The role of consciousness and intentionality in perception, semantics, representations, and rules

Renate Bartsch

ILLC, Department of Philosophy

This paper 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 are potential neural activation patterns and themselves 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.

Table of Contents

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 they are by themselves mostly not sufficient, since many such states are not accompanied by experiential awareness in consciousness. What kinds of relationships between such states result in consciousness is not known yet. With respect to states that usually go together with conscious 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 to see nothing, registrate differences of light frequencies and motions in their visual field, though they do not see colours or motions. Their ability to correctly assent or dissent to judgements of others about what colours and movements are presented before them is not perfect, but they make significantly better than random "guesses" about these. 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 orders them in form, colour, and motion aspects, and receiving them back in a more finely analysed form, recurrently after being processed in the specialised areas of the visual cortex, V3, V4, V5, 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.

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, potential or actual consciousness, 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, is selection from the categorised input under relevance and attention 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 object- and subject-related relationships, i.e. under all 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 onto maps in which generalisations take place. Processes of this kind have been modelled in connectionist models. These maps organised by similarity under certain perspectives, on the other hand, serve furtheron for categorisations of incoming data and are thus used for filtering the incoming data such that they can be sent in the so processed forms back on the episodic maps. These recurrent ways in the whole construction of the neuro-physiological background of our conscious data must require recurrent fibres, such that both kinds of maps, the contiguity preserving and the similarity preserving, i.e. categorising or generalising maps, are linked sufficiently.

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 ones. We can, for example, not imagine a concept. Rather we only can imagine series of situations that exemplify the concept, which may originate from the episodic memory and also are constructed in imagination analogously to situations of the episodic maps.

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

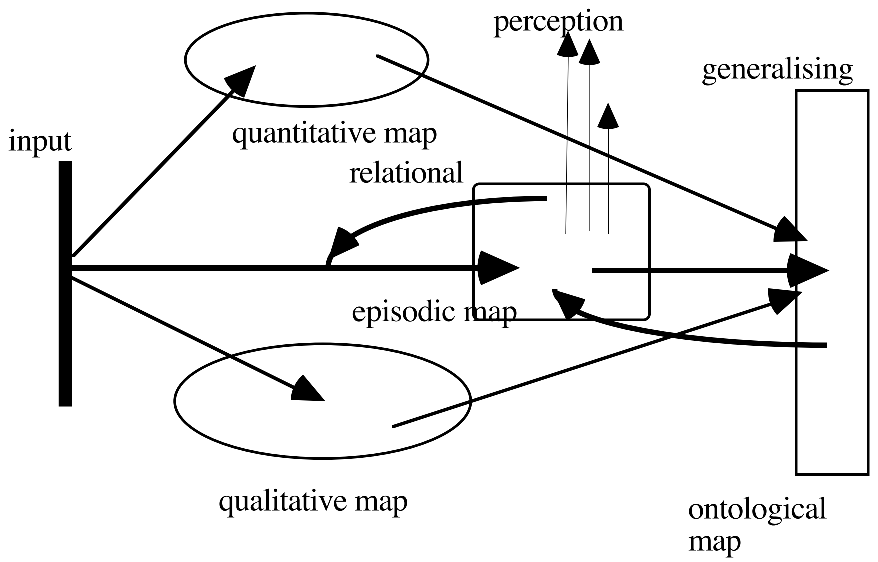

Sensoric and motoric input is being processed in maps specialised for qualitative and quantitative categorising these kinds of input. Parallel it goes to a relational episodic map, which represents space and time relationships. The input vector from the sensory units is concatenated by recurrent fibres with a vector from the episodic map representing the previous input. This concatenation establishes temporal, local and other contiguity ordering on the episodic map by strengthening linking connections or by geometrical closeness. The episode thus gets a marked place in a contiguity net of episodes. The episodic activation is linked to the parallel analysis of the sensory input on the qualitatively and quantitatively categorising map. This parallel episodic and categorising activation combined in one vector could be 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 episodic map. Taking into account that certain brain areas have been identified as processing form and motion, we additionally can assume a map for processing object-internal relationships amounting to form, and a map for processing object-external relationships amounting to relative object position and movement. These maps should be connected to the episodic map and the generalising ontological map by recurrent 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 then are presented, remembered, or imagined. This 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 such and such a situation or object, and only then we can locate a situation or an object with respect to others, i.e. understand it as being in the world. There cannot be an episodic memory without some degree of categorisation.

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 will be part of denotational semantics in the sense that it is the causal background of our semantic knowledge, consisting of stimulation by sensory input, neuro-physiological reactions on the input, and possible motoric output. I shall show the difference between denotational and causal semantics by means of thought experiments based on the blindsight cases.

1. Consciousness is necessary for denotational semantics.

In the case of blindsight, a person reporting not to see anything can react to things located in his visual field, especially to movements. 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, without touching it. However, he has no visual experience. The conclusion is that neural visual registration must have become directly connected, in a previous process of learning language, 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 actualised and potential knowledge, feelings, and attitudes towards surroundings and actions, including the operations by which the individual makes connections between these. The following thought experiments make no claim of being realistic; they merely serve the purpose of philosophical argumentation.

Imagine now that somebody has a defect causing blindsight from birth on. In such a case he cannot have 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 certain colours. This might be done by telling him to raise his left arm, when we see the colours 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 become 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 do you not tell me to lift my left arm? I feel like doing this." These are the occasions when we show him certain colours. 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, beyond his control the disposition to these reactions at certain moments. 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, 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.

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 of the colour he learns the name as a name of the type of his 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 arm, for example. 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 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 front of him 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 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 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 not perceived visual stimuli.

Our question now is whether a person such as has been described above, 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 the colour expressions for him would 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 blindsight person the semantics of our colour terms would rather be a denotational semantics of making certain movements connected with certain terms, and about this he could make correctness statements. If he had not learnt making certain arm movements in satisfaction situations for certain colour terms, but would have learned the colour terms previous to his blindsightedness, his semantics for our colour terms would now simply be a causal one, especially if he cannot anymore remember colours in imagination. 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. Having a concept, 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.

Note that I argue, against a current version of Wittgenstein's anti-private-language argument, 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 indeed refers to aspects of situations of application of a term which are present to him in consciousness. These acts of remembering or imagining aspects of situations of use for a term, or referring to some such situations outside, are additional 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 regularity under a 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, 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 whole language community. 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 to be confident in that he uses his linguistic means correctly, 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 blindsight person may also feel assured about the reliability of his use of 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 to 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 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.

We can compare the situation of the blindsight 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. This whole picture 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 perception and motor activity might just function like the correlation between unconscious colour effects in the brain and feeling compelled to use certain colour terms in a person who became blindsighted later in life.

For us, the question now arises what is the extra benefit of having available a representation which one can be conscious of. The answer will be that 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 compelled to do X, utter a certain expression or to move his arm, is regular. He misses the comparison with the associated colours. He only can take into account his behaviour or disposition, but not to what it is correlated 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 makes us understand learned behaviour as regular and makes it possible to sometimes abstain from the behaviour, because a rule or regular behaviour can be deliberately broken. It further makes it possible to imagine situations and plan behaviour according to these. 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. The term 'perception' as it is used here includes conscious presentation of what is seen. Intentionality relates the imagined or represented situation to the world in a direction of fit, which according to Searle 1983 can be world-to-mind or mind-to-world, in belief and perception or in desire and planning, respectively.

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. Rather we merely assume representations in our mind as perceptual presentations or images of linguistic utterances and of real and imagined situations, in which the respective term is applied correctly. Our concepts are equivalence classes of stabilising sequences of growing sets of such situations. Representatives of these equivalence classes are sequences or sets of situations, subsets of which we can call up in sequences in our mind, i.e. we can come up with these consciously, and can refer to them in the world. These sequences of examples of our concepts consist of presentations and representations of situations or images. Of course, there are also 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.

Although our blindsight person naturally only has a causal semantics of colours, like a robot endowed with a trained connectionist network can have, I have pointed out that a denotational semantics can be achieved if the blindsight 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, and is correctly, or incorrectly, used on the background of the sets of contrasts within each network, which correspond to each other, and by virtue of which a term can be judged to be correctly associated with a movement. The point of this second case is 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 the use of the expressions and the making of the movements can be unconscious, as we assume it to be in neural nets and robots generally.

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. 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. 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. 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 blindsight person with normal perceivers on the one side and robots on the other.

Instead of imagining a certain interplay of colours our blindsight person can either entertain a set of situationally bound sentences he has come to hold true, by a mechanism totally unconscious to him, since he misses the criteria for holding true in these cases. Indeed, we may better not call this 'holding true'. It is merely feeling compelled to utter or assent to certain utterances; he has no idea what this assent amounts to in terms of being able to imagine a situation that fulfils the sentence assented to. He even might feel compelled to form new combinations of recurring parts of the expressions he was trained to. According to the other set up of our thought experiment, when we experience a play of colours, he can imagine or experience a kind of motoric ballet with his finger, hand, and arm reactions he has been trained to in association with unperceived, but well-received, colour stimuli. Both these possibilities are structurally not different from what we are experiencing. With respect to quality, there certainly is a difference. However, emotionally the difference may or may not be very important; who knows how enjoyable such a motoric ballet is, compared with a ballet of colours which we perceive. We might even speak of a structural identity between our colour system and his system of movements with his left arm, hand and fingers, even though there is a qualitative difference. He even might enjoy feeling compelled to combine parts of previously encountered expressions into new expressions.

Our blindsight person would behave rationally with respect to colour and form judgements, i.e. draw implications like we do. But he would 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 of the colour ballet. And he might learn the same about his linguistic combinations.

In what way would our blindsight person be handicapped because of his merely indirect causal relationship with visual experiences of others? I have already mentioned the involved procedures necessary for verifying statements about colour and form of objects. The same is true for recognising situations. Also here quite involved procedures are necessary in order to identify the objects in them, their colours and forms. These procedures are a problem for recognising properties and events or actions that merely occur for a short time span. Here the separations between things and their surroundings are not possible while the events are taking place, or they are not possible without disturbing the events taking place. For our blindsight person, the registration of short term properties of things is not merely troublesome, but it is a problem in principle. The short-term properties of things can only be experienced as properties of the situation confronted with, but they cannot be attributed to certain things on a merely receptual basis. Imagine that in the situation there are objects A and B, and our person receives a certain colour stimulus from out the situation, which compels him to a certain movement or to uttering the colour word. How then can he now whether the colour is part of A or of B? In order to decide that, he would need to synthesise object- and colour-stimulus in one object, and for that he would need imagination in which he remembers what it looks like when an object has a certain colour, and for that he would need previous conscious visual experience with which the new experience can be compared. Only by available theoretical knowledge, a situation could in normal cases be sorted out in the way of short-term properties of things and relationships between things.

Let us suppose again that our blindsight 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. He then would have no visual imagination and also no motoric aesthetic substitute. Because he lacks any such imagination, he never could design anything 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 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.

A robot has no sense experience at all in which he consciously registers 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 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. This device then could be used to teach others, without knowing what it does.

The correctness judgements of such a monitor are unconscious and unintentional. The monitor does not know what correctness is, though he judges correctness 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 the monitor might be an effective teacher for a certain task. Even though he does not have a denotational semantics he can produce correctness statements. But they are correctness statements only in the eyes of the program designer and user. The monitor 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 monitor 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', as 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. What is missing is the development of the semantic ontology, by which input stimuli can be synthesised by the power of imagination, as Kant has emphasised, in order to become historical concepts, such as concepts of individual objects, events, and acts, and general concepts, Kant's empirical schemata, of objects, events, acts, and abstract objects. However, we might assume a bias towards our coarse grained ontology, 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, although we factually only realise partiality, which is a the root of our notion of reality and truth.

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 phenomenal 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 to 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.

2. Consciousness is necessary for knowledge-relevant 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 viewed as conceptual maps (Bartsch 1996). There they are relationships between 'contents' which are causally bound to perceptual input in learning processes. The 'contents' are potential activation patterns due to the adjustment of connection strength between 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 potential activation patterns, encoded in the connection strengths between units on the net and between input units and these units is established in the learning process and re-activated for each new occurrence of similar input.
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, 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. The potential activation patterns are our dispositions to establish the causal links between input and the respective activated patterns. Activation patterns for term perception 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 as a relationship between terms and 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 makes 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 connections to action routines and by connections to public language use. Connections to 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 highlight, 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 patterns of publicly relevant objects and situations, as encoded in our public language, get distinguished on the level of 'content' as it appears on the neural nets. These considerations show how our public denotational and intentional semantics and causal semantics depend on each other. The first is not possible without the second, and the second in its specific structure depends on 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.

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, namely, 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 as described above.

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. 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 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.

Let us change our thought experiment to supposing 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. 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, though 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 qualitative differences in colour perception as an example of an (empirically) meaningless statement.

3. Consciousness is necessary for creation and improvement of language

Probably 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. Another example of unconscious training would be a robot who has learnt to react to his surroundings, or the crab of Churchland's (1986) article, 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 and exclusion of potential activation patterns, which mostly may not evolve into consciousness, and therefore cannot be expressed in language.

We cannot by ourselves develop a language 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 properties occurs. We just would not now why we suddenly use these expressions. They therefore would not really be language for us. The training would have to be performed by some external teacher who is consciously aware of these external properties, or who can consciously perceive our neural reactions and has devised the terms for these properties. After such a training we would be able to use expressions like a robot does who has been trained to certain input- output correlations.

To create, develop, and improve a language by oneself, however, does require 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. This capacity to develop a language 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 as know that, but merely 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 making the statement true or false.

Non-experienced registrations, i.e. reception, can give rise to a non-perceptive and non-cognitive adaptation to the world and to 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 neural effects does not belong to 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 can be made conscious by just somehow lifting it to a conscious level. Only scientifically it could be reconstructed by a physiological and especially neural 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. As Wittgenstein has argued, internal rules and representations could 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 makes 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 semantic. The origin and the reality of denotational semantics is a public affair. It is social, but it can in principle also be simply public, at least outside of us, such that we can be conscious 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 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 neural effects. Therefore rules constructed by the linguistic researcher could at most be idealisations of the processes that in reality take place without being governed by rules. There also will always be situations where no action is possible according to the rules constructed. A small deviance in the preconditions of action is already sufficient for the break down of the 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 neural networks to be the only presently available models of acquiring basic orientation in a non-conscious way.

4. Linguistic ability is different from linguistic analysis:

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 actually or potentially 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.

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 attempting 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 conclude from the examples to some definition, which is a certain partial representation of a concept in our public language or some language devised publicly for scientific purposes.

Linguistic ability as a faculty within the language using individual is a brain state which stands in a causal relationship to the sets of input by which it has been formed 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 one 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 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 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.

On the other hand, linguistic research comes up with rules, and conceptual analysis provides at least partial definitions of concepts we use. In linguistic analysis we find systematisations or structures in the sets of linguistic data, by applying structuralist methods such as 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, and we can even improve and develop our language further by institutional language planning performed in academies and schools.

These 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 accepting 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 become part of our mind. These are real and imagined situations of 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 potential 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 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 reasonable arguments have been given to postulate it nevertheless. We experience theories merely as 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 right, 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 speech situations, have memories of these, and imagine new ones. All these 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 subcognitive background of our mind as the semantic, syntactic and phonological faculty, which is a set of dispositions encoded distributively in the connection strength 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, for example 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 neural connectedness in the brain explain the causal individual aspects of language use, while public rules and representations describe the normative aspects.

There is also a notion of unconscious mind in Psychoanalysis, according to which we have hidden beliefs and desires, which can be brought to consciousness by therapeutic communication. However, these and other not explicit beliefs are just dispositions of the brain to get activated under certain germane 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. If the activation is effective enough, 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, but only can detect them by ordering the data in systematic linguistic analysis, and we can afterwards recall them, 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 are part of our potential mind like any other belief or knowledge we have acquired. As such they 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 which are the kernel of our linguistic ability. If ability regulating rules formulated in some hidden unknown symbolic medium would exist at all, 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 since we do not need such hidden rules for explaining our language ability there is no need for assuming a hidden mind besides our mind in which we would have stored such rules in some hidden code of representation. In the following paragraphs, I shall elaborate the arguments against the assumption of rules and representations in 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, and propositional attitudes require a medium of representation of the propositional content, which such mental language provides. The first argument is weakened since the growing success of connectionist and probability based models of learning, which do not require that linguistic rules are learned (Bartsch 1996). The second argument is weak, since implicit propositional attitudes have contents which can be thought of in terms of dispositions, i.e. neural activation patterns causally connected via the previous learning processes to objects, events and actions, and explicit propositional contents are represented by the internalised public language. Neither for learning language, nor for propositional attitudes a mental language is necessary as a means of representation. 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 in forming activation patterns conforming to objects and events that are enhanced and delineated by their occurence in activity and action routines in which the learning individuals are involved. But this is not a mental or innate language. Rather every public language developed in human history will be constructed in parallel with such basic action, event, and object patterns. Piaget based the logical abilities of individuals on a process of internalising external operations on actions, objects, and situations (Piaget 1953).

Linguistic ability has been reconstructed on a philosophical, logical or 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, or, if non-mentalistic, have reconstructed sets of linguistic expressions, fragments of a language, as sets (re-) producable or characterisable by sets of rules. However, 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 be caused.

More recently, the claim that competence means rule following has been weakened (Matthews 1991): Utterances are not built in following the rules, rather in accordance with the rules. The last 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 subsymbolic or subconceptual connectionist processes. The level of connectionist modelling he views as a theoretical construct between the neural and the conceptual or symbolic level. It certainly is sub-cognitive and subsymbolic, and represents on a very abstract mathematical level some relevant properties of the neural network structure of the brain. I do not agree with Smolensky's use of the term 'individual knowledge', because also individual knowledge is representational and thus part of the cognitive level, the conscious mind. Rather his term 'individual knowledge' should be replaced by the term 'know how'. It comprises adaptation to surroundings and language use.

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 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 we have in mind. It is conscious or it is potentially conscious, which simply means that there is a brain disposition, a potential activation pattern, 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 architecture 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 its products, of which we then say that we have them in our mind.

Giving another twist to the computer metaphor, 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, or they play 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.

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 an relationship within a single ontological realm. The term is 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. The mind is at any moment what is conscious and it is extended by the potential mind of the results of brain processes that can come into the mind.

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 working as a public constraint 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 looks at rules as norms of the forms of linguistic products and their use 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. This architecture makes substitution of equivalent expressions in contexts possible, like it makes 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 recursion, requires recurrent architecture in neural networks, which is not specific to the language faculty but serves also in perception, imagination, and action. Architecture and adjustments of connections to 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 such that they can be described in terms of linguistic rules. But rules as registering regularities, and idealisations of these, are merely found in our products. Our brain does not follow these rules in speaking or interpreting. Later in this paper, 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 (Bartsch 1996, Bod 1995), the competence-performance distinction has to be reconsidered, and 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 production of utterances; our brain works differently, much more flexible and far less restrictive than the rule systems could do which have been supposed 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 our linguistic products and, furtheron, normative stipulations as far as we are norm authorities, being 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 will be more or less adjusted to the products that conform to these norms, and in this way speakers have built up a competence that is, by itself, not a set of rules, but 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 principles might be due to our general cognitive abilities, based on innate brain architecture and biases in neural nets due to phylogenetic adaptation to orderings in the world, and additionally 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. But all these biases and established states of neural net work connections are not the 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 potential activation patterns in the neural nets. 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. This can also be examples he had previously produced himself. These will be adapted somewhat more to the model input, but the perception of the model input will already be influenced by the state the neural network is in when perceiving. 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.

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 exists 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 to consciously construct some examples or to highlight the structure of a series of examples, which 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.

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 Chomsky's 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. These preferences concerning the quality 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 abilities and to the regularities in our surroundings, which also must have supported the formation of certain cognitive abilities during evolution.

When we construe a rule system that should re-produce, 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. Then we are able to form consciously new or old utterances by performing a synthesis based on the rules 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 flowers, 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, and it is also here that linguistic analysis itself 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 is not very informative, and introspection 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 aware of, but which are 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 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 creating a language, designing a language, 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. Without consciousness, one would not even be aware of the fact that one uses a language. However, also computers can use language. Using language there is based on algorithms programmed in, or also on a learning process, where not rules are learned, but behaviour that conforms to the rules, without being produced by the rules.

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. Knowledge about language is only possible after language analysis has been performed. Before, there is merely knowing how to use language. There is no once and for all act of identifying the conditions of application of an expression, because many possible conditions might be part of the previous situations of linguistic use, some of which we are aware of and others not. In the analysis of our 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, can bedevil our intuitions and give rise to wrong linguistic or logical reconstructions, which work in many instances by providing appropriate output, but at other occasions the rule 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 of these 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 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.

We have explicated the difference between the two different 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 set 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. Or we can construct programs that follow the rules on a computer. 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. On the one hand side we have described linguistic ability as an ability to language use, on the other hand, we have described the ability to linguistic analysis, to constrain 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 to function in establishing 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 we find in denotational and structural semantics. These rules we can 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, which would be required for conscious 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 performing tasks which are conform to the rule which describes the regularity in the behaviour or tasks.

5. 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 stimulate it, we can always expect new kinds of data produced by language users 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. 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 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 require a certain 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.

The different ways of continuing similarity sets of old expressions by new ones, parallel with our semantic structuring of similarity sets of the situations that form our world, may correspond with the flexibility of possibilities in the neural base, which has been, and continuously gets adapted to the world. There may be a flexible choice between several potential activation patterns built up in the course of learning, 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*]]. 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: 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.

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 and from the global and distributed organisation of activation patterns on the neural-physiological basis. The basic principles of substitution and analogy, which somehow must be encoded in the architecture of neural networks, admit different structurings on the sets of data. Though linguistics cannot do more than take seriously the linguistic data available in language use and develop theories that explicitly take care of these, it can profit from insights in neural-physiological models by being open to kinds of flexibility, which can be integrated into cognitive and especially 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), which combine structuralist methods with probability measures on sets of data.

On the other hand, theoretical linguistics has always striven for uniqueness of a given structural analysis, and much of argumentation in linguistics 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 our normative intuitions as competent speakers of a language. 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 Chomsky school that they describe the language faculty which is part of the mind of speakers. I have argued with Searle (1992) that the latter is not correct. The mind is all that is in principle conscious, one time or other. We may adhere consciously to certain norms when speaking, but not to deep systematisations of sets of norms. Here I add to the argumentation against Chomsky's position that 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 much idealised and abstract linguistic theory, which 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. The results of theoretical linguistics are impressive products of our will to rationalise, i.e. explain and understand, and even guide, our linguistic behaviour by rules. However, grammars never will be predictively adequate of natural language use. Rather they will more likely function in a normative way, and then not so much for natural language users, but for technical purposes, as part of systems of traditional artificial intelligence. In this point I agree with Schnelle (1996: 114 f.), who 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". In this perspective I have discussed the role of grammars.

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. They 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 be attractive merely under the view of simplicity and elegance of theories which are the best systematisations of normative grammars of such 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 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 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 to the same 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, the second side guides the conceptual built-up and basic syntax, reflecting the functionality of language with respect to our world, and 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 under 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 that is possible according to the structurings laid on a growing set of data is allowed. Much is not continued, or does not even come under attention, at a certain stage or never. Much linguistic output, created by individuals or by neural nets 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, frees this sub-cognitively available potential from the normative pressure enacted in normal circumstances, and gives rise to language change.

6. 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. 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.

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, on the cognitive level. 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 following 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 become 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 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 perceiving them, the neural activation patterns and processes that support these examples best become activated and modified according to the input examples. The neural pattern which is strongest activated by the constructed examples as input, supports best 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. 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 in our design. And in learning the products which are selected by 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 itself, possibly 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 by others and by oneself, namely on expressions, symbols, pictures, perceptions, and imaginations, and by playing back the results in receiving them as new input. Weakly associated products, or those which only become possible by certain additional contextual input, can be selected and in design strengthened in a conscious manner. This enables an individual, or a neural net, 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 the awareness of the products and behaviour which become, in consequence, analysed in different settings, and consciousness is involved in constructing new products and behaviour step by step, by selection of parts and new compositions according to rules formed or consciously adopted, 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, the level of consciousness, in the style 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 taking place on the mind level 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.

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 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. 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. Thus also a theory of mind has to be product oriented; the so-called mind-processes can only be postulated but cannot be experienced in the mind. They therefore do not belong to our mind. They have a reality only as subcognitive brain processes. They could only be part of the mind if they would be perceptions, memories of perceptions, or imaginations of behavioural data, and processes as sequences of steps in overt behaviour, like, for example following an algorithm step by step, where the steps can be seen because they are noted down in a series of products in imagination or on paper, for example written expressions, which are followed with the eyes.

Although mind is somehow emerging from the brain, mind and brain are ontologically totally different. We have to distinguish the reception of data by 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 further we can consciously operate on these mind phenomena. Internal activation and internal triggering caused by the external sources or generated by itself, need not be part of the mind. Its products, though, appear as imagined or as utterances or other behaviour in the mind. They get expressed in consciousness and thereby 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. Then it seems accidental that such a product comes up in the 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 consciously 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.

It seems safe to conclude that language use can be learned, as Wittgenstein suggested, just by training, and does not require consciousness as long as the semantics of such language use is merely a causal one. But intentional denotational semantics, i.e. being aware of the relationship between language and world, as well as deliberate change of a language, deliberate change of knowledge, and with this of theoretical concepts, explication and setting of norms and codifying these, as well as all linguistic analysis and rule formation, require consciousness in dealing with the linguistic products or data available, and in forming new ones according to rules designed. It follows that an unconscious system of artificial intelligence never could know that it uses language and never could design a system of language and language use or for that matter a system of artificial intelligence.

The level of consciousness is the perceptive and cognitive level on which we also formulate rules, and when it gets described and 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, finally 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 product of the mind, 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. It is a truly artificial intelligence, which is a design product of conscious human design activity, an activity in accordance with principles of rationality. This design product can be seen as an extension of the human mind, but only working on its superficial level according to the kind of operations that are the conscious operations performed on the mind phenomena, including all 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. It has an artificial intelligence which humans rationally have developed, like all other rule systems formulated on a cognitive level. It presupposes consciousness in design and setting it up, such that it conforms to all kinds of normative notions, semantic correctness, goal adequacy, true beliefs and preferences to be taken into account, 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 would create another artificial intelligence by observing and analysing its own products, simply because it cannot 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.

7. The primacy of consciousness and intentionality in cognition

One of the metaphors about the mind discussed above sees it as the monitor or screen of the brain, another sees the mind is 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. A third 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 incorrect, because the mind not exist independently from the brain, which suggests to speak of an internal relationship between mind and brain, and it explains why some authors, for example Chomsky, writes about the mind/brain, which is more obscure than that it clarifies anything.

There are arguments of principle against the possibility of reducing consciousness to brain states. These 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 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. 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 can talk about their relationships to each other, but we can not say, in terms of mind phenomena, such as brain states as known objects, what the mind is in which these objects are. We cannot say more than that space is the condition of the possibility of spatial objects, and 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 the mind is. It points out that there is always an epistemological primacy of the mind. 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).

Mind, the whole of conscious phenomena, is epistemologically primary to any correlation between states of consciousness and brain states. The latter only could be identified as states probably evoking consciousness by correlating them with conscious states. By themselves they cannot have any observable property to which consciousness can be reduced. Such a property could only be identified by correlations to a state of consciousness. But a correlation is not a reduction or explanation by which one could understand the correlation. Furthermore, 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.

Another way of explaining 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 further metaphor for the mind is the one 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 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 or in pictures restricted in their interpretation 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 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 has been extended to self-consciousness because the above mentioned improvements by judgements 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. 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 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 this way could be perceived, i.e. experienced. It has got a pictorial, procedural, conceptual, and furtheron even linguistic representations. Our perception, phylogenetically, is adapted to chunks of middle size, which we need to avoid or 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 we learn fastest what is most favoured by biases in our neural organisation. We have become endowed with sensors measuring input in certain dimensions or similarity spaces, and neural nets which favour reaction towards middle-sized objects and events, and especially 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 make more fine grained partitions in the world by performing an analysis on the middle-sized chunks, by which we order them into similarity sets under certain perspectives. 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 identification as well as discrimination of objects. We then build up a semantics according to the fine-grainedness of analysis or conceptualisation of objects and actions. In a process of synthesis we can rebuild old and build new objects.

Our consciousness of an object goes as far as the analysis goes, and 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. All the steps taken in a synthesis of different aspects or parts into a whole we can endow with intentionality, i.e. view them as intentionally taken steps of construction, and in this way we can reach finer grained levels of intentionality in rationalising a new product as a construction which has been deliberately set up, although it might be that the product was not constructed intentionally, but rather is a product of merely unconscious neural processes without an expression in the mind. The same happens with the syntactic built-up of utterances. We can view every syntactic part of the expression as intentionally used in its context. But this 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. 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 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 nets. What we see in an algorithm for performing a certain task or action, or generating a certain perception, is based on a rational analysis of intentional behaviour, whereby the intentionality which normally is restricted to the larger chunks, is now directed towards small parts which are normally not produced or perceived intentionally. For these 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 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 the normal level of 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 massive basic brain processes caused by certain inputs. In controlling activity performed by a self-monitoring process, they afterwards may get rationalised as being intentionally constructed. 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 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 our mind produced intentionally through refining the analysis of the products emerging in our mind into smaller parts, the artificial intelligence is cognition at all, namely representing our knowledge and performing procedures and tasks in a meaningful way. Our knowledge on the intentional level makes that the deliberately fine-grained representations of it, which are made 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 into another code. Only because there is the level of intentionality and consciousness, on which we represent and perceive 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.

References

Bartsch, R. 1987. *Norms of Language*. London and New York.

Bartsch, R. 1996. The Relationship between Connectionist Models and a Dynamic Data-Oriented Theory of concept Formation. In: *Synthese* 108: 421-454.

Bod, R. 1995. *Enriching Linguistics with Statistics: Performance Models of Natural Language*. ILLC Dissertation Series 1995-4. Institute for language, Logic, and Computation. University of Amsterdam.

Chomsky, N. 1986. *Knowledge of Language: Its Nature, Origin, and Use*. New York-Westport-London: Praeger.

Chomsky, N. 1995. Language and Nature. In: *Mind* Vol. 104. no. 413. Jan. 1995: 1-62.

Churchland, P. M. 1986. Some Reductive Strategies in Cognitive Neurobiology. In: *Mind* Vol. 95, no.379, July 1986: 279-309.

Dennett, D. 1991. *Consciousness Explained.* Boston: Little, Brown, and Company.

Fodor, J. 1976. *The Language of Thought*. Hassocks: The Harvester Press.

Fodor, J. 1987. *Psychosemantics. The Problem of Meaning in the Philosophy of Mind*. Cambridge, Mass.: MIT Press.

Kripke, S. 1982. *Wittgenstein: On Rules and Private Language*. Oxford: Basil Blackwell.

Matthews, R.J. 1991. Psychological Reality of Grammars. In: A. Kasher (ed.), *The Chomskyan Turn*. Oxford: Basil Blackwell: 182-199.

McGinn, C. 1993. Consciousness. In:*Problems of Philosophy. The Limits of Inquiry*. Oxford: Basil Blackwell: 27-45.

Peacocke, Ch. 1992. *A Study of Concepts*. Cambridge, Mass.: MIT Press.

Piaget, J. 1953. *Logic and Psychology*. Reprinted 1965. Manchester: Manchester University Press.

Putnam. H. 1981. Brains in the vat. In: *Reason, Truth, and History*. Cambridge University Press: 1-21.

Quine, W.V.O. 1974. The Nature of Natural Knowledge. In: Guttenplan, S. (ed.), *Language and Mind*. Oxford: Clarendon Press.

Schnelle, H. 1996. Reflections on Chomsky's "Language and Thought". In: *Theoretical Linguistics*. Vol. 22, No. 1/2: 105-124.

Searle, J. 1983. *Intentionality*. Cambrideg: Cambridge University Press.

Searle, J. 1992. *The Rediscovery of Mind*. Cambridge, Mass.: MIT Press.

Wittgenstein, L. 1960. *Philosophische Untersuchungen*. (Engl.: *Philosophical Investigations*). Frankfurt a. Main: Suhrkamp Verlag.

Smolensky, P. 1988. On the Proper Treatment of Connectionism. In: *Behavioral and Brain Sciences* 11, 1-23.

Van der Berg, M., Bod, R. and Scha, R. 1994. A Corpus-based Approach to Semantic Interpretation. In: *Proceedings of the 9 the Amsterdam Colloquium*. IILC. Institute for Language, Logic, and Computation. University of Amsterdam.

Weiskrantz, L. 1986. *Blindsight: a Case Study and Implications*. Oxford: Oxford University Press, Clarendon Press.
Zeki, S. 1992. The Visual Image in Mind and Brain. In: *Scientific American*, September 1992, Special Issue "Mind and Brain": 43-50.