to the implicit episodic memory and are constructed in imagination analogously to perceived situations of the episodic maps. Somehow, the binding of the different aspects of analysis, the contiguity and the similarity aspects, should occur in the episodic maps, where the characterisation and identification of an episode, situation or object take place. Characterisation by categorisation through connection to generalising maps, and identification by space-time and other contiguity relationships are united in the representation of an episode, situation or individual. Episodic maps have to register temporal and other relational properties.

A connectionist architecture of the following kind seems to be appropriate for explaining why consciousness is mainly episodic:

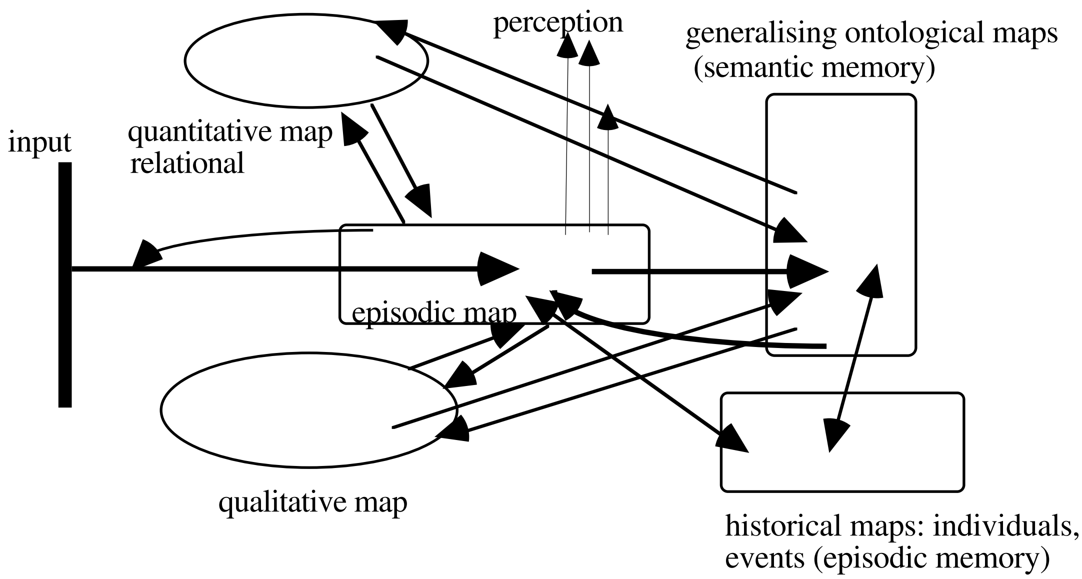

Sensoric and motoric input is being processed in maps specialised for qualitative and quantitative categorising these kinds of input. It is received there after going through the base of the episodic map, consisting of a primary sensoric map and a contiguity preserving map, which provide for the qualitative and the quantitative phenomenal material represented in space and time relationships. The current input vector from the sensory units is concatenated by recurrent fibres with a vector from the episodic map representing the previous input on the primary sensoric map, which forms the base of the episodic map. This repeated concatenation establishes the temporal ordering of incoming data on the episodic map by temporarily re-activating connections, and possibly causing a geometrical closeness of the received input on the episodic map. The episode thus gets a marked place in a contiguity net for ordering episodes. The episodic activation is the result of the sensory input and the analysis of the input on the qualitatively and quantitatively categorising maps. This parallel episodic contiguity and categorising activation is forwarded to the generalising maps for being further categorised in the coarse-grainedness we need for orientation in our surroundings. This amounts to an analysis of the input situation in our established ontology. From there the activation goes back by recurrent fibres to the primary modality specific episodic maps, which hereby, together in parallel and possibly also in interaction, constitute the episodic map of situations, which are understood in terms of our historical entities and our normal ontological categorisations.

Quantitative maps process object-internal relationships amounting to form, and process object-external relationships representing relative object positions and movement. These maps are connected to the episodic maps and to the generalising ontological map by circuits of fibres. On the episodic map we then get activation patterns which for us have their expression in terms of objects and situations, events and actions, as we experience them in consciousness, and especially in perception. They are presented, remembered, or imagined.

Such a model in principle explains why we can represent our concepts merely by series of episodes, i.e. real or imagined situations and objects. And, on the other hand, it explains why we can remember something consciously only when it has been categorised. Only then we can understand a situation or an object as different from others, only then we can remember it as a certain situation or object, and only then we can locate a situation or an object with respect to others, i.e. identify it and understand it as being in the world. There cannot be an episodic memory without some degree of categorisation. And as far as categorisation is a classification of individuals and events, and not just of situational impressions, the normal ontology of general concepts presupposes historical concepts, namely concepts of individuals and historical events, which have become part of the episodic memory. In chapter 3, the connectionist architecture of episodic and conceptual maps will be elaborated in more detail. In 3.4, the nature of the episodic memory will be elaborated as consisting of historically individuating contiguity patterns of indices with circuit connections to the characterising conceptual, including relational, categorisations.

Consciousness is restricted to perception, imagination, and episodic recollection. Even in so-called abstract thinking, sentence inscriptions or utterances or imaginations of these are the 'episodes' we deal with consciously. Episodes are perceived or imagined and they are characterised via the connections of the episodic maps to the categorising maps. They are identified by being represented on an episodic map by indicating their place in space and time and in other contiguity relationships with other episodes or situations and individuals that participate in them. In blindsight cases the primary sensoric area of vision is partly damaged. Since it is the basis of the visual episodic map, no representation and especially no percept can be formed there, which would have a conscious expression as a perceived situation. Representation on episodic maps makes use of the sensor related capacities of their sensoric base areas, as it also makes use of the motor related capacities of pre-motoric maps. The neurones dedicated to certain effects caused from the sensors must, at the same time, provide the phenomenal potential of the sensoric areas. In cases of blindsight, this phenomenal capacity is missing.

In order to shed some light on the role of consciousness in language use I now want to clarify the difference between a merely causal semantics and a denotational, i.e. referential-intentional semantics, and show what role consciousness plays for the latter kind of semantics. Causal semantics is part of denotational semantics in the sense that it is the causal background of our semantic knowledge. This background is causally established in learning and consists of stimulation by sensory input, neuro-physiological reactions on the input, and possible motoric output. In what follows I shall show the difference between denotational and causal semantics with respect to the role consciousness plays for semantics. I shall do this by designing thought experiments based on the blindsight cases.

2. Consciousness is necessary for denotational semantics and following rules.

In the case of blindsight, a person reporting not to see anything can react to things located in his blind visual field. With high probability of success he is able to fetch things before him, when asked to do so, and he can make good guesses about what was placed in front of him, when he is confronted with making a choice between suggestions about what objects are possibly there. However, he has no visual experience about these objects. The conclusion is that neural visual registration must have become directly connected, in a previous process of learning, to the neural registration of linguistic expressions as well as to the motor capacities. Missing is, in these cases, the visual experience, i.e. actual consciousness or the awareness of seeing something. This means that part of the mind is missing, if we understand a mind as the whole of an individuals current beliefs, feelings, and attitudes towards surroundings and actions, including the operations by which the individual makes connections between these. The following thought experiments, which radicalise the blindsight case, make no claim of being realistic; they merely serve the purpose of philosophical argumentation for the distinction between causal and intentional denotative semantics, between having dispositions and learning and following rules, and generally for the distinction between unconscious processes in the brain and their conscious results, which constitute our mind.

The first thought experiment:

Imagine that somebody has a defect causing complete blindsight from birth on. In such a case he cannot have instances of colour concepts and visual form concepts in his mind. They cannot play a role in his cognitive capacities, though he might be trained to show certain bodily responses to visual input. For example, lifting his left arm might, in a series of training examples, get indirectly associated to showing him instances of a certain colour. This might be done by telling him to raise his left arm, when we see the colour presented in front of him. Raising his left arm then not only gets linked to our telling him so, but also to the accompanying colour input he cannot consciously experience, but which nevertheless becomes associated by neural connection strength with the effects of the colour reception in his brain. He cannot see the colours, i.e. does not experience their effects on him, except that he somehow suddenly feels compelled to do certain things, namely to lift his left arm, or to ask his trainer "Why don't you tell me to lift my left arm? I feel like doing this." These are the occasions when we show him certain colour samples. But he does not know that we do so. Though his behaviour has become conditioned to these colours, we cannot say that he has learned a rule to the effect that he has to react in certain ways to certain colours. He cannot explain his behaviour as following a rule, because he never can state the rule; and he also can never follow that rule because he cannot perceive whether the conditions are fulfilled which require the reactions. He just feels at certain moments, beyond his control, the disposition towards these reactions. If there are no reasons against, he simply will react in the way he is trained to do. But we cannot say, he will react in the way he has learnt in the sense that he has learnt that he ought to behave in a certain way when a certain condition pertains. He also cannot experience that he shows a regular reaction towards certain colours because he cannot experience one of the two things that are connected in the associative training. Thus we conclude that following rules requires awareness of perceptual input and consciousness of a regularity, though a robot without consciousness can produce behaviour that can be described by rules.

An interesting question now is whether we can say that our blindsight person, in some sense, knows the colours which he has been trained to react to.

Second thought experiment:

He might learn to say that a certain colour or form is in front of him if he feels the bodily reaction which the experimenter induced in association with the colour in the training set. Instead of learning the name as the name of the colour he learns the name as a name of the type of his motoric reaction. He can be told that when he feels an urge to realise the satisfaction conditions for the name, something we call colour X is in front of him. When we then tell him that at a certain place somewhere away there is colour X, he understands this in his way: At that place he will feel that he has to lift his left arm. He nevertheless does not feel, i.e. experience, anything like colours. He can learn, to some extent, to make and to verify colour statements, and form statements if he also would be trained for these in the way described. He can by himself determine truth values for statements about colours and forms in his surrounding. In order to find out whether a certain thing which he can touch has a certain colour, or whether something else behind it or next to it has the colour, he has to perform experiments in the way of removing it and putting it at other places in order to see whether he still feels disposed to reaction X. This is rather troublesome, but in principle he would be able to learn to make the right statements about the colours of things, and also to state what forms things have without touching the things. For him these would not be colour- or form statements, but they would work pretty well in the social co-ordination of judgements. We could say that he knows in some way what colours things have, and that he knows in another way what forms things have because he can also classify them by touching, besides classifying them via his typical causal reactions to the visual input he was trained to, and of which he is unaware. He cannot see the real thing but he can classify it by symptoms, which are his kinds of reactions to received and partly processed, but unperceived, visual stimuli. We also can think of our own colour perceptions as symptoms of certain light frequencies, although they are based on innate dispositions, while the arm movement is learned.

Our question now is whether such a totally blindsighted person has a semantics for colour statements or form statements. He certainly has some causal semantics, since he assents or dissents to colour statements, according to the training he has gone through. This is a causal semantics we also find in neural networks trained towards perception patterns. These nets can learn to connect these patterns with expressions by building up associations between activation patterns of the object perceptions and the expression perceptions. They thus can learn to classify input by attaching names to it. As long as a person is not aware of any bodily sensations that occur to him in cases were we would see colours, he just would be caused to endorse statements about things having certain colours as we do, but he would not be aware of what causes him to assent. He would never be able to learn a rule that connects certain colour perceptions with certain colour labels, though he would learn to apply colour labels correctly. He would be able to judge about correctness in the sense that he could detect a difference between cases of correct and of incorrect applications of the terms by testing them against his dispositions towards certain arm movements in these cases. He remembers approved situations of correct application of terms in opposition to experienced or imagined situations of wrong application, in which he did not feel compelled to lift his arm. The denotations of our colour expressions for him would then be certain arm movements. Such a person would have a criterion of correctness of term application, but for him it would be correctness of applying movement terms. He would have a denotational arm movement semantics and not a denotational colour semantics. Rather we can attest him a merely causal semantics for colour terms. A denotational semantics, namely, requires consciousness about criteria of correctness, the application conditions of the labels, and in our example, the conscious awareness of colours in perception, as that to which the ascription of the labels is conditioned to. For our blindsighted person the denotational semantics of our colour terms is rather a denotational semantics of making certain arm movements connected with certain terms, and about this he could make correctness statements.

Third thought experiment:

If the completely blindsighted person had not learnt to make certain arm movements in satisfaction situations for certain colour terms, but would have learned the colour terms previous to his total blindsightedness, his semantics for our colour terms would now simply be a causal one, especially if he cannot anymore remember colours in imagination. - And this is very likely because recollection and imagination of qualitative experiences probably will make use of the same neuro-physiological processes that give sensory reception and processing an expression as a certain perception in consciousness. - Whenever he would try to imagine what the difference is between a situation of correct application and of incorrect application of a certain colour term, he would just register the difference between a situation in which he feels compelled and one in which he does not feel compelled to utter the respective term. But he could not correlate his feeling compelled with a certain aspect of the situation at hand, as we can, when we relate our utterance of *red* to something red in a situation. For us, having a concept which is expressed by a term, is not just feeling compelled to utter the term in some situations judged publicly as appropriate. Rather having a linguistically expressed concept implies to relate the term to a certain aspect of the situations in which we feel compelled to use it. To distinguish aspects in situations and to realise how situations are built up from such aspects requires perceptual awareness, i.e. consciousness about the situations, and not just feeling compelled to use certain terms. Consciousness about the relationship of reference and satisfaction, which is established between aspects of a situation and the linguistic expression used for these aspects, is typical for having a denotational semantics. We intend the relationship to hold for the expression to be true or satisfied.

Denotational semantics includes a conscious realisation of the correspondence between the conscious results of two causal processes, the image or utterance of the linguistic expression caused by the situational, here visual, input, and the percept or image of the visual input caused by the situational input. The completely blindsighted person trained to make or select the appropriate linguistic utterance, the completely blindsighted person trained to make the appropriate arm movement and to make or select the appropriate utterance, and the normal person, who has a denotational colour semantics can be represented and contrasted by the following three schemata:

1. The completely blindsighted person without arm movement
(after the training process):

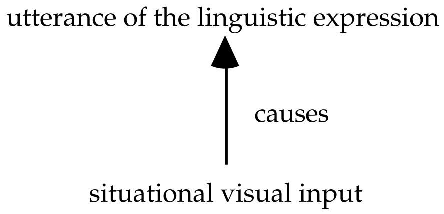

2. The completely blindsighted person with arm movement
(after the training process):

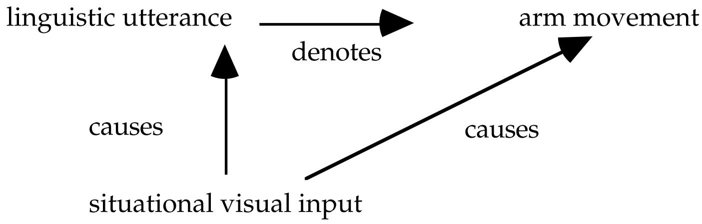

3. The normal person (after the training process):

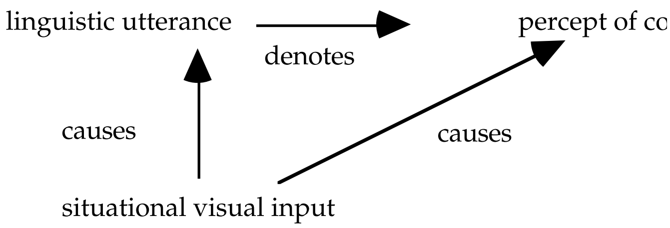

Against a current version of Wittgenstein's anti-private-language argument, I argue that an individual, all by himself, can make correctness judgements, if he is conscious about applicability conditions of a term in the sense that he simply remembers, imagines, or refers to aspects of situations of application of a term. These then are present to him in consciousness. Such acts of remembering or imagining aspects of situations of use for a term, or referring to some such situations outside, are conscious acts of control with regard to coherence and regularity of the use of a term, additional to just feeling compelled to use the term. The several cases of feeling compelled to use a term are compared to each other with respect to a regularity under the perspective or frame provided by a certain kind of practice. Hereby, the use of a term is placed within sequences of actual, remembered, and imagined situations, which have to show some convergence towards a regularity, i.e. an internal similarity or aspect of the situations that approaches what would be a set of the applicability conditions, if they could be formulated. In such a procedure of comparing situations of use, each of us makes correctness judgements, of course, without ever being able to exclude the possibility that one was wrong in the judgement by failing memory. And this possibility cannot be excluded even for the language community, which provides the criterion of correctness for public language. We only can say that we normally are quite reliable in these judgements, measured against the likewise, normally, quite reliable correctness judgements of other speakers. Finally, the relevant reliability is measured not merely by previous and actual consensus with other speakers, but by the undisturbedness of our worldly practices that rely on the regular use of the respective terms. Because of this aim of undisturbedness, regularities in the world have to be taken into account by appropriate regularities in the use of our terms. This undisturbedness is reason for a single language user, all by himself, to be confident in his use of linguistic means, even if no other language users are around to support or correct him. The same holds for a whole linguistic community. The last criterion of correctness of our linguistic usage is the undisturbedness of our practices and the experience of regularities in our surroundings in correspondence with regularities in our linguistic utterances.

A blindsighted person may also feel assured about the reliability of his use of our expressions of visual perception and about his correctness judgements, if he has been trained to the additional arm movements, or something similar. But without that training for some associated movements or other behaviour he can be consciously aware of, he would not have any means to understand what correctness of his use of the visual perception terms amounts to, other than that his judgements agree with those of others. He would, though, realise undisturbed practices. When he feels compelled to utter a certain expression, say *a tree before me*, he would at the same time feel compelled to take a detour and would be pretty safe in not bumping into the object located on his path. But all that does not provide him with the direct awareness that he uses the term in accordance with certain conditions of application of which he could become aware by the regularity in imagining situations. When he is sitting he would not feel compelled to evade the tree, only when he is walking, and it is on his path. And when he is just sitting there, he might not feel compelled to utter the expression he has been trained to, because he is occupied by something else, or has no urge to communicate. Thus the situations of application of the term he can be aware of, namely taking a detour or feeling compelled to a certain utterance, are much more dependent on situational factors and they do not show a common regularity, though they sometimes factually overlap temporally, without being similar to each other in any way: the situation of feeling compelled to take a detour, and the situation of feeling compelled to utter a certain expression overlap in time but are totally unrelated to each other. For a person with visual awareness there would be one situation which he can see, and with it certain actions are connected that might be appropriate, such as uttering a descriptive term and evading the obstacle seen ahead. A similar case is the following: When we ride our car while thinking about quite different things, we still are disposed to evade obstacles without being consciously aware of them. We have not seen them, though we have got visual input and automatically processed it such that the appropriate connections to the motoric reactions are made sufficiently.

We can compare the situation of the totally blindsighted person having merely a causal semantics for colour terms with Churchland's (1986) model of the connection between the crab's vision and its arm movements in catching a prey. There is a direct transformation between the sensory topographic map and the motor topographic map. The visual input causes a pattern on the sensory map, which is transduced to the motor map, whereby it gets transformed such that a motoric pattern is activated which causes the arms of the crab to reach for the prey it has seen in the sense that it has received visual input. We do not know whether the crab has really experienced seeing the prey. - That would, according to the present theory (cf. chap. 3), mean that the crab would experience episodes, and this would require that there are conceptual maps, which are in circuit interaction with the primary sensory areas. - The whole crab situation suggests a simple causal semantics, whereby no conscious awareness in the sense of having a conscious representation or picture of the prey is required. The crab's correlation between visual reception and motoric activity might just function like the correlation between unconscious colour effects in the brain and the use certain colour or form terms by a totally blindsight person, except that the person consciously perceives the terms he uses, or at least feels being compelled to use the terms.

What is the additional benefit of having available a representation which one can be conscious of? Consciousness is a prerequisite for intentionality and for learning and applying rules, instead of just showing regular behaviour achieved by training. The observer can see that the behaviour is regular with respect to colours shown, but the blindsight person himself cannot know whether his behaviour or his feeling of being compelled to do X, utter a certain expression or move his arm, is regular. He misses the comparison with the associated colours. He only can take into account his behaviour or disposition, but he cannot relate to what it is correlated with in a regular way. He cannot understand his behaviour towards colours as regular and thus he has no rule for doing X, for example using certain terms or making certain arm movements. Having available a rule implies understanding our learned behaviour as regular. It also makes it possible to sometimes abstain from the behaviour, since a rule or regular behaviour can be deliberately broken by intentional inhibition of the regular response.

Consciousness further makes it possible to imagine situations and plan behaviour accordingly. Planning requires imagination and intentionality, and imagination presupposes previous perceptional experience of situations, which is conscious, i.e. real perception and not just reception of input with some processing. -- Perception, as the term is used here, includes conscious presentation of what is seen, while reception merely includes visual input and part of its processing without coming to conscious awareness. Intentionality relates the imagined or represented situation to the world in a direction of fit which can be world-to-mind in belief and perception, or mind-to-world in desire and planning (Searle 1983).

In this context, representations should not be taken as syntactic constructions out of basic concepts in the way they are thought of in a mentalistic semantics, like the one of Fodor's (1976) language of thought. Of course, there are linguistic representations in our natural or designed formal languages, but they are not necessarily involved in propositional attitudes like belief and desire, or in perception and action. They only play a role in representations of these in reports. The only representations in our mind are the perceptual presentations and the recollected or construed images of linguistic utterances and inscriptions, of real or imagined situations, and of the relationship of approval between an expression and a situation in which the expression is used correctly. Our linguistically labelled public concepts are equivalence classes of stabilising sequences of growing sets of such satisfaction situations for an expression. Representatives of these equivalence classes are available to every language user. Each representative of such a concept or equivalence class is a sequence or set of situations, subsets of which we can call up in our mind as sequences of examples for the concept. We can come up consciously with such sequences of imagined examples, and we can refer to such sequences of examples in the world. These sequences of examples of our concepts consist of presentations and representations of situations or images. We conclude that our totally blindsighted person has no colour concepts, since he cannot imagine or design examples. Of course, he must have assemblies of neurones activated with the reception of certain colours and forms, which connect his arm movements or the endorsed colour statements to actual colours and forms he does not perceive. Maybe that in our completely blindsighted person some received input can go directly to V2 and the categorising visual areas, bypassing the damaged primary sensoric area V1.

Although our blindsighted person, like a robot, naturally merely has a causal semantics of colour terms, I have pointed out that a denotational semantics can be achieved if the blindsighted person is trained to both, making certain movements for different colours and also using our colour terms. Then he can combine in correctness judgements both external facts he is conscious of, the movements and the use of terms, and achieve a denotational semantics by connecting the terms with the respective movements, although he is not able to use any colour experience as a missing link. He now can imagine situations of false use of a certain term by just imagining a contrasting movement that does not belong to this but to another colour term. But note that these terms in his denotational semantics are in fact not colour terms but movement terms; he is conscious of the fact that he feels compelled to use a certain term, our colour term, in case he feels compelled to a certain movement. The term, as well as the movement, is located in a network of contrasting terms and movements, respectively. The term is correctly or incorrectly used on the background of the sets of contrasts within each network, which correspond to each other. By virtue of these contrasts a term can be judged to be correctly associated with a movement. The point of this second thought experiment was that the person is conscious of his use of expressions in connection with certain movements, of which he is also conscious, and he feels compelled to use a certain term in case he feels compelled to make a certain movement. He had been trained to connect two external things he was conscious about, expressions and movements in situations in which others applied these terms to colours, of which he is not conscious, but to which his movements and use of expressions have been causally connected by a training process. For the causal semantics between expressions and unconscious colour stimulation, and for the causal semantics between movements and unconscious colour stimulation, no consciousness of the visual uptake is necessary, and even using the expressions and making the movements can be unconscious, as we assume it to be in neural nets and robots generally. To have a causal semantics does not require consciousness. Therefore to point out that robots have a sensory uptake which is connected to a motoric output by complex processes of information processing is no argument for assuming consciousness in intelligently programmed computers and robots, as Chalmers (1998) is inclined to do.

On the other hand, for a denotational semantics, which is a semantics based on correctness, especially truth judgements, two consciously available media are necessary, which have to be related in a structure preserving manner. This relationship is intended by the language user. To intend this relationship of reference, satisfaction , and truth means to understand language as language, i.e. as a system of symbols, which stand for something. To understand a language as a means of representation requires to be consciously aware of the signs and of what they represent.

Denotational semantics has been learned by training on external examples one has been conscious of. The structures are defined by identities and oppositions between the cases of expression use on the one side, and by the identities and oppositions between the instances of the corresponding partial extensions or denotations, on the other. In their structural fit to each other, the individual's correctness judgements are founded. In this way he could have a private language with correctness judgements as far as the fit of the structures goes. In fact, such a private language is learnable and teachable via the external examples on which the structures are erected, and thus is in principle a public language. In order to learn a public language, the elements of the structures have to be related to external things and practices about which we are conscious, and to which we refer consciously in our use of expressions. This relation between the two structures and their external aspects is not one between two separate realms, rather the external denotations are the data on which the structures are built by operations of ordering data according to factual contiguity and to similarity under certain perspectives. The relationship between both structured realms of data is induced in establishing the structures on them in association between expressions used in satisfaction situations and situational data present in these (cf. Bartsch 1998). Thus the relationship between both structures is an internal correspondence rather than an external one between two independent settings. In the process of learning and understanding language and world, linguistic expressions and facts satisfying them are necessarily taken to be partially identical in structure. The linguistic expression selects from the set of orderings into which the situation fits, conceptual orderings depending on how far the expression is analysed in the available language. This selection is what is highlighted in a situation when we realise that it supports a certain fact. The situation is then represented under this respect as a situation described by the expression.

What now are the major shortcomings of getting information in a setting of a merely causal semantics, compared to our conscious manner of knowing and perceiving, in which our denotational and intentional semantics is rooted? In order to give an answer to this question let us push our thought experiments somewhat further and compare our imagined blindsighted person with normal perceivers on the one side and robots on the other.

Our totally blindsighted person behaves rationally with respect to colour and form judgements, i.e. draws the same implications we do. But he does not have any visual aesthetic judgements, since he has no colour experiences and visual form experiences, and therefore no visual imagination of situations. The main difference with us is that he has no visual imagination. He may have motoric aesthetic judgements, or aesthetic judgements about certain sentences he feels compelled to generate by some cause unknown to him. The aesthetic judgements may or may not coincide with our judgements about visual experiences, and he might learn that it is culturally appropriate to call certain of his motoric aesthetic experiences beautiful, sublime, or ugly, following our judgements about, for example, a colour ballet. And he might learn the same about his linguistic combinations.

Let us suppose again that our completely blindsighted person had not been conditioned to making motoric reactions to colours, forms, and movements, but merely had been trained to utter and accept certain statements in situations where he receives non-perceived visual stimuli, as described in the third thought experiment. He then would have no visual imagination and also no motoric aesthetic substitute. Because he lacks any such imagination, he never could design anything visual and he could not anticipate situations, except that he could form sets of statements which he knows are coherent in the sense that he has experienced that they belong together somehow because they occurred together before. He thus can only imagine feeling compelled to make certain linguistic combinations. He and others, together, somehow have used them and could use them as descriptions of situations he could not see. His only reactions towards these situations, he is otherwise unaware of, is linguistic, i.e. a disposition to utter certain sentences. The coherence of these sentences can not be felt in any way, except in a structural similarity or analogy to previous discourse. He can, of course, acquire the notion of a situation only via his other senses, as certain units of smell, touch and sound he is conscious of. On this basis he could make designs, using his imagination of previously experienced situations and modifying them. But this is not at issue here.

We can imagine the extreme case in which he can only perceive linguistic utterances and other sounds. A zombie or robot is a further radicalisation of this case, where no perception is available but merely reception of sensory input, processing and behavioural output. A robot has no sense experience at all in which he could consciously register situations. The robot also cannot feel compelled to make certain linguistic expressions, and therefore he could not imagine situations of feeling compelled to produce such expressions. He gets expressions as input but does not perceive them. He just is caused to give expressions as output triggered by a certain input. In fact, he cannot imagine any situation. Rather he just could come up randomly with new linguistic combinations, possibly guided or restricted by previous probability assignments to certain expressions or structures, closeness of expressions in previous examples of discourse, or he could produce expressions by being set to execute rules programmed or hard wired in. Yet a robot with sensors, but without perceptual consciousness, could be constructed to memorise and generalise input and output patterns, for example linguistic output associated to sensory input. Some correlations could be kept in the memory of a controlling device. In this way he even could mimic correctness judgements and conscious comparison between good and bad cases of term application. Such a monitoring device could replace our conscious monitoring and be developed into an ability to give automatic correctness judgements, including truth judgements. This device then could be used to teach others, without knowing what it does.

The correctness judgements of such a controlling and monitoring device are unconscious and unintentional. In fact, they are not real judgements for the monitor, because for him they are not about anything. The controlling and monitoring computer does not know what correctness or truth is, though he 'judges' correctness and truth rightly. But he could not intentionally break his own procedure and accept an incorrect example, knowing that it is incorrect by the standards. He could not decide voluntarily to follow or not to follow his own rules. He does not even know that there are any rules. Still this computer might be an effective teacher for a certain task. Even though he does not have a denotational semantics, he still can produce correctness statements. But they are correctness or truth statements only in the eyes of the program designer and user. The monitoring and controlling device itself does not know what correctness of a task performance amounts to. He blindly registers and checks the task performance of his pupil against his own learned task performance, or against the rules and representations programmed in. A teacher who understands what he does could ask himself whether he himself was quite correct in performing the task. A monitoring and controlling robot could not do that, though he could be programmed to, now and then, repeat a training session for himself on an old or new training set. But he would not know that he trains himself in order to uphold his own ability to task performance and correctness judgements on others' performances. This taking in of a training set can be seen as the causal semantics of the term 'learning task X' or 'retraining task X', like his performance of the control and correction task is the causal semantics for the term 'giving correctness judgements about performance of task X'. Thus he has a second order causal semantics, but this does not make his first order semantics denotational. Our advantage as conscious beings is that we are always ahead of the robot. He only can take the next step of meta-performance if we first program him accordingly. The point is that he does not understand his own behaviour. He does not know what he does. He just performs.

What is missing in a robot is the development of the semantic ontology, by which input stimuli can be synthesised by the power of imagination, as Kant has emphasised. The synthesis of series of situational impressions in space and time gives rise to historical concepts, such as concepts of individual objects, events, and acts, and it gives rise to general concepts. They are comparable to Kant's empirical schemata of classes of objects, events, acts, and abstract objects. In these schemata concepts are not merely sets of features, but these features are aspects of things or situations ordered in forms of space and time by the power of imagination. These aspects are not all explicated like features would be in traditional definitions of concepts. This same power of imagination is part of every act of perception and experience generally. Transposing these insights into today's approaches to cognition we may assume that a bias towards our coarse grained ontology is built into our neural net organisation in the course of evolution. The network then will be able to synthesise incoming data into the evolutionary selected ontological frame in a way which is causally efficient with respect to our behaviour within our surroundings. This all does not really require consciousness. But still, using historical concepts and general concepts in understanding and in recognising objects and categorising them requires consciousness of the unity of an historical object and of the unity of a class as a general concept. We have to recognise parts of space-time slices of an individual while assuming, in space and time, the continuous whole, i.e. the total individual, of which they are parts, and we have to perceive the unity of the class through the examples of the concept. It is the assumption of totality or completeness, although we factually only realise partiality, which is at the root of our notion of reality and truth. These are not part of our experience, but regulative ideas of reason, as Kant and Husserl have characterised them. However, it seems possible that such regulative ideas could be built into a robot as procedures towards coherence and completion of partial general and individual concepts. Such procedures would correspond to our intention towards truth for statements and reality for referents of our partial situational concepts and individual concepts.

Denotational semantics is built on the semantic notion of truth. The following argument shows that self-consciousness is a prerequisite for having such an objective notion of truth. In order to understand sentences about the past, *PAST(p)*, one must be able to think "If I had been a suitable observer in that past time I would have observed that p". In a causal semantics, such as a robot can have, a link can be laid between past experiences and the present in the following way: 'Yesterday it rained' is true if and only if in the memory on the location labelled as 'one day before now' a perceptually caused situational representation connected to the phrase 'it is raining' can be found. But note that this is not sufficient for a realistic notion of truth. It is a purely phenomenological or procedural relationship depending on the perceptive history and the memory of the robot. Rather a counterfactual relationship has to be added, namely "If I would have been a suitable observer the day before the day now, I would have observed that it was raining then", i.e. I would have had a situational perception fitting to the phrase 'it is raining'. Such a counterfactual statement requires consciousness and even self-consciousness. This statement can be generalised over the whole speech community as: If there would have been a suitable observer the day before now, he would have had a situational perception fitting to the phrase 'it is raining'. To have and to handle a notion of reality above mere phenomenal registrations presupposes self-consciousness. The realistic notion of truth requires consciousness about one’s own and others' possible and counterfactual positions as verifiers of a statement. The same is true about the realistic notion of an individual; to have the concept of a real individual, surpassing the knowledge of partial individual concepts, requires consciousness with respect to one’s own and others' possible and counterfactual positions as perceivers of the individual. Understanding simple statements about the past requires self-consciousness and consciousness about oneself as a possible observer among others. The same is true with respect to statements about the present, if they are about some situation at a distant spatial location.

3. Consciousness is necessary for having general knowledge and structural semantics

Structural semantics is a set of semantic networks in which a term is explicated by its place in the network. The semantic network of a term consists of relationships between this term and other terms, which are interpreted as semantic relationships, i.e. as holding between the extensions of the terms in real and possible worlds. These relationships make up the theoretical concept expressed by the term in question. Alternatively a theoretical concept can be expressed by a set of general sentences in which these relationships are formulated as holding between entities that belong to the extensions of the terms in question. A taxonomy, for example, is a special kind of semantic network which expresses the relationships 'is a', i.e. inclusion or subordination, hyponomy, and the relationship of exclusion, i.e. opposition between concepts. Such relationships can also be formulated symbolically in traditional lexical rules, meaning postulates. On the other hand, they can also be found, in a more flexible way, in structures on the sets of activation patterns in connectionist networks, where they are, so to speak, directly contained in the conceptual 'contents' as they appear as activation patterns on the nets of units which can be understood as conceptual maps (cf. chap. 2 and 3). The activation patterns can be taken as 'contents', or better as conceptual indicators, by virtue of their relationships in which they are causally bound to the perceptual input in learning processes. The activation patterns are due to input examples and the adjustment of connection strengths between neural units in learning, i.e. in generalising the examples presented as input. They are directly addressed, i.e. activated, by new perceptual input of similar kind. The causal semantics of activation patterns, encoded implicitly in the connection strengths between units on the neural net and the input units, is established in the learning process and reactivated for each new occurrence of similar input. The activation patterns therefore function as indicators of the concept or category of the instances.

Our structural semantics, established publicly in endorsing general statements and theory building generally, has on the connectionist ‘content’-level a counterpart which is built upon causal semantics while building up this ‘content’ level in learning. This basic structure of inclusion, overlapping, union, and exclusion of potential activation patterns is causally established in learning and does, by itself, not require consciousness. However, the symbolic formulation of relationships between concepts does. It would be an uninterpreted network between terms if the terms were not causally linked with the objects and situations in the world about which the relationships of structural semantics hold. This causal linking has been built up in a regular way as a stabilisation in the distribution of neuronal connection weights during learning processes. The current activation patterns are the results of our dispositions, implicit in the network architecture and the connection strengths, to activate the respective patterns in reaction to certain input. These activation patterns then cause a certain linguistic or other behavioural output. Activation patterns for the perception of linguistic expressions and activation patterns for situational perception are connected through the process of language learning, and in this way we connect the use of a term to situations.

The interpretation of expressions as a relationship between expressions and the world is a conscious one in intentional denotational semantics, which delineates objects and situations which are coarse enough to be the units or chunks in the world we have to reckon with, and which are publicly referred to and communicated about. Thus consciousness and intentionality in perception and action make us select the right size objects and situations to deal with, and our denotational semantics is directed toward these objects, which constitute our world, the world we know. Structural semantics is our conceptual ordering on this level of intentionality, and in this way is a derivative of denotational semantics, which constitutes the interpretation of its symbolic formulations. Only on our intentional level of coarseness of objects, actions, events and situations generally, we have knowledge, and according to this level of coarseness the activation patterns on the connectionist neural level get delineated and partitioned by causal relationships to action routines and to public language use. Causal relationships with these two aspects select and thereby delineate structures of activation patterns by enhancing them in joint stimulation. This happens in a consistent and precisifying way by increasing the connections' strengths built up in the association between conceptual activation patterns, on the one hand, due to perception of our factual surroundings and our actions in these, and, on the other hand, activation patterns due to the perception of linguistic expressions. Linguistic expressions, and also action routines, serve by means of these connections to emphasise, i.e. strengthen the parts of activation patterns which correspond regularly to the use of these expressions or to the performance of the respective actions. Hereby the patterns of publicly relevant objects and situations, as encoded in our public language, get distinguished on the level of 'contents' or conceptual indicators as they appear on the neural networks. These considerations show how our public denotational and intentional semantics and causal semantics depend on each other. The first can only be handled by means of the second, and the second in its specific causal structure depends on the first, in the sense that it is induced by the first.

I shall now show that for acquiring the semantics of certain terms, consciousness of a specific quality of perception is not required. However, there must be some perception of discriminating qualities we are conscious of, in order for a denotational and structural semantics to be possible. We have to be able to be aware of identities and differences necessary for building up structural semantics, or general knowledge. In order to see this, we now change our thought experiment in such a way that certain reactions to unperceived, though well received, visual stimuli are not brought about by training but rather are assumed to be innate.

Fourth thought experiment:

Imagine that someone has experiences connected with colours which are quite different from our experiences. For example, instead of having a colour experience like we have, he feels a certain kind of itching at a certain area of his body if a certain colour is present before him. In certain context types his concepts of certain colours would be structurally equivalent with the ones of other speakers.

But this would not be so in all contexts, especially not all our oppositions between the colours could be made. If, what is very unlikely, there would be a total equivalence in colour judgements between him and us in all satisfaction situations, then we have to say that his concepts are the same as ours. Suppose that with every colour term he has learned to associate a certain kind of itching. This serves for using the colour terms correctly with respect to satisfaction situations, except in cases where the itching is caused by something else than lightwave frequencies. He can report to others about colours in objective situations correctly. Still, the difference of his colour experience from those of others is statable. The subjective situation will be different; the others will say that for them colours do not cause itching anywhere. How far does the equivalence extend? Full equivalence would require that the itching would be located in the visual space. That is not possible. Only if a space can be formed in which different kinds of itching are related according to their location in space, his use of colour terms could be similar to ours. Otherwise he cannot locate colours at the place where they are. He only would know whether some colour is around. To locate them as belonging to certain objects would require the same complicated procedure of moving objects around, which has been described above in the context of the second thought experiment.

In our example, we have identical structural relationships in the conceptual space of colours and of itchings, but the external location, a property of the visual field, is not represented. The itching would have to be located outside the body in space. We would need something like a pain space outside us, in opposition to pain in and on our body. The itchings, though due to stimuli on our skin, have to be noticed by us as if projected outside our body such that we perceive them as properties of itchy objects outside. Instead of objects of different colours we would have objects of different itchiness. If objects outside would itch us by the light they reflect, and by certain qualities of the itchings, their intensity, and a kind of feeling of the angle under which the light reflections hit the skin, we would also get information about the place in space, the distance and direction in which we find the object which itches us by the light waves it emanates. This, of course, requires correlated motoric experiences about reaching out towards a place in space. Note that an auditive space is similar to this kind of construct.

Suppose now that somebody would merely see a certain black-white pattern instead of a colour, and for different colours there would be different patterns. They are placed in visual space corresponding to the colours we find there. Nevertheless, his notion of colour would differ from ours, and this difference can be described. His colour perception would be part of his form perception. He could not distinguish whether indeed there is the form pattern, or whether it is the colour that for him has become this pattern. He therefore would not have the same colour concepts we have. The difference is due to the fact that the colour oppositions and identities would be integrated into the form oppositions and identities and hereby get indistinguishable from these. This would also be so in the example of itchings, if he also would experience other itchings than those caused by colours.

The off-shot of these thought experiments is that perception of other qualities, replacing our visual perceptions, can result in partially identical semantic networks, i.e. in partially identical structural semantics, or alternatively expressed, in partially identical general knowledge. In principle the specific quality of sensory perception does not matter as long as it serves to make the publicly endorsed identities and differences. This is the old structuralist principle which holds not only for language but also for knowledge generally. All our knowledge, as far as it is publicly shared, is structural. Substance or quality matters only as far as it functions in stating identities and differences, and this can be done couched in different kinds of substances and qualities, or different kinds of phenomenal material, in an equivalent way.

Fifth thought experiment:

Let us change our thought experiment and suppose that, for example, somebody would see red as green, blue as yellow, etc. Then he would have the same notion of colour we have, and also his colour concepts could not be distinguished from ours. We would not even come to know that he perceives colours differently. We could not but say that he has the same colour concepts we have. This is due to the fact that all oppositions and identities are the same as ours, and we can communicate about these. However, he could connect colours differently to feelings from what we experience as normal. Suppose that our red is for him identical with our green feeling; then what he calls 'red', would calm him, and our green, which he experiences as our red feeling, would make him nervous or stressed. This we could interpret as a sign that he has a different colour perception from ours. However, we also could explain the difference by supposing that his colour perception is the same as ours, but that the connection to other feelings is different for him, which makes him nervous in sight of a green meadow and makes him relaxed when seeing a room decorated red all over.

The result of this last thought experiment is that by an interchange of colour qualities the structural semantics of colour terms would just remain the same, as, alternatively formulated, would our general knowledge about colours. We would not even detect whether there is such a difference in colour perception between people, and by this very reason logical empiricists have used the statement that different people might have a qualitatively different colour perception as an example of an (empirically) meaningless statement.

4. Consciousness is necessary for an intentional creation and improvement of language

There are many physiological registrations of internal and external processes, reactions on surroundings, and inside and outside effects on our body and its nervous system which we are not consciously aware of because we do not experience them consciously. For these bodily registrations also neural registrations and classifications could be built up, very much comparable to the ones of which we have experience. Of course orientation of our body can be achieved by conditioning to these influences, similar to the orientation of the man who reacted on what he "saw", without seeing it consciously. This leaves us with a lot of regular and irregular behaviour, unaccounted for in terms of rationality and intentionality, because in explanation and understanding we can only take recourse to causes and motives we are conscious of.

We might react regularly to sounds of frequencies beyond those we can be aware of in hearing. This we could detect if we would monitor our behaviour while carrying with us a frequency meter, which we routinely check. An example of unconscious training would be a robot who has learnt to react to his surroundings, or Churchland's (1986) crab, if we assume that such low animals do not have consciousness, and that their reactions are not fully hardwired in to begin with, but require some learning. There is implicit in the trained neural maps a rich structural semantics in terms of inclusion, overlap, union, and exclusion of concept indicating activation patterns, which mostly may not evolve into consciousness, and therefore cannot be expressed in a language which, also by the language user, is understood as a language, i.e. as a symbolic system.

We cannot by ourselves invent symbols for registering unconscious processes that are reactions to properties of our surroundings, but we could be trained to feel compelled to use certain expressions whenever the unconscious stimulation of our body by these external properties occurs. The training would have to be performed by some external teacher who is consciously aware of these external properties as an observer. He can consciously perceive our neural reactions and has devised the terms for these properties. He would call out the term the moment he registers the activated neuronal pattern. After such a training we would be able to use expressions like a robot does who has been trained to certain input- output correlations. We just would not now why we suddenly use these expressions. They therefore would not really be language for us, though they would be part of the observer's or teacher's language.

In order to create, develop, and improve a language we need consciousness of that about which the language is going to be used, and this is something a robot without consciousness cannot do. To design or develop a language presupposes some knowledge that certain things exist and events take place, such that one can make statements about these, and these things can fulfil the statements, i.e. can make them true. It presupposes the intention to use certain objects as signs, namely as symbols for something else which one is also conscious of. This capacity to develop a language which oneself understands as being a language or a symbol system, might be an essential distinction between human and artificial intelligence. In consequence of being trained within an unconscious state, a robot cannot have any knowledge which is a 'know that', but merely has a 'know how'. Even if he would utter expressions of statements, these would be statements, i.e. expressions of knowledge, only from the standpoint of the designer or user of the robot. An expression can only express knowledge if it is understood as being about something, which hereby is intended as that which makes the statement true or false.

We can imagine a computer run by a program in which a procedure has been programmed that randomly selects a term out of a set of possible terms in correspondence with a certain internal activation pattern, which has come up regularly in connection with a sensory input. If similar inputs occur the same term association is made. Another machine can adapt itself to this association by a learning process. For the outside observer it looks as if these machines create a language all by themselves. In fact, these artificial devices for language creation do not know that here a language is created. They use signs for outside objects via the causal semantics developed internally according to the program and they seem to communicate with each other. However, they do not know that they use a language and communicate by means of this language. Only the designer and observer can understand and describe the situation as one of language use and communication. In doing this he sees this situation in line with situations of human language use and communication, even though the machines miss the intentionality and consciousness about what is at hand.

Non-experienced registrations, i.e. reception, can give rise to a non-perceptive and non-cognitive adaptation to the world in terms of internal states and processes. Such adaptation is not knowledge that this or that is the case, or must be done. These registrations, and the connections among them, are part of our 'know how', abilities, patterns of reaction to natural and social conditions. Non-perceptive and non-cognitive orientation based on merely sensory reception and its effects on neuronal connections, which collectively embody dispositions towards certain activations, are not part of our consciousness. We are only aware of its consequences, namely that we, for example, walk pretty safely along a road, and that we can bicycle (cf. Searle 1983 and 1992 about the Background). There is also no reason to assume that this mostly non-perceptively but merely receptively acquired non-cognitive background, or ability to perform certain procedures, can be made conscious by just somehow lifting it to a conscious level. It can be reconstructed scientifically by a neuro-physiological theory about how a system adjusts to its surrounding. Searle, rightly, points out that our background, although being unconscious, is necessary for using any rules and representations, such as linguistic expressions and pictures, because only via the background these can be interpreted, i.e. be related to the world. This happens in an automatic causal way. Our causal semantics is such a know how, an ability which makes understanding and interpretation possible.

Causal semantics is a pre-condition for employing denotational semantics. It has to be acquired in training. As Wittgenstein has argued, internal rules and representations by themselves can not serve as a medium for relating public rules and representations, for example linguistic expressions, to their cases of application. Such internal mental objects would need interpretation themselves in order to be utilisable in interpretation of public expressions. This means that the mental objects must already be understood themselves as examples of the rule or as examples of application of the expression. Likewise consciously comparing new cases of use with recollections of old cases of use of the expression cannot be the solution for the problem of regulating the correct use of the expression. The recollections of the old cases would have to be judged as correct examples beforehand and for this we would already require the criterion of correct use which we try to construct as similarity with these recollected examples. Rather a causal, unconscious relationship established in learning from cases of publicly approved language use, establishes for each of us the link between previous cases of use of an expression and new cases. Causal semantics for each of us is the condition for the possibility of having a denotational semantics. The origin and the reality of denotational semantics is a public affair. It is in fact social, but it can in principle also be public in the weaker sense that it is simply outside us, such that we can be conscious in perception of the two sides of the relationship of interpretation, language and world. Consciousness is not a condition for causal semantics per se, but it is a condition for the specific denotational structures imprinted on it during the process of learning the public language.

Our ‘know how’ capabilities are not processes governed by rules, though the regularities in their results may be described, in idealisation, by rules. Firstly, learning rules as rules and conscious application of rules require data one is aware of. Secondly, representing these processes by context free rules, taking into account the kind of input and output in form of descriptions, proves to be incomplete because symbolisation of input and output cannot take into account all the specifics of these data and their surroundings which might have causal influences on reception and neuronal effects. Therefore rules constructed by the linguistic researcher could at most be idealisations of the relationships between the results of the processes that in reality take place without being governed by rules. There also will always be situations where no action is possible in accordance to the rules constructed. A small deviance in the preconditions of an operation is already sufficient for the break down of an assumed linguistic algorithm. Our behaviour is rather caused by sub-cognitive processes, which, as far as they remain unconscious, can not be understood in terms of rules, a point made by Searle (1992). Thinking of the background as providing basic orientation in our surroundings and as not being a set of rules, makes neuronal networks to be the only presently available models of acquiring basic orientation in a non-conscious way.

5. Linguistic ability is not a generative system of rules

We now address the question whether, and in how far, language is part of the 'know how' background of knowledge and action, and whether, and in what sense, it is part of cognition and knowledge as that what is conscious and thus part of our mind. I shall argue that both is the case: Our language ability is part of the subcognitive background, while language as a set of rules and representations is the result of linguistic analysis and does not describe our language ability in kind, but merely is about structural properties of the results produced by our language faculty. Hereby a system of rules and representations indirectly suggests some structural properties of our language ability as it must be encoded in our language faculty. However, these structural properties are not the results of hidden rules that would guide our language ability. Rather they are idealising descriptions of globally emerging conscious phenomena that come about by cumulative effects of local micro-processes of quite different kind.

We speak correctly without being able to cite the rules which might have regulated our speech; we use terms with meaning, i.e. concepts, without being able to define these concepts. We, apparently, do not have representations in our mind of external meanings such as concepts in the traditional sense, still advocated in mentalistic positions. We are merely able to present, identify, and imagine examples or cases of use of respective terms, or, even without using terms, we are able to arrange situations in classes, in alignments according to similarity, and in orders of contiguity. This is an ability, a know how, as is rightly stressed by Wittgenstein (1960), Putnam (1981), and Peacocke (1992). A concept is a certain ability, according to these philosophers, not a representation of something common to all the cases of application, like the mentalistic notion of a concept would have it. There are no concepts or representations of meaning to inspect, when we attempt to analyse a concept and define it. Rather we can only scan through what we are consciously aware of, namely real and imagined situations of language use, which are the external proofs or products of our linguistic ability, and which, for us, constitute what linguistic ability means. In our scientific endeavour of linguistic analysis we can only inductively and hypothetically conclude from the examples to some definition of their kind, which is a certain partial representation of a concept in our public language, or in some formal language devised publicly for scientific purposes.

Linguistic ability as a faculty within the language using individual is a brain state (not a mind-state) which stands in a causal relationship to the sets of input by which it has been formed and stabilised in learning processes, and it stands in a potential causal relationship to new data which activate those groups of neurones and structures of groups strongest which have been formed by previous input which was most similar to the new data under the perspectives of measurement available. Language as the individual's language faculty or ability is part of the background of our mind, and not part of our mind, and as such ability our language does not consist of rules and representations in our mind. The background can be thought of as a state of our neural networks, possibly with some innate biases in its architecture, in which the connections between the neurones have got and still get adjusted such that the neural nets come up with certain activation patterns under externally or internally caused input. The potential patterns are directly addressed, i.e. activated, by certain perceptual input or indirectly via other patterns, that got activated first. Input, hereby, is classified and associated with possible context types in which it has been learned. This association happens by activating directly and indirectly the respective patterns of units. Certain results of these subsymbolic processes appear as conscious perception and cognition in the mind. In chapter 4 and 5 it will be discussed what kinds of results have an expression in consciousness.

On the other hand, linguistic research comes up with rules, and conceptual analysis provides at least partial definitions of concepts we use, or rather explicates partially our conceptual ability. In linguistic analysis we find systematisations or structures in the sets of linguistic data, by applying structuralist methods such as syntactic-semantic substitution tests, elimination tests, movement tests, and paraphrase tests. Data are pairs of linguistic utterances with their situations of use, especially the satisfaction situations of the utterances. We have a normative awareness of what counts as correct or incorrect linguistic forms and correct or incorrect linguistic usage. This awareness of correctness is called 'linguistic intuition'. We can also improve our language use by consciously monitoring it in comparing it to normative models, and we can even improve and develop our language further by institutional language planning performed in academies and schools.

These quite different facts, the unconscious workings of linguistic ability on the one hand, and the conscious handling of language on the other, can be reconciled by realising that our linguistic ability belongs to the non-cognitive and unconscious background, and merely our linguistic products and their understanding and interpretation in the world are part of our mind. These are real and imagined situations of the use of expressions, and they are the data of our analysis, rather than some unknown internal conceptual representation, which, after all, does not exist in our mind. The concept indicating activation patterns, encoded in the weights of the connections between neurones as the result of learning, are not conscious and, therefore, are not part of our mind. Only what, at a certain time in our life history, is or can be experienced, and this implies that we can be consciously aware of it, is or can be part of our mind and can be treated rationally in contexts of knowledge, rules, and explicit preferences.

It is unnecessary to assume an unconscious part of our mind in which hidden theories and rules can have a place and where they could be found by introspection or reasoning. No introspection ever has lead to such a part of mind and no convincing arguments have been given to postulate it nevertheless, neither by McGinn 1991 in his argumentation for a hidden consciousness based on the blindsight case, which has been discussed above, nor by Chomsky who thinks that a hidden mind would be the best explanation of our linguistic intuitions. Our theories about the world and our knowledge about language is not something hidden in our mind. To have knowledge means that we deal with coherent sets of general sentences held true, related to each other and to experiences of situations by inferences and interpretation (Quine 1974). We experience, for example, what peace is in the use of the word *peace* in linguistic and practical contexts. In our mind we only can recur to such experiences, not to any hidden rules of application of terms. There is no rule or definition used in the application of the word *peace*. This inaccessibility of that which makes us use the expressions of our language in a correct manner, points to the fact that our linguistic ability belongs to a sub-cognitive background. This implies that language belongs to the mind only as far as we perceive utterances and speech situations, have memories of these, and imagine new ones. All these perceptions and images of linguistic usage are the data we are aware of and which constitute the results or products of our language faculty and what linguistic ability means for us. We can analyse them in comparison with each other, and we can compose constructions out of the parts gained in analysis. Linguistic ability is grounded in the sub-cognitive background of our mind as the semantic, syntactic and phonological faculty, which is a set of dispositions encoded distributively in the connection strengths between neural units. Our conscious awareness of our linguistic ability is nothing but a growing set of data of acceptable language use.

In mentalistic linguistics and philosophy, exemplified in Chomsky's and Fodor's work, it is assumed that there is an unconscious part of our mind in which linguistic rules preside and from where they secretly guide our behaviour. Such a notion of an unconscious mind does not make sense and moreover is totally superfluous for explaining linguistic behaviour, since connectionist network models of important structural properties of the neuronal connectedness in the brain explain the causal individual aspects of language use, while public rules and representations describe the normative aspects. There is no hidden mind in between these two sides of language use, as I shall argue for in the next paragraphs.

A notion of unconscious mind is used in Psychoanalysis, according to which we have hidden beliefs and desires, which can be brought to consciousness by therapeutic communication. However, these and other non explicit beliefs are just dispositions of the brain to get activated under certain kinds of stimulations. In contrast to the hidden linguistic rules and representations of the assumed unconscious mind, the unconscious beliefs and preferences assumed in Psychoanalysis are potential activation patterns which can have conscious expression in the therapeutic dialogue. If the activation is effective enough and inhibition sufficiently diminished, we can experience that we have these beliefs and desires in imagining cases in which they clearly play a role, and in the therapeutic sessions they become finally represented in our mind in natural language formulations. But the same is not true for linguistic rules. We never experience these rules in our data. Rather we only can detect them by ordering the data in systematic linguistic analysis, and we can recall them afterwards, if they have been formulated before, like we can recall any representation, which we have experienced, whether it is a picture or a sentence in some public language.

After having been explicitly formulated, rules can be part of our mind like any other belief or knowledge we have acquired. As such they actually come to our mind in the formulations of linguistic science. What we then have in our mind as knowledge are only representations of rules we find as observable regularities in our data, but these are not rules on which our linguistic ability thrives. Even if rules regulating our ability were formulated in some hidden unknown symbolic medium and would exist in this form, they could not be part of our mind, also not of our potential mind, because our mind can only comprise what we can be aware of. And furthermore, since we do not need such hidden rules for explaining our language ability, there is no need for assuming a hidden mind in which we would have stored such rules in some internal code of representation. In the following paragraphs, I shall elaborate the arguments against the assumption of rules and representations as internally guiding and organising linguistic ability.

The cognitive tradition following Fodor (1976, 1987) assumes a hidden language, Mentalese, somehow similar to logical languages. The arguments have been that learning applicability conditions for lexical items, and learning other linguistic rules presupposes such a mental language in which the rules get stated. It has also been argued that propositional attitudes require a medium of representation for their propositional content, which such mental language provides. The first argument is weakened since the growing success of connectionist and probability based models of language learning, which do not require that linguistic rules are learned. It is not anymore used in Fodor's (1998) position of informational atomism, where concepts are supposed to be atoms no longer defined in a language of thought. The second argument is weak, since implicit propositional attitudes have contents which can be thought of in terms of dispositions to come up with neuronal activation patterns, causally connected via the previous learning processes to objects, events and actions, while explicit current propositional contents are represented by the internalised public language or situational images. Neither for learning language, nor for propositional attitudes an innate mental language is necessary as a means of representation (cf. chap. 4). Therefore such a symbolism is useless, and in consequence there cannot have been any evolutionary pressure towards its development. Rather we can assume some evolutionary selected neural biases for forming circuits of activation patterns (cf. chap. 3) conforming to objects and events that are enhanced and delineated by their occurrence in activity and action routines in which the learning individuals are involved. But this is not a mental or innate language. Any conceptual system and any public language developed in human history will be constructed in parallel with such basic action, event, and object patterns. Piaget, for example, based the logical abilities of individuals on a process of internalising external operations on actions, objects, and situations (Piaget 1953). The capability to use such operations is partly determined by innately programmed neural growth, which enables the young child to learn certain operations at certain developmental stages.

6. Linguistic theory formation is rational reconstruction

Linguistic ability has been reconstructed on a logico-philosophical and on a linguistic level. The dominant brands of theoretical linguistics, the mentalistic Chomsky school and the non-mentalistic logical grammarians following Montague, have reconstructed linguistic ability as an ability to form linguistic expressions according to rules. However, logical and linguistic rule systems do not fit very well to real utterances, and even break down, when they are applied in the analysis of these. For a long time it has been assumed, for example by Chomsky and his followers, that a great many utterances in daily communication are corrupted, because they suffer from intervening performance factors, and that the competence of a speaker, the rule system by itself, would produce perfectly correct utterances if no disturbances would intervene.

More recently, the claim that competence means rule following has been weakened (Matthews 1991): Utterances are not built by following the rules, rather in accordance with the rules. This position leaves open how utterances are built. It merely says that they are constrained by the rules in some (unknown) manner. But still recently Chomsky writes that someone "follows the rules of T", the language, "because that is the way he is constructed" (Chomsky 1995: 36). What rule following means remains unclear, like the status of the "inner language" remains very ambiguous, mainly due to the ambiguity of Chomsky's mind/brain notion, a combination of two ontologically totally different notions that is not understandable. Schnelle (1996: 111), among others, points out that the working of the language faculty in the brain is a property of the implicit memory embodied in the dynamic neural connectivities, and not a property of declarative memory, and only for the latter a symbolically computational model would be adequate. This means that to view the language faculty as adequately described by algorithms of rules is far away from what it really is. Also Smolensky (1988) argues in this line when he points out that cultural knowledge and conscious rule interpretation, which have the properties of public access, reliability, and formality, is the domain of symbolic processes, i.e. is part of the conceptual level, while individual knowledge, skill and intuition are the domains of sub-symbolic or sub-conceptual connectionist processes. He views the level of connectionist modelling as a theoretical construct between the neural and the conceptual or symbolic level. It certainly is sub-cognitive and sub-symbolic, and represents on a very abstract mathematical level some relevant properties of the neuronal network structure of the brain. I do not agree with Smolensky's use of the term 'individual knowledge', because I take individual knowledge to be representational too, and thus part of the cognitive level, the conscious mind. Rather I would replace his term 'individual knowledge' by the term 'ability' or 'know how'. It comprises adaptation to surroundings and language use of others. The connections on the neuronal level, trained in learning processes, implicitly contain dispositions or tendencies towards regular linguistic behaviour, which gets idealised in the form of rules in linguistic analysis and theory formation. Here rule formulations are set up and linguistic behaviour is understood in a rationalising fashion as being guided by these rules, although it is merely conform to the rules to some degree.

According to Chomsky (1995), the mind is a computational state of the brain and the language faculty is part of the brain and brings about a cognitive state, i.e. a mind state, which is the I-language, our internal language. In this model it must be the language faculty, i.e. a brain part, which follows the rules according to a programme that is installed in it, basically innate and adjusted in learning. The installed program then must be part of the mind, as well as all states of execution it can be in. If we take the computer metaphor seriously, the program must be formulated symbolically and the states of its execution as such must be likewise formulated symbolically. And when the program is installed and runs it determines the states of the computer. Unfortunately, here the metaphor breaks down: the mind, or the installed I-language does not run the brain, it rather is a state of the brain and thus the program would have to be hardwired within the computer. Apparently in the hidden mind, brain states and hereby parts of the program appear like they can appear on a monitoring computer screen. However, a screen can be observed and the program can be read. A program is public and not in principle hidden. But the problem for the I-language is that we are unconscious of all this, and cannot become aware of our internal language. And therefore it cannot be in the mind.

If we know anything, then it is something that can come to our mind. It is conscious or it can be made conscious, which simply means that there is a brain disposition, a potential towards activation patterns, that can result in conscious awareness of knowing something, for example, knowing sentences of the public language as products of our language ability. There cannot be any, in principle, unconscious knowledge in the mind. However, like the workings of the language faculty, there are many processes and states in the brain which we are never conscious of. But that is no reason to claim an additional hidden part of the mind for these. They are just in the brain, and that is enough. There is no need to declare certain global brain states and architectures as showing up in the hidden mind or as being the hidden mind. A cognitive state, a mind state which would be our internal language, is totally superfluous and thus does not make sense in a scientific linguistic theory. We are never conscious of any brain state or process itself, as such a brain state or process. Rather we are conscious of some of its products, of which we then say rightly that we have them in our mind.

Giving another twist to the computer metaphor, which fits the facts better, the mind can be seen as a kind of interface between brain and world. On the level of this interface we deal with coarse-grained representations of parts of world, and all our mind phenomena are of the same kind, similar to our perceptions of parts of world, playing a perceivable role on this level. If we stress the computer metaphor somewhat further in this direction, we might be tempted to say that the mind is something like the interface running on a screen or monitor, where the input into the computer and its output are presented to us in the same format and can be worked with together in relation to each other. Mind phenomena are all of this kind: Perceptions and images of situations and of sentences, or other phenomena that are used as signs. They are all of the type of perceptions and motoric procedures applied to these phenomena, of which we are consciously aware. We can use, in our mind, images of sentences and other phenomena to operate on by procedures such as substitution, insertion, movement of forms to other places. These operations happen freely or according to rules. In this way we can consciously control the results of the unconscious brain processes that result into thought products in the form of representations such as series of images of situations or sentences.

On the logical and linguistic level we manipulate sentences according to publicly endorsed schemata, and on the level of experiential concepts we manipulate concepts by addressing sequences of examples which represent the concepts. We integrate new data into these representations of concepts salva stabilitate of their internal similarity degree or the internal coherence by contiguity. In rationalist philosophy and cognitive science on the other hand, it is assumed that input and output get related by means of a rationalising construct of rules which are taken to be applied to the input and result in the output. The steps in between are in principle of the same syntactic format as the format assigned to input and output. In such a model it makes some sense to talk of a hidden mind as an extension of the open mind which contains the same kind of syntactic objects and operations on these. The problem with this view is that it is just a construct of rationalisation which is useful for designing some AI-products, but which is not a model of how our brain works.

I basically agree with Searle's connection thesis (Searle 1992: 132) which says that only what can in principle be conscious can be part of the mind, whereby the mind phenomena are somehow caused by the brain. I think that the term 'caused' is problematic in this context, because it presupposes that we have here a relationship within a single ontological realm. The term 'cause' is here used metaphorically and only expresses that consciousness does not exist independent of the respective brain states that, so to speak, carry it. It might be a better metaphor to say that consciousness is an 'expression' of certain brain states or is 'brought about' by these brain states. The mind is at any moment what is conscious and it can be thought of being extended by the potential mind which consists of the results of brain processes and states that can come into the mind under stimulating conditions. The 'potential mind' is merely a theoretical construct which comprises all which can come into the actual mind. It is not a space in which mind products are located and mental operations take place, since it does not really exist. It is just a logical space of possibility of the products of brain activity that can come up into consciousness. The possibility of these mind products is implicitly contained in the connection strengths between neurones of the brain.

Instead of talking about the brain or the mind as following rules, we can better understand the impact of linguistic rules when we view the system of rules as the scientific representation of public constraints on the linguistic output by way of controlling and possibly correcting the much more generous linguistic products our language faculty comes up with. This view takes rules to be norms of the forms and the use of linguistic products which secure easy and co-ordinated uptake of utterances in communication, as opposed to viewing them as an internal algorithm. Nevertheless, that our linguistic norms are largely so regular that it makes sense to formulate linguistic rules comprising these norms in systematic ways is, besides being induced by the semantic function of language and thus bound to regularities in the world, partly caused by a certain architecture of the neural nets in our brain, which have evolved during evolution. It is also partly caused by a tendency towards global regularity which results from accumulating locally steered processes. The network architecture makes certain operations possible which result in substitution of equivalent expressions in contexts, by which we get categories of syntactic or semantic kind. Likewise substitution makes it possible to understand that parts of an object or situation can be substituted by other parts, and that certain parts can be extended in certain ways, without effecting their role in the whole in which they function. This general ability for substitution, insertion, and hereby also recursion, requires certain architectures of recurrent neural networks. This is not specific to the language faculty but serves also in perception, imagination, and action. Architecture and adjustments of connections to the stimuli from the outside world in situational perception and language learning gives rise to regularities in the output, and these structurings on sets of linguistic output get more systematised and idealised by analysis such that they can be described in terms of linguistic rules. However, rules which register regularities, and are idealisations of these, are merely found in our products, not in a hidden mind.

Our brain does not follow these rules in speaking or interpreting. It does not operate according to these rules. Rather the rules are idealisations of regularities that result globally on the basis of local micro-processes and of regularities in input data. -- Later in this chapter, I shall discuss how normative rules, which are external to the language faculty, can perform the function of constraining linguistic output. Idealised descriptions of linguistic output can be emphasised and become a norm, such that only products that conform to them are accepted as correct. By such normativity the creativity of our language faculty gets constrained in processes of learning from accepted examples.

In the light of new developments in cognitive science, especially in the field of connectionist en probability based models (discussed a.o. in Scha 1990, Bartsch 1996, developed for syntactic-semantic parsing in Bod 1995 and 1998), the competence-performance distinction has to be reconsidered (Scha 1990). The gulf between deficient utterance forms and the products of rule systems has to be understood in a different way. The rule systems are idealisations of the output, not models of our competence of language use, and especially not of our production of utterances; our brain works differently, much more flexible and far less restrictive than the rule systems could do which have been assumed to model competence. Utterances, whether conform or not conform to the rules, are rather the products of our subcognitive linguistic competence. Linguistic capacity is organised such that also so-called mistakes are output, which are proper products of the neural network activation patterns which constitute our competence at a certain time and which can change over time due to new input. On the other hand, our linguistic intuitions formulated in correctness judgements, and the standards presupposed by them, are not primarily provided by our individual internal language competence, as Chomsky (1986) has thought; rather they are idealisations and rationalisations of the linguistic products we find in public normatively endorsed language use. They are normative stipulations as far as speakers are norm authorities, members of subgroups that provide models for other members of the language community (Bartsch 1987). These standards, as standards, are linguistic products of these speakers, which can be consciously controlled by comparing them with products of other norm providing speakers. The standards are reproduced by those that have adapted to the norms, and they have to be reproduced by those who still have to adjust to these speakers in learning. The internal language faculty of speakers is not a separate module or organ; rather it is interspersed with, or even consists of perceptual, motoric, emotional, and cognitive functions. It will be more or less adjusted to the products that conform to the public linguistic norms. In this way speakers have built up a competence which is, by itself, not a set of rules, but still gives the output required within broad margins.

The normatively effective output of subgroups of speakers can be formulated in codified norms, and in more general and deeper ways in terms of rules, in which we can realise, in doing linguistic research, even some universal principles. These universal principles might be due to our general cognitive abilities, based on innate brain architecture, stabilisation processes, and even biases in neural nets due to phylogenetic adaptation to orderings in the world. Additionally, universality is induced in ontogenetic development by our common semantics, which is formed in learning to adjust to our surroundings and in acquiring knowledge about the world and the public language. These learning processes work according to principles of similarity and contiguity in sets of data and lead to stabilised clusterings of data, which represent concepts. A concept is the equivalence class of such a data-cluster or data-set comprised by similarity or held together by contiguity relationships. Two data-clusters or sets are equivalent if the can be united without destabilisation, i.e. while keeping intact their internal similarity measure or their coherence by contiguity relationships. In this way there is a process of concept formation going on (cf. Bartsch 1998), whereby public language use guides the individual learners concept formation.

The biases and established states of neural network connections are not rules or principles themselves, nor can they be understood as algorithmic programs. The neural system, in the state it is in at some stage, even if it is stabilised, is much more flexible and dynamic than any algorithm of symbol manipulating rules can be. It will always be able to produce utterances which are not quite according to the rules formulated by linguistic science. These deviations, when publicly accepted, are one source of language change, whereby the systematic orderings are changed, i.e. the rules we can devise for a language at a certain stage. Still, we can say that we show linguistic behaviour which more or less is conform to the rules that describe it under idealisation, but we cannot uphold that we follow the rules in producing our behaviour. In the light of these considerations and the new developments in connectionist network modelling, the traditional cognitivist positions of Chomsky and Fodor are no more convincing. In order to understand our language ability, we need not postulate an internal language as a set of rules and representations, formulated in some hidden mind language.

When speakers are generating data especially for linguistic analysis, they are normatively aware of their linguistic products in performing correctness judgements. They imagine situations of use of linguistic expressions and linguistic products and judge their correctness by simply testing whether they feel compelled to agree to them or to produce them in imagined linguistic situations. Though this method of testing against linguistic intuition cannot serve to test the correctness of one’s own linguistic ability, because for this it would be totally circular, it is a way of carefully controlling and monitoring data with respect to internal consistency and coherence, and of getting material for linguistic analysis in order to describe linguistic structures and devise linguistic rules, which are then not more than an idealised expression of ones own linguistic ability, as it is at a certain moment, monitored by a normative attitude. One’s own linguistic ability can be tested as correct by comparing its products to those of speakers who are accepted as providing models in the speech community, and if one belongs to these speakers one can only test for conformity with the others.

The rules, which are the result of linguistic analysis, provide a systematisation of a standard, exemplified by the linguistic products of speakers who have become adjusted to the norms exemplified by the models provided by norm providing speakers. The rules have to be correct with respect to these speakers' judgements of correctness about linguistic products. These model-speakers themselves need no rules or explicit application conditions for producing and using their linguistic expressions. However, when speaking in a formal style they can do this in a normative attitude, monitoring their speech and hereby adjusting it consciously to other products of this style. Also then no one needs the rules formulated by linguists; all we need is to focus on previous examples under certain points of view, like pronunciation, choice of register, etc. Rules merely belong to a meta-level, on which linguistic products and their situations of use are the objects of conscious systematising, as it is done by linguists.

Speakers produce their expressions from out their background ability, and they automatically compare their products with those of others for conformity, whereby some provide the models, and others follow. And speakers also naturally make their products such that they conform to their own previous products, if these have not been corrected. This all is not done by conscious comparison of new products to old ones, rather this comparison or adaptation happens automatically by making products such that they fit best with the previously established activation patterns in the neural nets. We can formulate it such: The new product is supported by the activation patterns of the previous examples to which it is similar. And when the new product is perceived, as it is in control, it recurrently triggers these activation patterns and at the same time strengthens the connection strengths between the neurones involved in bringing about the activation patterns. This leads to more or less adaptation to models of correct speech, because the input from these sources gets perceived by means of the neural nets of the learner, where it activates those patterns that have been caused by examples which, according to his measurements, are most similar to the input. These can also be examples he had previously produced himself. They will be adapted somewhat more to the model input, but on the other hand the perception of the model input will already be influenced by the state the neural network is in when receiving the input. We all have come across funny examples in which a language learner repeats his own wrong pronunciations even after having been corrected several times. He does not seem to hear the difference between his own linguistic products and the pronunciation of the model. His perception, i.e. what he hears is the product of his auditive input with the state of the auditive network in which the input effects activation.

The problem of completeness and of correctness of a system of rules for the representation of a language exists then in the same way as it would exist for representing how to bicycle by means of an algorithm of rules, or for representing our co-ordination of vision and motoric activity. Only a trained linguist, language teacher, or student will be able to compare his and others' linguistic products to rules and application conditions. He has to be trained in linguistic analysis and its terminology. And also he, who understands the rules, will not use them to construct linguistic products in real conversation. However, the rules can help for consciously construing some examples or for highlighting the structure of a series of examples. This is equivalent to putting a set of examples in line and continuing the sequence in an exactly specified fashion. What is then learned is not the rule itself, rather the examples and constraints of modifications on these are learned. Rules can help to present examples in an efficient way in a learning process by alignment and opposition. We can learn from the examples which are provided by means of rules, but not from the rules directly. We can learn rules as *know that*, namely that expressions of certain forms are correct and are used in situations of a certain kind, and that certain structures are found in language, but we do not learn rules as a *know how*, as a competence of how to speak. And we never learn directly from rules. If rules are given a role in learning it can only be an indirect one, by providing examples in an ordered manner and emphasising certain aspects of these. And these examples then are used as the data for learning.

A system of rules is correct if it produces the right output, especially the right behaviour, and behaviour might furtheron be judged as correct by stating that it conforms to the rules. But there might never be a complete system of linguistic rules. To view language, along with the tradition of the Chomsky school, as a device for generating a potentially infinite set of sentences from a (small) finite set of syntactic-morphological rules and a closed lexicon suggests a wrong way of looking at linguistic ability and activity, namely seeing it as a program or algorithm for symbol manipulation on a computer. New insights provided by modelling sub-symbolic processing of input and of learning makes it reasonable to discard the old picture of language competence as a system of rules. However, sub-symbolic linguistic capacity is not something that stands by itself. It is only possible, and can arise merely in learning processes in which a public language is learned, and this language consists of public products, utterances in their semantic and pragmatic context, which are, for a good deal, systematically related to each other, due to requirements of efficiency and orderliness, which make it possible to capture these products, to some extent, by linguistic rule systems. The creation of norms is not a random business, rather norms are worse or better, depending on how far they fit with what is already there. New forms will preferably be modifications of old ones, formed by analogy and continuing what is already created, in a regular way. Analogies are created by substitution operations. These preferences concerning the systematicity of norms gives rise to also deeper linguistic rules that themselves can be understood as systematisations over the set of norms (Bartsch 1987). The preferences, of course, are due to our innate general cognitive ability of adaptation of our brain and our behaviour to our surroundings and tasks and to the regularities in our surroundings which are relevant for us and which also must have supported the formation of certain cognitive abilities during evolution.

When we construe a rule system that should reproduce, or even 'explain' a set of utterances, and permit us to make new ones of the same kind, we are performing on the level of consciousness, and in the degree of discreteness achieved there, an analysis of existing or especially elicited linguistic products. Using these rules, we would be able to form consciously new or old utterances by performing a synthesis based on the rules which were devised in this research activity. This new production of utterances on a meta-linguistic level is quite different from how the utterances were produced in the first place, on the sub-conscious level. It is like replacing a natural process of growing flowers by cutting the results, the flowers, into well discernible parts and putting them together again in order to reconstruct the old flowers, or getting new flowers by new combinations of the parts. Analysing the flowers into parts, classifying the parts and putting them together according to rules declared over the classes achieved by analysis, gives new results, new flowers. But this is not what growing flowers means. It is a method of producing artificial flowers. Nevertheless, we might learn a lot about flowers under the taxonomic point of view, like we learn a lot about sentences in linguistic analysis. But we do not learn everything about how flowers grow, and certainly not how to grow flowers. For that, we have to do genetic studies, and studies about the influence of the surroundings, soil, other plants, animals, and the weather.

Likewise, our philosophical and linguistic reconstructions of linguistic usage and interpretation are based on systematically observing our linguistic behaviour, including our judgements of correctness of language use in inferences. We make assumptions about kinds of orderings on the data and the guiding principles of language use. Additionally, mathematical and logical tools can be used, as they are available at the time of research, or as they can be developed for these purposes. Language use is embedded in action and interaction, and in how we operate on our surroundings and understand and describe our world. A theory of language use therefore takes into account the built-up of actions and situations and how we represent parts of world. Use also shapes the linguistic products which are used, and thus can explain some of the linguistic forms and their semantics, especially a basic logical compositionality in syntax and semantics. The main profit of this kind of linguistic analysis focused on language use is to gain insight in the relationship between language and the matters in which it serves. On the basis of this relationship, we are able to construct formal languages for certain purposes. It is here where linguistic analysis generally is useful.

Theoretical Linguistics provides a rational reconstruction of linguistic behaviour, including its products, the utterances, and their production. Linguistic analysis uses certain kinds of operations and relationships, such as forming substitution classes in context-types, testing for dependencies, and paraphrasing. These first two methods are also applied in the analysis of the situations, actions, and objects that satisfy expressions. There is no way to observe in an interesting way the neural workings steering linguistic behaviour. Observing which areas of the brain are activated in which tasks and making contrast analyses by contrasting tasks or input data and contrasting activated neuronal patterns in a parallel fashion, or finding out in which way connections are strengthened by forming stronger synapses between the axon and dendrite fibres involved, in a biochemical way for short term memory, and in an anatomical way for long term memory, all these recent research results are of greatest interest to neurologists or neurophysiologists. For the philosopher, the linguist, and the cognitive psychologist, they show that the connectionist cognitive model is in principle the right one. But all these findings do nothing for linguistics itself, except that these general cognitive conditions suggests to take into account flexibility in language use for the construction of grammars, and to take the dynamics of concept formation and grammar formation into account for the explanation of language change. There must be structural parallels on a certain global level of description.

Introspection into our cognitive functioning is impossible. What is often called introspection is intuition about ones dispositions towards linguistic behaviour and correctness judgements. *Intuition* is awareness of certain dispositions, under conditions one is not fully aware of, but which are nevertheless operative in experienced or imagined situations of use one is aware of. These dispositions answer to conditions of use of expressions, which are partly found out in semantic and pragmatic analysis and can be seen as theoretical entities in semantic theory. Application conditions for expressions are such theoretical entities. Intuition about regularities in application of expressions and reactions towards them requires that one becomes aware of both, and this requires some training as a language researcher, and of course, consciousness about using language. Consciousness about language use is required for deliberately creating language, for designing a language, for scientific language research, and for making normative statements about language use, especially judging its correctness and devising rules, but it is not required for using language itself. Suitably programmed computers can use and analyse language. Their use of language is based on algorithms programmed in, or also on a learning process, where not rules are learned, but behaviour is induced that conforms to the rules, without being produced by the rules. Also detecting syntactic structures in sets of data and parsing new data on the basis of the structures detected so far can be done by computing probability based procedures. Such automatic analysis, for example in Bod 1998, does not require consciousness. However, the design of such programs does. Without consciousness, one would not even be aware of the fact that one uses or analyses a language.

Linguistic intuitions are in the first place replays of previous language use, and in this replay one becomes aware of the use and its conditions, and on the basis of this, one imagines new situations of language use, especially what one would say appropriately in this or that situation. These are our linguistic intuitions. Knowledge about language is only possible when its symbolic function has been understood and then some language analysis has been performed. Before, there is merely knowing how to use language. We do not consciously identify the conditions of application of an expression. Many possible conditions might be part of the previous situations of linguistic use, some of which we are consciously aware of and others not. In the analysis of our linguistic intuitions we rely on those linguistic products and situations of use of which we have been aware, or we make new combinations of parts of these in our imagination. Mistakes of rationalisation afterwards, i.e. after language use, especially in linguistic theory formation can bedevil our intuitions and give rise to wrong linguistic or logical rule reconstructions, which work in many instances by providing appropriate output, but at other occasions they might fail. Then we are inclined to speak of exceptions to the rule. The alternative would be to do away with the rule altogether and try to find another one. Norm-formulations are based on such intuitions and reconstructions by inventing rules. These norms then may get into conflict with language use on some occasions, which are thereby judged to be incorrect linguistic usage. For a descriptive attitude, it is rather wise to take into account the role of hidden factors and the fact that we never consciously realise all factors involved. Therefore, probability statements about cases of application, which are continuously revisable in the light of new data, are a good second choice solution instead of formulating fixed applicability conditions and rules.

We have explicated the difference between the two ways of looking at language, namely as a capacity or language faculty that steers our linguistic behaviour from out the sub-cognitive background, and as sets of rules and representations, which we come up with in linguistic analysis and which we also can, if we want to, deliberately follow by consciously following rules presented before us, and thus constructing new examples. Likewise we can construe programs that follow the rules. These are ways to implement linguistic norms.

The opposition between causal and intentional-denotational semantics treated in the first section, is a special aspect of the opposition between linguistic ability and linguistic normativity. Intentional-denotational semantics is a product of linguistic and situational analysis and of the normative attitude with respect to its results, typical for human language users. We have described linguistic ability as an ability to language use and, on the other hand, we have described the ability to linguistic analysis, to constrain deliberate linguistic use, to formulate and devise rules, and to consciously follow rules. The results of the scientific activity are implemented in the automatic execution of rules in traditional cognitive systems. We have pointed out that causal semantics, as it is established in our neural nets, belongs to our linguistic faculty and is also the essential background for our conscious denotational intentional semantics which explicates the relationship between language and world. The experience that linguistic expressions have a denotational semantics is the conscious result of our causal semantic processes. Our experience of what expressions denote is partially analysed and formulated in sets of semantic rules which we find in denotational and structural semantics. These rules we could then follow consciously, if we would like to. But, as Wittgenstein has pointed out, all such rule following presupposes the basic language ability to which we are trained. By this training process our causal semantics is established, and only if we are geared to the world in this way, can we interpret any representations or any rules. Such interpretation would be a prerequisite for consciously following of a rule. By training we do not achieve the ability to follow a rule unconsciously, rather we are trained to regular behaviour and to performing tasks which are, to a high degree, conform to the rule which describes the regularity in the behaviour or tasks in rational reconstruction.

7. Linguistic structuring is flexible and is not fixed by a single set of rules

Since language use of individuals is steered from out their neural basis and external factors, we can always expect that language users produce new kinds of data which do not quite fit with what linguists have figured out about previous language use, when ordering masses of data under methodological principles, and by tools we are aware of, such as substitution, movement, paraphrase and elimination tests. Even more, language use can be expected to be deviant in comparison with theoretical grammars which linguists have constructed on the basis of their own linguistic intuitions about correctness of linguistic forms. Deviant data are not just accidents of corrupted speech, but instances that in principle are in agreement with the physiological neural network organisation of linguistic ability at a certain stage of development. What looks unsystematic under one point of view can be systematic under another. Clashes between assumed norms and some particular language use will be resolved one way or the other, largely depending on social factors. Furthermore, what we have experienced in previous language use does not necessarily lead to a certain unique ordering and no other. There can be alternatives, i.e. alternative structurings on the sets of data. In principle all the alternative structurings may be taken into account, because they all can give rise to subsequent linguistic behaviour as continuation of one of these structures.

There are different ways of continuing similarity sets of old expressions by new ones, parallel with our semantic structuring of the growing set of experienced situations into similarity sets. This flexibility corresponds to different possibilities of activation patterns coming up under slightly different contextual conditions. They are delineated by contrasts on our neuronal networks, which continuously get adapted to accepted new input. There may be a flexible choice between several potential activation patterns built up in the course of learning in slightly different contexts, into which the new data fit more or less. The choice may depend on contextual factors and perspectives, we are even not all aware of. Examples of such alternative structurings, which do not exclude each other, are the different possibilities of syntactic coherence of prepositions with their neighbourhoods: the phrase *sit on the bench* can be structured by the constituents [[*sit on*][*the bench*]] or [[*sit*][*on the bench*]]. They come up by different ways of substituting and contrasting expressions. The first gives rise to new constructions like *the bench I sit on* parallel with *the apple I eat*. The second gives rise to constructions like *On the bench, I sit*, parallel with *This apple, I eat*. We have, in principle, two options in linguistic analysis: either admitting flexibility in rule formulation and admitting several rules that partly capture the same input, or rigidification, which allows only one rule and hereby one structure. The choice for algorithms means rigidification, which creates a wide gulf between natural and formal languages.

In the tradition of Generative Grammar always one of the possible structurings was chosen as the right one, while the others were discarded as being wrong. A lot of spurious argumentation was devoted on showing that a certain structure or rule was the right one within the system of rules that was assumed to be the right one, too. Likewise, the NP-VP, or Subject-Predicate dichotomy was declared to be the right universal structure of sentences, although a structure with the verb as head and flexibly arranged NPs are much more adequate for languages with a rich case system or adpositions (i.e. prepositions or postpositions) for expressing grammatical roles. The flexible positions of sentence constituent NP's were treated as results of additional movement transformations applied to the basic clause structure. In alternative theories of grammar such as Dependency Grammars and Flexible Categorial Grammars (Bartsch 1995) the flexibility of the NP-positions has been taken into account right away in formulating one rule-schema which generates all positions for an NP such that there was no need of transformations. The free syntactic order is much more in agreement with the free semantic order of the participants with respect to the action they take part in. Secondarily there are then some soft constraints for word order due to the preference for prototypical sentence forms and due to pragmatic and contextual factors.

It is advisable for theoretical linguistics to admit organising linguistic data in a flexible way, i.e. admit growing sets of data and alternative structures erected on them. Such flexibility follows the flexibility that can be expected from the context dependency of sentence forms and from the global and distributed organisation of activation patterns on the neural-physiological basis. The basic principles of substitution and analogy, which also must be encoded in the architecture of neural networks (cf. chap. 3), admit different structurings on the sets of data. Linguistics has to take seriously the linguistic data available in language use and develop theories that explicitly take care of these. It can profit from insights in neurological models by being open to kinds of flexibility, which have to be integrated into cognitive theory and linguistic theory. Linguistics should not push aside flexibility as a matter of performance, or even as performance errors. Especially systematic errors or deviancies have been part of the empirical basis of linguistic theories, since they indicate properties of the internal organisation of our language faculty. Flexibility of description has been inherent in data-based structuralism. Traditional descriptive linguistics, in the European and American versions of structuralism, has always been an endeavour of structuring sets of data. And there are now theories and programs of data-oriented parsing developed by Van der Berg, Bod, and Scha (1994) and Bod (1995, 1998), which combine structuralist methods with probability measures on sets of data.

As has been pointed out, traditional theoretical linguistics, assuming that there is a single system of representations and rules in the human mind as his internal language, has striven for uniqueness of a given structural analysis, and much of argumentation has been devoted to show that a certain structural description must be correct, while another is not. The data used as basis of the analyses were founded on the normative intuitions of competent speakers of a language. The intuitions had to be intuitions of natives speakers, especially of academic speakers, in formal situations of testing while the language use outside these test situations and language use of non-native speakers and speakers less exposed to the norm of the official standardised language shows more flexibility and leaves more room for the possibilities which are inherent in our natural linguistic abilities. The goal was to systematise the data by a unique set of most general and deep rules, of which it has been claimed by the Chomskyan school that they describe the language capacity which is part of the mind of speakers. I have argued, in agreement with Searle (1992), that the latter is not correct. There is no hidden mind as the place of deep rules. The mind rather is all that is in principle conscious, at one time or other. We may adhere consciously to certain norms when speaking, but not to deep systematisations of sets of norms. Systematising is an on-going process which is the result of unconscious brain activity of the language user on the one side, and it is the result of conscious operations of the linguist, such as substitution, insertion, and generalising over detected syntactic and semantic structures, on the other. Flexibility is a fact of language and language ability which is not compatible with the view that a language should be a unique and fixed set of rules and representations.

The flexible systematisations in structuring growing sets of linguistic data are an intermediate level between the flexible neural competence basis and the idealised and abstract linguistic theory. This theory is not so much a theory of linguistic competence as it is a product of rationalisation of conscious linguistic behaviour produced in eliciting intuitions of linguistic correctness in speakers well-trained in the official standardised language. The results of theoretical linguistics are impressive products of our will to rationalise, i.e. to explain and understand, and even guide, our linguistic behaviour by rules. It is part of the century-old endeavour to understand man as the rational animal. This attitude is inherent in rationalising all our behaviour, especially our linguistic behaviour. It always comes down to a very selective argumentation and explanation afterwards. This rationalising strategy repeatedly breaks down when used predictively. Also grammars never will be predictively adequate of natural language use. Rather they at most will function in a normative way for language users in formal contexts, and in a productive way for technical purposes in systems of traditional artificial intelligence. In this connection Schnelle (1996: 114) argues, extending his Leibniz-interpretation, that the language faculty itself as a brain function, "narrowly related to the general theory of cognitive units (i.e. monadology)", is describable by "numerical methods (the infinitesimal calculus as a basis of a dynamic theory)", while "symbolic calculi (formal logic) are necessary for external control and distinct reasoning". The latter serve as "external tools of cognition" and "discursive control of argument and deliberation". This perspective on the role of grammars gives a good picture of the use they can have for us.

Linguistic theories, which propose grammars that are unique sets of rules, are never descriptively adequate with respect to language use, and even farer away from explanatory adequacy in the sense of explaining how language is acquired and produced. They always are partial and fragmentary. Those grammarians that took into account our intuitions of inferences in natural language produced fragments of natural language grammar, for example Montague Grammars, which are semantically interesting. These are logical grammars, which formalise the relationship between linguistic expressions and our models of worlds. Grammars of natural languages, proposed by theoretical linguists and semanticists, can be understood as normative grammars presented in a very general, or deep formalisation. In them, well-formedness is defined by rules designed according to what we, as educated speakers, accept as linguistic forms when we look at our language, and what we rationally accept as inferences, given the information we have about the world. But these ways of looking at language do not very well represent language in all its practical purposes. They might become attractive under the view of simplicity and elegance of theories which are the best systematisations of normative grammars of such idealised linguistic means that serve for straight forward information transfer. Because they cannot fully take into account the empirical basis, i.e. language use, they give a merely partial outlook on language, namely a view on a normalised grammatical core which has in fact a lot of linguistic activity around it which is not captured by the rules proposed. These grammars are mathematical constructs of idealised or formal languages of information transfer, rather than theories about language, as it is used by individual speakers in a speech community.

Language is situated between three constraining factors, which themselves are not quite stable. Firstly there are the biological foundations of behaviour, perception, and cognition, and the physiological-neural registration within the neuronal connectivity in learning processes. Secondly there is the world around us which has to be represented by language in its relevant aspects. And thirdly there is the normativity of the language community, which is motivated by the need to conform more or less to the same standardised linguistic behaviour over very large groups of people across space and time. Linguistic usage is influenced from these three sides. The first side, as far as it consists of the neural network effects of learning processes, permits much flexibility, and more or less stability.

The second side, consisting of relevant regularities in the world, guides the conceptual built-up and basic syntax, reflecting the functionality of language with respect to representing situations we are engaged in. The third, the normative side, places constraints on what is acceptable. The first permanently gets adapted to the second and the third. But also the world for us is not stable. It gets understood by means of a conceptual capacity which is always developing further. Also our norms change under influence of the language abilities of individuals, when some of them present new models of linguistic usage, which get adopted in subgroups and larger groups. Not everything is allowed that is possible according to the structurings we lay on the growing set of data. Much is not continued, or does not come under attention at a certain stage or never. Much linguistic output, created by individuals or by neural networks that after training behave quite agreeable, is not accepted. Normativity, in judging phonetic, syntactic, semantic, and pragmatic appropriateness, cuts down on the possibilities which at each point in general and individual language development are available to language users. On the other hand, the admittance of creative language use at certain occasions and in certain contexts, for example in poetry or in the speech of subgroup leaders that become models of behaviour, frees this sub-cognitively available potential from the normative pressure enacted in normal circumstances, and gives rise to language change.

8. Learning from rules is possible via the construction of examples

An interesting question is how we can learn from our own mind's designs, such as rule formulations and norm codifications, and how we can teach others by confronting them with linguistically formulated rules. The problem is how norm formulations, explicit rules, and designs can have an influence on neural coding, and how norms of language can cause change of linguistic behaviour. Norms can have such influence, because they are exemplified in normative linguistic behaviour, which enforces linguistic behaviour via a learning process. Normative force influences behaviour by correction, by giving models or examples of accepted speech, and by providing training. Persons who issue, promote and enforce norms, offer models for correct linguistic usage. They provide data, i.e. linguistic products and behaviour, according to their own neural background, and others adjust via registering these data and being trained to accommodate their linguistic usage to these models. Even without models provided by other people, a person can construct examples following a rule step by step, and then he can learn from the examples, and hereby gain the ability for coming up with further similar examples without using the explicit rule.

An individual who is compentent in the language in which the rule is formulated, understands the rule formulation. In the process of understanding the rule formulation, the rule content is construed as the rule's satisfaction conditions. This happens on the situational episodic maps via the conceptual indicators on the conceptual maps which determine activation on the sensorimotor maps. This results in imagining the possible satisfaction situations of the rule, and thus in selecting the relevant features or aspects in the perception of the data given as positive examples for the rule. Likewise it facilitates seeing why the data of negative contrast are not conform to the rule. This control of the rule against the data presented speeds up learning the praxis of acting conform to the rule.

There can even be a creation of norms from the scratch. Whole languages have been designed and then learned, for example Esperanto, or formal languages like logical ones and programming languages. Large or small sets of newly designed norms have been constructed theoretically, by conscious deliberation. For designing a language we need imagination, based on analytic awareness and theoretical understanding of the built up of one or more languages. We have to imagine new combinations and new forms in analogy with linguistic experiences we have had. A design is a product of imagination, guided by previous experiences and rationality with respect to certain preferences and goals. How can then, furtheron, behaviour be produced which follows a design? This seems a difficult question for everybody who thinks that all our behaviour is the result of purely bodily states and processes, without assuming guiding actions from out the conscious level itself. An answer to this question may be the following:

Like we can bake a cake following a recipe, we can make linguistic products also by consciously following linguistic rules step by step. These are external operations. We then have externally built these products and view them as examples of our rules, i.e. structured according to the rules. Making designs and following designs, especially rules, is partially a conscious process. Though we might not know how we come up with certain aspects of the design, these aspects have to be conscious because we have to control and evaluate them with respect to their agreement with the whole design and their functionality or aesthetic quality. Because of their very form and structural properties the structures that are experienced, permit to make overt projections, modifications, and continuations, which lead to new products that may be evaluated. We can construct products consciously following a design more or less freely. In this activity of making examples by construction and then perceiving them, the neuronal assemblies and processes that support these examples best, become activated and modified according to the input examples. The neuronal pattern which is strongest activated by the constructed examples as input, supports these and every further extension of the series of examples. It will be strengthened and modified in order to fit better to the incoming data that were constructed following the design or rule. Now examples can be produced that are no more constructed following the design or rule. Nevertheless, they conform to it to a higher or lower degree. They can be checked in control by generating descriptions of the examples and finding how and whether they fit with the general rule formulation by seeing whether the example descriptions can be derived from the rule formulations. The rigid process of construction is replaced by a soft, flexible process of performing in a reasonable agreement with the design or rule.

Also imagination mediates between rules and behaviour. In imagination, internal construction of examples according to designed rules can be performed, and we experience these internal products like external ones and can learn from them. Such imagination can be exemplified additionally in externally realised products which also can be perceived. The design is realised by constructing and imagining examples following the design step by step, and then we adjust our behaviour to these examples. Thus the consciously constructed examples serve as input into a neural base which hereby learns how to produce these and similar examples directly. Thus, the products that have been constructed consciously by applying rules step by step, can be learned by the same underlying neural system, which has made these constructions and additional modifications of these, from which we then consciously have made selections and new combinations in our design. And in learning the products which are made according to the design, the system acquires a new ability, or an extension of an old one. This is the only way in which a mind that is thought of as merely being emerging experience of certain large structures of neural processes and states, can on its part influence what happens on its neural basis.

The constructive ability of the mind, as far as it is an ability to design by way of selecting from modifications of products that associatively emerge from the neural basis, seems to be an ability of the neural basis to monitor and control itself and to improve itself. This control and improvement is achieved by operating on one's conscious, i.e. perceived or imagined, products, namely on expressions, symbols, pictures, perceptions, and imaginations, and by playing back the results in receiving them as new input data. Also associated products, which only become possible by certain additional contextual input, can be selected and in design strengthened in a conscious manner. This enables an individual to do more than merely react on what it has received as input from the outside. It can transform its own output products by making new analyses, selections and syntheses of these on the conscious level, and it can use these newly constructed products as new input. Learning this new input can further the individual's own abilities in strengthening weak and only contextually available associative activation patterns. New concepts can be formed in this manner and put to use in planning, designing, or analysing new situations. Consciousness is involved in being aware of the products and behaviour which become, in consequence, analysed in different settings, and consciousness is involved in constructing new products and new behaviour step by step, by selection of parts and new compositions according to rules formed or adopted consciously, or according to structural patterns recognised and laid out consciously.

When our brain expresses certain results on the level of consciousness, it can, so to speak, see its own products on a par with perceptions of things from the outside world, and it can deal with its products like it deals in acting on outside objects and situations it has perceived. It gets its products, especially behaviour, represented in the same manner and degree of coarseness as perceived objects and situations are represented, and thus it can deal with them on this level, the intentional level. On this level of representation and understanding we can do the relevant substitutions of patterns within patterns, smaller parts within larger parts, getting new objects, new situations, and new sentences by modifying old ones via substitution and analogy. The products that appear in consciousness can be analysed, rebuilt, and changed like we do with things and actions in the world. Of course, our brain does this work; in a strict sense our mind does nothing at all because it has no brain-independent existence. But the brain can perceive some of its products in the mind, on the level of conscious expression. This happens in the style in which it experiences outside objects and thus can compare them with these and evaluate them. The comparison must be a recurrent play-back, whereby examples in imagination, constructed ones and simply imagined ones, are treated like perceptions of situations and objects. Substitution performed in the mind must be based on associations between parts of products with other parts that fit into the same context types. For example, in the situation in which we perceive or imagine a cat on the mat, the cat can be replaced by a dog, or a parcel, or a toy car, whatever gets mostly activated. All these fit into the context 'on the mat' by having an appropriate size with respect to the mat. By such replacement we construct a new situation, and likewise a new sentence: From *the cat is on the mat* we get *the parcel is on the mat,* etc.

To begin with, only properties in the world which affect us causally such that we can perceive them, can be part of our knowledge and thus be captured conceptually. Therefore also all our linguistic theory building is about external objects, which we can perceive, i.e. which can come into our mind. These are real or imagined linguistic utterances, examples produced on the basis of our intuitions, and their cases of use with respect to satisfaction situations and speech situations. As objects grasped in our mind they can be observed and seen in relation to each other. Our theory of language is a theory of the linguistic products, not of the production. We are conscious only of the products, not of the production process; and our mind, when we try to analyse it, is found to contain products, perceptions of linguistic utterances and situations, or imaginations of these. A theory of mind has to be product oriented; the so-called mind-processes are only postulated but cannot be experienced in the mind. They therefore do not belong to our mind. They, at most, are theoretical constructs of global properties of subcognitive brain processes. These constructs serve as steps in a rational reconstruction of a process that leads to the overt mind products. However, the constructs are superfluous when we are able to directly describe the brain processes that lead to the products. They could only be part of the mind if they would be perceptions, memories of perceptions, or imaginations of behavioural data, or processes which are sequences of conscious steps in overt behaviour, like, for example following an algorithm step by step. These steps can be seen and controlled because they are noted down in a series of products in imagination, or on paper, for example as written expressions, which are visually perceived.

Although our mind is somehow emerging from the brain, mind and brain are ontologically totally different. We have to distinguish the reception of data resulting in activation patterns in neural networks, as subcognitive processes, from experientially and linguistically achieved information, the latter being conscious knowledge. The first do not belong to the mind, rather only to the brain phenomena. Some of them have an expression on the mind level. Situational or linguistic triggering from outside is conscious as far as it is perceived, and not merely received by the organism. In perception it is part of the mind. The basic and prominent mode of the mind is perception, and furtheron imagination. We can perform conscious operations on these mind phenomena. Internal activation and internal triggering caused by the external sources, or also generated by internal stimuli, need not be part of the mind. Its products, though, appear as perceived or imagined, as utterances or other behaviour, in our minds. They are expressed in consciousness and thereby they are part of the mind. Products that come to mind may also be generated by subcognitive activity for which we see no external causes, though they may have causes nevertheless, even if these do not come up in our mind. In those cases it seems accidental that such a product comes up in our mind, and it is taken as an expression of freedom of the mind to form such products or entertain such ideas. Always, if we think about some imagined or abstract topic, we deal with products, namely with unspoken linguistic expressions, with experiences or imaginations of situations, and with series of modifications of utterances and of situations in our imagination (cf. also chapter 4).

Language use can be learned just by training, and it does not require consciousness as long as the semantics of such language use is merely a causal one. But consciousness is required for intentional denotational semantics, i.e. for being aware of the relationships between language and world, as it is required for deliberate change of a language, for deliberate change of knowledge, and with this of theoretical concepts, for explication and setting of norms and codifying these, and for all linguistic analysis and rule formation. In all these cases we deal with the linguistic products or data available, and form new ones according to rules designed. An unconscious system of artificial intelligence never can know that it uses language and never can design a system of language and language use, or for that matter a system of artificial intelligence itself. Even if a computer is programmed to create a simple language for perceived patterns, he just relates patterns of one domain with those from another domain without knowing that on the one side he refers to things which he uses as signs for the objects from the other domain. Because he has no conscious experience of both domains he can not intend that one kind of thing must represent another kind of thing. He thus has no understanding of what he does. The intentional act of symbolising something by something else can only take place in a conscious mind.

The level of consciousness is the perceptive and cognitive level on which we also formulate rules. When the rules get described and procedurally implemented on a computer, this implementation amounts to a rule-driven extension or reduplication of operations designed and performed on the cognitive level by the computer scientist and programmer, which finally have been transformed in series of smallest steps. The construction of new products according to rules can be done by a machine and program designed up to smallest steps of constructive action, conform to the physical basis used in digital computing. We then deal with an artificial system as a product of the mind of the designer, whereby it does not matter who runs the execution of the design, the mind or the computer. The latter will do this in a more perfect way than the human mind can do. Traditional Artificial Intelligence is a truly artificial intelligence, which is a design product of conscious human design activity, an activity in accordance with principles of rationality related to certain goals which the programs have to achieve. This design product can be seen as an extension of the human mind, merely working on its superficial level according to the kind of operations that are the conscious operations performed on the mind phenomena, including operations on known objects, actions, and situations. But this design product is not at all a model of the workings of the basic human intelligence, which lies on a sub-conscious, sub-cognitive level of brain activity. The design product has an artificial intelligence which humans have developed by rationally reconstructing their own overt behaviour, task performance and products, by means of rule systems formulated on a cognitive level, i.e. on the mind level. It presupposes consciousness in design and setting it up, such that it conforms to all kinds of normative notions, semantic correctness, goal directedness and goal adequacy, true beliefs and preferences, which all require in their application a conscious awareness of the data to which they are applied. All this suggests that an artificial intelligence, as long as it is unconscious of everything, never could by itself create another artificial intelligence by observing and analysing its own products, simply because it is not in the position to observe these. Observation implies conscious understanding on the intentional level in its manner and coarseness provided by our language's conceptual apparatus. This is not only sensory reception transmitted onto the neural networks where it activates strongest those parts that previously have been built up through activation by similar observational data. Rather it is having an expression of such activation in form of a conscious experience as the result of the received input. Of course, an artificial system can also be designed to take up its own products as data and do certain things with them, for example according to certain conditions of fitness to certain goals. But this presupposes that a conscious being, a human mind, has been one step ahead such that he has designed the system to perform in this way.

9. Consciousness and intentionality are prerequisites for cognition and knowledge

One of the metaphors about the mind discussed above sees it as the monitor or screen of the brain, or as the conscious interface between the brain on the one side, and the natural and social world on the other, whereby the mind is structured in conformity to the middle-sized objects and events, to which our behaviour is adjusted and among which all our dealings and activities take place. Another metaphor is the one that says that the mind is an 'expression' of certain brain states. All these metaphors are insufficient for really understanding the mind-brain relationship. Everybody nowadays seems to agree about the fact that there is no mind without a functioning brain. But to formulate this relationship as such a supervenience relationship does not add anything to understanding it. Even to speak of a relationship is already problematic, because the mind does not exist independently from the brain, which rather suggests to speak of an internal relationship between mind and brain. It explains why some authors, for example Chomsky, writes about the mind/brain, which even is more obscure than that it clarifies anything.

In the philosophical discussion arguments have been given against the possibility of reducing consciousness to brain states. The arguments are based on the totally different ontology of both, or they are based on the self-reference, which occurs because all knowledge, also theoretical knowledge about brain states and phenomenological knowledge about our mind is developed in our mind, and is therefore itself a mind phenomenon. Brain states are scientific objects outside our mind, but as scientifically known objects they are at the same time in our mind, and are known in the way our mind can know something which is outside itself. On the other hand my mind knows by self-experience what is on my mind. Here is an analogy: We can only talk about objects in space and their spatial relationships to each other and cannot say what space itself is in terms of speaking about spatial objects. Likewise, we can only talk about mind phenomena, including objects of scientific knowledge, and we can talk about their relationships to each other, but we can not say, in terms of mind phenomena of outside things, such as brain states as known objects, what the mind is in which knowledge of these objects is situated. Without spatial objects space seems to be nothing. We cannot say more than that space is the condition of the possibility of spatial objects. Likewise mind is nothing without mind phenomena, i.e. consciousness about objects, feelings, etc. We cannot say what mind is except that mind is the condition of the possibility of mind phenomena and all our knowledge about them, including knowledge about the world. But that does not explain what consciousness or what the mind is. There is always an epistemological primacy of the mind as the condition of all knowledge.

There is no explanation possible about what the transformation from brain states to consciousness is like, in terms of the kind of scientific and philosophical explanation we are capable of (McGinn 1993). This is the case because in doing so we would need to stand outside our mind in order to see and to describe its relationship to the brain. McGinn's argument is that our mind is 'cognitively closed' to the relationship between mind states and brain states. This means that our concept-forming procedures cannot be extended to grasp this relationship, because introspection cannot look into the brain to see the brain properties of the relationship and neurophysiological research cannot see the mind properties of the relationship in the brain. Thus there can never be seen more than a correlation between both aspects of the relationship, without any understanding of the relationship itself. Flanagan (1991) does not accept that the cognitive closure is one of principle; he thinks it is possible that there might be developed a theory that explains the relationship between mind states and corresponding brain states, although it is not based on really seeing the pairs of states that form the relationship. According to him, such a theory must be based on abduction, i.e. it must be a theory that provides the best explanation. However, this does not seem possible as long as the notion of explanation used in this context is not understandable. It would only be understandable if the steps from the brain part of the correlation to the mind part could be formulated in detail. And that is not imaginable, because there will always be the point where we get stuck with just a correlation between physical and mental phenomena.

Mind, the whole of conscious phenomena, is epistemologically primary to any statement of the correlation between states of consciousness and brain states. The latter only can be identified as states evoking consciousness, by correlating them with conscious states. Since not all brain states have an expression in consciousness there must be some observable property in consciousness creating states to which we can attribute the generation of consciousness. Such a property could only be identified by correlations to a state of consciousness. But a correlation by itself is not a reduction or explanation by which one could understand the correlation. Furthermore, it cannot be excluded that even different types of brain states might be correlated to about the same type of state of consciousness in different stages of the history of a learning brain, or even within a short time span, because different contexts might be involved, which cause different brain states to occur, though we are not conscious of these different contexts.

Another way of trying to explain consciousness is pointing out what consciousness is good for and that it therefore must have been an advantage with survival value by which it got supported in emerging in the course of evolution. A metaphor for the mind, based on this view, is the one of the mind as a mirror of the world. This mirror reflects according to its own dimensions and perspectives only what is relevant for the individuals and the natural sort. Taking a Darwinian stance, as Dennett (1995) has done, we can understand that in our natural history we, and all higher animals, got adapted to the middle size objects and events of which we have to take notice in order to survive. Consciousness about these, extended to include self-consciousness, makes rational explanation and prediction of our and others behaviour possible, and thus it makes behaviour to be action, and makes it possible to plan our dealings with the objects around us. The method of rational explanation presupposes the development of a linguistic medium of public and private representation, in which we can clearly delineate situations and beliefs in order to deal with them in organised and flexible ways. Intentionality is primarily consciousness in taking account of the outside world in perception and in producing behaviour as rational action, i.e. in a goal directed and goal-adequate manner. Self-consciousness makes it possible to judge about our doings as correct or incorrect, which, when applied on a broad scale, presupposes sufficient representation of our actions in language, in imagination, or in pictures restricted in their interpretation by aspectualisation, especially by language. Judgement of this kind gives us a way of conscious improvement of our activities and beliefs such that we can achieve growing social and practical abilities and knowledge. These advanced kinds of improvement are outside the reach of other animals. All this is very well, and explaining the emergence of consciousness in this way is at most a historical explanation about why the emergence of consciousness was supported by evolution.

If consciousness has come about in natural history and includes self-consciousness because the above mentioned improvements by judgements about own and other conduct have a survival value, then we can argue, against Dennett (1989), that intentionality is *intrinsic* in the human mind, and to some extent in higher animals as a natural result of the adaptation of our brains to the middle-sized world around us and our place among them. That intentionality is intrinsic to the human mind, and probably to higher animals generally, is Searle's (1983) claim, and he seems to be right about this, since this is what we experience and what we believe about other animals, with whom we make contact and who are also able to perceive their world and seem to have some knowledge about their own doings. That we have intrinsic intentionality is the best explanation for our ability to take the intentional stance towards our own and others behaviour, and towards all kinds of so-called intelligent machinery. It seems that the human mind can take an intentional stance precisely because it is intrinsically endowed with it, and because it is conscious about this in reflecting about his own and others condition in the world.

We can think of stages in the genesis of consciousness and intentionality. Firstly, only that can be conscious what is externally manifested and in this way can be perceived, i.e. experienced. It has got pictorial, procedural, conceptual, and furtheron even linguistic representations by being placed into the orderings of our experiences that are typical for these modes of representation. Our perception, phylogenetically, is adapted to chunks of middle size, some of which we need to avoid, others we can deal with. Positive and negative bodily encounters, as joy and pain, desire and avoidance become consciously experienced, as emotions which are primarily directed towards objects and events. Ontogenetically, our sort is equipped to learn fastest what is most favoured by biases in our neural organisation that have been formed in evolution. We have become endowed with sensors measuring input in certain dimensions or similarity spaces, and neural nets which favour reaction towards the middle-sized objects and events of our natural ontology, and especially favour emotionally activating objects, which become the centre of attention. These objects are the goals of our intentions, and our intentionality is structured by the contents that in some way represent these objects.

Secondly, we can then make more fine grained or more global partitions in the world. On the one hand, we can perform an analysis on the middle-sized chunks, by which we order them into similarity sets under certain perspectives. On the other hand, we can look for general relationships and properties in theory building. Under a perspective we order the chunks according to similarity and opposition. Thus we get under a perspective several concepts which are in opposition to each other and serve for characterisation as well as discrimination of objects. Over these generalisations are obtained. Achieving general knowledge we build up a conceptual semantics according to the fine-grainedness or coarseness of analysis or conceptualisation of objects and actions. In a process of analysis and synthesis we can rebuild old objects, build new objects, and envisage new situations.

Our consciousness of an object goes as far as the analysis goes by embedding it into sets of similar situations and objects salva stability of their similarity measures and thus characterising the object. By embedding it into sets of situations related by contiguity relationships salva stability or coherence we identify it. Our imagination about new objects depends on modifications and new combinations in accordance with the fine-grainedness of the analysis or conceptualisation at a given point. In principle, all the steps taken in a synthesis of different aspects or parts into a whole we can endow with intentionality. This means that we view them as steps of construction taken intentionally, and in this way we can reach finer grained levels of intentionality in rationalising a new product as a construction which seems to have been set up deliberately, although it might be that the product in fact was not constructed intentionally. It rather may be a product of merely unconscious neural processes without a conscious expression in the mind.

The same happens with the syntactic built-up of utterances. We can view every syntactic part of the expression retrospectively as intentionally used in its context. However, taking an intentional stance is only a rationalisation afterwards. It is not the normal way of producing or interpreting utterances. Likewise we can even go down to algorithmic constructions of expressions or actions and view these as intentionally built up according to the algorithm. But this does not mean at all, that in the human mind these intentional steps have been taken in producing the utterance or the action. They are merely intentional actions by the designer and programmer of the algorithms, designed in order to transform our symbolically formulated knowledge into the likewise symbolic formulations of programming languages and finally machine language, physically implemented into the machine's circuits. In the human mind nothing like this normally takes place, though such algorithms can be devised and followed in special scientific and design activities. Rather quite different processes take place in the parallel working neural networks. What we see in a logically built algorithm for performing a certain task or action, or for generating a certain perception, is based on a rational analysis of intentional behaviour, whereby the intentionality which normally is restricted to the larger chunks we are consciously aware of, is now directed towards small parts which are normally not produced or perceived intentionally. For these smaller steps no claim should be made as to their functioning as part of the human mind, as also Searle (1992) has argued. They have no function there, though they have a function in artificial intelligence as algorithmic levels by which the programmer bridges the gulf between a symbolic description of our intentional level of perception and behaviour, and the level of states and processes in a digital computer.

Also on our normal level of perception and representation of objects and actions these are not all produced intentionally. They can be constructed intentionally, but most of the time they just pop up as the products of many parallel brain processes caused by certain inputs. They afterwards may get rationalised as being intentionally constructed in control, performed in a self-monitoring process. Being able to deal with these products in such an intentional manner makes correction and improvements possible in an effective way. When we judge about our behaviour and make sense of it generally, and predict others' behaviour, we consciously relate representations of such behaviour to the world, i.e. we make use of our ability to entertain an intentional referential semantics. Our designs of artificial minds furtheron extend this level of conscious representations by devising finer grained representations of the same symbolic kind, down to finest symbolic representations and processes on the machine language level. And only by being an extension of what our mind produced intentionally through refining the analysis of the products emerging in our mind into smaller parts, the artificial intelligence is cognition, namely representation of our knowledge and a performance of procedures and tasks in a meaningful way. Our knowledge on the intentional level makes that the deliberately fine-grained representations of it, which are designed according to the requirements of the digital computing machine, can be viewed in large chunks corresponding to the objects on the knowledge level. The fine-grained representations can be seen as knowledge representations merely because they are transformations or translations from the knowledge level into another code developed according to a certain technical goal. Only because there is this level of intentionality and consciousness, on which we perceive, represent and know the world, the masses of very fine-grained representations encoded in the machine can be understood by delineating the large chunks, predetermined on the knowledge level, as information about the world. This also holds with respect to neural activation patterns. They only can be understood as representing knowledge by getting delineated according to observed correlations with conscious judgements about what is perceived, what is believed, what is thought, and what is done.

10. There is conscious thought without language

Carruthers (1996) defends the view that all our conscious thoughts are in natural language, and with this our thoughts are dependent on background culture and theories. Thus our mind consists in all the thoughts we can think in our acquired language and is universal as far as all natural languages have common properties. It is culture specific as far as a language is developed within a certain culture and contains concepts specific for this culture. Conscious states are higher-order discriminations of perceptual states and as conscious experiences they are available to acts of thinking. Such acts of thinking presuppose reflexivity because in thinking one has to refer to ones perceptions and thoughts. Reflexive thinking is only possible by means of a language which can represent ones experiences.

In his review Allen (1998) objects that many people think in pictures, not just in language. The generality of thought is then achieved through sequences of pictures which correspond to one thought, for example to the belief that grass is green. Carruthers' thesis implies that young infants and animals have no conscious thoughts, something the author accepts but others might not agree with. Allen concludes that the centrality Carruthers gives to thought in linguistic form in consciousness is questionable.

Allen's suggestion about the generality of thought which is not expressed linguistically but by sequences of pictures agrees well with the view about propositional contents and concepts in my theory of concept formation and understanding, presented as Dynamic Conceptual Semantics (DCS) in Bartsch (1998). Here general experiential concepts need not be linguistically framed. They are represented by growing similarity sets of experienced and imagined situations. Such similarity sets form equivalence classes, i.e. (quasi-)concepts, according to measures of internal similarity. Pre-linguistic quasi-concepts get furtheron stabilised and adjusted to public language in linguistically guided concept formation and in this way become concepts. Historical concepts, i.e. individual concepts and concepts of specific events, are formed as equivalence classes of sets of situations with internal space, time, and causal contiguity and also partial similarity relationships. Each member set is a partial individual concept, whereby the equivalence consists in that all identify the same individual, i.e. they are part of the same complete individual concept.

Basic thoughts, or propositional concepts, are construed as smallest situations (i.e. **B**-intersections, cf. Ch. 2) that fit into general and historical concepts salva stability of these concepts. By fitting into these they are analysed or internally discriminated, and identified. They can be linguistically expressed but need not be. Thus, having a basic thought or imagining a basic situation is construing a smallest situation as being integrated salva stability into similarity sets and contiguity sets, formed on the growing set of previously experienced or imagined situations. This integration requirement for smallest situations constitutes the satisfaction conditions or truth conditions of the basic thought. In the same manner, perceiving a situation means understanding it by integrating it, or fitting it salva conceptual stability into an experientially formed conceptual system, i.e. a system of such similarity and contiguity sets.

Characterisation of a situation happens by embedding it salva stability into general concepts, i.e. into similarity sets. Identification happens by embedding salva stability into contiguity sets, especially into partial individual concepts. This integration process can be expressed in consciousness by imagining an example of a situation that fits into previously established similarity and contiguity sets, which represent general and historical concepts, respectively. Thus, a basic thought does not need to be couched in language. It simply is this process of synthesising into smallest situations the constraints of integratability into historical and general concepts. It can serve as the content of a propositional attitude, such as a belief, or a desire. The integration process, and each play-back of it, defines the type of satisfaction situation for the attitude. Such contents are non-linguistic thoughts. They are not expressed in language, but they have structures that are isomorphic to syntactic-semantic structures of basic sentences. These contents or thoughts are close to language in its basic structures, though they are not real language because they miss the representational medium which we call a language. When they get linguistically represented these non-linguistic thoughts are the satisfaction conditions of the expressing sentences.

In speaking of sets of situational experiences, ordered according to similarity and contiguity, we speak about phenomena in consciousness. They are backed up by brain processes in neural networks which have been adapted to the surroundings in learning. The brain processes automatically perform orderings equivalent to measurements on which our conscious judgements about similarity and contiguity are based. Though there are no unconscious thoughts, there are the brain processes which activate those processes and neural activation patterns that give rise to conscious thought on episodic maps, as I shall argue in chapter 3. If the activation of the episodic maps is missing we still might have the brain activation that otherwise runs with a specific conscious thought, but now happens without consciousness. This brain activation can trigger further activation such that a certain behaviour comes up which in retrospect rationalisation of behaviour will be explained by assuming that the individual had an intention with a certain propositional content. Could we then say that the individual has an unconscious intention or an unconscious thought which triggered his behaviour? Such terminology would mix up brain with mind. Rather we should say that the brain activity was there that in other cases accompanies the specific thought or intention.

A brain activity by itself is not a thought. A thought rather is an equivalence class of stabilising sets of situations that make the thought true. Different people at a certain time entertain different representatives of this equivalence set, and at different times a person can entertain different representatives of the same thought, which are all representatives of the same thought or situational concept. The neural activity pattern of a thought is the result of a disposition implicitly contained in the neuronal connection strength after a process of learning. It is merely a causal effect of learning, and in a specific situation it is activated by a linguistic expression or experience of some aspect of a kind of situation or thing which had been regularly connected to the situations of learning that have been involved in creating this neural activation pattern or parts of it. Parallel with the parts of a thought and their combinations there must be neuronal activation patterns and relationships between these, which physically support the thought. A thought consists of parts which belong to our conscious experiences of situations in learning. The brain activation pattern of a thought, on the other hand, must contain a run-through of brain patterns corresponding to the parts of the thought and something corresponding to their compositions. These situation-induced activation patterns by themselves are not parts of thoughts. Only by way of being causally connected to our sensors and the stimulations that were and are received by them, they can function as the material carriers of our basic thoughts about situations which we can be conscious of in our mind. Without these connections to our sensory-motoric apparatus and the stimulations received by situations and behaviour outside and public, the neuronal activation patterns would not be carriers of thoughts and would not have a conscious expression in terms of thoughts. A brain activation pattern of a thought by itself does not contain the parts of a thought. Rather it just contains physiological states and processes which support the constituents of the thought. The activation patterns are not the aspects of the situations we have experienced, have thought of, and think of. The experienced situations, more or less indirectly, have caused these stable, recurrent activation patterns.

Consciousness is necessary in thought, but part of the underlying brain processes go through without consciousness, when they don't hit episodic maps. Such unconscious processes can work in producing behaviour which can be interpreted by observers in a rationalising manner as actions, although they might lack expression in thought and intention. Underlying brain processes by themselves are never thoughts, and there is nothing we can call an unconscious thought. Each thought consists of basic thoughts and they are all in terms of situational aspects we are conscious of. They need not per se be couched in language. Of course, more complex thoughts that are not just about basic situations, but about constellations of situations and logically structured knowledge and theory formation need language as a means of representation.

We conclude that neither a Fodorian language of thought built up from innate atoms nor a natural language is necessary as a vehicle of a thought, though, of course, much of what we think is couched in our natural languages.

11. Free will has to be learned on the level of consciousness

Action routines, in which activation conditions and execution of an action are specified, are procedures consisting of complexes of concepts related to each other by contiguity relationships, and bound to (pre-)motoric maps. As such they are represented in architectures of connectionist nets with recurrent fibres, in order to capture temporal sequences of sub-actions. They are indicated in the brain by potential activation patterns which are results of implicit dispositions in the set of neuronal connections, due to a process of training in the course of learning. Several action routines may be innate, as they are part of the natural way of securing survival of the sort. They can be further specified in a culturally determined way in additional learning processes. An example in which natural and cultural features are combined would be the action of eating something. When certain input directly addresses neuronal patterns which are strongly connected to certain other neuronal patterns which cause and direct the appropriate action, then the action is executed if there is no additional inhibition. This then is very much a deterministic picture of action. Action seems to be determined by the causal chain of stimulus and response. Free will seems to be excluded by determinism of action. I now shall show why this deterministic picture is not correct and in what sense we can speak of free will.

Hebb (1980) in his *Essay on Mind* has pointed out that although assemblies of activated neural cells in the brain cause action, there is still free will in the sense that many activation patterns are involved, which results in a lot of indeterminism. There is not just one sequence of stimulus and response, such that a series of actions will be realised in a causal chain of events. The mechanical picture of a cause resulting in an effect, like one ball hitting another and the other starting to roll, is not adequate. Many factors can be taken into account by the individual, and higher order strategies of invoking additional circuit activation, promoting or inhibiting activation patterns can be acquired. I now want to elaborate this more involved picture and also show why the assumption of a free will is necessary in order to promote learning processes which result in more freedom of action for the individual. Free will is a matter of degree and it is located between two poles of a spectrum of possibilities. One is the freedom of choice, the other is the acceptance of the necessity of the moral and public law as guiding principle.

A simple example of free choice is that I can choose out of several numbers one or the other. I can say, 'well, I take this', and then, 'no, I rather take another' and 'well, maybe I just take that one' and then possibly just decide 'well, why all this fuss, I will stick to the first, after all'. Here I had no special reason for taking one or the other. In normal cases of action we have reasons, desires, short time and long time preferences. The reasons and long term preferences are given by public norms, moral norms and social norms, and accepted values. We take reasons into account in our decisions to act in certain ways. To decide to act according to the necessities of the moral law has been described as freedom of action by Kant, in opposition to a mechanical or natural drift of an action coming about as a natural response to a stimulus. Our consciousness of the law makes possible our freedom of action against the determinism of action according to animal instincts.

Emile Zola in his grand series of novels *Les Rougon-Macquart. Histoire naturelle e sociale d'une famille sous le Second Empire*, written between 1870 and 1893 illustrated his 'scientific' theory that all doing of people was determined by inherited strengths, weaknesses, and instincts, and by the learned patterns of behaviour, determined by their cultural and social surroundings. In any way, the individual is a product of nature and nurture which reacts to accidents and social conditions within its surroundings. Zola's work provides an illustration and elaboration of determinism, as it was shaped by Darwinism and by the great impression the workings of mechanical inventions had on psychological theorising.

In connectionist network models there is no relation of mechanical causation between stimuli and effects. Rather there is a probability based relationship in which many factors are involved that cannot all be spelled out. Very little effects on the micro level may have or not have, depending on further conditions, a large effect on the macro level of activation of neuronal assemblies which are relevant for causing certain actions to take place. Anyhow, the model of mechanical causality is not adequate for understanding the causal basis of action, and the indeterministic character of probability based models of activities of cell assemblies seems to give new openings to the possibility of free will. However, that kind of indeterminism by itself is not a guarantee for the possibility of free will. It just seems to be a property of our models of description, and the possible levels of description of brain activities that make for indeterminism or determinism. The argument would be that if we could describe the processes on the micro level of cell activity, and if we would be able to take into account all determining factors, the description would look as if our actions were causally determined. And the argument could proceed that on the even finer grained micro level of quantum mechanics, indeterminacy would again be part of the descriptions of the goings on in the brain which are the causes of our actions. We conclude that the property of determinism or the property of probability induced indeterminism of our alternative descriptive models generally has no bearing on the question of whether free will exists, because we can have the one or the other, determinism or indeterminism, depending on the relative macro- or micro-levels of description.

However, the argument above is not to the point, because only indeterminism in the prediction of actions on the level of our ontology of situations, objects, and actions, and its parallel macro-level of neuronal assemblies in the brain has some relevance to the question whether free will is possible. This we may conclude if we assume that conscious phenomena, such as actions and felt dispositions towards certain actions, are expressions of certain actualised neuronal activation patterns. I shall point out below in what sense indeterminism is possible on this level of macro phenomena which constitute the contents of intentionality and action. Especially I want to show that to have a free will is a question of development and learning. To put it crudely: an individual has to learn to have a free will.

Free will means that there is the possibility of choice between courses of action. It may very well be that the choice is not free in the sense that an individual is caused by brain processes to choose one option. This then does not seem to be a free choice but a forced one; the individual is forced by internal brain processes towards one of the options. This, in fact, is always so. The individual's mind and behaviour, all the individual is conscious of, and even the conscious individual itself is an expression of brain processes. I want to show that nevertheless this fact is compatible with having a free will. This is so, because on the other hand, the individual can gain insight in several options to choose from, and for each of these reasons can be gathered, desires and long term preferences can be considered, and in this way many brain processes take place that lead to different activation patterns causing the different possible courses of action. By being able to make explicit possibilities of choice, different activation patterns get actualised and the whole process becomes one of competition between activation circuits and patterns for strength. In the whole process of deliberating between options by consciously taking into account reasons and long term preferences, one activation pattern will come out strongest and thus one option will be realised. By consciously considering alternatives, choices are introduced and the chance is provided that not the activation pattern which is strongest by nature or by habit will win automatically, but one of the other potential activation patterns established in learning values and norms gets sufficiently enhanced by reasoning, consideration, and emotional activation of publicly accepted values. Then an alternative action pattern, or a positively valued natural drift or acquired habit can be the winner. In such a case we would say that free will has gained a victory over a natural drift or habit which is socially unwelcome.

I shall now elaborate on how such a situation can be achieved. For that let us imagine a young criminal who has the tendency to commit criminal act B in situations of type A. We say that he has a criminal habit, i.e. he has internalised an action routine A->B. What is necessary to restore in him a free will with respect to actions in situations of type A?

Punishment following realisation of B in situation A is one means. By creating an adverse emotion with respect to B, it might establish just enough inhibition of his tendency towards A->B in future cases to give room for the introduction of a situation of choice between possible alternative actions in situation A. In order to get a choice, the young criminal has to be trained to formulate alternative action concepts for situation type A. He furthermore needs to have a training to evaluate the alternative actions. Learning the ability of evaluation implies that he has to learn to formulate and take into account his short term preferences, his long term preferences, public norms and values. The publicly acceptable preferences, norms, and values have to be attached to favourable emotions in training. He also has to be trained to deliberate between the evaluations of the alternative actions and to attach a positive emotional value to outcomes which he has learned to be publicly approved of. The strength of enforcement of public norms and values, and the public emotional valuation of these, has to be exemplified in training processes and thus be learned from examples. This internalisation gives strength to certain potential activation patterns in the brain which favour a socially acceptable outcome of the choice between actions in situations of type A. This means that education has to take place in the way described, such that learning processes give rise to further potential activation patterns in the brain which promote choice towards a publicly acceptable action.

The considerations above imply that free will is taught and learned in planting inhibition by punishment, by teaching conscious reasoning and consideration, and by teaching emotional, publicly accepted, evaluation of actions by examples. What comes out in each situation of type A cannot be predicted for certain. Choice of action is not deterministic on the levels on which we describe situations and actions or, for that matter, global brain activation patterns established in learning about these situations and actions. There are too many factors involved concerning internal and external stimulation and strength of potential activation patterns and their internal connections. They make that all our predictions will be a matter of chance. But anyhow, there are at least chances for realising alternative actions C, D, E, or even new surprising ones, instead of just doing B according to the strong primary tendency A->B, we started out with. To have these chances means to have free will of action in situations of type A. This implies that the employment of free will can be established by guided processes of learning, and this is in principle part of the goals of imprisonment and social training of persons with criminal habits.

To those philosophers who would object that what is described here is not free will, we may answer in the line of a Wittgensteinian approach to meaning that this is what we call 'free will' in our social practices, and that there is no other meaning for the term than the one established in these practices of law enforcement, of learning socially acceptable behaviour, and of pointing out to individuals the possibilities of choice and the ways of reasoning about alternatives. Scepticism with respect to the existence of free will then can be seen as meaningless as is scepticism about the nature of knowledge, namely whether there can ever be truth of beliefs. Like we have a clear everyday notion of truth with respect to sentences and texts, we have a clear everyday notion of free will with respect to actions. And like telling the truth can be learned, also acting according to free will can be learned. We have to be content with these everyday notions of free will and truth without searching for absolute notions, which have been assumed to be stable, ideal, and context free.

The discussion about consciousness up to this point justifies the assumption that there is, like in connectionist models, no central control unit for cognitive and emotive processes. However, there is control on the results of brain processes, and this control happens in consciousness and takes into account comparison of the resulting imaginations of situations and actions, represented on our episodic and pre-motoric maps, with external and especially intersubjective data, i.e. with other situations and actions, and with correctness judgements, in theoretical as well as in practical affairs.

Chapter 2

Formal Constraints on the Relationship between Understanding and Connectionist Neural Network Activity

In the previous chapter the relationship between denotational semantics and causal semantics was at issue in the discussion about the roles of consciousness. Denotational semantics, a semantics from a realistic point of view and with normative constraints is the hart of interpretation of linguistic utterances. From a cognitive point of view not primarily interpretation, rather understanding is the topic. Our goal now is to make some formal assumptions about the mapping of the cognitive structures of understanding on processes and states of the brain, such that there we can speak of structures with a degree of coarseness that is appropriate for discussing the causal backbone of cognition. The connectionist neural networks discussed in this chapter are fairly abstract models of some structurally relevant aspects of brain processes. We formulate desiderata for these, which then can be transferred onto the brain, leading to hypotheses about global brain structures.

In this chapter a formal relationship shall be postulated and spelled out between, on the one hand side, understanding of situations and basic sentences, and on the other hand, global neural network activity that models the causal background of understanding in an appropriate coarseness in the brain. Understanding is couched in the model of Dynamic Conceptual Semantics (Bartsch 1998), which will be explained presently. Especially we will try to constrain the relationship between the level of experiential concepts in Dynamic Conceptual Semantics (**DCS**) and concept formation in distributively representing connectionist neural models (**CNM**). Both approaches model learning of concepts, especially linguistically expressed concepts, and at the same time they imply theories of understanding situations in perception, and of understanding simple sentences that describe situations. **DCS** gives a logico-philosophical reconstruction of concept formation and understanding on the cognitive level of conscious phenomena, namely on growing sets of data which are experienced situations. **CNM** tries to model on an abstract level of neuronal connections and activation patterns the learning of classifications and associative connections of contiguity in perception and behaviour, also linguistic behaviour, as generalisations and associations over sequences of examples. Special attention shall be given to the problem of syntactic systematicity, which is a notorious problem in **CNM**.

The formal relationship proposed between both models of concept formation and understanding will be a postulate motivated by semantic considerations. Both positions ground language learning in perceptual and behavioural experiences, on which the operation of selecting aspects of situational experiences under perspectives leads to generalisation, while the structuralist operation of forming substitution classes of items within identical contexts leads to paratactic classes of contrasting parts which take part in the same syntagmatic and contiguity relationships. These form syntagmatic and contiguity patterns.

The data of linguistically guided concept formation are pairs of situations and utterances, primarily simple sentences with their experienced satisfaction situations. Generalisation, on the one hand, yields situation concepts or types, especially action types, and on the other hand it yields utterance types, i.e. sentences. Substitution and contrasting within identical contexts leads to distinguishing parts of situations and parts of utterances, as well as parts of situation types and parts of basic sentences. Paradigmatic classes, short: paradigms, are formed as sets of situation parts which can be substituted for each other within situations in order to yield new situations. Conceptual paradigmatic classes are similarity sets, i.e. sets with an internal similarity between the members, due to an identical property or relationship. Likewise linguistic paradigms are formed as sets of expressions which can be substituted for each other within sentences in order to yield new sentences. Within situations, parts of situations stand to each other in contiguity relationships, and likewise sentence parts stand to each other in contiguity relationships, called syntagmatic relationships. These syntagmatic relationships are generalised over all members of two different paradigmatic classes which combined together constitute a new constituent of a situation or of a sentence, respectively. This is the common structuralist basis of our conceptual and linguistic structures, especially of basic syntactic structures. With these syntactic-morphological structures of basic sentences correspond the semantic roles of the sentence parts. They express a part of the possible satisfaction situation for the sentence and the role this part plays in building up this situation. The paradigmatic classes are the categories that are characterised by the syntactic role their members play in the construction of sentence constituents and sentences. In parallel with their syntactic role in sentences, they express types of actions and other relationships within situations, the role participants of these actions or relationships, and the circumstances of the actions. Basic sentences express a situational concept which characterises the possible satisfaction situations of the sentence. Relationships between basic situations and their circumstances are contiguity relationships of spatial, temporal, causal and other kinds. They give rise to sentence connectives of causal, adversative, concessive, temporal, and local kind, and in parallel to compound sentences with adverbials phrases, especially adverbial clauses describing the situational circumstances.

Generally, syntagmatic relationships express contiguity relationships within basic situations and between situations. Paradigmatic relationships express generalisations, i.e. similarity relationships based on identical properties or relationships. Contiguity relationships further play a typical role in forming historical concepts, such as concepts of individuals, actions and events. They can be expressed in texts, such as life histories and descriptions of actions, events, and their position and course in space and time. Over these, further generalisations can be formed leading to general concepts of kinds of individuals and other general historical concepts.

Concept formation starts out on the experiential level, which provides growing sets of data. The data, as far as we are conscious of them, change for us in the process of concept formation, because for us the data are always data which are understood in so far that they, in uptake, are embedded directly into sets of previous data by similarity and contiguity. On the experiential level, which consists of growing sets of situational experiences and their delineating linguistic expressions, we form two kinds of experiential concepts, general concepts and historical, especially individual concepts. Our general concepts are represented by series of examples consisting of linguistic experiences accompanied by the satisfying situational experiences, whereby each member of the series is seen in the light of the previous members as similar to these under certain selecting perspectives. In each stadium of the process of gathering more experiences, such a representation of a general concept as a sequence of examples is partial with respect to the full set of possible examples. Partiality works as a constraint: Each partial representation of a general concept has to be embeddible into the complete representation salva stability of the internal similarity of this complete set of situations, which is the extension of the realistic and publicly endorsed concept.

A *general concept*, strictly speaking, is the equivalence class of all partial representations that are equivalent by being embeddible into the complete representation salva stability of internal similarity. Each member of this equivalence class represents the concept. Since we are conscious only of these sequences of examples, these are the representations of a concept we, in fact, have at a given moment, and this concept as we have it is partial with respect to the complete concept. *Historical concepts*, especially individual concepts, are sequences of situations formed by contiguity associations in space and time connecting each member situation with previous ones. Such a sequence is a partial representation of the complete individual concept, which would be the full life history of the individual. Also here, partiality works as a constraint: The partial individual concept has to be embeddible into the complete one salva stability of the contiguity and coherence of the complete individual concept. The *individual concept* is the equivalence class of all the partial life histories which are part of the same individual life history. Also here all the partial representations are equivalent in that they are parts of a complete individual concept. Our individual concepts are always such partial representations, i.e. parts of life histories of individuals.

On top of this situationally based linguistic stratum, which is grounded in our perception of situations, performance of actions, and reporting about these, we get the logical stratum consisting of a vocabulary of logical connectives, negation, quantifiers, and modal operators of different kind. The logical vocabulary gets its meaning from the ways we organise our information in forming generalisation and existentialisation over particular information by means of general and existential quantifiers, by seeing information in conjunction or disjunction, in conditionalisation, and by seeing it as reality, as possibility, or as necessity. This stratum of the organisation of information, which makes it possible to keep information at a distance from sensory and motoric input and output, on a level of knowledge which is not directly manipulated by situational contingencies, is not captured in the discussion below, because it lies outside, or above, the treatment provided by typical connectionist modelling.

It has already been pointed out (Bartsch 1996, 1998) that connectionist models of concept formation share certain welcome properties with Dynamic Conceptual Semantics:

1. Perception is the core and basis of concept formation.

2. Sensory input is directly embedded into classificatory and contiguity content and thus understood as a perception of a situation. This amounts to direct content addressability by input. A situation is immediately perceived within the conceptual system available at a certain moment, which means that it is perceived as categorised and identified within the generalising classificatory, and the individualising contiguity structure established so far on the set of experienced, imagined, and reported situations. Likewise linguistic utterances are understood by directly addressing the linguistic type in contiguity relation with the type of the given context. Hereby also the semantic content is addressed, due to the association of situation parts with the parts of the expression in the respective contexts, and according to their syntactic contiguity connections, i.e. their syntagmatic relationships, within the expression. There are efforts to work out the corresponding operations of syntactic systematicity in connectionist frameworks by architectures with recurrent fibres using vector concatenation, by vector addition, and by tensor products (Horgan and Tienson 1992), some of which will be discussed below.

3. The structures stabilise in the course of learning from experience (conceptual stabilisation).

4. The structures are flexible in accommodating new and even deviant input (conceptual flexibility).

5. Conceptual growth and change are part of the system (conceptual dynamics).

6. Traditional cognitive representations of mental contents by means of an innate language of thought are avoided. In **DCS**, representations are always in the form of expressions of some natural or formal public language, or in the form of series of pictures and examples. In **CNM**, representations are causally effected neuronal activation patterns, which I rather shall call "indicators of concepts or categories", in order to reserve the term “representation” for representations in the strict sense.

As Wittgenstein has pointed out, representations always presuppose understanding, because they have to be related to something they represent, and thus they function as representations. Therefore they cannot function in establishing understanding and interpretation. They are what has to be interpreted. Interpretation is the relationship between expressions or pictures and their extensions in real and possible worlds. Understanding is the cognitive operation which the interpreter has to perform in order to be able to establish interpretation. In **DCS**, understanding basically is embedding a situation or utterance into sets or sequences of previously experienced situations or utterances salva stability of their internal similarity and contiguity relationships. In this way the conceptual semantics of representations, like sentences or pictures, is part of the cognitive level of perceptual and behavioural conscious phenomena. These consciously experienced phenomena and embedding constraints are publicly co-ordinated in the intersubjective acts of acceptance of satisfaction situations for utterances.

On the other hand, states and processes of neuronal connectivity and activation are the causal background for the conscious results of understanding. These results are representations. Our ability to handle a publicly co-ordinated conceptual semantics is implicitly encoded by the causally effected results in learning. They are encoded within the connection strengths between units in neuronal networks. These connection strengths embody a set of dispositions towards producing sets of neuronal activation patterns within the network. Their realisation in the network depends on appropriate input stimulation. The neuronal system also includes the neuronal units in the sensory and motoric organs. The activation patterns arising from dispositions of a neuronal net, can be assigned to types of input which triggers these patterns, given a certain state of the network. On the output side they can be seen as dispositions towards reactions to perceived situations and to situations which satisfy or dissatisfy utterances, and furtheron they can be seen as dispositions towards actions and utterances connected to these situations in contiguity relationships. With every input a vast cascade of reactive neuronal patterns is generated and it is difficult, though not impossible, to isolate by contrast analysis the activation pattern that primarily corresponds to an expression. Anyhow, all the typical contextual effects of the expression are also activated. They thus belong to the content of the expression represented on a neuronal net, and they can possibly be isolated in contrast analysis by being strengthened or inhibited.

7. In **DCS** semantic taxonomic systematicity is taken care of by inclusion and overlap relationships between sets of data. Based on this, we would like to require in **CNM** a reversed structure of inclusion of activation patterns in neuronal nets. This is motivated by the fact that we can see an activation pattern as corresponding to a set of examples which trigger this pattern when they are input, and some of these examples have been in the learning set that has resulted in convergence towards this pattern for this kind of input. Some sub-assemblies of activated neurones should correspond to semantic features of the examples. These current internal feature indicators have been created in the stabilising process of learning from examples. Such sub-patterns of activated neurones are

not real features, though they must correspond to features of the input examples; and such assemblies or sub-patterns must play a role in bringing about our conscious awareness of similarity between the input examples. They must be the causal basis for the internal similarity we find in the course of growing experiences in the sets of examples which correspond to this pattern. We recognise certain input data as examples of a concept due to that activation pattern. This neural pattern then works as an internal cause for classifying something as being an instance of a concept, and thus can be called a current network "representation", or better "indicator" of the concept. Thus the internal similarity in a set of examples is not just brought about by the sensory measurement of the input examples by the sensory units, which serve as micro-features or dimensions of the input vector. Rather the similarity is established by an reaction of the trained internal neuronal nets, whereby sub-assemblies of internal units identify features in the perception of the data. This also implies that in the course of learning our perception of the data changes due to perceiving other and new similarities between data. This model also leaves open the possibility that we measure similarity by internally indicated features which have no expression in consciousness, i.e. we have no phenomenal awareness of the identities that determine measurements. We might judge two data as similar without knowing on the basis of which identical features we do this.

In standard connectionist networks, the internal units have connections to all input units of the sensor. This means that an internal unit can be activated by quite different combinations of activity from out the input units. In the learning process some connections get strengthened, others inhibited. But still an internal unit is multifunctional, because its activation in combination with other units' activation can indicate quite different input; it can be activated from out different input units. This makes that it cannot be seen as a constant internal feature indicator. It at most can be seen as a secondary context-dependent micro-feature of an input example: only in the context of the activation of certain input units, which are strongly connected to the internal unit such that enough activation is transmitted, it can indicate an internal feature, namely the combination of the activation of these specific input units. The multi-functionality has as a consequence that a neurone in different contexts of other activated neurones can indicate different things. This hinders a neat correspondence between semantical structure and pattern structure in terms of sets of activated neurones and set theoretic operations on these sets. Intersection between two patterns does not always mean similarity of the input. Single neurones in the intersection may indicate nothing in terms of semantic similarity. But the larger the amount of neurones in the intersection the higher the chance that this indicates similarity. In order to make the set of activation patterns of a network indicative for semantic similarity and inclusion, we have at least to require that certain large enough assemblies of neurones have structural properties of inclusion and intersection. These assemblies must be the ones corresponding to features.

It would not be a good idea to take as a concept representation such sets of sub-assemblies of a distributed pattern which correspond to the features included in the concept. On such sets we could apply the set theoretic operations, like we can on sets of features. However, the disadvantage of representing a concept by a set of features, like it is done in linguistics, would be the loss of contiguity relationships between the features, and the loss of those features we are not aware of and have not expressed in our language, but which still might play a role in measuring similarity. All that which makes a neural network's concept indicator correspond to the richness of a gestalt-schema of a concept would get lost in such a set up of representing a concept by a set of features.

Because parts of internal neuronal activation patterns function in our perception of data we get the following correspondence: With forming hypernyms by forming unions of sets of examples for concepts, i.e. unions of concept representations in **DCS**, there should correspond intersections of neuronal activation patterns in **CNM**, and with inclusions of concepts in **DCS** there should correspond reversed inclusion. The formal relationship between both models should then be an injective homomorphism which maps a set of examples of a concept, i.e. a concept representation in **DCS**, onto a corresponding neuronal activation pattern in **CNM**, maps unions of concept representations on intersections, and reverses inclusions. We shall see presently that this requirement is not generally fulfilled.

8. Contiguity relationships in **DCS** can be represented in **CNM** by closeness and association of activation patterns. But this certainly is not specific enough for the different kinds of contiguity relationships, for which special conceptual maps for recognising local, temporal, causal, means-ends, action-result relationships would have to be trained additionally.

9. Syntactic structuring of the process of understanding amounts to adding up specifications of the set of possible satisfaction situations, whereby the binding of the specifications into a basic situational concept is is essential. This binding is crucial for semantic-syntactic systematicity. The final result is a representation which functions as such merely by the way in which it is built up.

Syntactic systematicity on the level of simple sentences is displayed in substitution based on contrast analysis between sets of data, stating identity and difference between simple sentences on the one hand and between situations on the other. Structural hierarchies in conceptual semantics, according to the semantically relevant syntactic hierarchies in a sentence, are realised in the constructive steps in the process of construing the set of situations into which the possible satisfaction situation of a sentence must be integratable salva stability. These constructive steps of a progressing specification of the possible satisfaction situations can be represented in what is called **B**-intersection ("binding intersection") of two concepts (cf. Bartsch 1998). These are intersections of two sets of situations, i.e. intersections of the two concept representations, such that the intersection contains the basic situation into which the two concepts are synthesised. This means that an intersection of two concepts is a **B**-intersection if and only if it contains the basic, i.e. the smallest situation which fits into both concepts salva stability. This means that in concept composition the basic situation fits into both concept representations salva stability. By performing **B**-intersections between the concepts in their syntactic-semantic role, a stepwise specification of constraints on the set of possible satisfaction situations for the sentence is achieved in the process of understanding (for details see 2.2.1 of this chapter). In order to find a formal correspondence between dynamic conceptual semantics and neural network activities, it has to be decided what operation on neuronal nets might correspond to **B**-intersections. It cannot be a simple union of neuronal patterns, or vector addition if the net is represented in the form of a vector. There must be an additional synthetic operation in which forming a union plays a role, like intersection plays a role in the synthetic operation of **B**-intersection. This point will be discussed in section 2.2.

This requirement of parallel and reversed properties between **DCS** and **CNM** is based on a postulated strong structural similarity between the ways in which concepts are represented in both models. Furthermore both models share a common background assumption, namely a non-classical position in cognitive science. From the semantic properties of understanding, explicated in **DCS**, which are based on principles of interpretation and truth evaluation in semantic models, semantically motivated conditions can be derived for constraining the structures which stabilised neuronal networks must exhibit, will they be semantically successful. I finally shall show that understanding a sentence in terms of a resulting conceptual semantic representation can be paralleled by circuits of activation in an appropriate connectionist architecture of neural maps. The process of activating certain neuronal patterns, which are causally related to input from previous experiences in learning processes, in partially fixed syntactic orders is significant as the neuronal basis of sentence understanding. It guides the interpretation of expressions in the world by providing the referential relationships to parts of the world in action and perception. Architectures in neuronal network models which let the process run in some syntactic order, have to be postulated as a neural basis for the process of understanding situations in perception and action, and for the parallel process of understanding sentence utterances. It is questionable how far the postulated formal relationship and requirements on neuronal nets can be realised successfully in current connectionist models which still are restricted in their capacity by their implementation on serial computers. However, at least some beginnings have been made.

2.1 The non-classical cognitive model

The classical model of cognition is designed following the standard computer set up, consisting of a central executive module and peripheral modules for output and input. It has as the centre a unit of rational reasoning, including the execution of rule application on representations in some formal language, corresponding to what has been referred to as the language of thought. Around the centre, peripheral modules are located, namely the perceptual and the motoric module. Information comes in via the perceptual module and output goes out via the motoric module, having trespassed the central unit of rational processing, which connects the peripheral units. This model strokes well with the ideas of rationalist philosophy, which saw the human mind as the rational core of man, getting input and securing output via peripheral sensoric and motoric modules, under influence of an emotional module, mostly competing with, and sometimes supporting the rational processing.

Both models of concept formation treated in this chapter, however, assume a non-classical cognitive model. The non-classical model is rather based on empiricist philosophy, taking perception and motoric behaviour as central for cognitive ability and deriving ideas, i.e. concepts and their combinations in thinking, from perception and memory by operations of similarity and analogy in processes of perspective dependent selection and substitution, and by operations of contiguity in registering factual relationships. Speaking in terms of centre and periphery, the sensoric and motoric faculties and the conceptual orderings established on their processing results on conceptual maps form the centre of primary cognitive activity. The rational faculties form the periphery, which consists in controlling the conscious results delivered from the central basis by applying to them rational and public constraints and making use of the general goal directed and logical procedures in analysis and in composition according to rules.

It has been shown in neuro-biological brain research that the activated neuronal assemblies for performing certain actions, for example for grasping certain kinds of things, and the activated assemblies for perceiving these actions, also when performed by others, largely overlap. Furthermore, perceiving the kind of things that regularly are grasped in this way additionally activates the neuronal assemblies for the respective actions of grasping (Fadiga and Gallese 1997). Thus not only the classification of the actions is represented neurally, associating first and third person perspectives on the actions and representing their perceptual and motoric aspects, but also the contiguity relationships of the actions with the type of things used in these actions are associated in the neural activation patterns. In fact, the whole frame or schema of an action is represented including its types of role participants and its typical circumstances. For example, starting out from seeing a certain kind of things, which we are used to grasp, grasping them is an action possibility which is offered straight away by the neuronal assemblies. Of course, such association comprises much of what is rationally reconstructed in practical reasoning and it follows the ways of performing practices while using certain things. This reminds us on Wittgenstein's example about the use of language by masons in handling different kinds of stones. The stones are what they are by the way they are used, and likewise, we can add, our neural indicators include in the activation pattern for a certain kind of stone the activation pattern of the way in which we normally use it. What appears as practical reasoning in rational reconstruction is not just some rational faculty innate in individuals, but it is guided in a context dependent way and organised by great amounts of practical experience, and by being trained according to public norms and conventions. Therefore practical reasoning is for a large part learned in education within certain types of situations.

In the way sketched above, perception, motoric, and emotional ability form the core of the non-classical model, while rational reasoning rather is the controlling periphery, possibly partly an externalised and explicated continuation of inborn preferences for stability, conservativity, monotonicity, which are publicly endorsed in education. Reasoning is induced mainly through public normative and controlling behaviour, directed towards the goal of getting true, reliable information. Reasoning works on representations, especially in public language, and trained in schooling. This implies that rational reasoning is to some part natural to human minds, but also very much a culture bound and enforced set of constraints for control concerning reliability on what otherwise is produced directly as output from the centrally encoded relationships between sensory input, emotional disposition, and motoric activity. The possibility of control is provided by our faculty of consciousness, i.e. by our awareness of others' and own behaviour, of emotional evaluation, and by our awareness of previous and current perception in memory.

A non-classical model is proposed, for example, by Smolensky (1989), who takes connectionist neural networks as models of the sub-symbolic centre of cognition, while rules and representations are the publicly induced and endorsed ideal rational constraints, imposed from out the periphery. The latter require symbolic and especially linguistic processing on a level of logical vocabulary. This is compatible with the possibility that specific practical and logical constraints are not merely publicly induced in the symbolic mode, but emerge also from sub-symbolic processes resulting into generalisations over neural indicators of concepts, and association (union), overlap, inclusion between such indicators. Parallel we have relationships in consciousness as operations on concept representations, which are here available as stabilising series of instances of a concept.

Concept formation is primarily based on the experiential level of perception and motoric activity, though it is furtheron developed on the level of linguistically explicated theoretical concepts, constrained by public language and theory-formation. These latter concepts are constraints on possible perceptions, actions, and expectations from out the periphery of cognition, which works via operations on symbolic representations of such perceptions, actions, and expectations. These linguistically explicated, i.e. analysed, concepts form the second level of concept formation. They are theoretical concepts in a broad sense which are explicated by coherent sets of general sentences held true. Linguistic labelling on the experiential first level of concept formation is a prerequisite for the linguistic explication of concepts in theories. Experientially introduced concepts can be elaborated, or even be redefined as theoretical concepts, i.e. as linguistically explicated concepts. And in different occupations, technologies, sciences and cultural achievements, there are different theoretical elaborations of experientially introduced concepts, according to their different perspectives, goals, and contexts in these different cultural fields. Hereby a single experiential concept might have several corresponding theoretical concepts as elaborations in different language registers and varieties.

In this chapter I shall concentrate on the experiential level of concept formation, because there we can hope to formulate a relationship between connectionist neural network models and the data-based dynamic model of concept formation. Connectionism, as it is developed today, cannot say much about theory formation, in as far as it is more than generalising over associative relationships. Therefore the genesis of theoretically and formally explicated concepts will not be treated here. If we presuppose that a theory has already been built up on the discursive level, the similarities between theoretical concepts, their opposition and their inclusion relationships can be represented on connectionist networks, because linguistic distributions, i.e. sets of contexts of an expression in normalised sentences can be represented in these networks. Therefore the semantically characteristic syntagmatic distribution of a term X, roughly the set of sentential contexts of the universally quantified term *all X, every X, X in general*, which form a true sentence together with the universally quantified term, can be ordered such that expressions are semantically more similar the more semantically characteristic sentential contexts they have in common. Comparing the semantically characteristic distribution of two terms results in a similarity degree according to how much their distributions overlap. A subset of such a semantically characteristic distribution is a feature, which can also be part of the semantically characteristic distribution of another expression. This way, terms have common semantic features and they can be classified by the features created as sets of sentential contexts on the theoretical level. Hereby similarity, opposition, and inclusion between linguistically explicated concepts is represented in conceptual analysis. By applying these linguistically explicated concepts to objects these will be classified according to the explicit features, and thus will be judged similar according to common features explicated in the general sentences held true.

2.2 Projecting Structures of Dynamic Conceptual Semantics onto Structures on Activation Patterns in Connectionism

In the DCS-model of concept formation and understanding, linguistically guided concept formation on the experiential level consists in creating structures on growing sets of data. Data here are pairs of utterances of simple sentences such as *The boy eats an apple, Peter drinks milk, Fred runs*, together with their satisfaction situations, which are situations that make the respective sentence true. Two types of structures are established. On the one hand there are similarity structures or classifications under perspectives, i.e. similarity sets of situations formed, to begin with under certain basic perspectives such as colour, form, taste, touch, sound, form, motoric behaviour. On the other hand there are contiguity structures. These are contiguity relationships within situations, such as the ones between actions and participants and circumstances. And there are sets of situations which are connected to each other by relevant contiguity relationships. These are factual relationships in space and time.

Similarity structures and *contiguity structures* on the growing sets of experienced situations are compatible in that relationships under one structure do not destroy relationships under the other structure. A situation is determined as that what it is under both structures together. A situation is always an experienced situation the borders of which are not fixed and vary with the focus of attention, which implies making a selection of aspects by the relevant perspective induced, for example, by a desire, an action context, or a question asked on the basis of previous knowledge. A situation is not determined once and for all. Rather the result of applying the methods of structuring under one or more perspectives gives an analysis of the situation. The situation in this form can be used further in applying more structuring, and the results of these further steps can be used as a new data in applying the methods of structuring again. This recurrent structuring is a kind of hermeneutic circle applied in perception. In this way, our set of data does not only grow, but the data also change in the manner they are understood, and in this way old data can become new data. In principle this spiral like process of understanding a situation can go on, and will practically terminate differently in different settings and different states of development, i.e. according to the different structurings available under the relevant perspectives, created by different practical and theoretical contexts, and available at a certain moment in the history of an understanding individual. Results of the process of understanding a situation, whether momentary or final, become conscious in perception and in pictorial or linguistic representation, imagined or uttered. After such a classified and identified item is represented in consciousness, it can come up again in current memory, and the recollected item, or an imagined item of the same kind, can become a new input in a new round of the process of understanding. In what follows I shall formally characterise the two kinds of structures we lay on our growing sets of data, and I shall formulate what their compatibility consists in.

Parallel to the conceptual structure we assume a formal structure on the set of partitions on a neural network which are formed by stabilising activation patterns, actualised with suitable input data. An activation pattern corresponding to a concept in a certain stage of learning and in a certain context is the partition of the neuronal net that is actual whenever the net gets input by examples of the concept or is activated via the use of its linguistic expression. The effects of certain contexts may be extrapolated by using contrasting contexts for the same object, and by keeping the same context but using contrasting objects as data. The activation patterns are results of dispositions of the network to react to input by actualising these activation patterns consisting of assemblies of activated units or neurones. The dispositions can be looked at as functions built up in learning, analogous to semantical intensions, which have internal extensions on the sets of neurones by causing their activation under suitable input into the net, and which have as external extensions the sets of examples which are suitable inputs for getting these activation patterns. This set of dispositions is encoded implicitly in the connection strengths between the units, which has been built up during the learning process. The basic idea about building up connection strengths or weights was Hebb's 1947 conjecture that units which are repeatedly activated together strengthen the connections between them (elaborated in Hebb 1980). The formal relationships between potential activation patterns or net partitions are union, overlap, and especially inclusion.

A problem is the delineation of a pattern because the pattern for a certain type of object, for example a chair, cannot easily be separated from all those activated parts of the net that are due to typical contiguity associations with chairs, such as sitting, table, sitting persons, properties of typical materials used, etc. However, some delineation can be achieved by contrast activation, for example by contrasting the patterns activated by perceiving standing persons, tables alone, chairs alone, whereby the parts activated by what is visually present must be activated stronger than the other parts of the patterns which correspond to what is not present in the input, but is merely associated. Also the associative correspondence with linguistic expressions uttered helps to single out a strongly activated part from other parts of an activation pattern in a contrast analysis of the patterns elicited by appropriate input. Hereby the important point is that the stabilisation of activation patterns and the relationships between them are induced in the course of learning from data, namely from the linguistic utterances together with their respective satisfaction situations. Each pattern can be represented by a vector with the neurones of the net as the vector's dimensions and the activation values as the values in the dimensions.

We need not assume that all members of a speech community show the same activation patterns for the same kinds of objects, nor that for a single person these are identical across time. The only requirement is that the formal relationships between these patterns remain stable in correspondence with the relationships between situations and sets of situations that are causally bound to the patterns by the individual's history of learning. For an individual using language also the relationships to the linguistic expressions of the concepts have to become stable during his history of the learning.

Of course, also pre-linguistically experienced situations are ordered into similarity and contiguity structures, but language learning fixes or modifies these in accordance with the public vocabulary learned. For example, a common name strengthens the similarity relationships between its cases of application in satisfaction situations; a proper name strengthens similarity and contiguity relationships between the different occurrences of an individual object: syntagmatic relationships in recurrent small sentences, the so-called lexical solidarities, strengthen contiguity relationships, for example the ones between chair and sitting, dog and barking, etc. All this strengthening of class-identity or contiguity identity, for example the identity of an object, happens because in principle the use of the same expression in several satisfaction situations is an indicator of a recurrent identity. The linguistic utterances themselves are likewise partitioned and classified paradigmatically under similarity and contrast, and thus become utterance types or expressions, and they get at the same time represented in their contiguity relationships, i.e. syntagmatically within larger expressions. The linguistic and the situational systematisations on the growing sets of data constrain each other by correspondence within the data-pairs, such that situational information helps linguistic analysis, and linguistic information helps delineating linguistically expressed situational concepts.

Between the formalised structures in both types of models, **DCS** and **CNM**, there must exist a mapping from one to the other, which must be structure preserving with respect to the linguistically explicated semantic structures. Because of this mapping we may speak of a "representation" of our conceptual structures in the set of potential activation patterns, or corresponding dispositions towards producing these patterns. These dispositions are implicit within the state of the connection strengths of a network. As far as these networks are designed to model cognitively relevant structural aspects of the brain and succeed in this task, we can speak of implicit representations of our conceptual structures in the brain, which become explicit as internal extensions in a set of activation patterns, or net partitions, which we here shall call indicators of concepts. The structural relationships between activation patterns must have a formal relationship to conceptual semantics. In requiring such a formal relationship we are able, in principle, to capture formally the bridging relationship between the linguistic data and the perceptional data, as it must exist in the causal semantics of the understanding individual. The neuronal network's activation patterns form the bridge per occasion between the linguistic and situational data that are publicly co-ordinated and agreed upon. In the learning process they impose constraints on a person's structural relationships between possible brain activation patterns. The latter internal structures get adjusted in learning to the external linguistic and semantic structures by adjusting the connection strengths between neurones in such a way that the network embodies dispositions towards those activation patterns in which the external semantic relationships are preserved in the mapping from the external onto the internal structures.

2.2.1 The structure of Dynamic Conceptual Semantics

A model of **DCS** is a growing set of data on which two kinds of structures, classificatory and contiguity structures, are established in the process of concept formation. The structures are compatible. Each new datum is understood by being placed into these two kinds of structures in a way that preserves the stability of the structures. This holds for the structures on the linguistic data, the utterances, as well as for the corresponding situational data, the perceived situations. What stability amounts to will be explained below. These two kinds of structuring our data by embedding can be employed recurrently on different ontological levels of categorisation and identification, in discriminating and identifying qualities within situational impressions, situational impressions within (the life histories of) individuals, individuals within situations and constellations of individuals and situations.

The data are pairs (u, s), whereby u is an utterance and s its satisfaction situation. On the set of utterances U types of utterances are formed, and types of their parts are formed by substitution of parts in contrast analysis between utterance types that are partly identical. These types are sets of utterances and sets of parts of utterances, namely sentences and parts of sentences, and sets of sentences and sets of parts of sentences, which are the syntactic categories. They are formed as substitution classes of utterances and of their parts that can be substituted for each other in certain sentential contexts and context types. Categories are thus similarity sets of utterances or parts of these, whereby similarity is relational, i.e. given by an identity of sentential context. Categories are linguistic concepts, i.e. concepts of linguistic and semantic forms. Likewise types or similarity sets, i.e. concepts, are formed on the set of situations. These are situational concepts, or situation types. And as far as linguistically guided situational concept formation goes, types or concepts corresponding to linguistic parts are formed by a parallel substitution in contrast analysis on sets of situations. This implies that for non-negated parts of simple sentences we have corresponding sets of satisfaction situations, namely the satisfaction situations of those sentences in which the expression part occurs in a non-negated fashion. Thus for the phrase *a/the dog* we have a set of dog situations, namely those situations that satisfy sentences in which this phrase occurs.

To keep a concept *stable* means that its internal similarity degree under the relevant perspective is not diminished by adding new examples. - The similarity degree is reconstructible in a quantitative way, or by stabilisation on neural networks.- An evolving concept, a quasi-concept, expressed by *e* stabilises by growing towards a *conceptually complete* set of satisfaction situations for *e*. Such a set is conceptually complete in the sense that its internal similarity degree does not diminish any more by adding new examples of satisfaction situations of *e* under the relevant perspective. A conceptually complete set of satisfaction situations is a representative of the concept expressed by *e*, and the concept itself is the equivalence class of such conceptually complete sets. In dealing with concepts in Dynamic Conceptual Semantics we just take such representative sets as concept representations. The members of a speech community entertain the same concept expressed by *e* if and only if their concepts expressed by *e* can be thrown together into one equivalence class, which is the public concept expressed by *e*. This means that they accept each others' examples for the concept.

On the set of experienced situations, and especially on the growing set of satisfaction situations for sentences, and likewise on the set of the corresponding linguistic utterances, similarity based and contiguity based structures are established, which are compatible. Situations are categorised and identified by their place in both structures, i.e. on the one hand these are categorised by similarity to other situations, especially by identical parts in their internal contiguity, whereby parts get distinguished under identity and contrast with other situations, and on the other hand they are identified by contiguity relationships with other situations in space and time.

On a set S of situations a set structure is induced by forming the set of subsets of S on which the normal operations such as union, intersection, inclusion, complement $\{\cup, \cap, \subset, \backslash\}$ are declared, with the normal set theoretic axioms:

$A \cap B \subset A$, $A \cap B \subset B$,

$A \subset A \cup B$, $B \subset A \cup B$,

$A \cap B = B \cap A$, $A \cup B = B \cup A$,

$A \subset B \ \& \ B \subset C \rightarrow A \subset C$,

$A \cup (S\backslash A) = S$.

From all subsets those are selected that are similarity sets, i.e. have an internal similarity relationship under some perspective among their elements, and also those are selected that have contiguity relationships among their members. Hereby we get a similarity based structure and a contiguity based structure on the set of situations and on the set of utterances.

Similarity based structure:

On the set of experienced situations we form similarity sets as subsets of S. Σ is the set of similarity sets on S. Similarity is a relationship with the following axioms:

x similar x;

x similar y $\rightarrow$ y similar x.

Similarity is graded: x can be more similar to y than to z.

A *similarity set* is a set S such that $\forall xy$ ($x, y \in S \rightarrow$ x similar y).

A similarity set has an internal similarity degree D which is the maximal similarity among all the members of the set. If $D > 0$, then there is a *perspective* **P** under which a transitive similarity relation holds between the members of the set by virtue of an identical feature or identical relationship that comes into view under the perspective under which identities and differences are discerned. *Stability* of a similarity set or class means that its internal similarity is kept constant in adding new members.

> A *class* is a similarity set with transitivity of the similarity relationship which defines the similarity set as a class: (1) $\forall x,y,z \in C$ (x similar y & y similar z $\rightarrow$ x similar z)

A *class* is the extension of a *classical concept*. Here all elements have a common property or feature which defines the similarity as the identity of the class in opposition to other classes under the same perspective, for example under the perspective Colour, or Shape, or Behaviour, or a composition of these. A perspective provides a selection of aspects of identity and difference which come into attention in determining similarity and contrast under the perspective. To view something under a certain perspective does not require that one is conscious of the perspective or has a concept of the perspective, and it does not presuppose that one already has concepts of the aspects that play a role in establishing categories under the perspective. In being directed towards situations under a perspective, we can become conscious of the aspects constituting identity and difference under the perspective. Hereby they can become concepts which function as features in a conceptual analysis. Although a perspective can become conscious and be conceptualised itself, it basically is a way of bodily orientation in activities and action by which certain channels of perception, understanding, or analysis are opened or selected. In advanced levels of gathering knowledge and in theory formation, perspectives can be constructed and chosen by making conscious decisions to attend to these or that questions and aspects. A perspective can be thought of in extensional terms as a set of classes of examples which stand in opposition to each other under this perspective.

> X is *in opposition* to Y: $X,Y \in \mathbf{P}$ and $D(X) > D(X \cup Y)$ and $D(Y) > D(X \cup Y)$.

This means that the classes which are in opposition under a perspective have each a stronger internal similarity than their members have with other members outside the class under the same perspective. We also say that X *contrasts* with Y under **P**. The concepts under a perspective are formed by identities and contrasts.

Complex concepts (a term borrowed from Vygotsky 1986) are sets of satisfaction situations for an expression which can be seen as unions of classical concepts that are related to each other by similarity or contiguity between their members under a common perspective. Complex concepts are similarity sets without a class-defining identity, i.e. they are not defined by an identity which must be present in all members and must be absent in all non-members. A special kind of complex concept is one that is based on similarity only. Such a complex concept with at least a family relationship among the members is Wittgenstein's example of German *Spiel* ('play', 'game'). Different plays/games still have a common property, for example they are all activities. However, this property is not sufficient to define the whole similarity set, because it does not discriminate this set from other contrasting ones, such as reading or working.

There are even weaker organised sets of satisfaction situations for an expression which are based on similarity or contiguity (or both), but where there is no general perspective under which similarity is seen as an overall common identity. These have a much looser internal organisation by similarity than complex concepts; they cannot be understood as unions of classical concepts under a single perspective. Rather they merely are complexes of concepts with metaphoric or metonymic relationships between these concepts. Metaphorical relationships between concepts expressed by a single expression typically require a change in perspective under which the similarity is seen. Such complexes of concepts are *polysemic complexes*. They are chain and heap complexes (Vygotsky 1986); the concepts in them are related to each other as members of a chain or a heap by common similarity aspects or by contiguity between the elements of two related concepts.

The more regular complexes of concepts are *centralised heap complexes*, whereby each concept in the complex has common similarities or contiguities with a central concept. If a centralised complex of concepts in union forms a similarity set, then it is a prototypically organised complex concept, whereby the central concept is the prototype to which the others are each related by some common identities. The complex concept contains sub-concepts which each overlap, under a certain perspective, by similarity with the central or prototypical sub-concept. Let now Σ be the set of concepts.

> A *complex of concepts* is a subset Σ' of Σ, such that there is an $A \in \Sigma'$ such that for all $B \in \Sigma'$ there is a chain of concepts $X \in \Sigma'$ from B to A whereby each two neighbouring concepts in the chain are connected by similarity or contiguity, i.e. if X and Y are neighbours in the chain then for $\forall x \in X, \forall y \in Y$: x similar y or x contiguous with y in a constant contiguity relationship.

> A *centralised complex of concepts* contains a central class C such that $\forall X \in \Sigma'$: ($\forall x \in X, \forall y \in C$: x similar y) or (there is a contiguity relationship R such that $\forall x \in X, \forall y \in C$: x R y).

A *prototypically organised complex* is a centralised complex of concepts which is based on similarity relationships under a governing perspective and a couple of more specific subperspectives. This centralised complex is not based on contiguity relationships between the members of its member concepts. Its central member is a class that furthermore has to each other member a stronger similarity than each other member has to all the other members together.

> A *prototypically organised complex* of concepts thus contains a central member which is a class C such that $\forall X, Y \in \Sigma'$: ($\forall x \in X, \forall y \in C$: x

similar to y to degree D, and if x is similar to degree D' to all z, with $z \in Y$ and $z \notin C$, then $D' \leq D$).

The subperspectives under which the different member categories are formed also contain the prototypical member class such that their elements are comparable with the prototypical examples under the respective perspectives. This central member class is the prototypical concept. The union of the members of this kind of complex of concepts forms a *prototypically organised (complex) concept*, where the prototypical set is the central subset with which all the other subsets share a stronger similarity than a possible overall similarity across the whole complex.

C is a *prototype* in Σ' if and only if $\forall X, Y \in \Sigma', \forall x \in X, \forall z \in Y, \forall y \in C$: x similar to y to degree D, and if x is similar to degree D' to all z, $z \in Y$ and $z \notin C$, than $D' \leq D$.

Contiguity based structure:

There is a set **R** of contiguity relationships between situations, for example spatial and temporal overlap and adjacency, which define spatial and temporal connectedness, and also distance, and ordering in spatial and temporal space, cause-effect, action-result, means-goal, and other factual relationships. Our basic *ontology* of continuants is due to the most general spatial and temporal contiguity structure laid on sets of experienced situations.

m is *contiguous with* n = there is a contiguity relationship $R \in \mathbf{R}$ such that mRn.

There are chains of situations formed by space-time and causal connectedness. The space time connectedness is provided by overlap and adjacency.

C is a *chain* = $\forall s \in C$: $\exists s'$ such that s connected with s', and $\forall s,t$: $s \leq t$ or $s \geq t$ or s overlaps with t.

A *continuant* can be a patch of colour lasting for some time, or a portion of stuff, or an individual, or an event which is a process or non-permanent state or an action. It has internal and relational properties which define its identity through time. An *individual* is a special kind of continuant such that there is some continuous property during its existence and other properties which are changing slowly. Often the property which defines the identity is not known, it is rather merely assumed that there is such a property. However, the individual is identified by its slowly changing properties by means of which it can be followed in space and time as a continuous bundle of states changing only slowly, and by some constant relationships to other individuals. An individual in this model is its life history, and as a complete individual concept it is the complete knowledge of this life history, while the experienced individual is merely a partial individual concept, namely a part of this life history. We can have an individual concept of someone which is partially incorrect. If such incorrectness becomes evident it has to be corrected in a way that the partial individual concept can, in a consistent and coherent way, be part of the complete individual concept, the actual individual. Reference to an individual via a partial individual concept means that the partial concept is assumed to fit coherently into series of steps towards completion into the complete individual concept. *Stability* of a partial individual concept amounts to its growth into further partial individual concepts which are consistent and coherent within one possible completion, the complete individual concept.

An *event* can be point-like or can be a continuant. As a continuant it is a closed process with a temporal structure consisting of a beginning state, a couple of intermediary states, and an end state. An *action* is an event brought about by an agent due to motives and goals. For a thorough analysis of these ontological notions a large philosophical literature is available. To take it into account here would take us too far off. What is important rather is that our ontology arises by establishing a combination of similarity and contiguity based structures on sets of experienced situations. The *compatibility conditions* for these two structures imply that in historical concepts, i.e. in individual concepts and concepts of other spatio-temporal entities, both structurings co-incide: Each situation is *categorised* by embedding the situation into similarity sets, i.e. general concepts, and it is *identified* as a specific situation by embedding it into one or more individual concepts, partial life histories, or into relational contiguity networks of other spatial-temporal entities. Both structures intersect in determining, i.e. classifying and identifying, situations and more complex contiguity based entities in space and time, such as individual concepts and other historical concepts. Both kinds of embedding a situation have to be done into a stabilising system, or salva stabilitate within an already stabilised system of general and individual concepts.

In order to understand or analyse a situation, also as far as this takes place in experiencing a situation, we place it into these two kinds of structure, the similarity based and the contiguity based structure, by integrating it into similarity sets and contiguity sets salva stability. That both structures are *compatible* means that situations, as that what they are for us, i.e. as experienced situations, are determined by both structures at the same time. Both structures get developed on growing sets of data, i.e. develop with growing sets of experiences of situations, by ordering each new situation by similarity and contiguity into those subsets of situations which are already built up in the previous states of experience. In the course of development, the understanding of situations changes. Old experiences can be re-analysed in the light of new experiences, which give rise to new categorisations or new continuants into which they can be included.

Our sets of data grow. New data are added to old similarity sets and to old contiguity sets in a way that keeps these stable or coherent. To keep old structures stable implies to extend them by adding new similarity and contiguity sets for those data that do not fit into old sets salva stabilitate of their internal structures. This also holds for data that are nevertheless referred to by the same expressions as the old data with which they do not fit in a stability preserving manner. Such new language use is metaphoric or metonymic use of linguistic expressions. These newly created similarity sets partly overlap with old sets, on the account of similarity or contiguity. In correction of conceptual systems, on the other hand, old structures are partly cancelled in favour of new ones.

Conceptual growth comes about with conceptual change by extending and elaborating a conceptual system. Here complexes of concepts get extended by forming new similarity and contiguity sets from old and new data that do not fit together into one of the already existing sets salva stability. A preference for keeping old sets stable in their internal similarity and contiguity structure leads to forming new sets, based on some members of the old sets to which the new data are added because of their similarity or contiguity to these members, often under new perspectives. Thus similarity or contiguity under a new perspective leads to metaphoric or metonymic transfer of an old expression onto new data of a different kind. If the perspective remains constant in the transfer of a term to new examples and the internal similarity degree decreases, we deal with a case of concept broadening. If the labels of the old sets are used also for these new sets, either a polysemic complex of concepts is metaphorically or metonymically extended under a new perspective, or a concept is broadened under the old perspective. If new names are devised for these new sets of examples, the vocabulary is extended. - This all is extensively treated in Bartsch 1998. -

Basic syntactic structure: *The combination of similarity and contiguity structures in perceiving and understanding expressions and situations*

Our basic ontology, to start with, consists of experienced situations which are structured according to the impressions we have of the situations, our situational experiences. With growing numbers of such experiences they become analysed more and more by establishing more and larger similarity sets on this growing set of data. The ontology gets revised by embedding situations into both structures, similarity and contiguity sets. Contiguities give rise to further similarities, especially relationally based ones. Thus a banana is classified together with apples and pears as fruit because of its similarity to these due to qualitative identities of their contiguity relationships to eating, tasting, touching these things, and, with further experience about these things, to growing on trees and on other plants. Rice in a bowl, for example, is identified and classified together with rice as a plant by its contiguity relations with processing of rice seeds and growing rice plants. In this way we come to see situations and objects as belonging to certain classes and to the life histories of individuals or other historical entities.

Individuals, in this situation-based ontology, are sets of situations structured by several similarity and contiguity relationships. An individual is a completion of partial individual concepts, which are sets of situations embedded into a contiguity order of situations that typically holds within life histories of the special kind of continuant which this kind of individual is. Furtheron, similarity sets are formed over individual concepts, which have been established in this way. These are classes of individuals as there are the natural kinds and other general concepts over individuals. After having formed individual concepts, a situation can be seen as a part of an individual concept and thus can be perceived as a state, an action, or an activity in which the individual is engaged as an agent or in an other participant role. On this higher level of analysis and understanding, reversing the order of antecedency in concept formation, situations are now perceived not as basic, rather as being analysed according to our ontology of individuals, events, actions, states, and processes, which we have built up by combining similarity and contiguity structures in our uptake of situations. Since perception consists in embedding sensor-motoric data into these structures, the data as perceived data are already understood as such and such situations in the act of perception. With the dynamics of our conceptual systems also our perception of situations has changed over time, and thus our world as a perceived and understood world has changed.

The two types of structures, similarity and contiguity structures, in combination with certain restrictions, naturally or conventionally imposed for different kinds of entities, set up our ontology. The two types of structures are not only compatible but they together establish the coherent network of our historical and general concepts, as far as it is based directly on experience. Therefore a perceived situation, an individual, event, action, state, or process does not just consist of parts perceived due to sensoric stimulation within certain motoric contexts. Rather these stimuli are automatically embedded into similarity and contiguity sets and thus are analysed or understood in the process of perception. The perceived object contains additional aspects which are induced by the embedding of the incoming sensor-motoric data into general concepts in categorisation, and into historical, especially individual, concepts and in identification of the object. We perceive a car, although we do not really see the backside of the car. And we perceive the whole front of a car as a gestalt, although in rapid eye-movement we merely repeatedly scan over the front side, connecting stimuli of its parts in recurrent processes. By embedding the perceived data into contiguity and similarity relationships, the whole object is the perceived object, whereby what is perceived is hereby ontologically categorised and identified. Perception of something, i.e. what we perceive consciously, does already contain cognition of the object. This is a Kantian and Husserlian notion of perception, which also is implied in **DCS**, and also in **CNM**. What is perceived is not just a bunch of received sense data without categorisation and identification. In perception, data are always already embedded into previous organisations of data. Though at first we might merely have qualitative and quantitative categorisation and some space-time contiguity of situational impressions received via our senses, as soon as an ontological order is established on our growing sets of data, perception does not anymore consist simply of situational impressions within some qualitative and quantitative categorisations by similarity and contrast, and within some space-time contiguity relations. Rather by combining similarity and contiguity structuring in an ontological categorisation and identification, situations are then understood as built up of individuals and actions, events, states, and processes, in which the individuals are involved as participants.

The primary data are basic sentence utterances and situations. From a set of basic sentences with, in parallel, their satisfaction situations, whereby sentences and likewise situations are related to each other by identities and contrasts, concepts of the parts of sentences and concepts of the parts of situations are constructed in their relationships to the other parts of the sentences, and to the other parts of the situations respectively. Hereby also the relationship between the morpho-syntactically expressed (relational) linguistic case concepts with the semantic role-concepts are constructed in a context-dependent way; i.e. the cases can express different participant roles in the context of different verbs. An individual in a situation has a certain participant role, which is morpho-syntactically expressed by case marking, by pre- or postpositions, or by the syntactic position of the respective nominal term within the sentence.

Construing and understanding new basic sentences and new basic situations goes hand in hand: the structure of this process is recoverable in the structure of the result, the sentence and its satisfaction situation. In understanding a situation in perception, we can construe in parallel a sentence to describe the situation, and in parallel with understanding a sentence as a linguistic entity we understand its content by construing its possible satisfaction situation in a hierarchically organised process. The steps of the construction can be represented by **B**-intersections (Bartsch 1998), i.e. binding intersections, namely intersections between concept representations such that the intersection contains the smallest situation. This is the situational concept in which the combining concepts are bound together into a basic situation.

> A ***B**-intersection* of two concepts A and B is the intersection of a representative of A, i.e. a partial extension of A, with a representative of B, i.e. a partial extension of B, whereby the intersection is formed by adding the smallest situation s to both concept representatives salva stabilitate.

This means that we construe a smallest, i.e. basic situation s which must fit into A and B without impairing the stability of A and B. For general concepts, this implies that the addition of s does not diminish the internal similarity within the maximal, i.e. conceptually complete or stabilised, similarity sets under the relevant perspectives. For individual concepts, this implies that there is a part of it to which we can add s such that the coherence of the partial individual concept is not impaired by adding s. This is enough for understanding s as being part of a certain individual. For taking s as true more would be required, namely that situation s is embeddible into the individual, i.e. into its whole life history, preserving stability, which means that it is embeddible into the complete individual concept in a coherent way. Construing a smallest situation s in the intersection of two concepts is what in Kantian epistemology has been called the synthesis of concepts or categories by the power of imagination It is not simply the logical operation of conjunction or intersection. Rather both concepts are bound together into a smallest situation. This smallest situation is aspectualised to the composing concepts, i.e. its only aspects are the composing concepts. Note that the smallest situation defines the similarity measure of a concept, i.e. determines the similarity of a conceptually complete set of situations. It follows that any set of situations that contains a basic situation with this aspectualisation as a member is a conceptually complete set, i.e. a representation of the composed concept.

In imagination we can consciously represent the smallest situation, or the construed situational concept, by a series of imagined basic situations all containing the smallest situation as their aspectualisation. The concept construed by the **B**-intersection of two concepts can then be represented in consciousness by a set of situations each of which contains the smallest situation of type s as its aspectualisation. For connectionist networks, on the other hand, the difficult task is to devise an operation on the networks that will mirror this synthetic operation of binding and which might be considered as the kind of real neuronal network operation which has in our consciousness the result of the act of imagining basic situations which combine two, or more concepts. This problem will be tackled in section 2.2.2 and in chapter 3.

We shall presently see that the smallest satisfaction situation of a sentence is not construable without a minimal hierarchical structure of sentence construction. The result of understanding a sentence, and parallel of understanding a situation, is the construction of a smallest situation which is located as an element in a hierarchical architecture of sets of situations which are situational concept representations. Some sets are classes, others are selections from individual concepts. For example, the general concept Agent is represented by a conceptually complete class of situations in which someone is an agent, and the individual concept John is represented by a set of situations which are part of John's life history. Then a set of situations representing the combined concept John as Agent, a partial individual concept, is a subset of a partial individual concept of John, such that in these situations he is an agent. This is a subset of the set of situations that is John's life history. Within the selection expressed by *John as Agent* a step of a sentential hierarchy, for example of *John beats Paul*, is conceptually encoded. This requires that in understanding the sentence *John beats Paul*, the set of *John*-situations and the set of *Agent*-situations are **B**-intersected first before this whole is **B-**intersected with the set of *beat Paul*-situations. The last is construed by first **B-**intersecting the set of *Paul*-situations with the set of *Patient*-situations and then **B**-intersecting it with the set of *beat*-situations. It would be semantically equivalent to first construe *John beat*-situations and then form the **B**-intersection with *Paul as Patient*-situations. Likewise, but redundantly, we also could **B**-intersect the set of *John beat*-situations with the set of *beat Paul*-situations. Either way, the result is the same: The possible satisfaction situation for our sentence has to be a member of the set of situations that each contain the smallest situation construed by the above procedure, which is syntactically guided, though different syntactic routes lead to the same result. Of course, this procedure distinguishes *Paul beats John* from *John beats Paul.* The result of construing the smallest situation that has to be contained by all possible satisfaction situations of a sentence is dependent on syntactic structure.

In consciousness, relationships, especially actions, get classified by substitution operations employed on representations, whereby a variation of the participants of a relationship can be achieved on a set of representations of sentences and situations. We form constant one-place predicates by variation of the first place term, by subject variation, we get *x eat(s) bread* from sentences like *I eat bread*, *John eats bread*, *a girl eats bread*, etc. By variation of the second term, object variation, we get *I eat y* from sentences like *I eat bread*, *I eat pudding, I eat an apple*, etc. By variation of subject and object term we get the relationship *x eat y*. The relationship in fact is a growing set of sentences with their respective satisfaction situations, whereby the sentences and likewise the situations are related to each other by such substitutions of terms in subject or in object position, or by substitutions of the respective Agent-participants and Patient-participants in the respective situations. The internal similarity of the set is given by what remains constant under substitution.

B-intersections can be performed by substitution. In order to form the **B**-intersection between the partial individual concept John as Agent, i.e. a set of situations in which John is involved as agent, and the relational concept Beat, i.e. a set of beat-situations, we extend both with the smallest situation s that fits into both, salva stabilitate. This means we construe a smallest situation s that fits into the set John as Agent and into the set Beat, salva stabilitate. This now can be a Beat-situation construed from some other Beat-situation by substituting John for the Agent-participant, parallel with construing the sentence by substituting *John* for an other subject term in, for example, the sentence *Peter beats Paul.* An analogous step is taken to get the **B**-intersection between the set of John Beat-situations and the set of Paul as Patient-situations, and in this way the smallest situation in which John beats Paul is construed, parallel with substituting *Paul* for possibly some other Patient-term in, for example, *John beats Greta.* We can construe new complex expressions from old complex expressions by substitution, or variation, a method also called 'projection' from old to new expressions (cf. Ziff 1960). The operation of **B**-intersection is a parallel operation to substitution in sentences, now on the level of situational concepts, which are represented by stabilising or already stable, but still growing, sets of situations. Substituting Paul for Greta in a situation type or situational concept such that we get the situational concept John Beats Paul amounts to forming the **B**-intersection between the concept John Beats and the partial individual concept of Paul as Patient. Such substitution is the same as a unification of Paul as Patient with the Patient-aspect in the action concept Beat.

On the level of **DCS**, the sentence *John beats Paul* is understood by construing **B**-intersections between sets of situations representing the involved concepts:

A: John, i.e. a set of John-situations that form the partial individual concept of John;

A': John as Agent, i.e. a set that is the **B**-intersection between the concepts John and Agent, which is the set of all situations that contain the smallest situation in which John is Agent;

B: John Beats, i.e. a set of situations in which John beats somebody. It is the **B**-intersection between John as Agent and the relational concept Beat, which is a set of situations that contain the smallest situation in which John is Agent and the relationship Beat is realised. In construing the smallest situation the Agent-aspect of John as Agent is unified with the Agent aspect that is included in all Beat situations. That the agent aspect is included in all Beat situations means set-theoretically that the set of all Beat-situations is included in the set of all situations in which there is an agent, i.e. in all activity situations.

L: Beat, i.e. a set of Beat-situations

C: Paul, i.e. a set of Paul-situations that form the individual concept of Paul;

C': Paul as Patient, i.e. a set that is the **B**-intersection between the concepts Paul and Patient, which is a set of situations that contain the smallest situation in which Paul is Patient;

D: Beats Paul, i.e. a set of situations in which someone beats Paul. It is the **B**-intersection between the relational concept Beats and Paul as patient, which is a set of situations that contain the smallest situation in which the relationship Beat is realised and Paul is Patient.

E: John beats Paul, i.e. a set of situations that contain the smallest situation in which John beats Paul. It can be construed as either the **B**-intersection between B and C' or between A' and D.

E can also be construed as **B**-intersection between A' and C' followed by **B**-intersecting the result with L:

F: John as Agent and Paul as Patient, i.e. a set of situations each of which contains as a constituent the smallest situation in which John is Agent and Paul is Patient. This set then can be cut down further by **B**-intersection with the relation-concept Beat, which gives us a representative set of the situational concept of John beats Paul.

This account of alternative ways of understanding the simple sentence *John beats Paul*, by construing the situational concept it expresses in forming **B**-intersections, makes that the situational concept of John Beats Paul implies several set theoretic inclusions between the constituting concepts. The order relationships in the series of inclusions represent the possible alternative hierarchies of syntactic construction that all lead to the same result, a representative set of the situational concept of John beating Paul:

$A \supset A' \supset B \supset E$ i.e. E = **B**(**B**(**B**(John, Agent), Beat), **B**(Paul, Patient)))

$C \supset C' \supset D \supset E$ i.e. E = **B**(**B**(**B**(Paul, Patient), Beat), **B**(John, Agent)))

$L \supset B \supset E$ i.e. E = **B**(**B**(Beat, **B**(John, Agent)), **B**(Paul, Patient)))

$L \supset D \supset E$ i.e. E = **B**(**B**(Beat, **B**(Paul, Patient)), **B**(John, Agent)))

$A \supset A' \supset F \supset E$ i.e. E = **B**(**B**(**B**(John, Agent), **B**(Paul, Patient)), Beat)

$C \supset C' \supset F \supset E$ i.e. E = **B**(**B**(**B**(Paul, Patient), **B**(John, Agent)), Beat)

We see that the process of construing the complex situational concept is a gathering of constraints on the set of possible satisfaction situations for the sentence at issue; in understanding the sentence step by step, more and more constraints are placed on the set of possible satisfaction situations, by which the set becomes smaller. The process of understanding a sentence is thus a specification, which involves some fixed hierarchies, namely the combination of the role-concepts with the participant individual concepts. The concepts construed by **B**-intersection between two concepts are not empty, because either the smallest situation construed is already a member of both concept representing sets, or it is added to both sets, salva stabilitate.

We have seen that understanding a sentence follows hierarchies of constituents in performing the syntactic-semantic **B**-operations. For a simple sentence the combination between noun phrase concepts and participant role concepts has to be performed first. After that the formation of constituents can be done either way. This means that within the sentence the combinations of the sentence constituents are free, in other words, the main sentence constituent structure is flat. The traditional VP is not a necessary constituent in the built-up. The nomen-agent constituent can be connected first to the verb, and after that the nomen-patient constituent can be added. Even the nomen-agent and the nomen-patient constituent can be connected first in **B**-intersection, and then this new concept, John as Agent and Paul as Patient in an Action, can be specified further by **B**-intersecting with the Beat concept. Hereby the concept Beat is unified with the general action concept. Sentence constituent order is free and can be treated differently in different languages, without a difference in the conceptual semantics of the sentences. If we would also take into account extended sentences with adnominal phrases and adverbial phrases, we would have fixed hierarchies in the comprising nominal and verbal phrases, within which the adnominal has first to be connected to the noun or noun phrase, and the adverbial has first to be connected to the verb or extended verb phrase, according to the scope it has.

Since not all the participant roles in basic sentences are analysable by means of the general Agent and Patient roles, we can as well use the most specific roles for a given verb in the syntactic-semantic structure of basic sentences. In our example this would mean that we form the concepts of John as Beater and Paul as Beaten One by the binding operation and then bind these two concepts together into a smallest situation: **B**(**B**(John, Beater), **B**(Paul, Beaten One)).

In what follows we shall first formulate desiderata for any model of activation of neuronal connections and assemblies which also connectionist networks should have if they would be semantically adequate. This means that they should fit with the conceptual structures on the level of consciousness, as they are expressed here with respect to the dynamic conceptual semantics of basic sentences. Later on we shall evaluate some methods proposed in connectionist approaches for treating syntactic combinations in simple sentences with respect to their semantic adequacy. After showing them to be unsatisfactory, we shall make some suggestions about architectures of conceptual and episodic maps which provide richer models for how real neuronal networks in the brain could solve the problem of binding and of syntactic systematicity, and what role consciousness plays for understanding.

2.2.2 The desired structural properties of neuronal network processing.

Here I first formulate three desiderata for semantically adequate neuronal networks. Although existing connectionist modelling cannot live up to the second and especially the third of the desiderata, I shall suggest that human neuronal networks with some expression in consciousness can and must fulfil the desiderata.

First desideratum:

All the inclusions and non-empty intersections on the growing sets of experiences and construed situations in **DCS** must be reflected in **CNM** in the relationships between activation patterns that actually come up on conceptual maps in processing a sentence and in perception of suitable situations. Also there we have relationships of inclusion, union, and overlap.

A conceptually indicative activation pattern can be widely distributed over a compound network consisting of several parallel conceptual maps. Such maps can be maps devoted each to a different aspect of visual processing and maps representing categorisation of input to other sensory modalities. For example, all green things activate patterns that have at least in common the activation of the pattern that is indicative for green. Activation patterns due to reception of input data of green apples have also an intersection in common with activation patterns for input data of other apples, and fruit in general, and an intersection with activation patterns for other input of round things, and of green things.

With perception we find widely distributed activation patterns. The activation for seeing a car, for example, is spread over maps for form, movement, colour, and pre-motor maps for the action of driving a car. The activation of such parallel lower level maps being areas of perceptive modalities, of motoric activities, and other proprioceptic feelings can be composed together in a compound map. A state of activation of the compound map can be represented by a vector consisting of a concatenation of the vectors representing the activation states of the submaps. This can be a new input vector for higher level maps, where then a more compact activation pattern corresponding to cars, or to green apples can be found, which is merely locally distributed. Each submap, for example, can be understood according to the model of Kohonen maps, the compound map as a compound of Kohonen maps, and the higher maps can be again organised according to the Kohonen-model (Bartsch 1996, 1998): On a non-compound higher map we don't have intersections or inclusions between activation patterns corresponding to concept union and hyponomy. The concept indicating patterns there are not widely distributive, rather locally distributive over several neurones. Each concept indicating assembly of neurones is centred around a neurone which matches strongest with the kind of input, the best matching unit (BMU). Thus for roses and for tulips different assemblies would be indicative, though with the strong activation of the rose-BMU and its surrounding also the tulip-BMU and its surrounding might be weakly activated due to the general flower properties which roses and tulips have in common. The general concept of a flower would be 'represented' on a higher order map in which the overall flower input is registered in contrast with tree-input, bush-input and other classes that are given as input in contrast with flowers. All this input can be contrasted with animal-input and thus on a higher map we get contrasting BMU-surroundings indicative for the category of plants and for the category of animals. Thus more or less deep taxonomies might get formed by the availability of conceptual maps, whereby the compound input from the lower maps leads to classifications of the next higher map. The lower and more compound the maps are, i.e. the closer the maps are to the perceptive and proprioceptive maps, including motoric maps, the wider is the distribution of the activation patterns that are indicative for certain inputs. The wider the distribution, the better it is possible that similarity between concepts can be indicated by overlap of patters and that implicational relationships, hypo-relationships in taxonomies, can be indicated by pattern inclusion.

We require of conceptual maps that with the activation by data input from apples also data input from fruit generally is activated, and that on the lower compound maps the activation patterns for apples and pears overlap more than those for apples and bananas. On the other hand the higher level concepts have a function in that they create a stronger coherence between certain features by tying these together into a more general concept. Without this, activation would spread in almost equal strength to all sides; this would result in masses of unwelcome associations of lower concepts or features, because every two feature indicators that are activated together repeatedly get connected and the activation spreads from one to the other. A higher level concept A can bind features {a,b,c} on a lower level together by a circuit of activation between the A-indicator and the a,b, and c indicators (cf. chapter 3). Hereby the indicators for a,b, and c get stronger activated, highlighted or concentrated, among the other possible associative activations going from a to d, d to e, b to f, f to h and c to k and k to l and l to m, for example. This concentration also makes it possible to associate a linguistic item to {a,b,c} via concept indicator A, because hereby this set of features is delineated and highlighted among other feature indicating assemblies which are activated weaker than {a,b,c}. In chapter 3, I shall show how this principle of circuit forming with indicators on higher level plays a role in assigning lexical items to concept indicators and to processing constituents of basic sentences. A layered architecture of conceptual maps with circuit activation will enable us to combine both requirements, the one of having available logical relationships between concept indicators, and the one of singling out aspects of conceptual representation from a mass of possible associations and relating this selection to associated lexical items in forming a describing sentence.

Generally, intersections in **DCS** have to be mapped into unions in **CNM**, unions have to be mapped onto intersections and inclusions have to be inverted. In this way we should find potential activation patterns in **CNM** corresponding with the similarity based structures in **DCS**.

The mapping from **DCS** into **CNM** should be:

(1) $f(A \cap B) = f(A) \cup f(B) = A' \cup B'$, for example: for green round things we should find the activation pattern for green and the one for round activated, i.e. the patterns for green round things all should include the union of these two patterns.

(2*) $f(A \cup B) = f(A) \cap f(B) = A' \cap B'$, for example, for something that might be a donkey or a horse, is very similar to both, we should find the intersection of the horse and the donkey pattern activated.

(3*) $f(A \supset B) = f(A) \subset f(B) = A' \subset B'$, for example, if roses are flowers, then we should with the rose pattern also find the flower pattern activated.

If we represent the activation pattern on the assembly of units in a net by a vector, union there is vector addition and intersection is simple vector multiplication. The workings of the network should be such that the projection f from semantic conceptual structures to activation patterns on connectionist network structures is satisfied. But this requirement is too strong in this form: It holds for intersections of sets of satisfaction situations (see 1) because with growing specialisation the set of features grows that characterise the situation type, and these features somehow must be represented in the assemblies of activated neurones the net comes up with on the basis of the previous learning process. But it does not generally hold for unions of sets of satisfaction situations (see 2*): If we here intersect the features, too little remains to restrict the set of alternative satisfaction situations. Take, for example, 'apple or banana'. The intersection of features is not more than the features for 'fruit'. These would mean that at best the activated assembly in the intersection would accept any satisfaction situation in which a fruit is prominent. Therefore we need in the indicator of the alternative 'apple or banana' both indicator assemblies to be preserved, but in a way that this assembly does not imply one of these, except if the other one has already been excluded. If the net somehow should represent the satisfaction conditions for the expression, the net would have to represent a test procedure of the form: Test whether the apple-activation pattern is activated. Test whether the banana activation pattern is activated. If at least one is yes, then signal true, i.e. approve of the expression 'apple or banana'. The 'or' is not represented on the net, rather it means that a test operation has to be executed over the net. We thus have to exclude (2*), except for disjunctions in recognising things which are very similar, i.e. more similar to each other than to other kinds of things, such as apples and some fairly roundish kind of pears.

Also (3*) is problematic: A widely distributed pattern for rose, should comprise the one for flower. In this, (3*) his is right. But if the pattern for John as Agent should comprise the one for Agent, we get into problems. Our semantically motivated desideratum would have the consequence that with the activation of the Agent-pattern also the Agent-part in the indicator of Paul, who has the patient role in the sentence *John beats Paul*, would be activated. This would make it impossible to solve the problem of syntactic systematicity, namely that the network distinguished "John beats Paul" from "Paul beats John". The reversal of the inclusion relationship becomes problematic as soon as individual concepts are involved, which at different times in their life history have different properties and different roles which are incompatible at one point in time. In their indicator the different roles they have taken are represented as possibilities. This could be overcome if we would use an architecture of several maps in which the Agent-concept would also be represented on a higher-level map instead of merely being distributed over parts of all individual concept indicators that contain an Agent aspect and all indicators of verbs that have a frame with an Agent-participant. The higher level Agent-indicator then has to activate the agent-part in the John-indicator additionally in a recurrent circuit such that we get here a stronger activation than for the agent-part in the Paul-indicator, and the patient part in the John-indicator. The architecture can be such that it separates stages in the process of understanding according to major constituent structures, which are finally bound together in an episodic map, which has an expression on the level of consciousness. In chapter 3 such an architecture will be discussed and proposed as a solution to the problems of binding and systematicity.

Second desideratum :

Contiguity relationships generally might best be represented by sequential associations of activation patterns involving recurrent stimulation in temporal order. A special kind of contiguity is contiguity within a sentence and within a situation. Contiguity within sentences are the syntagmatic relationships which involve syntactic systematicity. This problem for **CNM** has been discussed by a.o. Fodor 1997. In this chapter I shall concentrate on the problem of syntactic systematicity in **CNM**, because it is insufficiently dealt with in these models.

Third desideratum:

Syntactic systematicity should be reflected in **CNM** such that the semantic properties, especially implications, of sentences are preserved. Before reporting about how that problem has been tackled, I shall first attend to the following questions:

1. What should the relationships between activation patterns look like in the process of syntactic-semantic understanding?

2. Should there be an activation pattern which is the result of the process of understanding and thus would that be something like a network internal semantic representation of a sentence?

3. What should the process of understanding a sentence be in terms of network architecture and mathematical operations?

I shall deal with the questions in that order. Conceptual Semantics, as it has been sketched in the previous section, places restrictions on what should happen in neural networks, will they have any chance of performing their tasks in the process of relating linguistic expressions to external semantics, and generally of being attuned to external situations and tasks. The process of understanding in **DCS** was one of gathering constraints on possible satisfaction situations, represented by **B**-intersection on concept representing sets of situations, which amounts to cutting down on the sets in a process of specification by intersecting expansions of these sets which are stability preserving, and by binding concepts together into basic situations by unification in an act of synthesis. At least some hierarchies in constituent formation were fixed. The fixed, endocentric, constituents have to be processed first, before they as wholes can be connected to other constituents in exocentric constructions. Except for these fixed hierarchies, the processing just can proceed in any order of major exocentric sentence constituents, whereby the concepts corresponding to these major constituents being processed first, are taken as units which get bound together in a concept of a situation, or a situation-type. How could such process of sentence understanding be mirrored in connectionist networks? We shall focus on the problem of binding concepts together and on the problem of syntactic systematicity.

Ad question 1: The formation of **B**-intersections should have the formation of unions of activation patterns as correspondent, because the activation patterns can indicate collections of features, established in the process of learning on the internal net, and mostly unknown to the language user. Concept indicating assemblies get activated in causal dependence of certain input features we may or may not be conscious of. They are internal indicators of outside features which we experience according to the way in which they have been processed or 'understood' by the network on the basis of its previous learning history, and according to the processing of the context available in the situation of experience. An internal activation pattern, due to an input and its effect on the established connection weights, can be represented by a vector which has as dimension the units of the internal net. In a real neuronal network we can think of such potential activation patterns as restricting and enriching our perception of possible external semantic correspondents, which are the input of the activation pattern. These internal feature indicators are learned and unconscious. They are activated by an input example according to the connection weights built up in the learning process. Assemblies as feature indicators, i.e. groups of activated neurones, must play a role in our conscious perceptions of input examples. They must have some indirect expression in consciousness by determining how we perceive sensory input in terms of consciously perceived features and relationships in and between the input examples; they are causally involved in recognising semantic properties and relationships. This causal connection induces the formulation of semantical adequacy for processes and activated patterns on a neural network.

Under a semantic point of view we desire the following structural properties realised by fully distributive 'representations': The greater the activation assembly or area the more restrictions or features it may represent in principle. The pattern for Flower should be the intersection of the patterns for the different kinds of flowers, and every activation pattern for a specific individual flower should comprise the pattern indicating its kind. Thus the pattern for roses comprises that for flowers generally. The Rose pattern should contain the flower pattern and also encode additionally more specific Rose features. Likewise John as Agent should be indicated by a potential activation pattern that contains the one for John and the one for Agent. The processing thus has to be such that potential activation patterns are added to each other. For example, first the pattern for John should be activated and then, according to the case marking or syntactic position of the phrase and the kind of verb, the Agent pattern should be activated. Likewise the pattern for Paul should be activated, and to this the pattern for Patient should be added. Furthermore, the pattern for Beat should be activated additionally. Internally, it always has some parts that indicate a generalisation of the Agent-participant, and some other parts that indicate a generalisation of the Patient-participant. The activation of the Beat pattern will again provide some activation for the Agent pattern that had already been activated additionally to the John pattern. Likewise the activation of the Beat pattern will imply some additional activation of the Patient pattern that is also activated additionally to the Paul pattern. This overlapping of activation can be seen as a unification: the Agent part within the Beat pattern gets unified with the especially activated Agent part in the John pattern. But this simple pattern union on a (compound) neural map does not work for syntactic systematicity because with activating the Agent-pattern in connection with the John-pattern it automatically also will be activated as being a part of the Paul-pattern, and with the activation of the Patient-pattern in connection with the Paul-pattern, it automatically will be activated as part of the John -pattern. Without taking into account the temporal order of processing we are stuck with the problem of syntactic systematicity; 'John beats Paul' could not be semantically distinguished from 'Paul beats John'. We shall see in chapter 3 that it does not just require temporal order; rather an architecture of neural maps and circuits of activation between such maps should do the job for binding and syntactic systematicity.

Ad question 2: Patterns of distributed 'representation' or indication of concepts are distributed over several parallel submaps, which together form a compound map. There should, though, be a unity that connects the John pattern and the Agent pattern by binding them to each other. Such binding operation, in principle, should hold for the whole sentence: the concepts that constitute a basic sentence should be bound together in the representation of a basic situation. But note that a simple union of all the involved concept indicating patterns would not do, because it neither can bind the concepts in the way required nor can it solve the systematicity problem.

A 'representation' of the satisfaction situations for a sentence which is sensible to structure must be such that the semantically relevant processing steps are somhow still recognisable in the overall pattern. Thus, we must

(1) either claim that a resulting final 'representation' of the conceptual semantics of a sentence by means of a potential activation pattern does not exist in a meaningful way, and that the syntactically structured procedure of connecting activation patterns is all there is to sentence understanding: Then the path to the goal would be the goal itself, and the resulting activation pattern on a net would be of no interest. Instead of a final 'representation' of a proposition there then is merely a series of activation patterns, as in type (4) connectionist networks according to Ramsey (1992). But note that in this series the binding requirement must be satisfied: the concepts must be bound together in the right way such that the perception of a situation in which John beats Paul is distinguished from a situation in which John does something else and Peter is likewise present and he beats Paul. In the perception of both situations the patterns for John and Beat get activated, but only in the first the Beat-concept is bound to the John-concept. Only the first situation may be recognised by the network as a satisfaction situation for the sentence *John beats Paul*.

(2) Or, if we want to cling to representations of propositions on a connectionist network, we must see to it that there are connections between successive patterns which are not just the union of two patterns, but imply a modification of one by the other, as has been proposed by Elman (1990, 1993, Garson 1994). Methods for trying to do this, as there are Elman's vector concatenation in recurrent processes, and Smolensky's forming of tensor products of vectors, will be discussed in the following sections 2.3.1 and 2.3.3. The problem with these proposals will shown to be that the relevant semantic properties are lost in these

'representations' and therefore they, by themselves, cannot be indicative for the possible satisfaction situation of a sentence. I shall show that such 'representations' are insufficient.

Our question right now is whether a finally resulting activation state that is typical, i.e. unique, for a certain proposition, is indeed necessary for understanding a sentence and relating it in interpretation to a model or piece of world. This requirement might be too strong. If we want to have some fixed representation for a proposition, we could best take an inscription of the sentence itself, or some semantically equivalent sentence, or some imagined or produced picture of the satisfaction situation. These would be representations in the full sense of the word, and they are part of our consciousness. Such a picture can be seen as an expression or articulation of the steps in the process of understanding the respective sentence, whereby a binding operation on the neuronal network expresses itself in seeing or imagining John as Agent and Paul as Patient in the Beat situation. There is no need for a final or single activation state of a connectionist network that should be the representation of a proposition, and which would be stored as such in some memory space. To require such a representation is a relict from the traditional cognitive model, where fixed mental representations were assumed to be construed and stored in memory as something that can be used and re-used in various contexts of processing. Instead of this, our memory either may be able to link up to a describing sentence presented, or to an approximate picture of a certain situation, or it may initiate forming a new representation of this kind. Thus, if we do not want to assume that the network somehow makes copies of resultant activation states and stores them somewhere in a big mental library, we have to think of a complex neural network (architecture) as being itself modified in its connection weights in such a way that experienced episodes are kept in it implicitly such that the respective patterns can be activated again if some parts of them get activated by some appropriate stimulus. Within the state of connection weight distribution in a network architecture, including the connections to the sensors, there is no representation to be found of a particular episode or of a particular sentence, although it must have had some effect in strengthening some connections and not others. The distribution of connection weights or strengths has been changed somewhat, such that with a current appropriate input the system is stimulated to bring about the representation.

In modelling understanding we shall restrict ourselves to those parts of the process of understanding in which brain activation patterns for the syntactic-morphological concepts and lexical concepts are run through constituentwise in connecting the sentence to the brain external public situations or to imagined situations. The sentence is in this way seen as a program or series of instructions for checking whether a situation is a possible satisfaction situation for the sentence. What is then kept as a representation in our mind is the image or percept of the sentence itself in its, in principle, public connection with situations. That connection must somehow be represented implicitly in the brain in a generalised manner as routines or procedures implicit as dispositions in a connectionist superstructure that works over the neuronal maps that indicate historical and general concepts. In connectionist processing the implicit procedures merely appear as an actualised sequence of steps in the concrete realised procedures, which serves in connecting the sentence with the external semantic conditions by measuring the sequences or rather circuits of activations, which it causes during the process of understanding, against the perceptual or imagined input which is hereby either accepted or rejected as a satisfaction situation. The syntactically organised processes select the relevant aspects in the perception of satisfaction situations in two counter-running processes of activation, namely by understanding the sentence and by perceiving the situation. these two processes are measured against each other. In case of fit the sentence is satisfied by the perceived situation. The procedure need not be frozen up in any resulting state of the neuronal network. Momentary results of the procedural steps may have as a conscious expression the sudden experience of having understood parts of the sentence or having imagined or perceived a respective satisfaction situation. The circuit architecture of neural maps in chapter 3 is designed to have this property.

Ad question 3: Connectionist networks typically work with operations on vectors, which are used to represent input and output, and also can be used to represent the activation of the internal net, such that all the hidden units are represented as dimensions of a large vector, which can have the values 1 or 0, depending on the state of being activated or not. An activation pattern then can be represented by such a vector which, of course, will be very long. In this purely theoretical discussion let us assume that lexical expressions, but also compound expressions, trigger activation patterns on the neuronal networks. The patterns have been formed as results of dispositions acquired in the course of becoming acquainted with the surroundings and in learning language associated with perceiving situations and performing activities. These dispositions are implicitly encoded in the connection strengths of the network. An activation pattern comes about according to the acquired dispositions by taking in and processing appropriate input, namely input of the kind that has served previously in building up and stabilising the occurence of this pattern in reaction to the appropriate input. What was discussed above already implies that for syntactically compound expressions certain operations should give a structure to the brain process of building up the resulting conceptual semantic representation in consciousness, which a resulting activation pattern by itself cannot express. The structure cannot be found back in such a resulting pattern; rather it has to be found in the process of producing and reproducing it following the syntactic form of the representing linguistic expression. In accordance with this view we can also assume that syntactically compound expressions which have been learned as wholes in situations, may correspond to a situationally caused activation pattern which can only then be an activation pattern indicative of an understood situation, if in its production a contiguity based, i.e. a syntagmatically organised situational built-up is followed. This also would mean that a situation can only be perceived in combining situational constituents according to contiguity relationships in space and time. - Of course, such a parallel between a syntagmatically organised understanding of sentences and of situations is a speciality of human perception and understanding, not found in speechless animals. But it leaves open the possibility that animals understand situations as organised by basic contiguity relationships and a by a restricted repertoire of general concepts and individual concepts. -

2.3 The problem of systematicity and binding in simple connectionist networks

The problem of syntactic systematicity arises because simple commutative and associative addition of patterns represented by vectors consisting of ones and zeros lead to a mix-up, such that sentences like *Paul beats John* and *John beats Paul* cannot be distinguished. We shall now see what operations on vectors can do, which have been proposed to deal with syntax in connectionist models. What network architecture can do will be illustrated in the next chapter.

1. Vector concatenation with recurrent fibres

Vector concatenation means that the vectors of two sentence parts in a sequence are fed into a net in such a way that first the vector of the first part is fed in with a certain activation result on the net. Then in the second circle, by recurrent fibres, a new vector consisting of this result concatenated with the vector of the second part is fed into the net. Like with all input vectors, the concatenated vector is per dimension multiplied with each of the connection weights. The sum of all these products per connection is measured against the threshold of the receiving unit of the connection in order to cause an activation of this unit or not. - See Garson 1994 and Elman 1990 for the details of representing information built up in time by recurrent networks. - If we have the vector for John on the net and then concatenate it with the one for Agent and feed this concatenation into the net we should get a modification of the John pattern whereby in contiguity with the John pattern some assemblies of units are activated due to the Agent-part of the vector concatenation that has been fed in. This seems to solve the problem of syntactic systematicity and provide for binding. In the course of time, i.e. if John has often been described or perceived as an agent, then with just activating the John-pattern also its associated Agent-assembly will be somewhat activated. This means that we expect him to be an agent, even if that is not explicitly stated. This method, if it can be refined enough to work in general, will result in networks in which contiguity in situations and between situations is represented by contiguity or nearness on network maps, or at least by association between activation patterns. These maps could be maps representing episodes and individual concepts. The representation of general concepts then would be distributed over individual concepts. On a second map the results of the first map could be generalised such that we there would receive activation patterns of locally concentrated assemblies of units as 'representations' or indicators of general concepts. We shall look at such a model later, in a more advanced architecture. The method of using recurrent fibres is essential for treating phrases in contexts and may be useful for treating recursion in linguistic structures as substitution of phrases into the context of other phrases, a problem not dealt with in this chapter, which concentrates merely on systematicity. How far this method in a more involved architecture of networks can also be semantically satisfactory, i.e. preserves implications and secures recognition of satisfaction situations will be discussed below, in chapter 3, where circuits of activation play a role in representing constituency.

2. Vector addition and vector multiplication

It is clear that simple pattern union, i.e. vector addition, gives rise to the systematicity problem. We shall now investigate whether we can use pattern intersection, i.e. vector multiplication, as an additional means to secure binding of concepts and by this also solve the systematicity problem. Let us assume that on the neural net we have individual concepts, i.e. concepts of individuals, 'represented' by locally contiguous activation patterns which indicate properties the individual has shown in the course of time. General concepts, such as Agent or Man are 'represented' globally distributed over all the activation patterns indicating individuals that have exemplified these general concepts.

Let us consider our sentence *John beats Paul*. The activation pattern caused by the utterance part *John* causes the corresponding activation pattern on the conceptual map. The activation pattern for the concept of John contains sub-patterns for Agent as well as Patient, because John has often been experienced as agent and often as patient. This, of course, gets generalised for all humans and also animals. By the subject marking of *John* in the context of the active voice of the verb, in English its pre-verbal position, the Agent pattern is additionally activated. In fact we here have an addition of the John activation pattern and the Agent activation pattern, which can be expressed by adding the two vectors that represent the activation by John and the activation by Agent. The vectors have as dimensions all the units on the conceptual map, and the values in these dimensions are the activation values. By performing vector addition, <JOHN> + <AGENT>, in each dimension the values are added to each other with as result a new vector which is the sum from the two. On the net, the addition results in additional activation of the Agent part in the John pattern, but unfortunately also in the activation of all other of the Agent-pattern's distributed parts, thus also the Agent-part in the Paul pattern. We rather want that only the Agent part of the John pattern gets activated. In order to achieve this we might want to form the intersection between the John-pattern and the Agent-pattern, which amounts to the simple vector product <JOHN><AGENT>. Hereby we have constructed the activation pattern for John as Agent which contains all the assemblies of units that are indicative for the features which John has as far as he is an agent. Since this would mean that we loose the features of his blue eyes and blond hair, because they have nothing to do with his being an agent, we have to add the vector <JOHN> to the product. This gives us <JOHN> + <JOHN><AGENT> as a representation for John as Agent. Proceeding like this we would get as a representation for the whole sentence: <JOHN> + <JOHN><AGENT> + <BEAT> + <JOHN><AGENT><BEAT> + <PAUL><PATIENT><BEAT> + <PAUL> + <PAUL><PATIENT>.

The resulting activation pattern is just the same as what we get by simple vector addition, i.e. pattern union, which is unsatisfactory. A difference can merely be detected in the process of construction. The overlapping parts expressed by vector multiplication get additionally activated. The Beat activation pattern generally has as sub-pattern a part that is due to the agents perceived in the set of Beat situations in the learning process, and it has a part due to the patients perceived in the same situations. By multiplying the vector for John as Agent with the vector for Beat, we get those assemblies which are indicative for the features John has as a beater. Likewise we get the part of the pattern which is indicative for the features of Paul as a beaten one. We in fact have a union of the pattern for John as Agent and the Beat pattern, whereby the overlapping parts are especially activated. There is a unification of the two patterns achieved by the overlapping part, namely the Agent pattern, which is identical in both and gets by repeated activation especially activated. Likewise the activation pattern for Paul as Patient is called up by multiplication of the vectors for Paul and the one for Patient, and it is united with the pattern for John Beats. Here the especially activated Patient subpattern in the pattern for Paul as Patient is additionally activated by vector multiplication with the Beat-pattern, where a Patient subpattern is included. The result should be a pattern for the situational concept of John Beats Paul. In fact, in the resulting pattern we cannot distinguish John Beats Paul from Paul Beats John, because it does not preserve the way in which it is constructed. But if we take into account the process, the systematic semantic distinction between John beating Paul and Paul beating John can be made. In vector addition and vector multiplication we have a semantically meaningful operation, whereby the process must account for syntactic systematicity. In the resulting pattern the semantic distinction made by syntactic systematicity is lost.

Note that the sentence *Paul, John beats* likewise is understood simply in the temporal order of its constituents. No transformation is required. But here first an overview of the sentence is necessary to see from the pre-verb position that *John* is the Subject. Since *John* is such marked as the Subject, *Paul* must be the Object. Then the pattern for Paul as Patient and for John as Agent are activated, and with the union of these, the pattern of Beat is united, whereby its Agent part additionally activates the already especially activated Agent subpattern of the John pattern in construing John as Agent and the Patient subpattern of the Paul pattern in construing Paul as Patient. Also here, the overall activation pattern finally resulting does not make the necessary conceptual distinctions, though the process of construction does.

If somebody would have mainly perceived beat-situations in which John does the beating, this would strengthen his concept of John as a notorious beater and his concept of Beat situations as situations in which typically John is actively involved. But if he never has seen John as a beater, there is a problem with the above representation in terms of a sequence of vector additions and vector multiplications: Some of the intersections can be empty, i.e. the vector product can be zero. If the network has never experienced John as an Agent the product <JOHN><AGENT> will be empty, if the network has never experienced John as a beater the product <JOHN><AGENT><BEAT> will be empty. Still, somebody who never has experienced John as a beater also must be able to understand our sentence. The above representation therefore cannot be achieved as an activation pattern in the process of understanding the sentence, rather it expresses conditions the network puts on possible satisfaction situations: a satisfaction situation for the sentences has to be such that its effect on the net, i.e. the pattern it causes there, contains at least what is expressed in the above representation. Not only must the products not be empty, but the situation-specific “John as Agent” features must include the Beat-agent features:<AGENT><BEAT> <AGENT><JOHN>. This requirement cannot be represented on the net as an operation on the activation patterns that are available in the process of understanding the sentence. The products could be empty. Rather the requirement has to be stated over the network: the network may accept an input situation as satisfaction situation only if the requirement is satisfied by the activation it gets from the input. Formulated in this manner, understanding a sentence means performing a process of generating requirements to be put on the net in case of perceiving a situation as satisfaction situation. The sentence then is the formulation of these requirements which can be activated in a sequence of instructions for checking on the activation of the net when it gets situational input which it has to accept or to reject as a satisfaction situation for the sentence. We shall see above how this can be achieved in a proper architecture.

The classical position in traditional cognitive science is that there must be a semantic representation, which must exist at some moment as a whole, and a copy of it can be kept and stored in an explicit memory. If we abandon this picture of a singular representation which can be stored in a memory like in a library, we still can be content with what is achieved, because we have pictures and sentence tokens as linguistic representations in external stores on paper, in notebooks, in computers, and in libraries, and we can call them up in reality and in imagination, consciously in our minds by stimulation through tokenings of some of their parts. This does not mean that there must be an unconscious part of our mind, as a kind of library, in which they would be stored in explicit forms. Rather they merely must be implicit in the connection strengths within our neuronal networks such that they can be called up by some appropriate linguistic or situational stimulus, resulting in a process of coming up with circuits of activation patterns with some articulation in consciousness, i.e. in a mental representation, which is a sentence token or a picture in our mind. - Note that our mind is constituted by what is conscious. Everything else is merely brain and the states the brain is in. - On our linguistic and pictorial representations we can do all kinds of explicit logical operations. We have an implicit memory in the connection strengths in our neural nets, which can come up with activation patterns in causally established connections with input from linguistic utterances or situations. Of course, the temporally ordered patterns occuring in procedures, like all other kinds of routines, must in principle be results of the capacities of our neural net architectures and must have been specified in learning processes. Then the presentation of an external linguistic or pictorial representation in perception, or the imagination of such a thing, can activate the relevant procedures of conceptual processing and episodic imagery.

3. Tensor product

Smolensky (1990), discussed in Horgan and Tienson (1992 and 1997), has made a proposal to treat sentences like the above ones by means of the tensor products between vectors, and not by simple addition. The reason for using this more involved operation is the systematicity problem, namely that the resulting conceptual activation patterns for a sentence like *John beats Paul* cannot be distinguished from those for *Paul beats John*, if we just use pattern union, i.e. vector addition.

In the tensor product, on the other hand, an n-place vector and an m-place vector are connected by multiplying each value of one vector with all the values of the other vector. The result is represented in a vector with n times m dimensions, or in form of a (n,m) matrix, having n lines and m columns. For vectors representing patterns on a net with n units we get as result vectors with n times n dimensions, or (n,n) matrices. For sentences like *John beats Paul*, Smolensky forms the tensor product of the vector for John and the vector for Subject, the tensor product of the vector for Paul and the vector for Object, and the vector additions of these with the vector of Beat as verb, which is also a tensor product. The result is something like

(*) [**v**(John) x **v**(subj)] + [**v**(beat) x **v**(verb)] + [**v**(Paul) x **v**(obj)].

Generally, the semantic vectors of words are tensor multiplied with a vector for the grammatical category they are used in. We see immediately that this combination of tensor product and vector addition does satisfy the requirement of syntactic systematicity, but does not satisfy the binding requirement. The participants are specified as agent and as patient, here as subject and object, but they are not bound to the verb, because vector addition cannot achieve this. A net with (*) as an activation pattern, which must be the criterion or indicator for accepting a possible satisfaction situation, would also accept a situation in which John beats Peter and Mary kicks Paul. The vector addition in (*) cannot secure that only a situation containing the basic situation, i.e. the smallest situation, in which John is agent, Paul is patient and the action is beating is accepted as a satisfaction situation.

Moreover there are some further problems for this 'representation' in that the tensor product does not make sense as a semantically relevant representation for John as subject, Paul as object, Beat as verb, if subject, object, verb are just concepts of linguistic theory. The activation value in each dimension of the vector for the individual concept of John is multiplied with the indicating vector of the grammatical category 'subject'. However, if a semantic concept of subject is meant, something like Agent or Carrier of a property, then each positive value of the John pattern will be multiplied with that. An assembly of such activated units may stand for a known feature of John. Assume that an assembly of activated units is indicative for John's blue eyes, then also the value of these units will be multiplied with all the values in the dimensions of the Agent-vector, and where these vector dimensions have the value zero, the John-values will be multiplied with zero and disappear altogether. That does not make sense. The result of the tensor product has no semantic significance. The pattern does not in a semantically effective way contain the one of John, nor the one of Agent as sub-patterns, and thus the activation pattern of John as an Agent does not include the activation pattern of John, nor that of Agent. It represents nothing meaningful because it does not have causal connections to John situations, nor to Agent situations. That we can retrieve the values for the Agent vector from values of the tensor product vector and the values of the John vector does not give any semantic relevance to the tensor product. The point is that the tensor product does not have semantic relevance to the causal effects of the activation patterns of the John concept or the Agent-concept in the recognition of John or of Agent-instances, which have been stabilised in learning. The net has to recognise the causal effects of input situations as being satisfaction situations of the sentence or not. For this it has to recognise the implications which the truth of the sentence has. The tensor product seems to be merely a technical device without any meaning.

Criticism of this kind, besides other points which I do not follow, has also been put forward by Fodor and McGlaughlin (1990) and Fodor (1997) against Smolensky's solution to the problem of systematicity: Fodor (1997) points out that vector composition by tensor product misses the causal connection to the possible referents of the expressions in composition. In short, semantic composition is not respected by tensor products. Compositionality of representations is not kept up in constituent combination by tensor product: BROWN COW does not partly overlap with BROWN TREE, as it should. The mental representation of Brown Cow has to include the one of Brown, as has the one of Brown Tree. That is not so in Smolensky's tensor product. Fodor (1997: 114) argues that a mental representation must be causally explanatory in the sense that "a tokening of BROWN COW will include the effects of a tokening of BROWN".

Horgan and Tienson, and also Garfield (1997), do not accept that semantic compositionality should be respected in the results of the syntactic composition of vectors and they support Smolensky's syntactic structures by claiming that they need not work compositional under the point of view of semantics and causality. It is enough, they argue, that the tensor product makes for a systematic change in the resultant pattern when we replace John as Agent by Paul as Agent, or Peter as Agent, etc. We therefore can abstract the Agent pattern by systematic substitution of John in the tensor product by Paul, by Peter, by Fred, and so on. Against this I want to point out, that if basic syntactic structure would not be semantically relevant and motivated, there would be no sense in having syntactic structures at all. In order to uphold semantic relevancy of vector composition by tensor product formation, we would have to assume that perceiving John as an Agent is already a syntactic process in which first a John-pattern and an Agent-pattern get activated, such that both corresponding concepts can be recognised as instantiated and then a syntactic step is taken such that also the result of the tensor product of both conceptual representations gets activated. This then should result in recognising the binding of both concepts as instantiated in the perceived satisfaction situation. Generally, the perception of a situation would then have to involve a syntactic process such that additionally to the result of the composition also the composing parts are preserved. This is required because the result of the syntactic composition alone does not contain these parts anymore, but we still should be able to automatically get all the semantic implications from a situation. This view on the perception of a situation as containing all parts and the syntactic-semantic processes is acceptable under the requirement of systematicity, since the process of composition distinguishes a situation of Peter beating Paul from Paul beating Peter. However, the result still has to contain all the conceptual parts in a semantically active way. The result of forming the tensor product alone cannot 'represent' the proposition expressed by the sentence; rather it has to be aided by also keeping intact the activation patterns of the composing concepts separately. This can be done in an architecture of several conceptual and sensory maps, as I shall argue below.

A formal problem of the tensor product is that the vector length is n times n, if both contributing vectors represent an activation pattern in the same net consisting of n units. A vector of the length n times n does not represent a pattern in the same net but has to be represented by a larger net with n times n units, which can be represented in form of a n times n matrix. This requires a special architecture of the network. A way out for a one-map architecture would be to have an operation of compression performed on the n times n matrix to reduce it to an n-place vector. A compression that preserves semantic compositionality would be preferred. Just taking the diagonal of the matrix would not do, because the diagonal is not more than the intersection between the composing patterns, i.e. the simple product of the vectors, which would cancel out too much. If, for example, on the assembly of activated units on which the blue of John's eyes is indicated nothing of the Agent-concept is indicated, then it is cancelled in the product John as Agent that John has blue eyes. The pattern has the value zero at these units or vector dimensions. The blue eyes of John could not be recovered from the pattern of John as Agent, even though in our perception of John as Agent this feature would remain stable. In this way, John as Agent may not be distinguishable from, for example, Peter as Agent, if such cancelling happens for several specific features of John. What remains of John in the product with the Agent-vector are merely the activated units that also are members of assemblies indicating Agent features, and of the Agent-features only those remain that are also found in the representation of John. The vector product then would amount to an intersection of features. This would be fine as a step in the syntactic-semantic process, where it not that in case John never has been experienced as an agent the agent features would be cancelled out. Whatever is the solution to this problem of forming the intersection between these activation patterns, in terms of satisfaction situations, the parallel to this intersection would mean that there is a union of John situations and Agent-situations. But such a union is much to broad, and thus is inadequate for indicating John as Agent. John as Agent is rather found in the intersection of the John-situations with the Agent-situations, which is indicated by a pattern union on neural nets. As has been pointed out this union alone would neither satisfy the binding, nor the systematicity problem.

Another semantically relevant way of reduction of the matrix would be to compress the matrix by adding up the rows of the matrix and by adding up its columns. By adding up the columns we would just get the vector of the first constituent, in our example the vector of John, multiplied with the number of ones appearing in the Agent-vector. By adding up the rows we would get the vector of the second constituent, i.e. the Agent-vector, multiplied with the number of ones in the John-vector. This would amount to an additional activation of either the John-vector or the Agent-vector. We can add up the sum of the columns with the sum of the rows, and maybe additionally with the diagonal. This way of compression of the tensor matrix would give us the activation of the pattern union with higher activation values for the activated units. This would be a semantically relevant matrix compression, since it contains both vectors of the composition, though with extra activation of the activated units. If we would normalise the positive values to 1, we get by addition of the columns simply the first vector of the tensor composition, and by doing the same for the rows we get the second vector. Adding these normalised sums of the columns and of the rows would simply come down to vector addition of the two composing vectors. With or without normalising the positive activation values to 1, these semantically adequate compressions give us a good result for John as Agent and Paul as Patient separately, but when adding up the compressed tensor products according to (*) above we loose again the syntactic systematicity. The whole boils down to pattern union and intersection, and in the non-normalised case with some more differences in the strength of activation of the positively activated units in the pattern union and intersection.

3. Conclusions with respect to a one-map architecture

The result of considering these methods is that vector addition or union of activation patterns makes sense semantically as far as the implications due to the composing parts are concerned. This also holds for adnominal constructions. Take, for example, *red ball*. The Red activation pattern contains parts that generalise over surfaces, because we mainly experience coloured surfaces. This is parallel to a so-called Surface-slot in a frame representation. The Ball-activation pattern does not contain a specific colour but a generalisation over colours on the surface generally, an aspect that in a frame representation would represent a Colour-slot for specifying a ball's surface. It further contains parts for representing a round surface, including parts that represent surface with colour in general. In the union or vector addition of these patterns of Red and of Ball, these common parts get double activation, namely by the generalised Surface-parts of the Red-pattern, and by the generalised Surface parts of the Ball-pattern. In this way there is a union between the Red- and the Ball-pattern by unification of the common subpattern for surfaces. The intersections in the union take care of the unification, which amounts to binding the concept representations together. The important point is that we take into account that activation patterns contain parts that represent typical contexts or situations of occurrence of the phenomenon to be represented. These get especially activated in syntactic and situational composition by way of unification or overlap like in the activation pattern of (John[Agent)Beat{Patient]Paul} or (Red[Surface)Ball], whereby Agent-features have been especially enhanced within the John representation and Patient-features within the Paul-representation by morpho-syntactic information that the Agent-concept applies to John and the Patient concept to Paul, and these agent and patient features have again been enhanced by intersecting with the Beat-pattern which contains generalisations over agent and patient patterns.

However, all these combinations lead to a mix up in the resulting pattern which appears as a whole on the net at a certain moment. The process of construction contains syntactic-semantic information that is lost in the resulting pattern. The resulting pattern is only relevant for keeping up information about the involved concepts, the semantic binding and syntactically expressed semantic systematicity are represented in the process, not simply in the resulting activation pattern, which therefore never can be a 'representation' of the proposition. Furthermore, the unification or binding by non-empty pattern intersection and inclusion is something that cannot be realised on one single net. In order to represent these requirements on the neural activation due to a possible satisfaction situation, and to preserve the information represented in the syntactic-semantic process we need a more involved architecture of nets. This will be the topic of the following chapter.

We can conclude that a mapping between **DCS** and connectionist networks representations in **CNM** can analyse the potential activation patterns on a net by delineating them in their parts and by analysing the processes of understanding by a network in terms of semantically relevant steps. Understanding a sentence should not differ in structure from perceiving a situation which satisfies this sentence. Thus, as far as understanding a basic sentence is concerned, syntactic structure has to be taken into account and its parallel semantic structure plays a role in perceiving and thereby understanding a situation. Perceiving implies performing a conceptual synthesis in which **B**-intersections are formed, i.e. smallest situations are construed in a sequence of specifying the situation by classification and identification. This process runs parallel to the syntactic combination of the concepts in forming the constituents of the sentences that describe the situation. Note that in the representation by **B**-intersections the result contains the syntactic built up. A situation is perceived by embedding the sensory input into a conceptual structure: we attend to the action that takes place and to the individuals which are involved in the participant roles. The process of achieving the result of perceiving a situation, or of basic sentence interpretation, must consist of these processing routines for situations, which are likewise represented by the morpho-syntactic form of the describing sentence itself and get activated by understanding the sentence. In connectionist modelling we need an architecture of maps, rather than a one-map network, in order to capture the process of understanding and make it part of the result of understanding basic sentences and situations.

Chapter 3

An architecture of episodic and conceptual maps for understanding basic sentences

In the previous section it was shown that connectionist models using only a single map for the generalisation of input data are not able to provide for semantic transparency, systematicity, and binding. It was suggested that an architecture of maps with interaction circuits between them would be able to fulfil the semantic desiderata. Also results in neurological-cognitive research suggest such an architecture. Thus Saffran and Sholl (1999) write that semantic features of a concept of an object are widely distributed over the different modalities of perception and motion, for example as visual perceptional features, kinaesthetic features, action oriented features, tactile features and auditory features. They suggest that they are interconnected via reciprocal activation. The different features of an object will not all be activated with the same strength. "Properties will be differently accessed as a function of context. So, for example, if the scenario has to do with moving, the weight of the piano is likely to be the object's most salient feature, rather than its musical properties." (1999: 263). They find that the distributed model is supported by evidence; the loss of perceptional properties or of functional properties makes a difference for which kinds of lexical items can be accessed and how partial the access is. The loss of functional properties leads to a loss of artefact concepts, the loss of perceptional properties leads to a loss of animal concepts. Saffran and Sholl cite evidence from brain damage patients and PET and fMRI scans during semantic processing tasks. Research in Neuro-cognition is relevant to modelling cognitive processes because models, which are constructed such that certain functions can be fulfilled, should at least be compatible with brain research, although they abstract from the biological aspects of the functioning of the brain.

3.1 Episodic and conceptual maps

The basic distinction we shall make is the one between one or more episodic maps on the one hand side, and several layers of conceptual maps, ordered according to the generality of concepts these maps achieve, on the other. On the episodic maps happens all that is conscious, namely perceptions of situations, including perception of linguistic utterances, and imaginations of situations and utterances. On these maps, sensory and conceptual activity meet in consciousness, namely in perception and imagination. On the other hand, everything that happens merely on the conceptual maps is unconscious, i.e. below the level of consciousness. Even thinking, as far as it emerges in consciousness, is couched in episodes, namely in the form of imagined situations or inscriptions of sentences, pictures or graphic schemata. Schemata are compositions of linguistic and pictorial representations, which serve to represent, to hold and to control in consciousness large networks of logical and conceptual relationships. The episodic maps are the ones that provide for phenomenal content due to their intricate relationship with the perceptual system. Everything that is conscious must have phenomenal content, i.e. it must be couched in some perceptual or proprioceptual medium.

On the episodic maps phenomenal content is provided by including parts of the perceptual and proprioceptual system, especially feeling of our own motor activity. The phenomenal contents, being shaped by conceptual activity, make up feeling, perception, and imagination. The phenomenal contents are rooted in the sensory capacities of the episodic maps and include to a lower or higher degree categorisation and identification which consists in an interaction between activation on the episodic maps and on the conceptual maps. I take it that phenomenal content never is without at least some form or conceptualisation. It contains categorisation through the involvement of generalising maps, by which a perception or feeling is shaped via the neural network by the effects of previously experienced similar episodes: because of the trained, i.e. weakened or strengthened, connections, activation from an activation from a new instance of sensory reception travels most easily towards those neurones on the episodic and conceptual maps which had been previously activated by similar episodes. Hereby the sensory input of the new episodes get classified together with those previous examples, and hereby becomes a perception, a perceived episode.

As we know from brain research, there are different maps or regions for our different senses or modalities of perception: vision, hearing, touching, smelling, tasting, and also balance, motoric and other proprioceptual experience. And some of the maps are compounds consisting of several maps; for example, vision consists of a primary region through which all visual input goes, and more specialised regions such as a region for colours, a region for shapes, spatial relationships, and a region for motion and some shape and locality aspects. The specialised regions are inter-connected with the primary region. We can assume that on the primary visual region visual episodes are momentarily represented by the activation pattern caused by the visual input and the processing results achieved by the specialised visual areas. These results are categorisations as to colour, shape and movement, and they are received back on the primary visual region, which in this recurrent process becomes to be the map of visual episodes. A sensoric map, i.e. an area receiving sensoric stimuli, in interaction with categorising maps, which generalise over the different aspects of the stimuli, is called an episodic map. Thus a sensoric map is not by itself an episodic map; rather it is an episodic map only as long as it stands in interaction with categorising, i.e. conceptual maps. These are firstly modality specific conceptual maps, creating modality specific episodes, such as visual episodes, auditory episodes, etc. And these are secondly higher conceptual maps which categorise and identify individuals of our normal ontology of situations and entities. The episodic map of representing situations and objects according to our normal ontology consists thus of at least two spaces, whereby the space of modality specific episodic maps is embedded into the space of the normal ontology map. The smaller circuits, which run between the primary sensoric areas and the conceptual maps categorising aspects of the sensoric input in a modality specific way, are embedded into the larger circuits between the layer of conceptual maps of individual concepts and kind concepts and the sensoric input. There is not just a one-way path from input categorised in a sensory specific way to the higher categorisation according to the concepts of normal ontology. Rather reciprocal fibres also make the higher conceptualisation play back to the modality specific conceptualisation by injecting into the perception of the received sensoric data also such features which are imagined due to the general and specific knowledge couched in our concepts of individuals and kinds. In this way we do not just perceive certain clearly sensed parts of the front side of a house, but we see the whole front of the house, which even gets completed more or less accurately to the whole house by our power of imagination.

In the figure below the architecture activated in the visual perception of a green ball in rest is presented. Note that the circuits of activation are strongest for those aspects which are really sensed. The activation pattern for the concept "ball", which is called up by the perception of the form properties also activates in a weaker way the possible motion a ball can make and the possible actions which are the standard ways of playing with a ball, which have been experienced before. Also how it feels to touch a ball, the felt roundness, gets secondarily activated from out the ball indicating activation pattern, although the ball has not been touched. Via these secondarily activated aspects, the perception of the ball contains imagined features on the primary visual, haptic and motoric maps. Since the actually sensed visual features are stronger activated because of the (repeated) input stimuli, the imagined features are not dominant. The movement of the ball cannot not be imagined at exactly the same moments in which the resting ball is perceived. Rather the potential movement of the ball can only be imagined in alteration with the percept of the ball in rest. The imaginations on maps of different modalities can be entertained at the same time. Thus, while seeing the ball we can imagine what it is like to touch it, or to play with it in a certain manner. If the child who is observing the green ball, has the desire to move around and play, then the areas for indicating play-concepts will be activated and herewith the corresponding pre-motoric areas. Thus the concept-indicators for ball-play will be additionally activated and a ball-play action will be meditated and prepared on the pre-motoric areas, such that the child will move towards the ball and begin to play with it.

Circuit activation in perception: seeing a green ball

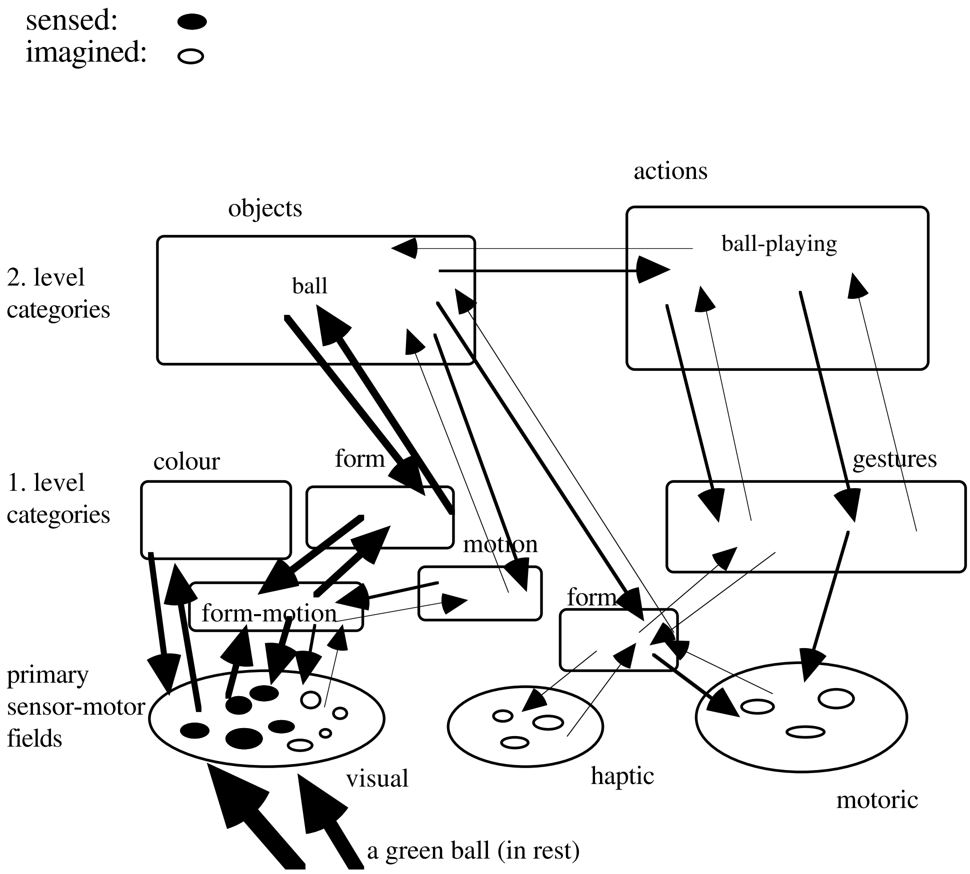

Similar to this architecture, models are developed in neuro-dynamical system research. Thus Grossberg, Mingolla and Ross (1994) have developed a neural theory of attentive visual search in which pre-attentive boundary segmentation, surface and figure ground perception, and spatial information activate certain object representations, which play back activation that helps to define (1994: 471) “the perceptual units that attentive spatial and object recognition mechanisms can further modulate and recognise during visual search.” Attentively organised search recursively applies to emergent groupings that may contain multiple items. Grossberg (1996) further explicates the object recognition mechanism by his Adaptive Resonance Theory (ART), where the bottom up activation from out stimulated sensorial maps activates a pre-learned concept representation, or rather concept indicator on a higher map, which then feeds back activation from the top down to the bottom. This top down activation, providing expectations about what is perceived, influences the perception of the received data by highlighting some features, suppressing others, and adding expected features. If the activated concept does not fit well enough with the data, the interaction circuit does not reach the synchronisation of the oscillatory activity of firing in which object recognition would result. Rather the mismatch causes an inhibition and hereby destabilisation of the circuit activity, while another competing concept takes over. It is also activated from out the data, and also possibly somewhat from out the first concept, which in some ways is similar to it. The second concept then has the chance to establish a new circuit of synchronised activity. Grossberg (1999) suggests that the circuits of resonant dynamics in between bottom up and top down activation processes plays a role in constituting conscious states in the brain; all conscious states are resonant states. He furthermore suggests that the ART-mechanisms, which have been shown in visual perception, are also operative in the other modalities. And, of course, the context-dependence of perception fits nicely into this model, because a context activates conceptual frames, which from top down influence the perception of the data presented in that context, i.e. interact with the bottom up activation from the reception of the data.

Another study in the same line of research has been done by Van Leeuwen, Steyvers, and Nooter (1997), who investigate perceptual segmentation by using a model of circuit activity organised by attractors, which are “stable limit sets of system states towards which nearby states converge”, while unstable limit states are called repellors (1997: 319). The model shows an interplay between stability and instability in chaotic behaviour. For calculating this interplay of non-linear change they use differential equations. Stability in these “discrete-time dynamic systems" means convergence towards attractors. An optimal segmentation of a perceived pattern of dots, for example, is one which is relatively stable. In one of their examples two sequences of dots form two lines which cross each other. Depending on from where one goes through the pattern, from the middle, from each one of the sides, or from the top or the bottom, the segmentation of the pattern may switch. The overall perception of the pattern as a crossing of two lines, or and X, is determined by the Gestalt-principles of proximity and good continuation. Hereby the starting points from which one takes up the pattern of dots serves in organising different attractors, between which one can switch such that “the two lines are desynchronised from each other “(1997: 337). Switching requires “an ability to disengage attention to a stable pattern by making it unstable and creating alternative stable ones” (1997: 320). Generally, switching between the interpretations of ambiguous figures is not organised by a control schedule, rather it occurs spontaneously. Perceptual destabilisation and switching are based on “deterministic chaos”, calculated by non-linear differential equations (1997: 322). They also remark, apparently following Crick and Koch (1990), that the synchronisation of neurone firing leading to oscillatory activity patterns, solves the binding problem in visual perception (1997: 322). Van Leeuwen et al. furtheron suggest that “two perceptual models could be assumed to operate in interaction: a perceptual segmentation system and a classification system. The classification system learns from the segmentation system and the segmentation system is corrected by the classification system. ...such a model would develop a bias for more stable perceptual organisation in the long run and combine this with flexibility in the short run" (1997: 342). This is a point also modelled in Grossberg et al. (1994) and in Grossberg (1996), and it also supports nicely the interaction between contiguity ordering and similarity ordering between data as the two ways of structuring growing sets of data in the process of concept formation and understanding (cf. chapter 2). Segmentation and classification in visual perception are ordering data by contiguity and similarity. In these dynamic models learning segmentation and object recognition takes place along synchronised circuit activation. Neuro-physiological and anatomical research by Kandel, Schwartz, and Jessell (1995) has supported this view, based on Hebb's conjecture, that units that are activated together strengthen their connections. Units that are connected in circuits of synchronised firing, strengthen their connections, which amounts to learning the combinations of input data, their forms and segmentations, and their classifications.

Also studies about deficiencies in perception caused by brain lesions show that binding of different aspects in perception is essential for object recognition. From experiments done with brain-damaged patients we know, for example, that there are patients who can see the colour of an object put before them, but not its form, and thus cannot identify the object by visual input. When they are allowed to touch it, though, they recognise the form by tactile input and are able to identify the object, for example, as a cup. This points to the fact that the specialised region for categorising visual form or shape is impaired.

It seems evident from case-studies by Sacks (1985) that also memory of scenes and faces, both requiring ordering of impressions by contiguity relations, has to result in activation on episodic maps, which have an expression in consciousness. An act of recollecting something from memory can be seen as similar to an act of perception or imagination. It is caused on generalising maps by some re-activation of parts of the neural net effects of previously experienced scenes. This includes maps that generalise, i.e. categorise, over spatial and temporal contiguity relationships between situations. From Sacks' case-descriptions it seems that episodic memory might function in parallel with the more analytic and abstract categorisations. It contains implicitly in connection strengths the partial orderings by contiguity relationships of the indicators of individual experienced situations. It can best be understood as the capability of ordering a new episode in space-, time-, and causal contiguity, while at the same time categorising it by comprising it together with previous episodes that are similar to it under certain perspectives. This two-fold ordering of the new with the old, keeping unimpaired coherence and stabilisation, is of course the basis for understanding in general. Parallel with the several levels of generalisation on the left side, we seem to have an implicit memory of schemata of episodes in their contiguity contexts and orderings between episodes. Parallel with the concepts on the several levels of generalisation, we have generalisations over sequences of episodes on maps of episodic memory, for example schema's of the kind "When it rains the ground becomes wet", "When a mushroom tastes bitter, I will not take it". Furthermore, parallel with the higher levels of generalisation, we have general concepts or schemata of spatial and temporal contiguity orderings, such as "in", "outside", "before", after", "at the same time", "next to", etc.

The episodic map for visual episodes is the primary visual region in interaction with the specialised regions, whereby the visual episodes consist in the primary sensoric activation of the primary region which is transmitted to the specialised regions categorising colour, shape, relative locality, and motion. From there the resulting activation patterns send back activity by reciprocal fibres to the primary map, where they cause a secondary activation, emphasising the features which have also been encountered in similar previous examples. The results of visual categorising, corresponding to features such as certain colours, forms, spatial relationships, movements must be bound together as belonging to one object or situation. I shall suggest that this happens by binding the repeated primary sensoric activation together with the categorising secondary activation in a small circuit that runs between the primary, sensoric map and the specialised, categorising maps of the respective modality. The sensoric map has to contain collections of specially dedicated neurones. They must be dedicated such that they contribute the phenomenal qualities as the medium in which all our experiences are expressed.

Crick and Koch (1990, 1992) have proposed that 40 hertz oscillations in the visual cortex solve the binding problem in the way that the common frequency of oscillation of neurone firing binds the visual aspects together, as they are distributively processed on the specialised visual maps. They even claim that this unifying frequency of oscillation must be the causal base of consciousness in perception. Of course, here as elders, the criticism by Chalmers (1996: 116) applies, namely that Crick and Koch cannot answer the key-question of why these oscillations, and hereby why binding, should be accompanied by conscious experience. However, I shall pick up the idea of binding by unifying frequencies, in proposing an architecture of constituent structures of sentences and situational perceptions by circuits of activation between several episodic and conceptual maps. The circuits might be individualised by certain rhythms, based on certain interferences of frequencies. Thus small circuits with certain frequencies in the firing of the neurones, which are involved within a larger circuit, determine the rhythm in the firing of the neurones in the larger circuit. This, of course, is speculation, but it would make sense.

For each sensory and proprioceptual modality we can assume a similar compound structure with a primary map and specialised categorising maps. Thus we have an episodic map for each of the modalities; we have visual, auditory, olfactory, tast+touch/texture episodes, arm and hand movement+touch, leg and foot movement+touch episodes. These specialised episodic maps are further bound together by activation interaction, via connections, with conceptual maps for identifying individuals and categorising objects and events into kinds of objects, events, activities, or actions. The interaction between the parallel episodic maps in a compound can happen via trained connections between these compound maps, or it can go via the conceptual maps in which groups of modality specific episodes are bound together as aspects of one object or one situation. The categorising activations on the higher conceptual maps are played back over the lower conceptual maps, to the primary sensoric maps. Such a large circuit of interaction headed by an activation on a conceptual map, defines an episode on the compound of sensoric maps. If the activation pattern on the higher level conceptual map indicates an individual concept, a kind concept, and a situational concept such as an action concept, we have an episode of an individual, being of a certain kind, involved in a certain action. The compound of sensoric maps in interaction with higher conceptual maps functions as an episodic map on the level of our normal ontology of individuals and situations. It is non-specific as to modality, but includes different modalities. In what follows I shall call this a normal episodic map.

The episodes on the episodic maps, caused there in recollection stimulated by activation from the generalising maps, are ordered in space and time via the connections to those generalising maps that show the activation effects of spatial and temporal contiguity relationships. They consist of compounds of parallel visual, auditory, touch and motoric maps and a higher level conceptual map categorising the compound input into spatial and temporal relationships between objects and events. Space and time are then for us instantiations of such spatial and temporal relationships on episodic maps. We experience space and time only in episodes involving spatial and temporal relationships. There are first and second order maps for storing the effects of contiguity relationships, which have been found in previously encountered constellations of situations.

Recognition of contiguity can be impaired by certain kinds of brain damage. Such cases have been reported by Sacks (1985). Somebody who is blind, wholly or for one half of his visual field, because of a certain kind of damage of his visual regions in the brain, cannot remember any scene which originally had been perceived in the part of the visual field now blind. Thus a patient who was asked to tell what buildings he passed on his way could only report about the buildings on his right side, and when returning on the same way he could report about the buildings on the other side, which were then on his right side. The left side scenes he could not remember, even if he had seen them before the damage of the visual regions in the right hemisphere. He was able to see and recognise only basic geometrical forms and patterns, and separate features. In so far he could report about the buildings, but he could not connect features in contiguity, and therefore he neither saw nor remembered faces or scenes, nor could he detect any emotional expression of faces. At most he could identify people by a single clear geometric feature, a very prominent form of the head, or the nose, or the chin, or a very perspicuous mark on the cheek. Because his right side visual region was impaired by a tumour, he could not perceive, imagine, or remember faces; and when he was asked to take his hat he grasped for his wife's head, as something of likewise round shape. He compensated his inability to perceive and imagine visual contiguity relationships in objects, scenes, and between scenes by getting things and actions into order through organising them in contiguity by means of singing tunes and songs, which he more or less did continually. His world was organised by musical contiguity relationships, replacing visual contiguity. Auditory contiguity got apparently connected with motoric contiguity of hand and foot movement, and thus he could organise his actions. On his auditory episodic maps he employed an especially rich gamut of contiguity relationships in which he embedded sequences of actions he had to do regularly, like washing and clothing himself. We can explain this behaviour by assuming that conscious recollection of sequences of actions and situations requires appropriate activation on episodic maps of perception and imagination in contiguity, which apparently are to be found mainly on the right hemisphere. Since the patient in question could find his way walking in his familiar surroundings, we can assume that he still had implicit memory of previously seen things and scenes, possibly on the right hemisphere, which does not come to consciousness on the visual episodic maps, which are now damaged. As far as his space-time orientation was conscious it was ordered sequentially in songs accompanying his activities. Also other cases reported by Sacks point to the conclusion that the right hemisphere is especially suited for perceiving and imagining contiguity within scenes, and in history, while the left side is stronger equipped for conceptual order, higher order concepts, and abstract thinking.

Visual episodes get connected by repeated co-stimulation with episodes processed at the same time in other sensory regions, such as those of sound, touch, smell, taste, motoric experience. Hereby these experiences are coherently bound into the unity of an object or situation. Thus, when seeing a plane in the sky we expect to hear typical aeroplane-noise, and the other way around, when we hear a typical noise in the sky from a certain direction we look up and expect to see there the aeroplane which causes the noise. A little boy I know, of one year and three month, who speaks no other word than 'ae-o' is able to connect the sight of an aeroplane and the noise of an aeroplane by using the word 'ae-o' when he sees or hears an aeroplane. Combining the several sensory perspectives of perceiving an object by establishing and strengthening connections between neurones of the different sensor-specialised episodic maps is part of the way to achieve the objectivity of an episode.

Objectivity is not achieved by just one sensory impression, but rather amounts to a connection of different kinds of sensory impressions bound together into a single object, linked by contiguity to other aspects of the situation and other episodes, and categorised together with other contiguous sensory impressions of different kinds on superimposed or 'higher' maps. Also these categorisations of different sensory kind must be bound together by neuronal connections established and strengthened in experience. This must be the case because we are able to draw implications from concepts like 'aeroplane' to the kind of noise, kind of form, kind of movement, kind of material typical for an aeroplane.

In what follows I shall simplify this very involved picture quite a lot in order to be able to design a map architecture of the relationship between situational episodes and the descriptions of these episodes in language. For the sake of representational simplicity, I throw together into one episodic map the various episodic maps of different modality, the activations of which have as expressions in consciousness the sensory impressions we are aware of in our perceptions and imaginations.

The notions "different layers or spaces of maps", and "higher" and "lower" maps, are used as notions of logical and cognitive organisation. It does not mean that these have to be different layers of the cortex, though that certainly would be a possibility of physical realisation of some regions of the different layers of maps. Future improvements on MRI- and PET-scan technics, showing small differences in depth of cortex activity, might give information about the position of layers of maps in and on the cortex.

Activation patterns on layers of conceptual activity, by themselves, cannot to be understood as representations of concepts. Such patterns have a concept indicating value only by contrast relationships to other such patterns, and by their relationships to perceptions and imaginations due to stimuli which originate externally from examples, experienced in a history of learning. Across different learning individuals and even across the life history of one individual the identity of such a pattern is not relevant, and does not even exist. Merely on the sensoric areas we must assume pre-established collections of dedicated neurones which are fixed in their connections to certain measuring devices of the input sensors. Different patterns on conceptual maps, due to the different initial situations of learning systems, the difference in their learning history, and the changes due to progressive learning experiences, will have the same discriminatory value, i.e. the same conceptual value, if the established contrasts and the relationships to the kinds of input are constant. This conceptual indicator value of a pattern is given by these two aspects and not by itself, i.e. not by the set of activated neurones alone. The identity of a neural pattern on a conceptual map is not necessary for indicating the same concept. However, on the episodic maps, in their basic areas which are directly sensitive to sensoric input, there should be a strong relationship between the phenomenal quality aspects of perceptions and the involved fibres leading to specialised sensoric cells and the spatial constellations of these, in order to have coherent phenomenal qualities within the experiences of things. Thus, specific frequencies of light should consistently correspond to certain qualities of colours, and for that one needs groups of neurones that are especially dedicated to certain kinds of sensoric input by being fixed in their connection to the respective measuring devices of the sensors.

Contrast between patterns has originated from contrast between examples and thus provides a structural representation of contrast relationships between examples. Concepts are merely implicit in the distribution of the weights of neuronal connections and their relation to sensors. A concept can be seen as a disposition to accept and create examples as belonging to a set of examples representing a concept. It is thus a capacity of the neuronal system and its relationship to the sensory and motoric organs. Conceptual ability is the ability to recognise, order, and create examples. The only representations of concepts are sets of examples and sets of general sentences we hold true, and which we are aware of in the conscious expressions of activations of the situational and linguistic episodic maps. Activation patterns on conceptual maps are merely concept indicators; they are what they are merely by their causal relationships to certain input examples, stabilised by adjusting connection strengths in the course of learning.

Only neuronal activity which happens on the episodic maps has an expression in terms of phenomenal content in consciousness. But that does not mean that phenomenal content consists just out of phenomenal qualities. Rather such qualities are never experienced as such, they always are imbedded in reactions from the conceptual maps, even if these are not yet stabilised. What happens on the episodic maps is created either as perception by sensory stimulation in space-time contiguity together with conceptual categorisation provided by the connections with the higher conceptual maps, or it is created as imagination stimulated by some other input, by processes based on inclusion and association, that takes place on maps of general and individual concepts and contiguity orderings with respect to other episodes. Imagination is especially triggered by linguistic input via the process of understanding, and perception can be modified by such imagination. Perception is caused by reception of data via the sensors, combined with registering of contexts consisting in bodily movements, and in interaction with the categorisations of the data: The data are sent from the episodic maps to the generalising maps; then the categorised data from the generalising maps are sent back via reciprocal fibres to the episodic maps. The activation from out the generalising maps calls up additional imagination on the episodic maps by adding further phenomenal content triggered by the generalised aspects that normally come into the picture under other perspectives than those under which the data are received by the external stimuli in this case of perception. Thus every perception of a categorised object or situation automatically is enriched by aspects which are imagined, and which are got from the effects of previous inputs which were contiguous with input similar to the current input. Which aspects are triggered depends on context, as has been pointed out in the example by Saffran and Sholl (1999) about moving a piano, where its heavy weight is to be considered, but not its musical aspects.

Imagination can best be understood as a derivative of previous perceptions: New situations or objects can be imagined in analogy to old perceptions, or there can be imagined really new ones of a kind that has never been perceived before. The process of imagination modifies previous or current perceptions by taking parts from other previous or current perceptions. The modifications are achieved by substituting parts of perceptions or previous imaginations by parts from other perceptions or imaginations, keeping intact contiguity relationships as they are classified in general situational concepts and individual concepts. Substitution is an operation with direct conscious results, which can be performed also consciously on representations due to activation on situational or linguistic episodic maps. On the neuronal level it is a change in association by blocking certain connections serving for associating contiguous content and by opening up other connections by which new content is associated to the contextual content which is kept constant in substitution. The results of substitution are conscious, and thus they can be controlled and evaluated as normal, as comical, as surprising or as revealing some deeper truth. General and individual concepts respectively characterise and individualise the imagination by being bound together in the examples generated on the episodic maps in analogy to previous experiences. The analogy is construed on sets of examples by the operation of substitution mentioned above. Substitution is constrained by coherence, i.e. by the contiguity relationships found in general and individual concepts. Note that imagination in the arts often reaches special effects of surprise and pleasure by breaking these coherence constraints. This means making a new connection against the normal expectations and current associations. Thus we can find on surrealistic pictures the strangest creatures, such as pigs with wings and a trunk; or we find an "artistic statement" presented as a toilet bowl used as a vase.

3.2 Understanding situations and linguistic expressions

That syntactic processing is not done by a separate linguistic organ or module but rather by an interaction of different brain areas involved in neuro-cognitive -perceptive and -motoric processes, some with linguistic task specifications, seems to be a widely accepted position. Hagoort, Brown, and Oosterhout (1999: 305) write: "it is most likely that both parsing and grammatical encoding are based on a concerted action of a number of different brain areas with their own relative specialisation. These relative specialisations can be memory requirements for establishing long-distance structural relations, the retrieval of a verb's argument structure and other lexical-syntactic information, the use of implicit knowledge about the structural constraints in a particular language for grouping words into phrases, etc." In what follows, the process of understanding a linguistic expression, and the reversed process of describing a perceived situation, shall be analysed in its functional components by an architecture of episodic and conceptual maps presented below. In the text we shall provide more detailed descriptions of the architecture than our figure of the architecture of understanding depicts. The figure is kept simple in order not to destroy the clarity of a global overview.