LDV-FORUM

Forum der Gesellschaft für Linguistische Datenverarbeitung (GLDV)

LDV-Forum 17 (2000) 1/2

Zeitschrift für Computerlinguistik und Sprachtechnologie

GLDV-Journal for Computational Linguistics and Language Technology

Offizielles Organ der GLDV

Herausgeber

Prof. Dr. Gerhard Knorz; Gesellschaft für Linguistische Datenverarbeitung

Anschrift: Fachhochschule Darmstadt, Fachbereich Information und Dokumentation, Haardtring 100, D-64289 Darmstadt; Tel: (06151) 16-8499; Fax: (06151) 16-8980; e-mail: knorz@www.iud.fh-darmstadt.de

Redaktion

Gerhard Knorz

Editorial

Die Fachbeiträge des vorliegenden LDV-Forum-Doppelheftes gehen auf einen Workshop des GLDV-Arbeitskreises Generierung und Parsing in Morphologie, Syntax und Semantik am 15. Februar 2000 in Bonn zurück. Dem Herausgeber des Forums Gerd Knorz sei Dank gesagt für diese Publikationsmöglichkeit der Beiträge.

Der Workshop stand unter dem Thema *Communicating Agents* und versuchte Überlegungen zu den verschiedensten Aspekten des Entwurfs natürlichsprachlich interagierender Agenten aufeinander zu beziehen. Die hier publizierten Beiträge geben ein Bild davon, leider ohne die fruchtbaren Diskussionen des Workshops protokollieren zu können.

In einer ersten Abteilung des Heftes finden sich grundsätzliche Überlegungen zur Theorie der Begriffe (Bernd S. Müller) und der Syntax-Semantik-Schnittstelle (Roland Hausser), es folgt die Vorstellung eines Systemprototypen, der viele Aspekte natürlichsprachlicher, situierter Kommunikation erfolgreich meistert (Jan-Torsten Milde). Rodolfo Delmonte befasst sich mit der Implementierung und Evaluierung von Strategien bei der Syntaxanalyse des Italienischen im LFG-Rahmen. Schließlich folgt eine Gruppe von Beiträgen, die sich mit Fragen der Pragmatik-Semantik-Schnittstelle befassen (Henk Zeevat, Paul Piwek, Anton Benz).

➢ Prof. Dr. Ulrich Schmitz (e-mail: ulrich.schmitz@uni-essen.de)

Erscheinungsweise 2 Hefte im Jahr, halbjährlich zum 31. Mai und 31. Oktober. Preprints und redaktionelle Planungen sind laufend und aktuell unter der Adresse http://www.iud.fh-darmstadt.de/iud/wwwmeth/publ/ldvforum/menu1.htm einsehbar.

Bezugsbedingungen Für Mitglieder der GLDV ist der Bezugspreis des LDV-Forum im Jahresbeitrag mit eingeschlossen. Jahresabonnements können zum Preis von DEM40,- (incl. Versand), Einzelexemplare zum Preis von DEM20,- (zuzügl. Versandkosten) bestellt werden: LDV-Forum, c/o Dr. Bernhard Schröder, Poppelsdorfer Allee 47, 53115 Bonn

Fachbeiträge Unaufgefordert eingesandte Fachbeiträge werden vor Veröffentlichung von mindestens 2 ReferentInnen begutachtet. Manuskripte sollten deshalb möglichst frühzeitig eingereicht werden und bei Annahme zur Veröffentlichung in jedem Fall elektronisch und zusätzlich auf Papier übermittelt werden. Artikel sind bevorzugt einzureichen in den Formaten Microsoft Word für Windows® oder Word Perfect® für Windows. Eine Dokumentvorlage für Word für Windows® kann unter der Adresse ftp://www.iud.fh-darmstadt.de/iud/wwwmeth/publ/ldvforum/ldvforum.dot heruntergeladen werden. Sie enthält die wichtigsten Styles.

Rubriken Die namentlich gezeichneten Beiträge geben ausschließlich die Meinung der AutorInnen wieder. Einreichungen sind – wie bei Fachbeiträgen – an den Herausgeber zu übermitteln.

Druck und Vertrieb GLDV **Satz** Kurt Thomas/Alex Schulz/ Patrick Jülich, Bonn **Auflage** 400 Exemplare **Anschrift der GLDV** Prof. Dr. R. Hausser Universität Erlangen-Nürnberg Abteilung für Computerlinguistik Bismarkstraße 12 D-91054 Erlangen; e-mail: rrh@linguistik.uni-erlangen.de.

Sehr bedauerlich ist das (euphemistisch gesagt) späte Erscheinen des Heftes. Der Grund liegt in den nur partiell erfolgreichen Versuchen, für das LDV-Forum zumindest einen provisorischen Produktionsmodus zu finden, nachdem das Heft nicht mehr von IKS e.V. hergestellt werden konnte. Nun zeichnet sich dank der Bereitschaft von Dr. Christian Wolff, das Forum zukünftig in Leipzig herzustellen – synergetisch sinnvoll verbunden mit der Übernahme der Herausgeberschaft des Forums – glücklicherweise eine tragfähige Lösung, die das Forum bald wieder in einen regulären Rhythmus bringen soll.

Mit dieser Aussicht wünschen eine aufschlussreiche Lektüre

Bernhard Schröder
und Hans-Christian Schmitz

Remarks on Concept Processing for Cognitive Robotics

Bernd S. Müller
Biomimetic Autonomous Robots Team
Institute for Autonomous intelligent Systems (AiS)
GMD – German National Research Center for Information Technology
D-53754 Sankt Augustin

Abstract

We present preliminary considerations on the architecture of a CONCEPT *processing system for robots. With the help of* CONCEPTS *cognitive robots will be able to structure their sensory input, access their inner motivational states, and gain flexibility for reacting on changing circumstances. Concept is a theoretical construct in Cognitive Science for which considerable experimental evidence from psychological experiments exists. We propose that* CONCEPTS *should be considered to be pre-linguistic, language agreement processes making use of them in forming the* **CONCEPTS** *we "see through" natural language. As pre-linguistic concepts,* CONCEPTS *may also be generally assumed for living beings. A tentative comparison with Natural Language Semantics, concept theory in Cognitive Science, and Jackendoff's Conceptual Grammar is tried.*

– Es werden erste Überlegungen zu einem CONCEPT-Verarbeitungssystem für Roboter angestellt. Mit der Hilfe von CONCEPTS sollen kognitive Roboter in die Lage versetzt werden, ihren sensorischen Input zu strukturieren, auf ihre inneren Motivationszustände zugreifen zu können und Flexibilität beim Reagieren auf wechselnde Umstände zu erlangen. "Konzept" ist ein theoretisches Konstrukt in der Kognitionswissenschaft, für dessen Existenz einige Evidenz aus psychologischen Experimenten abgeleitet werden kann. Wir meinen, daß CONCEPTS vorsprachlich sind; mit ihrer Hilfe bilden Sprachvereinbarungsprozesse die **CONCEPTS**, die wir "durch die Sprache" sehen. Als vorsprachliche Konzepte können CONCEPTS bei allen Lebewesen angenommen werden. Ein vorläufiger Vergleich mit Aussagen aus der Semantik natürlicher Sprachen, aus der Konzepttheorie der Kognitionswissenschaft und aus der Conceptual Grammar Jackendoffs wird versucht.

1 Introduction

For purposes of Cognitive Robotics, we introduce a construct called CONCEPT. We think that CONCEPTS would be of great importance for Cognitive Robotics. Cognitive robots, most probably, will get at least part of their autonomy from the stable recognition of known situations with the help of concepts, and from the flexibility with which they can react on changing circumstances via modified or new concepts. We present preliminary considerations on the architecture of a concept processing system for robots. Concept is also a theoretical construct in Cognitive Science for which considerable experimental evidence mostly from linguistically oriented psychological experiments exists, and which can also look back on a long tradition of philosophical thinking. According to these considerations, concepts play an important role in representing circumstances accessible through the sensors of a cognitive agent like living beings are, and in accessing inner states necessary for an overall problem-solving by the cognitive agent.

In the first part of this paper, we give a broad introduction, concentrating on answering the following questions in a first attempt: What is the role of concepts in robot cognition, what is their structure and what are the processes working on that structure? We do not describe any implementation on a robot. In the second part, we raise more theoretical questions: What is the status of our approach in relation to the problem of the infinite regression when defining new concepts (and their attributes) using known concepts and attributes? We will propose two ways to overcome the problem at least partially. What is the relation of concepts to **concepts** we see "through" natural language? We will propose that concepts should be considered to be pre-linguistic, language agreement processes making use of them. At the end of the paper, a tentative comparison with Natural Language Semantics, concept theory in Cognitive Science and Jackendoff's Conceptual Grammar will be tried.

We are well aware that fundamentally different approaches exist, either philosophically different, like radical constructivism, or technically different like the dynamical systems approach (Bajscy&Large 1999), and the hope to solve everything by applying current function approximation methods like training of artificial neural nets; reinforcement learning, etc. (Müller 1997). But we think, our approach has the advantage of being conceptually nearer to solutions in natural cognitive agents.

Concept notation

As the reader may have already realized, we use the following notation for concepts when writing about them in meta-language (i.e. in the plain text): If we emphasize the pre-linguistic character of concepts, as we mostly do in this proposal, we write concepts; if we write about concepts emphasizing their provenience from natural language, we write **concepts**; if we write about concepts in more general terms we write: concepts. Similarly, in the following, our notation for examples of pre-linguistic concepts will be small capitals (our default) and for examples of natural language concepts bold small capitals.

2 Concepts and concept processes for cognitive robotics

CONCEPTS are structural elements of the memory of cognitive agents, natural and artificial. They play an important role in representing and accessing circumstances with the help of the sensors. CONCEPT processes also access inner states, as part of the general problem-solving necessities, of the (artificial) cognitive agent. CONCEPT processes are part of different cognitive processes. Typically, if, e.g. in a recognition process, parts of the sensory input match parts of the inner structure of a concept exemplar, the process can work with the whole concept as an hypothesis, speeding up the recognition, as not the rest of the whole inner structure must be checked. CONCEPTS are products of the CONCEPT formation process. As a first approximation, lacking the experience from experiments with artificial agents or the definite certainty from analyzing natural cognitive agents scientifically, we propose a concept system that is influenced by representational theories in cognitive science and artificial intelligence.

2.1 Structure of CONCEPTS

A CONCEPT is an entity that serves as an unit in cognitive processes. CONCEPTS have an inner structure which is accessible by these processes through CONCEPT processes. It is common to think of the inner structure of concepts as consisting of attributes that can take certain values: A set of attributes and their values together make up a specific CONCEPT. A CONCEPT may be part of another CONCEPT. Such aggregates of CONCEPTS may themselves be processed as wholes.

The exact nature of the inner structure of concepts is still under discussion in cognitive science. After we have introduced our view on concepts in this and the next section, we will have glimpse on this discussion in section 4.

Advantage of CONCEPTS in cognitive processing

CONCEPTS are important, if not the central, tools for a cognitive agent in handling the world by generalization and abstraction. Also, processing CONCEPTS as whole entities has advantages: As we have seen the recognition process is speeded up. Generally in cognitive processes, due to the CONCEPT structure, only parts of the input necessary for a certain cognitive process must match parts of the inner structure of a CONCEPT relevant for that process. The process can work with the whole CONCEPT as an hypothesis already in early states of processing. Only if this "early" hypothesis fails, (parts of) the rest of the whole inner structure of the CONCEPT must match.

CONCEPTS as entities

A CONCEPT becomes an entity, that serves as an unit in cognitive processes, during CONCEPT formation, CONCEPT verification and CONCEPT use as influenced by the general problem solving needs of the artificial cognitive agent, and taking into account the needs of the complete cognitive apparatus of the agent. A CONCEPT as a whole is referred to by its name.

Inner structure of CONCEPTS

The name of a CONCEPT is unique in a specific cognitive agent. This name is not decomposable, it has no meaning of its own. CONCEPTS consist of attributes that can take values, and, optionally, hierarchy markers that relate the CONCEPT to the CONCEPT aggregates it is part of. Attributes and hierarchy markers together "define" the CONCEPT. There are two types of attributes: features and propositional attributes. Features and propositional attributes take different types of values: the value of a feature is descriptional and qualitative in character; the value of a propositional attribute is a truth value telling if the proposition must be true or not if the CONCEPT has to be applied successfully. During CONCEPT formation, it is decided whether a CONCEPT is "defined" by using propositional attributes or features, or a mixture of both. In most cases, propositional attributes can be expressed as features, and vice versa. This decision must also be made when displaying CONCEPTS for inspection purposes during the implementation of cognitive agents.

Attributes are generated by the CONCEPT formation process. They can be understood as special kinds of CONCEPTS, having values as their attributes. Values are also generated by the CONCEPT formation process, in cases of observable values respecting the physical nature of the sensors.

Between the name of a CONCEPT (standing for the CONCEPT as a whole) and an attribute, a relational operator holds; similarly, relational operators connect an attribute with other attributes of a CONCEPT. Semantically, these relational operators are very elementary. In concept display they often are left implicit (like the relations between the elements of natural language expressions).

CONCEPTS contain pointers to the snapshot sequence they are "made of" (see paragraph on CONCEPTS and time below).

Examples of CONCEPTS can be found in the section 2.3 and 3. Our notation for attributes are small capitals, for values plain roman letters.

Aggregates of CONCEPTS, composition of concepts

CONCEPTS can be parts of aggregates of CONCEPTS that may themselves be processed in a similar way as CONCEPTS (i.e. as wholes). This hierarchical CONCEPT composition makes use of relational operators. CONCEPT composition typically takes place in a given situation, time-slot, or sensory context. While it is a fact that we find non-hierarchical composition of natural language concepts in natural language syntax, it is not so obvious if or how non-hierarchical composition of pre-linguistic concepts is possible. In Example 1 we use sentence-like constructs that are meant as descriptions of pre-linguistic circumstances, in the first place.

Recursive definition of CONCEPTS

If world facts allow, CONCEPTS that are made of CONCEPTS can be simplified by introducing recursion. If, e.g., a tree has to be conceptualized, there are different possibilities to conceptualize its parts: branches up to which branching depth do we conceptualize as **BOUGH**, which as **TWIG**, just the ones with leaves at their ends? A **BOUGH** may be a (main) **BRANCH**, a **TWIG** may be a **BRANCH**, too. Theoretically, all non-stem parts of a tree, except leaves and fruits, could become concepts of their own. While the choice has to do with the general cognitive goal governing the CONCEPT formation process, in some situations the recursive definition of the CONCEPT could also be a good optimization. A **BOUGH** is a **BRANCH** which branches from a **STEM** or from a **BRANCH** unless it has a leaf at its end.

CONCEPTS and time

CONCEPTS describe circumstances in time. There are no timeless CONCEPTS. Even a seemingly "static" CONCEPT like CHAIR is not timeless since **CHAIRS** are viewed at in time: from different perspectives, under changing light conditions, as moveable parts of a room, experienced in different situational contexts at different times, etc.; all these time-dependent aspects of a **CHAIR** are important for the concept formation process that generates the CONCEPT CHAIR which further on can be used in cognitive processes.

CONCEPTS describe changes: they may be compared to a sequence of snapshots. These changes over time that a CONCEPT represents can be traced by the cognitive processes. Introducing time dependency for CONCEPTS must be seen in the context of the ongoing discussion on representation (Bajscy&Large 1999).

Exemplars of CONCEPTS

While CONCEPTS are abstractions a cognitive agent makes from its outer and inner world guided by its needs and goals, exemplars of CONCEPTS are the results of analyzing a given situation with the help of the CONCEPTS, i.e. they are instances of CONCEPTS; as such they are abstractions drawn from a given situation with the help of memorized abstractions. Exemplars are products of the CONCEPT use process. An exemplar of a CONCEPT consist of its name and a hierarchy marker that refers it to its CONCEPT.

2.2 Processes working on concepts

Three types of processes can be discerned: concept use, concept formation, and concept verification. In *concept use* the different cognitive processes that access the outer world of an agent via sensors, or the inner world of an agent, e.g., its disposition of what to do next, make use of concepts for their process-internal considerations, inferences, and conclusions. *Concept formation* is the process that acquires new concepts as outer or inner circumstances of the cognitive agent change in a way that makes a restructuring or complementation of the concept set necessary. Like concept use, concept formation is a continuous process that never ends in an agent. Likewise, *concept verification* works continuously, confirming concept use decisions or refuting them as values for attributes of an already inferred concept come into the focus of the governing cognitive process. All these concept processes are governed by general cognitive processes (as vaguely defined in cognitive science like perceiving, categorizing, problem solving, remembering, learning, communicating with language, reasoning, imagining (see, for instance, Müller (1998)) and depend on the motivational and emotional systems and other parts of the cognitive system of the agents. The concept processing decision procedures, active in what we call concept processes, are not yet well understood.

concept **use**

In concept use, cognitive processes access the memory in order to interpret the sensory impressions of the outer world with the help of already formed and memorized concepts. The inner world of an agent is also accessed via concepts, e.g., its factual and episodic memory, its motivational disposition, or its action control.

Categorization is concept use applied to the sensory input of a cognitive agent. It can only be successful if memorized concept structures are consulted. There are also cases of concept use working completely "internally", i.e. only with already formed concepts and other parts of the cognitive system.

concept use usually works top-down: the concept is compared with the "facts" (sensory or inner) up to a degree that a "first or crude" hypothesis becomes "reasonable". At that moment cognitive control is handed back to the governing cognitive process. As soon as concept verification intervenes concept use has to refine its decision by consulting more attributes. This iterative process ends with a "final" decision of concept verification.

concept **formation**

Concept formation is activated when a general cognitive process that evaluates the overall performance of the cognitive agent signals a failure that is not ultimate, but can be circumvented. A failure of the cognitive agent in a certain situation can be caused by a too coarsely grained conceptual structure. Concept formation restructures the concept set, or adds new concepts. New concepts are formed when they are needed and allow new differentiations ("when they are useful"). See Example 2 (in paragraph 3.1) for an example.

CONCEPT verification

CONCEPT verification verifies hypotheses brought forward by the other CONCEPT processes, CONCEPT use and CONCEPT formation. In CONCEPT verification CONCEPT use decisions are confirmed, if values for attributes of an already inferred concept meet the expectations of the governing cognitive process, or are refuted, if they contradict them. Furthermore, CONCEPT verification results are used to improve the performance of the other CONCEPT processes.

Sensors and CONCEPT use and formation processes

The processing of the sensory input is a central problem in concept use and formation processes. CONCEPTS are most probably useful in deciding what to accept as useful "information" and what to regard as noise. One way of tackling the problem is a binary dichotomy of sensor scales and sensor fields: Dichotomy points (that divide the scale of linear sensor) or dichotomy lines (that divide a sensor field into areas) are chosen according to process needs.

Concept display

Concept display is not part of a conceptual system, but part of the support system for the description, construction, analysis, and evaluation of a conceptual system. Concept display we call the tools and structures that are used for displaying of a concept set, for instance, for introspection purposes.

2.3 A concept system in action: A small predator looks at a changing scene

The following example would better be a series of sketchy drawings, certainly closer to the "representational reality" in animal cognition than the natural language expression we used, which too easily allow a anthropomorphic interpretation. Using Example 1, we will try to make our terminology, concerning the structure of a concept system and the processes working on it, a little bit more clear. For purposes of illustration, we assume the following situation: A small hungry predator is sitting in the grass behind a small rock on a slope overlooking a meadow with brushes in its center. As time passes by, different circumstances must be conceptually handled by the small predator (see left column of Example 1; as concept display form we use propositions in natural language). By reading the

Example 1: A small predator looks at a changing scene

robin	known ROBIN, a BIRD, a harmless one, could be prey
a robin sings	hear a ROBIN SING;
the robin sings	see the ROBIN that SINGS, I am hidden by the rock, attack possible, as bird does not see me, but it is too far off
hawk	known HAWK, a bird, a dangerous one, minor attention, as I am hidden by the rock!
the hawk attacks robin	see THE HAWK ATTACK the ROBIN
the robin sings and the hawk attacks	THE ROBIN SINGS AND THE HAWK ATTACKS; ROBIN does not see the HAWK; danger for the ROBIN
many birds attack the hawk	see MANY BIRDS ATTACK THE HAWK
the hawk disappears	THE HAWK DISAPPEARS; attention no longer necessary?
birds sing	hear BIRDS SING; attention no longer necessary! BIRDS could be prey; attack possible?
two hawks appear	TWO HAWKS APPEAR; I am hidden by the rock, but increase attention!
two hawks attack me	TWO HAWKS ATTACK ME; alarm! I am not hidden from HAWKS by that rock
I must flee	I MUST FLEE; prepare quick flight action! I am not hidden from HAWKS by that rock! [modification of protection attributes in rock and hawk CONCEPTS]

propositions from top to bottom, we as meta-observers can infer what is happening. For the predator, as a cognitive agent, things are not as easy, because it has to identify objects as objects, even if they are moving, or if the view is deteriorated by that upcoming rain. The small predator has to analyze many details omitted in our description of the scene. In its action planning and memorizing, it has to omit them, too, and keep only details relevant for its survival and other central motives. That exactly is the main function of conceptualization. In the right column of Example 1 we verbally circumscribe the concept processes active in the small predator..

In Example 1 we see the interaction of concept use and the motivational system of an animal in a normal situation, though with a remarkable demand on perceptual processing, even under time constraints. There is no situational slot where the cognitive agent has to learn (form) a new concept or attribute of a concept: all the concepts necessary to solve the problems of the present situations have already been formed. During the last time slot, it has to modify an ATTRIBUTE. Purposefully, we choose an animal as actor in our example. So, we hope to have illustrated that concept processing is not a privilege of the highest evolved living beings.

3 Theoretical problems: attribute formation, language dependency of concepts

3.1 The attribute formation problem

Concept formation must make use of known attributes and their applicable values. To a certain degree, in acquiring a new concept, attributes from existing concepts may be used. Or, if this is not possible, the formation of a new attribute may be tried by applying a process similar to concept formation to form the attribute needed, using more elementary attributes. As can easily be seen, this leads to an infinite regression, letting the whole conceptual building of concepts "tumble". Unless, something like atomic attributes and values could be identified. Work of this type is currently in progress, trying to understand the basis (Roth&Menzel 1996) and the evolution of cognition (Lengeler/Müller/diPrimio 2000, Stewart 1996):

In a description of the neurobiological architecture of cognition, Roth and Menzel (1996) see as basic processes that must be distinguished from cognitive processes: "Precognitive processes like constancy processing (color and form constancy), simple perception processes like the differentiation of figure and background, or the automatic segmentation of complex scenes into "good gestalts", the detection of simple states of order, of patterns and objects (p. 539)". As we have argued in Lengeler/Müller/diPrimio (2000) and in diPrimio/Müller/Lengeler (2000) these order-analyzing processes are very elementary and cognitive processes that may well be found in earliest forms of life. May be that in their context atomic attributes and values can be found.

Let us look at an example: The concept BIRD is known. It has the attributes (features) shown in Example 2, part 1:

Example 2 (part 1): The CONCEPT

BIRD		
	FLIES	yes
	SINGS	yes
	LAYS EGGS	yes
	NESTS IN TREES	yes
	EATS INSECTS	yes
	SIZE	small

To introduce the concept ROBIN the new attribute HAS A RED BREAST has to be added to the description of the exemplar of a bird we conceptualize as ROBIN (see Example 2, part 2):

Example 2 (part 2): An exemplar of BIRD

ROBIN		
is a BIRD		
	HAS A RED BREAST	yes

Propositional attributes are only accessible for CONCEPT use "as a whole". If we want to break down the propositional attribute HAS A READ BREAST into observable features, we have to introduce several new attributes, that could be also important for the conceptual processing of BIRD in some later context. Note that a considerable CONCEPT formation effort, namely observing the details of a robin body and inferring the function of its parts, is necessary before the new attributes can be formed. The following could be an intermediate result (see Example 2, part 3):

Example 2 (part 3): New attributes to discern a ROBIN from other birds.

ROBIN					
IS A BIRD					
	BODY	BODY-PART	HEAD		
			BACK		
			FEET		
			TAIL	POSITION	
HORIZONTAL	back				
VERTICAL	middle				
LATERAL	no				
			BREAST	COLOUR	red
				POSITION	
HORIZONTAL	front				
VERTICAL	middle				
LATERAL	left-right				
	COLOUR	light-brown			

Attributes are displayed here as having a hierarchical structure for illustration (they can easily be transformed into a non-hierarchical form). As we know that ROBIN is a BIRD, to end up this intermediate attribute formation process, the new attributes must become part of the description of BIRD.

3.2 Concepts and language

It is a common understanding that language and concepts come together: no concepts without language. As may be inferred from the statements above, concepts can also be considered as pre-linguistic. We think that CONCEPTS play an important role in living beings not capable of language. CONCEPTS take part in, or are crucial building blocks of, cognition of animals, primates, and men; and once, hopefully of man-built artefacts. Pre-linguistic CONCEPTS are the basis of **CONCEPTS** as seen via language. Without pre-linguistic CONCEPTS, the mutual agreement process, that forms a common language, would not be possible. Language is a tool for communication between agents which have a common set of pre-linguistic CONCEPTS and agree upon the naming of them. The question if phenomena like "prototypicality" ("a robin is a typical bird"), as inferred from the analysis of **CONCEPTS** in language (Fodor 1998a/b, Smith&Medin 1981, Laurence&Margolis 1999), are "universal" phenomena also found in pre-linguistic CONCEPTS, cannot yet be answered (may be they are epiphenomena stemming from psycholinguistic tests or from language agreement).

Pre-linguistic CONCEPTS

Language and concepts are often seen as being inseparable: we think that language is a tool that allows to communicate about concepts which themselves have an important function in cognitive and cognition-like systems, independent from a language capacity. Evidence from the functionality of natural agents leads to the assumption that there must be structures very similar to the concepts known from cognitively highly developed beings that play an important role in many living beings, even those not capable of language. If natural agents can tackle the situations they are confronted with in their life, they must have command over an apparatus which in the most evolved living beings is associated with concepts. CONCEPTS simplify the problem-solving of living agents.

The substrate of CONCEPTS, or CONCEPT-like structures, and their exact form is not yet known. Our proposals concerning CONCEPTS in this article are working hypotheses. We hope to learn more about them when we implement CONCEPT processing features into robotic artefacts. CONCEPT processing is also part of joint efforts of biology and cognitive science in redefining cognition; efforts which we support and are already involved in (diPrimio/MüllerLengeler 2000, Lengeler/Müller/diPrimio 2000).

We assume that pre-linguistic CONCEPTS are the basis of **CONCEPTS** expressed in language.

CONCEPTS in language formation and learning

Language is a tool for communication between agents which have a common set of pre-linguistic concepts and agree upon how to name and arrange them. Without pre-linguistic CONCEPTS, this process of mutual agreement, the forming of a common language, would have no basis, and could never start.

If two or more agents independently from each other formed the CONCEPT CHAIR because, in solving problems, they needed it, they can agree upon naming it in the language formation process with the word "chair" or with other elements of their common language, having afterwards (after a certain experience in linguistically referring to CHAIR) a common **CONCEPT CHAIR.** For a cognitive agent, the **concept chair** refers to both language immanent experiences with **chair** (e.g., using **chair** in language expressions and experiencing **chair** being used in language expressions, including defining and redefining **chair** via language expressions) and to problem solving experiences with CHAIR (e.g., categorizing objects as CHAIRS when having the need to sit down, thereby integrating over complicated form and aspect differences).

In Example 3 we illustrate the assumption that there is no one-to-one relation between pre-linguistic CONCEPTS and natural language **CONCEPTS**.

Example 3: From CONCEPT to **CONCEPT**

conceptualized prelinguistically in Agent 1	*conceptualized prelinguistically n Agent 2*	*conceptualized linguistically in Agent 1 and Agent 2*	*commentary*
LEAF	LEAF	**LEAF**	as communication affords the naming of this shared CONCEPT
LEAF-WITH-HOLE	LEAF-WITH-HOLE	**LEAF**	as communication does not afford the naming of this shared CONCEPT
LEAF-WITH-YELLOW-TIP	LEAF	**LEAF**	as the agents do not share this CONCEPT

Concept phenomena and epiphenomena

Which of the **CONCEPT** phenomena can be found in CONCEPTS? Do, for instance, CONCEPTS show "prototype" effects ("a **ROBIN** is a typical bird")? One of the following may be the case, and we do not know enough about CONCEPTS to decide which one without further research and experimentation:

1. CONCEPTS do not show the prototype phenomenon as in a "mute" problem-solving context it is not advantageous to be able to refer to the most typical exemplar or subconcept of a CONCEPT, because would have the same advantage as to go to the next higher level in the CONCEPT hierarchy.

2. Prototypes are epiphenomena generated during the mutual agreement on a language (which can also be in learning one's mother language). Causes could be teaching preferences or cultural preferences; it might be that the first **CONCEPT** verbalized is a candidate for becoming the prototype of its class.

3. The prototype effect is a result of the psycholinguistic experimentation itself, because it very much depends on the choice of the features for the concepts that are used for the experiments (see Example 4 below. If, for instance, the attribute SIZE in BIRD would allow larger sizes, the prototypicality of robin could no longer be explained with that set of attributes.). And it depends on the psycholinguistic task posed and the psycholinguistic question to be answered: If the task is "Bring the **CONCEPTS** in an order, that shows what **CONCEPT** is more typical as compared to another **CONCEPT**!", you cannot help but getting a prototype. We do not want to criticize a long tradition of experiments and we know that this argument might be a bit unfair, as there is also evidence for prototypes that comes from quite different experiments, but we want to emphasize that it is worthwhile to look at the phenomena freshly with the new aspect of linguistically codified **CONCEPTS** and pre-linguistic CONCEPTS in mind.

Example 4: Prototypes — The concept **ROBIN** has the same attributes and values as **BIRD**, it is a typical **BIRD**.

	different "birdy" concepts			
attributes	BIRD	ROBIN	CHICKEN	VULTURE
FLIES	yes	yes	no	yes

SINGS	yes	yes	no	no
LAYS EGGS	yes	yes	yes	yes
NESTS IN TREES	yes	yes	no	yes
EATS INSECTS	yes	yes	no	no
SIZE	small	small	middle	large

4 Comparison with selected approaches in the theory of concepts

In this section we take a glimpse at the very old, very broad, and still ongoing discussion of the concept of concept. We picked out natural language semantics (de Swart 1998), Jackendoff's Conceptual Grammar (1994), and Laurence and Margolis' overview of concept theory in cognitive science (1999). The aim of this section is to show where our own assumptions are supported, and where they are decidedly different, or do not cover important phenomena.

Support may be drawn from de Swart, in her exposition of natural language semantics (1998, 1–7), when she argues, that language has content and this content is anchored to reality via some aboutness relation. We assume that the pre-linguistic concepts are this "aboutness anchor". When she says that communication is only successful if the idea the hearer gets is the same as what the speaker intended the hearer to get, she gives a formulation which is not only good for describing communication, but also for describing what is the goal of the language agreement processes. de Swart admits that the construction of complex concepts is a problem which has not yet been entirely understood by natural language semantics. Which is certainly by far more true for our approach: what we call composition of CONCEPTS is formally less elaborated than the theory of the compositionality of meaning in semantics; ours is only one aspect of the latter. This hints at the question of a semantics of pre-linguistic concepts, totally left open in our presentation above.

Our conception of pre-linguistic CONCEPTS is different from de Swart's point of view regarding the *direct interpretation* of natural language with respect to the outside world, which she characterizes as being not easy. Instead of direct interpretation, de Swart advocates the *indirect interpretation* by means of translation into formal languages, which capture increasingly more complex parts of the meaning of natural language, as she says: the exact meaning will be captured, according to her, as long as this translation is perfect. While being a common point of departure in semantics, we nevertheless think it is a problematic one: one problem (direct interpretation) is replaced by another unsolvable problem, the translating perfectly into a formal language. Certainly, this has also to do with a basic shyness of linguistics to tackle the interface between language and world, a sound shyness in many other respects. Pre-linguistic CONCEPTS deal with this interface, hopefully supported by evidence from circumstances in natural agents.

Many aspects in de Swart's description are not covered by our exposition above. This concerns the dichotomy of the scope of semantics in *lexical semantics* and *meaning at the sentence level*, mental states and imaginary worlds (we alluded to it by associating CONCEPTS with internal circumstances in an agent), and the role of natural language as both object language and metalanguage. A lexical semantics phenomenon we referred to is hyponymy (*isa*-hierarchies), but we did not cover ambiguity, synonymy, antonymy, semantic features, thematic roles (argument structure) of verbs). It is to question which of these lexical meaning aspects can be found at the CONCEPT level and which are only found with **CONCEPTS**. The same is true for the meaning at sentence level (coreferentiality, binding (reference of pronoun varies systematically with the choice of the individual determined by a quantifier), forms of semantic inference like presupposition and implicature).

Jackendoff, in his conceptual semantics, as de Swart (1998, p. 5) sees it, mixes ideas from predicate logic, theories about thematic roles (agent, patient), and psychological theories on (visual) perception. We share Jackendoff's interest in perception, as it supports our view that object recognition might be a key to pre-linguistic CONCEPTS, and inversely, CONCEPTS a key to object recognition. We would like to emphasize, with respect to Jackendoff's argument for a mental grammar and innate knowledge, that it is an innate capacity for problem-solving the CONCEPT system can recur to; which Jackendoff expresses as : "… the language capacity must have evolved from other capacities in the brains of our precursors (Jackendoff 1994, 160)". In the conclusions of his 1994 book (p. 203), Jackendoff summarizes his ideas as follows: "Our thoughts are built out of a finite set of unconscious patterns which give us the potential for thinking an infinite number of thoughts of indefinite complexity. … These patterns in turn are constructed from an innate Universal Grammar of concepts …" A more formal comparison with our approach is still to be done.

In the introductory chapter of Margolis & Laurence (1999), Laurence and Margolis give a detailed description of the state and the history of concept theory. They discern several types of concept theories which we will address and evaluate, one after another. Support of our own views comes from the *Classical Theory*, which states that most concepts (esp. lexical concepts) are structured mental representations that encode a set of necessary and sufficient conditions for their application, if possible, in sensory and perceptual terms. (Where this last condition is especially important for us). Important for the assumption of pre-linguistic concepts are also the following parts of their *Summary of Criticism of the Classical Theory* (the order of the arguments differ in Laurence & Margolis (1999): 1. “It is possible to have a concept in spite of massive ignorance and/or error, so concept possession can’t be a matter of knowing a definition.” As we interpret it, this is supporting the existence of CONCEPTS, and illustrates the fragility of the conversion of CONCEPTS into **CONCEPTS** in language. 2. “Concepts and categorization both admit a certain amount of indeterminacy (fuzziness), not possible in the Classical Theory.” As abstractions that are continuously verificated, rearranged and modified, concepts are “fuzzy” in themselves. The language agreement process adds to the indeterminacy, in itself and as it is also continuously active, language agreement results never being completed because accurate. 3. “Typicality effects can’t be expressed by classical models.” Here, we can only repeat our epiphenomenon suspicion with regard to typicality. The arguments 4 “There are few, if any, defined concepts” and 5 “Lexical concepts show no effects of definitional structure in psychological experiments” are not covered as we do not systematically treat concepts in a “definitional” role in our proposal.

When it comes to *Prototype Theory* Laurence and Margolis argue: “Most concepts (esp. lexical concepts) are structured mental representations that encode the properties that objects in their extension tend to possess”. If our assumptions concerning the concept processes are valid, they support the statistical nature of CONCEPTS and **CONCEPTS**. The argumentations of the *Theory-Theory* (“Children and scientists have the same method of exploring the world”) and of the *Neoclassical Theory* (“partial definitions are allowed”) are not covered by us, whereas we are inclined to support the *Conceptual Atomism Theory* in its coinage of Fodor, at least what regards the general statement that “lexical concepts are primitive”, while we think different concerning the statement “they have no structure”. It is certainly too general to say that pre-linguistic CONCEPTS are primitive, but in a certain sense they are “more primitive” than lexical concepts. An evaluation of Fodor’s concept theory (Fodor 1998a, 1998b) is on our agenda.

5 Outlook

In current robotics the concept processing problem has not yet been tackled. An experimental implementation is overdue. Solutions in the line of our argumentation would have a considerably impact on the re-engineering of biological solutions, and thus on the engineering of autonomous robots, and as a source of a tentative verification, on the theory of cognition.

Acknowledgements

My thanks go to the members of the former Cognitive Robotics group at GMD, to the members of the Biomimetic Autonomous Robots Team, its team leader, Frank Kirchner, and the institute's director, Thomas Christaller, both for discussing the concept of CONCEPTS with me and having patience with me when I insisted on pursuing my own ideas. This applies to a degree worth mentioning separately to Joachim Loosen. Bernhard Schröder (University of Bonn) kindly commented on a draft of this paper.

References

[Bajscy&Large 1999] Bajscy, Ruzena and Edward W. Large (1999) When and Where Will AI Meet Robotics? *AI MAGAZINE*, Fall, 57–65.

[Chater&Heyes] Chater, Nick; Heyes, Cecilia (1994) Animal Concepts – Content and Discontent. *Mind & Language* 9, 3, 209–246.

[deSwart 1998] de Swart, Henriette (1998) Introduction to Natural Language Semantics. Stanford, Cal.: CSLI Publications.

[diPrimio/Müller/Lengeler 2000] di Primio, Franco; Müller, Bernd S. & Joseph W. Lengeler (2000) Minimal Cognition in Unicellular Organisms. To appear in: Meyer, Berthoz, Floreano, Roitblat and Wilson (Eds.). *Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior (SAB2000) – Supplement Book*. Honolulu: International Society for Adaptive Behavior.

[Fodor 1998a] Fodor, Jerry (1998) When is a dog a DOG? Review of Sloman, Steven A. and Lance J. Rips (Eds.) Similarity and Symbols in Human Thinking, MIT Press 1998. *Nature* 396, 26 November, 325–327.

[Fodor 1998b] Fodor, Jerry (1998) *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon.

[Jackendoff 1994] Jackendoff, Ray (1994) *Patterns in the mind – Language and human nature*. New York: BasicBooks.

[Lengeler/Müller/diPrimio 2000] Lengeler, Joseph W., Müller, Bernd S. & Franco di Primio (2000) Neubewertung kognitiver Leistungen im Lichte der Fähigkeiten einzelliger Lebewesen. *Kognitionswissenschaft* 8.

[Laurence & Margolis] Laurence, Stephen and Eric Margolis: Concepts and Cognitive Science. Introductory chapter in: Margolis, Eric and Stephen Laurence: "Concepts: Core Readings ", Cambridge, Mass.: MIT, 1999, 1–81.

[Müller 1997] Müller, Bernd S. (1997) *Konzeptbildung und Kommunikation bei Robotern – Problemskizze und Vorschläge zum experimentellen Einstieg.* Arbeitspapiere der GMD Nr. 1080.

[Müller 1998] Müller Bernd S.(1998) *Identifikation elementarer kognitiver Leistungen*. GMD Report No. 17.

[Roth&Menzel 1996] Roth, G. & Menzel, R.. (1996) Neuronale Grundlagen kognitiver Leistungen. In: J. Dudel, R. Menzel & R. F. Schmidt (Hrsg.). *Neurowissenschaft – Vom Molekül zur Kognition* (S. 539-558). Berlin: Springer.

[Smith&Medin 1981] Smith, E.E. & D.L. Medin (1981) *Categories and Concepts.* Cambridge, Massachusetts: Harvard University Press.

[Stewart 1996] Stewart, J. (1996) Cognition = Life: Implications for higher-level cognition. *Behavioural Processes* 35, 311–326.

Natürlichsprachliche Kommunikation im Rahmen der Datenbanksemantik

Roland Hausser

14. September 2000

Abstract

Database semantics is a computational model of natural language communication based on the following innovations. One is a new data structure, cclled a wordbank, which stores content in the form of distributed, bidirectionally pointered feature structures; these code concatenated propositions and serve as the semantic representation for language interpretation and production. The other is the time-linear algorithm of left-associative grammar (LA-grammar); it is used for reading content coded in natural language into and out of the wordbank. The system is presented as a declarative model of a cognitive agent which functionally integrates the procedures of interpretation, conceptualization, production, query, and inference.

– Datenbanksemantik ist ein Computermodell der natürlichsprachlichen Kommunikation, das auf folgenden Innovationen beruht. Die eine ist eine neue Datenstruktur, *Wortbank* genannt, in der Inhalte als verteilte, bidirektional verzeigerte Merkmalstrukturen gespeichert werden; diese kodieren verkettete Propositionen und dienen als semantische Repräsentation zur Interpretation und Produktion von Sprache. Die andere ist der zeitlineare Algorithmus der linksassoziativen Grammatik (*LA-grammar*); dieser wird für das Ein- und Auslesen natürlichsprachlich kodierte Inhalte in die, bzw aus der, Wortbank verwendet. Das System wird als deklaratives Modell eines kognitiven Agenten präsentiert, das die Prozeduren der Interpretation, Konzeptualisierung, Produktion, Befragung und Inferenz funktional integriert.

1 Datenbankmetapher natürlichsprachlicher Kommunikation

Bei der Interaktion mit einem Standardcomputer ist das Verstehen der natürlichsprachlichen Zeichen fast ganz auf den Benutzer beschränkt. Wenn ein Benutzer z.B. in einer Datenbank nach einem roten Objekt sucht, wird das Wort **rot** vor der Eingabe in den Computer und später wieder bei der Ausgabe vom Benutzer verstanden. Innerhalb des Computers wird **rot** dagegen als ein Zeichen manipuliert, das bezüglich der denotierten Farbe uninterpretiert ist.

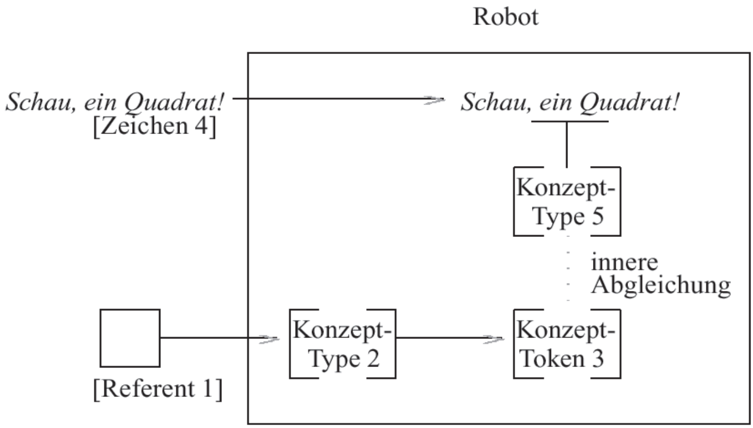

***Abb. 1:** Unmittelbare Referenz auf der Grundlage interner Abgleichung*

Was für einen Standardcomputer gilt, trifft jedoch auf die Mensch-Maschine-Kommunikation nicht allgemein zu. Die heutige Technologie ermöglicht den Bau eines autonomen Roboters, den man beauftragen kann, einen Gegenstand zu bringen, dessen Typ ihm bekannt ist, den er aber bisher noch nie gesehen hat (z.B. das rote Buch auf dem Schreibtisch im Nebenzimmer). Wenn diese Maschine in der Lage ist, spontan eine offene Menge verschiedener Aufgaben dieser Art auszuführen, hat sie ein Sprachverständnis, das auf einer bestimmten Abstraktionsebene mit den entsprechenden kognitiven Vorgängen beim Menschen als funktional äquivalent angesehen werden kann.

Die Interaktion eines solchen Roboters mit seiner Handlungsumgebung basiert auf Mustern, die auf die herein- und herausgehenden Parameterwerte seiner Erkennungs- und Handlungskomponenten abgeglichen werden. Diese Muster konstituieren eine prozedurale Definition von Konzepten, die neben ihrer Funktion im nichtsprachlichen Erkennen und Handeln eine zweite Funktion als Sprachbedeutungen übernehmen können.[1]

Die Sprachinterpretation des Roboters muß funktionieren, ohne daß dafür eine externe (metasprachliche) Relation zwischen dem Zeichen und dem Referenten postuliert wird. Dies wird über das Prinzip der *internen Abgleichung* geleistet, das im folgenden mit der sprachlichen Referenz auf ein Quadrat illustriert wird.

Der Roboter erkennt das Referenzobjekt (1), indem ein geeigneter Konzept-Type (2) auf die hereinkommenden Parameterwerte (hier *bitmap outline*) abgeglichen und das Referenzobjekt intern als ein Konzept-Token (3) instantiiert wird. Die Oberfläche des Sprachzeichens (4) wird entsprechend erkannt. Der Oberfläche **Quadrat** ordnet das Lexikon denselben Konzept-Type wie (2) als wörtliche Bedeutung (5) zu.

Die Beziehung zwischen der Sprachebene und der Ebene der nichtsprachlichen Kognition (auch Kontextebene genannt) wird über das Abgleichen zwischen dem lexikalischen Konzept-Type (5) und dem referentiellen Konzept-Token (3) etabliert. Eine lexikalisch als Type definierte wörtliche Bedeutung kann auf beliebig viele Referenz-Token abgeglichen werden. Dies modelliert die Flexibilität, die die natürlichen Sprachen von den Logik- und den Programmiersprachen unterscheidet.

Die in Abb. 1 illustrierte Art der Referenz wird *unmittelbar* genannt, weil Zeichen und Referent in der aktuellen Handlungsumgebung präsent sind. Dagegen wird die Referenz auf Objekte, die nicht in der aktuellen Handlungsumgebung enthalten sind, z.B. auf die Person J.S. Bach, *mittelbar* genannt. Bei der mittelbaren Referenz interagiert der Agent mit seiner Handlungsumgebung allein auf der Sprachebene, wobei als Referenten Objekte im Erinnerungsspeicher verwendet werden.

Die interne Abgleichung wird in Abb. 1 an dem einzelnen Wort **Quadrat** illustriert. Das Grundprinzip kann jedoch auf voll entwickelte natürlichsprachliche Kommunikation generalisiert werden, und zwar mit Hilfe der *Datenbankmetapher*. Diese geht davon aus, daß das Wissen der Sprecher-Hörer in der Form von Datenbanken repräsentiert ist. Kommunikation ist erfolgreich, wenn der Sprecher eine bestimmte Teilstruktur seiner Datenbank in Sprache kodiert und der Hörer diese in seiner Datenbank *analog* rekonstruiert, und zwar sowohl bzgl. der korrekten Dekodierung des Inhalts als auch bzgl. der korrekten Speicherung an einer entsprechenden Stelle.

Die Datenbankmetapher geht von vornherein davon aus, daß die kommunizierenden Datenbanken in kognitive Agenten (z.B. Roboter) eingebettet und dabei Schnittstellen zum sprachlichen und nichtsprachlichen Erkennen und Handeln realisiert sind. Deshalb kann die Datenbankmetapher auf Aussagen, Fragen und Aufforderungen gleichermaßen angewendet werde. Bei Aussagen speichert der Hörer den vom Sprecher kodierten Datenbankinhalt an einer entsprechenden Stelle seiner Datenbank. Bei Fragen bestimmt der Hörer, wo die vom Sprecher erbetenen Information gespeichert ist, bzw. gespeichert sein sollte. Bei Aufforderungen erstellt der Hörer eine kognitive Rekonstruktion der vom Sprecher erbetene Handlung.

Die programmiertechnische Umsetzung der Datenbankmetapher umfaßt folgende Aufgaben.

1.1 Realisierungsaufgaben der Datenbankmetapher

1. Kodierung von Datenbankinhalten in natürliche Sprache (Sprechermodus)
2. Dekodierung von Sprachzeichen in Datenbankinhalte (Hörermodus)
3. Befragung der Datenbank nach bestimmten Inhalten (Sprecher befragt den Hörer)
4. Inferenzen auf Datenbankinhalten zur Bestimmung der Identität zwischen Substantiven, der temporalen Relationen, der Interpretation nichtwörtlicher Verwendungen (z.B. Metaphern) etc.
5. Indizierung der Speicherstelle, von der die in Sprache kodierten Sprecherinhalte stammen, so daß der Hörer die entsprechenden Speicherstelle für die rekonstruierten Inhalte bestimmen kann.

Um zunächst die Natur des Indizierungsproblems zu erklären, vergleichen wir die Interaktion zwischen Benutzer und konventioneller Datenbank mit der zwischen Sprecher und Hörer.

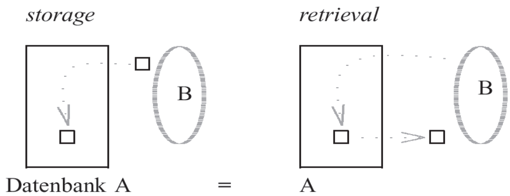

Abb. 2: *Interaktion mit einer konventionellen Datenbank*

Die Interaktion findet zwischen zwei verschiedenen Entitäten statt, dem Benutzer (Ovale) und der Datenbank (große Vierecke), wobei die kleinen Vierecke die Sprachzeichen repräsentieren, die als Ein- und Ausgabe dienen. Der Benutzer steuert die Speicher- und Abrufoperationen mit Hilfe einer Programmiersprache, deren Befehle auf der Datenbank als elektronische Prozeduren ausgeführt werden.

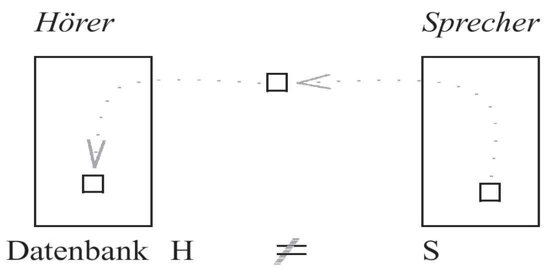

Abb. 3: *Interaktion zwischen Sprecher und Hörer*

Hier repräsentieren die großen Vierecke kognitive Agenten, die natürlich oder künstlich sein können. Es gibt keinen Benutzer. Vielmehr findet die Interaktion zwischen zwei gleichen und gleichberechtigten Agenten statt, die ihren Informationsfluß gegenseitig steuern, indem sie abwechselnd den Sprecher- und den Hörermodus übernehmen (*turn taking*). Der Sprecher steuert die Sprachproduktion als autonomer Agent. Die Interpretation des Hörers wird von den hereinkommenden Sprachzeichen gesteuert.

2 Inhaltsdarstellung: verkettete Propositionen

Die beiden ersten Aufgaben zur Realisierung der Datenbankmetapher, also die Kodierung und Dekodierung, erfordern eine allgemeine, Sprach- und Domänenunabhängige Methode zur Darstellung von Inhalten. Hierfür verwenden wir den klassischen Begriff einer Proposition. Im Einklang mit Aristoteles (384–322 v.Chr.) sind Propositionen einfache Zustandsdarstellungen. Sie sind so allgemein und abstrakt, daß sie gleichzeitig als Teilzustände realer oder möglicher Welten und als Bedeutung sprachlicher Sätze aufgefaßt worden sind.

Propositionen setzen sich aus drei Arten von Grundelementen zusammen, den *Argumenten*, den *Funktoren* und den *Modifikatoren*. Eine elementare Proposition besteht aus einem Funktor, der mit einer für ihn charakteristischen Anzahl von Argumenten kombiniert wird. Die Modifikatoren sind optional und können auf Funktoren oder Argumenten angewendet werden.

In der Welt sind die Argumente die *Objekte* (im allgemeinsten Sinn), die Funktoren die (intrapropositionalen) *Relationen* zwischen den Objekten und die Modifikatoren die *Eigenschaften* der Objekte oder Relationen. Mit diesen Grundelementen ist eine allgemeine Ontologie zur Darstellung kognitiver Zustände gegeben - im Unterschied zu spezielleren Ontologien, wie etwa zur Beschreibung der Welt aus Sicht der Physik (basierend auf Atomen, Masse etc.), der Biologie (basierend auf Stoffwechsel, Fortpflanzung etc.) oder der Ökonomie (basierend auf Angebot, Nachfrage, Geldmenge etc.).

In den natürlichen Sprachen sind die Argumente die *Substantive*, die Funktoren die *Verben* und die Modifikatoren die *Adjektive* (wobei Adnominale auf Substantive und Adverbiale auf Verben angewendet werden). Die Substantive, Verben und Adjektive umfassen die *Inhaltswörter* einer Sprache.

	Logik	*Welt*	*Sprache*
1.	Argument	Objekt	Substantiv
2.	Funktor	Relation	Verb
3.	Modifikator	Eigenschaft	Adjektiv

Abb. 4: *Die Bausteine elementarer Propositionen*

Elementare Propositionen können mit *Operatoren* verändert (z.B. Negation) oder miteinander verknüpft werden (z.B. Disjunktion). In den natürlichen Sprachen entsprechen den Operatoren die *Funktionswörter*. In der Welt entsprechen den Operatoren u.a. die *extrapropositionalen Relationen*.

Als Beispiel einer Inhaltsdarstellung mittels verketteter Propositionen betrachten wir einen Roboter, der gerade ein Dreieck und ein Viereck in Feld A2 seiner Handlungsumgebung findet und diese Situation automatisch analysiert. Wir gehen davon aus, daß die Konzepte **Feld**, **Dreieck**, **Viereck** und **enthalten** als Programme realisiert sind, die von der visuellen Wahrnehmung des Roboters gelieferte Bitmap-Umrisse analysieren.

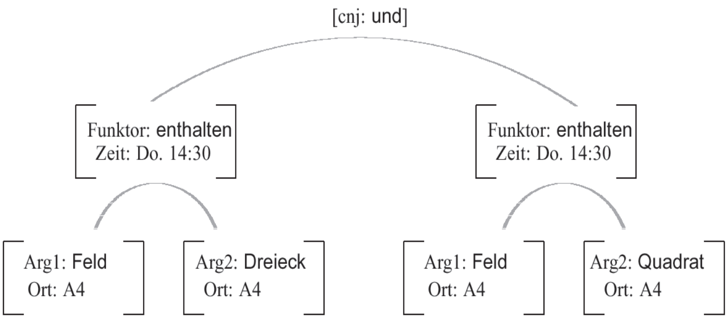

Abb. 5: *Verkettung zweier elementarer Propositionen*

Die elementaren I-Propositionen in Abb. 5 können mit **Feld enthält Dreieck** und **Feld enthält Viereck** paraphrasiert werden. Sie sind über die extrapropositionale Relation **und** miteinander verkettet.

Die Merkmalstrukturen für Argumente (Objekte, Substantive) bestimmen den Ort, an dem das System ihnen begegnet ist. Die Merkmalstrukturen für Funktoren (intrapropositionale Relationen, Verben) bestimmen den *Zeitpunkt*, an dem die Relation festgestellt wurde. Bei Modifikatoren (Eigenschaften, Adjektiven) stimmt die raum-zeitliche Bestimmung mit derjenigen der modiozierten Struktur überein.

3 Datenstruktur einer Wortbank

Das in Abb. 5 illustrierte Format verwendet *graphische* Mittel, um die Funktor-Argument-Struktur des Inhalts auszudrücken. Diese sind gut geeignet, um semantische Intuitionen auszudrücken weshalb Baumstrukturen in der Linguistik sehr beliebt sind. Für eine Implementierung der natürlichsprachlichen Kommunikation im Rahmen der Datenbankmetapher sind graphische Darstellungen jedoch ungeeignet. Denn Datenbanken kodieren Relationen abstrakt, wobei in Beziehung stehende Einheiten beliebig weit von einander gespeichert werden können.

Um eine einfache und effiziente Prozedur zum natürlichsprachlichen Einbetten und Extrahieren propositionaler Inhalte zu unterstützen, rekodieren wir das graphische Format von Abb. 5 als Tabelle aus Merkmalstrukturen, wobei die Funktor-Argumentstruktur elementarer Propositionen (intrapropositionale Relationen) und die Verkettung elementarer Propositionen (extrapropositionale Relationen) mit Hilfe von bidirektional verzeigerten Attributen ausgedrückt wird. Diese neuartige Datenstruktur heißt *Wortbank*.[2]

TYPES	PROPLETS	
Konzept: enthalten Rolle: Funktor	funk: enthalten ARG:Feld Dreieck cnj: 23 und 24 prn: 23	funk: enthalten ARG:Feld Quadrat cnj: 23 und 24 prn: 24
Konzept: Feld Rolle: Argument	arg: Feld FUNK: enthalten id: 7 prn: 23	arg: Feld FUNK: enthalten id:7 prn: 24
Konzept: Quadrat Rolle: Argument	arg: Quadrat FUNK: enthalten id: 9 prn: 24	
Konzept: Dreieck Rolle: Argument	arg: Dreieck FUNK: enthalten id: 8 prn: 23	

Abb. 6: *Rekodierung von Abb. 5 als Wortbank*

Die Daten sind als alphabetisch geordnete Listen von Types angeordnet, wobei jedem Type eine offene Anzahl von Token folgt. Die Types drücken die Grundeigenschaften der Inhaltswörter aus, wie z.B. ihre grammatische Rolle. Die Token, *Proplets*[3] genannt, fungieren als die Elemente der einzelnen Propositionen. Eine Zeile, bestehend aus einem Type gefolgt von einer offenen Anzahl von Proplets, wird eine *Tokenzeile* genannt.

Proplets, die zu derselben Proposition gehören, werden durch ihre gemeinsame Propositionsnummer **prn** zusammengehalten. Merkmale, die zusammenhängende Proplets über intra- und extrapropositionale Relationen bestimmen, heißen *Fortsetzungsprädikate*.

In intrapropositionalen Relationen bestimmt ein Argument (Objekt, Substantiv) den zugehörigen Funktor, ein Funktor (Relation, Verb) die zugehörigen Argumente und ein Modifikator (Eigenschaft, Adjektiv) das zugehörige Modifikandum und umgekehrt. In extrapropositionalen Relationen kann ein Argument die Identität oder Nicht-Identität mit anderen Argumenten, und ein Funktor konjunktionale Relationen (z.B. **und**, **dann**, **weil**) mit anderen Funktoren bestimmen.

Die Tokenzeilen einer Wortbank entsprechen der Struktur einer klassischen *Netzwerkdatenbank*. Eine Netzwerkdatenbank definiert eine 1:n Relation zwischen zwei Arten von Verbünden, den *owner*- und den *member*-Verbünden. In einer Wortbank fungieren die Types als *owner* und die zugehörigen Proplets als *member records*.

Die Kodierung propositionaler Inhalte als verteilte, bidirektional verzeigerte Merkmalstrukturen in der Databansemantik ist nicht nur eine technische Bequemlichkeit. Sie ist vielmehr entscheidend für eine einheitliche Behandlung der NL-Interpretation, Inferenz und NL-Produktion.

4 Intuitive Überlegungen zur Definition einer Datenbanksprache

Die Definition einer Datenbanksprache (DBS) zum automatischen Ein- und Auslesen von Inhalten in die Wortbank entspricht der Definition z.B. des Prädikatenkalküls. So wie der Prädikatenkalkül als künstliche Sprache entwickelt wurde, die bestimmte wahrheitskonditionale Eigenschaften der natürlichen Sprachen modelliert, so wird DBS dafür entwickelt, die automatischen Ein- und Ausleseprozeduren zu modellieren, die für die natürlichen Sprachen charakteristisch sind. Und so wie der Prädikatenkalkül bzgl. eines mengentheoretisch definierten Modells interpretiert wird, wird DBS bzgl. einer Datenbank (Wortbank) mit bidirektional verzeigerten Merkmalstrukturen (Proplets) interpretiert.

Wie ein Modell, ist eine Wortbank sprachunabhängig. Während jedoch ein Modell als Darstellung der Welt konzipiert ist, dient eine Wortbank zur Darstellung kognitiver Inhalte. Außerdem werden Elementarbedeutungen wie *rot*, *Dreieck* und *enthalten* in der logischen Semantik extensional behandelt, während sie in der Wortbank auf den Wahrnehmungs- und Handlungsprozeduren des kognitiven Agenten basieren, in den die Wortbank eingebettet ist.

Die semantische Interpretation eines DBS-Satzes besteht in der Konstruktion einer Propletmenge während der syntaktischen Analyse. Dies wird im folgenden mit der Analyse der Elementarproposition **Feld enthält Dreieck** illustriert.

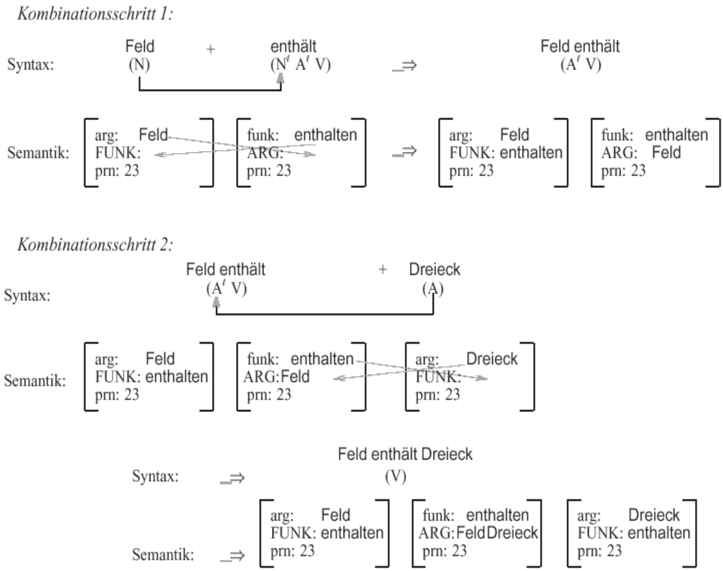

***Abb. 7:** DBS-Interpretation von Feld enthält Dreieck*

Die syntaktische Analyse besteht in dem zeitlinearen *Kürzen* von Valenzstellen durch Valenzfüller. Beim ersten Kombinationsschritt kürzt der Satzanfang **Feld** die Nominativvalenz N´ von **enthält**. Beim zweiten Kombinationsschritt kürzt das nächste Wort **Dreieck** die Akkusativvalenz A´ in dem Satzanfang **Feld enthält**.

Die semantische Interpretation besteht in dem *Kopieren* von Propletnamen in die Fortsetzungsattribute anderer Proplets. Hierfür wird jedem Wort lexikalisch eine geeignete Merkmalstruktur zugeordnet, die den Namen des Proplets, z.B.

Feld, und die intrapropositionalen Fortsetzungsprädikate (Großbuchstaben) mit dem Wert NIL (hier durch den leeren Wert dargestellt) enthält. Die Interpretation der ersten syntaktischen Komposition kopiert den Propletnamen **Feld** in das ARG-Merkmal von **enthalten**, und den Propletnamen **enthalten** in das FUNK-Merkmal von **Feld**. Die Interpretation der zweiten Komposition kopiert den Propletnamen **enthalten** in das FUNK-Merkmal von **Dreieck** und den Propletnamen **Dreieck** in das ARG-Merkmal von **enthalten**.

Das Ergebnis ist eine (ungeordnete) Menge der drei koindizierten Proplets **Feld**, **enthalten** und **Dreieck**. Diese Proplets sind autonome Entitäten, die gemeinsam die Funktor-Argument-Struktur ‘ihrer’ Proposition über ihre jeweiligen Fortsetzungsprädikate spezifizieren. Das Indizierungsproblem wird automatisch gelöst, indem jedes Proplet am Ende seiner alphabetisch geordneten Tokenzeile in der Wortbank gespeichert wird. Die Datenbanksemantik verwendet also die Propletnamen als Primärschlüssel.

Als nächstes betrachten wir die Produktion eines Satzes aus der Wortbank:

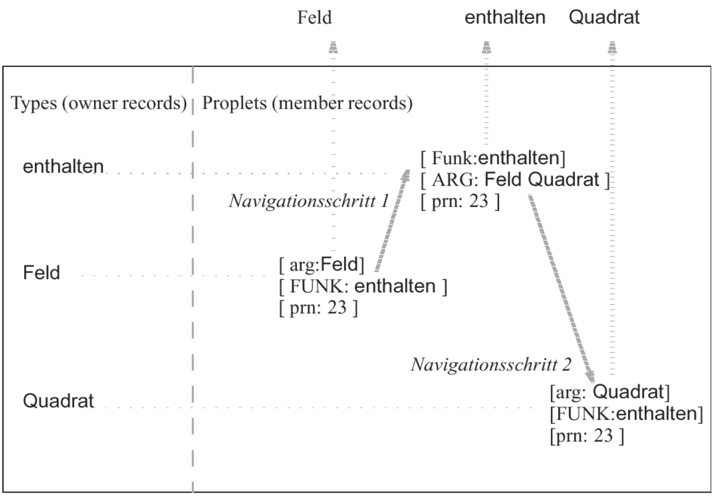

Abb. 8: *DBS-Produktion von* ***Feld enthält Quadrat***

Eine Wortbank geht über eine klassische Netzwerkdatenbank insofern hinaus, als sie im Rahmen ihre Verbundstrukturen *mögliche Fortsetzungen* definiert. Diese bilden die Grundlage für Operationen, die konventionelle Datenbanken nicht zur Verfügung stellen, nämlich eine autonome, zeitlineare Navigation durch die verketteten Propositionen der Datenbank.

Die Navigation beginnt z.B. beim Proplet **Feld** mit der Propositionsnummer (prn) 23. Aufgrund des Fortsetzungsprädikats [FUNK: **enthalten**] sucht der Navigationsalgorithmus alphabetisch nach der Tokenzeile von **enthalten** und bewegt sich vom Type zum Token mit der prn 23. Das 'nächste' Proplet **enthalten** hat das Fortsetzungsprädikat [ARG: **Feld Quadrat**]. Der erste Wert **Feld** bestätigt die vorangehende Navigation. Der zweite Wert **Quadrat** ist das neue 'nächste' Proplet. Es wird gefunden, indem der Navigationsalgorithmus alphabetisch nach der Tokenzeile von **Quadrat** sucht und vom Type zum Token mit der prn 23 geht. Auf diese Weise wird Proposition 23 vollständig traversiert. Dabei kann der Inhalt automatisch herausgelesen werden, indem der Wert des ersten Attributes der traversierten Merkmalstrukturen automatisch kopiert und geäußert wird – wie die oberste Zeile von Abb. 8 zeigt.

Nach einer intrapropositionalen Navigation kann als nächstes eine extrapropositionale zu einer anderen Proposition folgen, entweder auf der Grundlage einer Identität zwischen Argumenten (Objekten, Substantiven) oder einer Konjunktion zwischen Funktoren (Relationen, Verben). Auf diese Weise kann der Inhalt einer Wortbank traversiert werden, entweder zufällig im Sinne der freien Assoziation oder nach bestimmten Strategien oder Mustern. Diese Navigationen sind ein einfaches Model des Denkens. Wenn das System im Sprechermodus ist, äußert es Wörter, die den traversierten Knoten entsprechen.

Auf einer abstrakteren Ebene als Abb. 7 und 8 kann der zeitlineare Charakter der DBS-Interpretation und -Produktion folgendermaßen beschrieben werden:

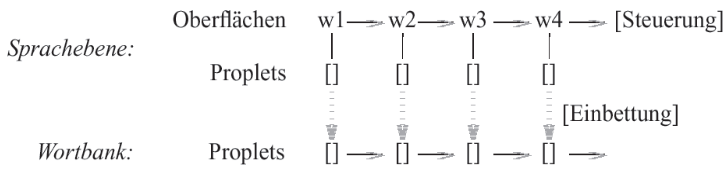

Abb. 9: *Schema der DBS-Interpretation (Analyse)*

Das System folgt den Oberflächen der Zeichen, bestimmt ihre Bedeutungen (Proplets) über das Lexikon, ergänzt die Werte der Fortsetzungsprädikate bei jedem zeitlinearen syntaktisch-semantischen Ableitungsschritt und ordnet die vervollständigten Proplets in die entsprechende Tokenzeile der Wortbank.

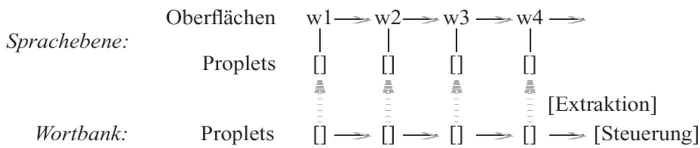

Abb. 10: *Schema der DBS-Produktion (Generierung)*

Das System navigiert durch die verketteten Propositionen der Wortbank, paßt auf jedes traversierte Proplet einen entsprechenden Lexikoneintrag und äußert dessen Oberfläche.

Die Schemata 9 und 10 genügen der natürlichen Auffassung, daß Interpretation (↓) und Produktion (↑) inverse vertikale Prozesse sind. Trotzdem haben Interpretation und Produktion ihre eigentliche Hauptrichtung gemeinsam, nämlich die Richtung der horizontalen zeitlinearen Grundstruktur (→).

5 Begriff des Fragments

Die bisherigen Überlegungen zur Interpretation und Produktion natürlicher Sprache können als ‘Fragment’ formuliert werden. Der Fragmentbegriff wurde von R. Montague für die vollkommen explizite Definition eines formales Systems zur Analyse natürlicher Sprachen geprägt. Ein Fragment ist vollständig bzgl. seiner Funktion oder Funktionen, kann aber zunächst nur eine kleine, repräsentative Untermenge der fraglichen natürlichen Sprache behandeln. Nachdem ein Fragment im Grundsatz funktioniert, soll es Schritt für Schritt erweitert werden, bis es schließlich die ganze Sprache abdeckt.

Montagues Ansatz, auch Montague-Grammatik genannt, erforderte die Neuerung einer expliziten Übersetzung analysierter natürlichsprachlicher Ausdrücke in Formeln der intensionalen Logik, die mit den bekannten Methoden der Modell-theorie wahrheitskonditional interpretiert werden können. Als Montagues Übersetzungsalgorithmus im Prinzip funktionierte, wurde er auf die Behandlung neuer Sprachkonstruktionen erweitert, die wiederum Innovationen bzgl. ihrer mengentheoretischen Interpretation erforderten.[4]

Während Montague eine maximal explizite und elegante Bedeutungsbeschreibung mit den Mitteln der Mengentheorie anstrebte, ist unser Ziel eine deklarative Definition, die für eine effiziente prozedurale Implementierung geeignet ist.

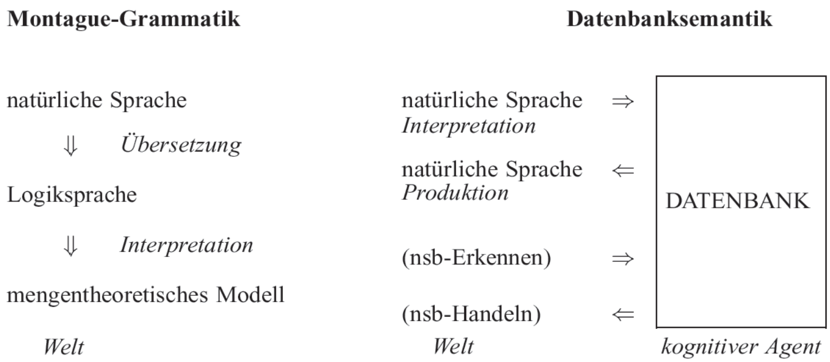

Abb. 11: *Montague-Grammatik und Datenbanksemantik im Vergleich*

Ontologisch behandelt Montague die natürlichsprachlichen Zeichen und das mengentheoretische Modell als Teil der 'Welt', die vergangene, gegenwärtige, zukünftige und mögliche Zustände umfaßt. Seine semantische Methode basiert auf einer *metasprachlichen* Definition der Sprachzeichen, des Modells und der Beziehung zwischen Sprachzeichen und Modell.

Das DBS Fragment behandelt Bedeutungen dagegen ontologisch als kognitive Inhalte. Die semantische Methode basiert auf *kognitiven Prozeduren*, mit denen die Beziehung zwischen der Sprache und den Objekten der Welt etabliert wird. Dabei wird die Welt nicht definiert, sondern als die reale Handlungsumgebung des Agenten vorausgesetzt.

6 Rolle der nichtsprachlichen Kognition

Die schematische Darstellung der Datenbanksemantik in Abb. 11 unterscheidet zwischen der sprachlichen Ebene und der Kontextebene (vgl. auch Abb. 1). Die sprachliche Ebene umfaßt Sprachinterpretation (Erkennen) und Sprachproduktion (Handeln). Die Kontextebene umfaßt nichtsprachbasiertes Erkennen und Handeln.

Entwicklungsgeschichtlich entstand zuerst die kontextuelle Kognition. Die sprachliche Ebene entwickelte sich später als Spezialisierung des nichtsprachbasierten Erkennens und Handelns. für eine computerbasierte Modellierung ist die Kontextebene jedoch schwieriger als die sprachliche Ebene.

Dies liegt daran, daß heutige Standardcomputer weitgehend auf den Sprachkanal beschränkt sind, wobei die Eingabe über die Tastatur und die Ausgabe über den Bildschirm erfolgt. Eine Modellierung der Kontextebene setzt jedoch einen kognitiven Agenten voraus, der seine Umwelt selbstständig erkennen und in ihr Handeln kann. Dies geht über die Möglichkeiten von Standardcomputern hinaus und erfordert stattdessen die Technologie von Robotern.

Wenn das DBS Fragment im folgenden für Standardcomputer definiert wird, so ist dies vor allem ein Zugeständnis an die heutige Technologie. Dabei muß man sich jedoch im Klaren sein, daß dieses Zugeständnis zu schwerwiegenden Einschränkungen bei der Modellierung kognitiver Agenten führt, und zwar sowohl in der Theorie als auch in der Praxis.

Praktisch wird es dem resultierenden Systems unmöglich, natürlichsprachliche Befehle als Handlungen auszuführen oder eigenständige Beobachtungen zu machen, die dann natürlichsprachlich mitgeteilt werden können. Theoretisch entzieht dieses Zugeständnis dem resultierenden System zwei wesentliche Fundamente.

Das eine ist die prozedurale Definition elementarer *Konzepte*, die einerseits zum eigenständigen Aufbau kontextueller Inhalte und andererseits als Sprachbedeutungen dienen (vgl. wiederum Abb. 1). Das andere ist die Grundlegung der kognitiven *Kohärenz* bei kontextuellen Inhalten.

Zum Beispiel kann ein Inhalt wie **Peter überquert die Straße** direkt über die unmittelbare Wahrnehmung oder indirekt über Datenträger, z.B. Sprache, aber auch Film etc., in die Wortbank eingelesen werden. Beim direkten Einlesen über die unmittelbare Wahrnehmung folgt die Kohärenz der Inhalte aus der Kohärenz der externen Welt – also der zeitlichen und räumlichen Abfolge von Ereignissen, den Teil-Ganzes-Beziehungen etc.

Beim indirekten Einlesen besteht dagegen die Möglichkeit, daß ein Autor die aus dem unmittelbaren Erkennen und Handeln vertrauten Elemente auf dem Datenträger neu gemischt und verkettet hat. Dies ist der Grund, warum datenträgerbasierte Inhalte inkohärent sein können.

Beispielsweise ist ein Schwimmer, der zuerst am Beckenrand steht, dann ins Wasser springt und schließlich unter dem aufbrausenden Wasser verschwindet, kohärent. Dagegen ist ein Schwimmer, der mit den Füßen zuerst aus dem aufbrausenden Wasser auftaucht und dann in hohem Bogen am Beckenrand landet, nicht kohärent – es sei denn, man macht die zusätzliche Annahme, daß dieser Inhalt mit Hilfe eines Datenträgers, z.B. einem rückwärtslaufenden Film, dargestellt wird.

Entsprechend sind Menschen, die sich miteinander auf Deutsch oder Englisch unterhalten, kohärent. Eine Biene, die sich mit einem Grashüpfer auf Deutsch unterhält, ist dagegen nicht kohärent – es sei denn, man macht die zusätzliche Annahme, daß dieser Inhalt mit Hilfe eines Datenträgers, z.B. Sprache, dargestellt wird.

Die Verwendung von Standardcomputern hat zur Folge, daß Inhalte nur indirekt über Datenträger, speziell Sprache, in die Wortbank eingelesen werden können. Deshalb liegt die Verantwortung für die Kohärenz der Datenbankinhalte allein bei dem Benutzer, der diese eingelesen hat und nicht beim System. Aus der Kohärenz bzw. Inkohärenz der Datenbankinhalte folgt wiederum die Kohärenz bzw. Inkohärenz der vom DBS Fragment produzierten Sprache. Denn die Konzeptualisierung bei Sprachproduktion beruht auf einer Navigation durch die vorhandenen Datenbankinhalte, die direkt in Sprache abgebildet wird (vgl. Abb. 8).

7 Simulierung der kontextuellen Kognition?

Anstatt auf die kontextuelle Kognition zu verzichten, um die Schwierigkeiten mit der Verwendung von Robotern zu vermeiden, könnte man eine Simulation der kontextuellen Kognition in Erwägung ziehen. Hierfür stehen die Möglichkeiten der kognitiven Modellierung (*cognitive modeling*) zur Verfügung.

Die kognitive Modellierung kreiert Trickfilm-Agenten (*animated agents*) in virtuellen Umgebungen auf Standardcomputern. Der Benutzer kann die Bewegung durch eine künstliche Welt steuern (*automated cinematography*, Funge, Tu, & Terzopoulos 1999) und mit künstlichen Agenten kommunizieren, die mit ihrer virtuellen Handlungsumgebung interagieren (Rickel & Johnson 1999, Loyall & Bates 1993).

Kognitive Modellierung und Datenbanksemantik haben entgegengesetzte Ziele: der eine Ansatz arbeitet an der Interaktion eines natürlichen Agenten (dem Benutzer) mit einer künstlichen Welt, der andere an der Interaktion eines künstlichen Agenten mit der natürlichen Welt. Daher ist für die kognitive Modellierung eine Kontrollstruktur als *message dispatcher* oder *blackboard architecture* wahrscheinlich besser geeignet als die Wortbanknavigation der Datenbanksemantik. Umgekehrt ist die Verwendung virtueller Realität für die Datenbanksemantik aus folgenden Gründen problematisch.

Erstens führt sie zu einer Fehlleitung des Aufwandes: anstatt Kohärenz in einer wirklichen Interaktion mit der externen Welt zu begründen, richtet sich fast die gesamte Programmierarbeit auf eine möglichst realistische Simulierung von 3-D Welten mit Trickfilm-Agenten. Zweitens ist die Interaktion mit künstlichen Welten nicht glaubhaft als Fundierung der Wortsemantik. Drittens trägt die Simulation der sprachbasierten Referenz auf virtuelle kontextuelle Objekte wenig zur Modellierung der Referenz auf wirkliche Objekte bei (vgl. Abb. 1).

Die Unterschiede und Gemeinsamkeiten zwischen Datenbanksemantik und kognitiver Modellierung können wie folgt zusammengefaßt werden. Aus Gründen der Allgemeinheit wird die modelltheoretische Semantik mit in den Vergleich aufgenommen.

Kognitive Modellierung und Datenbanksemantik haben die programmiertechnische Methode gemeinsam. Dagegen basiert die modelltheoretische Semantik auf metasprachlichen Definitionen, die für eine prozedurale Implementierung meist ungeeignet sind.

Modelltheoretische Semantik und kognitive Modellierung haben gemeinsam, daß sie mit Nachbildungen der Welt arbeiten. Datenbanksemantik setzt die externe Welt dagegen voraus und stützt sich auf eine prozedurale Interaktion mit ihr.

Datenbanksemantik und modelltheoretische Semantik haben gemeinsam, daß sie formale Definitionen des Gesamtsystems, sogenannte Fragmente, entwickelt haben. Dagegen ist die kognitive Modellierung möglicherweise zu anwendungsorientiert, um eine deklarative Speziokation – über die Definition von Modellierungssprachen wie VRML oder CML hinaus – zu erlauben.

8 Schematische Darstellung von DBS-FRAGMENT

Als Hintergrund zur schematischen Beschreibung eines datenbanksemantischen Fragments mit der Bezeichnung DBS-FRAGMENT, betrachten wir zunächst ein Fragment der Montague-Grammatik. Das in Montague 1974, Kapitel 8, auch PTQ genannte Fragment besteht aus folgenden Komponenten:

1. Eine rekursive Definition syntaktischer Kategorien und ihre intuitive Interpretation für eine endliche Untermenge.
2. Ein Lexikon, das diese Kategorien verwendet
3. Eine Kategorialgrammatik für die syntaktische Analyse der Sätze des Fragments.
4. Eine intensionale Logik mit einer modelltheoretischen Interpretation.
5. Eine Menge von Übersetzungsregeln, die syntaktisch analysierte Sätze in logische Formeln mit entsprechenden Wahrheitsbedingungen übersetzen.
6. Eine Menge von Bedeutungspostulaten.

Abb. 12: *Komponenten von Montagues PTQ-Fragment*

Dabei wird die Sprachproduktion als Generierungsprozedur der Kategorialgrammatik in Verbindung mit dem Lexikon behandelt. Die Sprachinterpretation hat das Ziel, die Wahrheitsbedingungen der generierten Ausdrücke zu charakterisieren.

DBL-FRAGMENT hat dagegen die Aufgabe, kognitive Inhalte von einem kognitiven Agenten zum anderen zu transportieren. Diese Aufgabe wird von einer Reihe von Motoralgorithmen übernommen – Komponenten, die in der Montaguegrammatik keine Entsprechung haben. Aus praktischen Gründen ist DBL-FRAGMENT für Standardcomputer definiert.

1. DBS-LEXIKON: Eine kleine Anzahl lexikalischer Merkmalstrukturen, die jeweils mit einer Menge zugehöriger Wortoberflächen assoziiert sind, mit denen die Merkmalstrukturen instantiiert werden.
2. WORTBANK: eine verbundbasierte Netzwerkdatenbank, die bidirektional verzeigerte Proplets enthält.

3. LA-INPUT: eine LA-Grammatik für eine natürliche Sprache, deren syntaktische Regeln semantisch interpretiert sind, und zwar in der Form von Kopierregeln, die in die Fortsetzungprädikate der lexikalischen Merkmalstrukturen Werte eintragen.
4. LA-MOTOR: eine LA-Grammatik für die Navigation durch die verketteten Propositionen der Wortbank. Die Traversionszähler der Wortbank und vier Spurprinzipien verhindern Rückfall und Spaltung.
5. LA-OUTPUT: eine LA-Grammatik, die die von LA-MOTOR produzierte Propletsequenz in sprachspeziosche Oberflächen abbildet, unter Berücksichtigung des Navigationstyps.
6. LA-QUERY: eine Menge von LA-Grammatiken, die Frageproplets auf die Tokenzeilen der Wortbank anwenden, um gewünschte Informationen abzurufen.
7. LA-INFERENCE: eine Menge von LA-Grammatiken, die aus dem Inhalt der Wortbank neue Propositionen ableiten.

Abb. 13: *Komponenten von DBS-FRAGMENT*

Die Motoralgorithmen sind die Komponenten 3–7. Sie interagieren wie folgt:

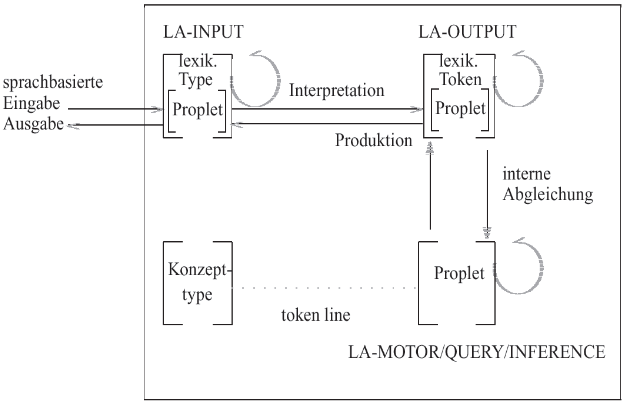

Abb. 14: *Schematische Struktur von DBS-FRAGMENT*

LA-INPUT wird von den hereinkommenden Sprachzeichen angetrieben und verwandelt sie in Propletmengen, die in der Wortbank gespeichert werden. LA-MOTOR, LA-QUERY, und LA-INFERENCE motorisieren gemeinsam die Wortbanknavigation. LA-OUTPUT wird von der autonomen Navigation angetrieben, und reflektiert sie mit natürlichsprachlichen Ausdrücken.

Unter Berücksichtigung der algebraischen Definition der LA-Grammatik (Hausser 1992) und der Standarddefinition klassischer Netzwerkdatenbanken (siehe z.B. Elmasri & Navate 1989) ist DBL-FRAGMENT formal ebenso stringent wie PTQ. Auf die explizite Darstellung der LA-Grammatiken von DBS-FRAGMENT muß hier allerdings verzichtet werden, da sie über den Rahmen dieses Aufsatzes hinausgehen würde. Sie findet sich in Hausser 1999 und 2001b.

9 Schlußwort

Ein Fragment der Datenbanksemantik besteht aus einem Lexikon, einer semantischen Repräsentation und formalen Grammatiken. Die Grammatiken haben die Namen LA-INPUT, LA-OUTPUT, LA-MOTOR, LA-QUERY und LA-INFERENCE, und sind für die Interpretation und Produktion natürlicher Sprache sowie für die Konzeptualisierung, Befragung und Inferenz, respektive, definiert.

Die Integration dieser verschiedenen Funktionen basiert auf der Navigation durch verkettete Propositionen in einer Datenbank (siehe 4.1 und 4.2). Die Navigation wird mit Hilfe des zeitlinearen Algorithmus der LA-Grammatik realisiert. Das 'Schienensystem' der Navigation liefert die Datenstruktur der Wortbank, die verkettete Propositionen mit Hilfe von distribuierten Merkmalstrukturen darstellt. Sie drücken intrapropositionale Funktor-Argument-Strukturen und extrapropositionale Relationen mittels bidirektional kodierter Fortsetzungsprädikate aus.

Als Zugeständnis an die heutige Technologie wurde das formale Fragment für Standardcomputer definiert, und nicht für Roboter. Die generelle Theorie der Datenbanksemantik basiert jedoch entscheidend auf kognitiven Agenten, die in der Lage sind, auf der nichtsprachbasierten Ebene wahrzunehmen und zu handeln.

Aus diesem allgemeinen Ansatz folgt, daß das Fragment strukturell für eine Übertragung von Standardcomputern auf Roboter geeignet ist. Wenn die zusätzlichen Möglichkeiten einer prozeduralen Fundierung der Wortsemantik, einer Fundierung der Kohärenz in der kontextuellen Kognition, und einer Umsetzung von Sprache in externe Handlungen und von externen Wahrnehmungen in Sprache einmal zur Verfügung stehen, können sie in das gegenwärtige Fragment integriert werden.

Zudem steht die Verwendung von Standardcomputern dem grundsätzlichen Funktionieren des Kommunikationsmodells nicht im Wege (mittelbare Referenz). Sie verhindert auch nicht die Möglichkeit, das bestehende Fragment zu erweitern. Eine Erweiterung sollte mit der Definition eines größeren Lexikons und einer Syntax beginnen, die zusätzliche Konstruktionen analysieren und interpretieren kann.[5] Um das integrierte Funktionieren des Gesamtsystems zu gewährleisten, muß eine solche Erweiterungen konsequent durch entsprechende Erweiterungen aller anderen Komponenten ergänzt werden

Als programmiertechnisches System ermöglicht die Datenbanksemantik eine leistungsfähige und bequeme Verifikation an großen Mengen realer Daten. Dabei müssen nicht nur die empirische Abdeckung von Lexikon und Syntax verifiziert werden, sondern auch die semantische Korrektheit der Interpretation, die Gültigkeit der Inferenzen, die Kohärenz der Konzeptualisierung und die rhetorische Korrektheit der Produktion.

Systeme der Datenbanksemantik mit ausreichender Abdeckung haben eine Fülle praktischer Anwendungen. Sie ermöglichen neue Methoden z.B. bei natürlichsprachlichen Schnittstellen mit klassischen Datenbanken, bei Indizierung und Abruf in textuellen Datenbanken, beim Modellieren der Informationsübertragung in Dialogen, bei der Interlingua-basierten maschinellen Übersetzung und bei der Unterstützung der Spracherkennung mit Grammatik- und Domänenwissen.

Bibliographie

Elmasri, R. & S.B. Navathe (1989) *Fundamentals of Datebase Systems*, Benjamin-Cummings Publishing Company, Inc., Redwood City, CA.

Funge, J., X. Tu, & D. Terzopoulos (1999) „Cognitive Modeling: Knowledge, Reasoning and Planning for Intelligent Characters," *SIGGRAPH 99*, Los Angeles, CA.

Hausser, R. (1992) „Complexity in Left-Associative Grammar," *Theoretical Computer Science*, Vol. 103, Elsevier.

Hausser, R. (1996) „A Database Interpretation of Natural Language," *Korean Journal of Linguistics*, Vol. 21, No. 1,2:29–55.

Hausser, R. (1999) *Foundations of Computational Linguistics*, Springer-Verlag, Berlin–New York.Hausser, R. (2001a) „The Four Basic Ontologies of Semantic Interpretation", in H. Kangassalo et al. (eds.) *Information Modeling and Knowledge Bases* XII, IOS Press Ohmsha, Amsterdam.

Hausser, R. (2001a) „The Four Basic Ontologies of Semantic Interpretation“, in H. Kangassalo et al. (eds.) *Information Modeling and Knowledge Bases* XII, IOS Press Ohmsha, Amsterdam.

Hausser, R. (2001b) „Database Semantics for Natural Language“, *Artificial Intelligence,* Vol. 129, No. 1.

Hausser, R. (2001c) „Spatio-Temporal Indexing in Database Semantics“, in A. Gelbukh (ed.) *Computational Linguistics and Intelligent Text Processing,* Lecture Notes in Computer Science Vol. 2004. Springer-Verlag, Berlin, New York.

Lappin, S. (ed.) (1996) *The Handbook of Contemporary Semantic Theory*, Blackwell Publis hers, Oxford, England.

Loyall, A. B. & J. Bates (1993) „Real-time Control of Animated Broad Agents“, *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Boulder, Colorado.

Montague, R. (1974) *Formal Philosophy*, edited by R. Thomason, Yale University Press, New Haven.

Nasr, A., O. Rambow, M. Palmer, & J. Rosenzweig (1997) „Enriching Lexical Transfer with Cross-linguistic Semantic Features“, *Proceedings of the Interlingua Workshop at the MT Summit*, San Diego, California, October.

Rickel, J., & Johnson, W.L. (1999) „Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control“, *Applied Artificial Intelligence* 13:343–382.

ENDNOTEN

[1] Die prozedurale Behandlung der Konzepte ist eine grundsätzliche Alternative zur logischen Semantik, die Bedeutungen metasprachlich über Wahrheitsbedingungen zu beschreiben sucht. Dieser Unterschied wird in Hausser 2001a ausführlich thematisiert.

[2] Ein möglicher Vorläufer ist die Verwendung der bidirektionalen Verzeigerungsmethode zur Beschleunigung beim Cashing (persönliche Mitteilung von Jaime Carbonell, CMU, April 2000).

[3] Der Begriff *Proplet* ist analog zu *droplet* gebildet und bezeichnet die Grundbausteine elementarer Proposition. Oberflächlich mögen Proplets den Merkmalstrukturen der HPSG ähneln. Letztere sind jedoch als Teil von Phrasenstrukturbäumen konzipiert, und nicht als Teil einer Datenbank; zudem fehlen die Kodierung der Funktor-Argumentstruktur über *bidirektionale* Zeiger, eine Verkettung von Propositionen über extrapropositionale Relationen, und somit die Basis für eine zeitlineare Navigation.

[4] Viele dieser Extensionen sind in S. Lappin (ed.) 1996 überblicksartig zusammengestellt.

[5] Eine LA-Grammatik, die einen maßgeblichen Umfang des Englischen abdeckt, findet sich in Hausser 1999, p. 458 f. Eine entsprechende LA-Grammatik für das Deutsche wird auf S. 364 f. definiert.

The instructible agent Lokutor

Jan-Torsten Milde,
Fakultät für Linguistik und Literaturwissenschaft,
Computerlinguistik und Texttechnologie,
Universität Bielefeld,
Universitätsstr. 25, 33501 Bielefeld,
email:milde@coli.uni-bielefeld.de

Abstract

In this paper we describe Lokutor, a virtual human. Lokutor is a partially autonomous agent, inhabiting a 3D virtual environment. The agent can be controlled via natural language directives of the user and by direct manipulation using a graphical user interface. Its control system consists of a behavior-based reactive layer, a deliberative control layer and a motivational subsystem. Lokutor's domain knowledge, its behavior modules and its communicative behavior is encoded in an integrated multimodal format, which has been implemented using XML. Currently Lokutor is used to present information about a car.

1 Introduction

In natural environments humans rely on their ability to communicate using language. Natural language encodes information in a very compact way. Nevertheless in most cases processing language is seen as an isolated cognitive capability of an artificial system. Language understanding is restricted to the mapping of natural language expressions into an internal semantic representation, whereas language production takes an explicit semantic representation as input, from which a natural language utterance is generated. This point of view is inadequate while interacting with an situated agent living in an environment (virtual or natural).

The design of Lokutor reflects the definition of Franklin and Graesser, who define an autonomous agent to be: "*... a system situated within and as part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future...*" (Franklin & Graesser, 1996).

Lokutor is an intelligent agent living in a simulated 3D environment. It has been designed to be as independent of the application domain as possible. To achieve this goal standard technologies have been used to implement the agent:

Figure 1: *Lokutor and the Opel Calibra. The agent is able to walk through the environment, pointing and looking at details in the scene, opening and closing doors, trunk and bonnet of the Calibra. It uses the XML-encoded background knowledge to explain part of the functionality of the car.*

1. Java/Java3D for the implementation of the distributed simulation engine[1]
2. VRML 2 as the basis for the geometric object description[2]
3. H-Anim as a basis for the geometric joint description[3]
4. XML as a basis for the representation of the domain knowledge[4] (Holzner, 1998)

Lokutors current task is to present a car (an Opel Calibra) to a human user. The agent is able to convey information about the functionality of the car (e.g. how to open a door, how to open the tank lit, what type of fuel the car takes, the size of the space in the trunk etc). The Opel Calibra model has been converted from CAD data of a desgin model which has been kindly provided (just as its user manual) by Opel Germany. The user is able to interact with the agent by natural language directives.

Lokutor will follow these instructions, while integrating situative non-linguistic information in the analysis and interpretation of the directives (see Milde & Ahlers, 1999). Information given by the agent will be presented to the user in synthetic spoken natural language. Currently Lokutor is able to "speak" in English and German. The content of the information is retrieved from an XML annotated version of the cars user manual.

2 Multiagent simulation system

Our first approach to implement a multi agent system was based on a combination of VRML and Java technology (Jung & Milde, 1999). Unfortunately, the VRML technology integrated into the standard WWW browsers has shown to be very unstable and slow. It has been difficult to control the visualization process. Either control was accurate, but slow, or the agents would move smoothly but almost uncontrollably. Still the basic principle of separating visualization, world representation and agent control into three different processes seemed to be promising. We therefore switched from VRML to Java3D, which now is sufficiently matured to implement a stable and fast system.

As Java3D is able to import geometric objects described in VRML, the reimplementation of the visualization was relatively easy. Lokutors geometric appearance is based on a standard HANIM 1.0 figure. The skeleton structure of HANIM 1.0 allows definition flexible articulated agents. The arms, legs, the upper body and the head can all be moved seperately, making it possible to define realistic animations of the agents.

The world server encodes most of the state information of the simulated virtual world. Position and orientation of the agents are represented. The same holds true for a number of parts of the car. Virtual objects, like sensors, are also represented here. Based on this information the world server is able to calculate the current perceptual status of the agent. Lokutor is equipped with visual and tactile sensors. It is also possible to request the current sensomotoric status of the agent, mainly the joint angles of the limbs and wether the hand is open, closed or in a pointing position. Any number of control clients can connect to the server. This allows to control a number of interacting agents in the simulated environment. For the presentation task we only implemented one control client.

3 A hybrid control architecture

The control architecture (see figure 2) contains a deliberative system, which models „higher'' cognitive competences, and a behavior-oriented[5] base system, which integrates language, perception and action on a lower level. This hybrid architecture (Milde, 1997) allows the distribution of the necessary competence and therefore the tasks of the whole system onto both subsystems. The deliberative system is responsible for the sequentialization of complex actions into simple basic actions and schedules the execution of these actions. However, because of the complexity and uncertainty of the real world, a fully detailed plan of the agents movements cannot be provided by the deliberative system. The behavior system on the other hand is embedded in the real world by means of sensors and actuators, which enable it to detect changes in the world and react to them immediately. It has the necessary competence for the autonomous execution of basic actions, but cannot aim at fulfilling given goals by itself. It is only through the coupling of the behavior system with a deliberative system, that the accomplishment of tasks can be realized by the interplay of goal-directed and reactive behavior.

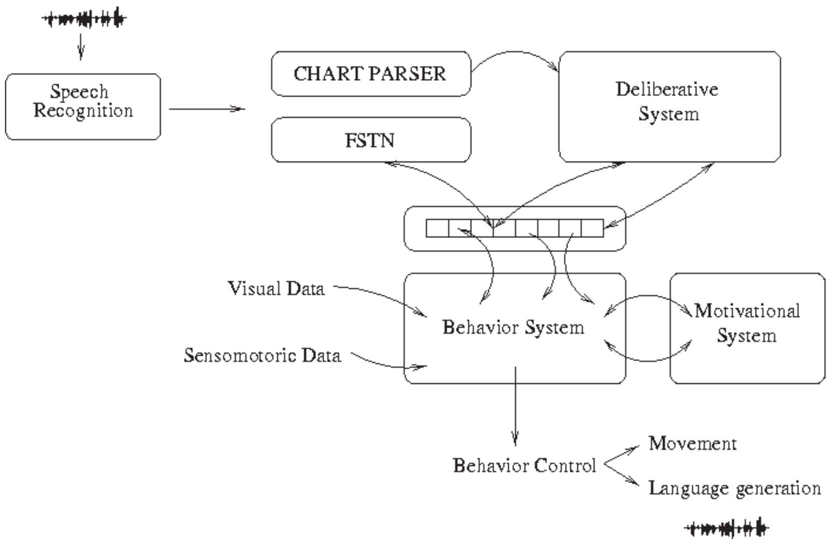

Figure 2: *The control architecture of Lokutor. The behavior system selects the next action of Lokutor based on sensoric date and linguistic information. The motivational system evaluates the internal state of the behavior system. The deliberative system controls longer term goals of the agent.*

Instructions cannot be processed by the base system directly. They provide resources for planning goals or action sequences and guide the planning process. First, instructions are parsed by the chart parser that builds up typed attribute-value pairs. The semantic part of those structures – based on the work of

Jackendoff (Jackendoff, 1990) – is passed on to the deliberative system, which is responsible for keeping track of long-time goals. The deliberative system uses this semantic part to initialize a corresponding action scheme (Lobin, 1999).

Action schemes contain explicit knowledge about the decomposition of higher level actions into basic actions. The resulting information blocks of an action scheme are mapped one after the other onto so-called internal sensors, which provide the basis for the communication between deliberative system and behavior system. Thus the necessary sequentialization of control parameters to initialize the corresponding basic actions in the behavior system can be produced. Suitable feedback from the base system allows the deliberative component to monitor the activity state of the behavior system and feed in the control parameters that are needed for the completion of the next subtask just in time.

The behavior system is partly autonomous. It can carry out basic actions and react to unexpected events without the help of an external control, but it is still controllable by a „higher“ system: The selection of basic actions is initialized by the deliberative system – as explained above – or through user interventions.

The behavior system consists of a hierarchy of behavior modules, each of it specialized for a certain task, which it can fulfil autonomously. In contrast to traditional, knowledge-based agent control, action related knowledge is not represented in a model of the world, but is distributed among all behavior modules. This leads to a reduction of complexity, because an error does not cause global replanning or the intervention of higher system components, but only local reactions inside the relevant behavior module. The modularization of the behavior system is motivated by the requirements concerning reactivity and autonomy of the behavior system on the one hand and by the expected user directives on the other hand: All possible linguistic sub-directives must be depictable to a – as small as possible – set of corresponding behavior modules.

Interventions refer to the direction and velocity of the agents movements and actions or to simple object-oriented actions. They are fed into the behavior system directly, thus allowing the immediate manipulation of ongoing behavior. The behavior system is responsible for the situated and time-adequate translation of sensor input into actuator output, treating information from the internal sensors just like any other sensor data. The integration of the different sensor inputs allows the situated interpretation of directives. As a consequence the processing of elliptical utterances like situation-dependent object references and indexical expressions, which can only be comprehended in the current context of sensing and acting, is made possible.

Language generation is based on textual information of the user manual. A text-to-phoneme system[6] transforms the textual information to a pre-speech level. The actual sound is then generated by the MBrola speech synthesizer[7] In our approach language instructions are treated as preprocessed sensoric input to the behavior system. Accordingly language generation will be issued by the behavior system as part of the action selection process. Suppose Lokutor is instructed as follows:

a) *Öffne den Tankdeckel!* `(Open the tank lid!)`

This instruction, once it has been processed by the language system, will be decomposed by the deliberative system into a moving towards the tank lid, a pointing at the object you are moving to and then to issue some explanation about the object you are pointing at action. Based on the current sensomotoric state, the explanation action filters the background knowlegde:

b) *Der Kraftstoffeinfüllstutzen mit Renkverschluß befindet sich an der rechten Wagenseite hinten. Tankdeckel aufschließen: Schlüssel einstecken und nach links drehen, den Deckel nach rechts ausrasten.* `(The tank lit is at right rear side of the car. Insert the key, turn left and open the tank lid to the right.)`

Most of the background knowledge is taken from the printed user manual, which is delivered to the customer with the car. It has been scanned and processed by OCR software to produce an electronic version.

An XML DTD has been developed which allows to annotate the text in a structured way. Also a number of attributes are defined in the DTD, which allows for the integratin of metaknowledge into this textual database:

```
<!ELEMENT agent (geom,scene,knowb)>
<!ELEMENT geom (#PCDATA)>
<!ELEMENT scene (#PCDATA)>
<!ELEMENT knowb (entry+)>
<!ELEMENT entry (name, desc)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT desc (sit, action)>
<!ELEMENT sit (cond*)>
<!ELEMENT cond (att,val)>
```

```
<!ELEMENT att (#PCDATA)>
<!ELEMENT val (#PCDATA)>
<!ELEMENT action (move*, say*)>
<!ELEMENT move (basic*)>
<!ELEMENT basic (att,val)>
<!ELEMENT say (utt*)>
<!ELEMENT utt (#PCDATA)>
```

4 Outlook: Talking to each other

Currently the work on Lokutor has reached the level of a stable prototype implementation. It is possible to quickly set up the simulation system, define the basic animations of the agents including their sensoric abilities and connect the control system to the simulation. Lokutor runs a number of platforms without any difficulties, performing best under Win9X with a fast graphics adapter.

The hybrid control architecture has shown to produce believable behavior for the agent. Lokutor can be instructed, is able to follow longer term goals, while still being reactive. The agent is able to process natural language directives on different levels of complexity.

The next step is to populate the environment with a number of different Lokutors. Here Lokutor can be used as a means to do research on communicative agents. From a computational linguistic research standpoint this could lead to a better understanding of an extended speech act theory, which allows to model complex human communication. The experimental scenario will consists of two Lokutors standing face to face near a table on which a number of Baufix parts are lying. A set three experiments will be performed:

1) **Sorting game**: The parts are scattered over the table. Each Lokutor is selecting a part on the table and tries to grasp it and move it to its side of the table. Picking the parts has to be coordinated, such that each Lokutor is taking a piece at a time. If conflicts occur, these have to be resolved using natural language. The game ends, if all the parts are distributed amongst the agents.
2) **Maximizing game**: Each Lokutor starts sorting its pieces. First it selects a sensoric feature (color, size or type), then it will try to move every appropriate part to one side. Once all the pieces have been sorted, it will start asking the other Lokutor for more pieces. This is done as long, as either no more pieces will be exchanged, or the Lokutor is satisfied with the result, e.g. has reached

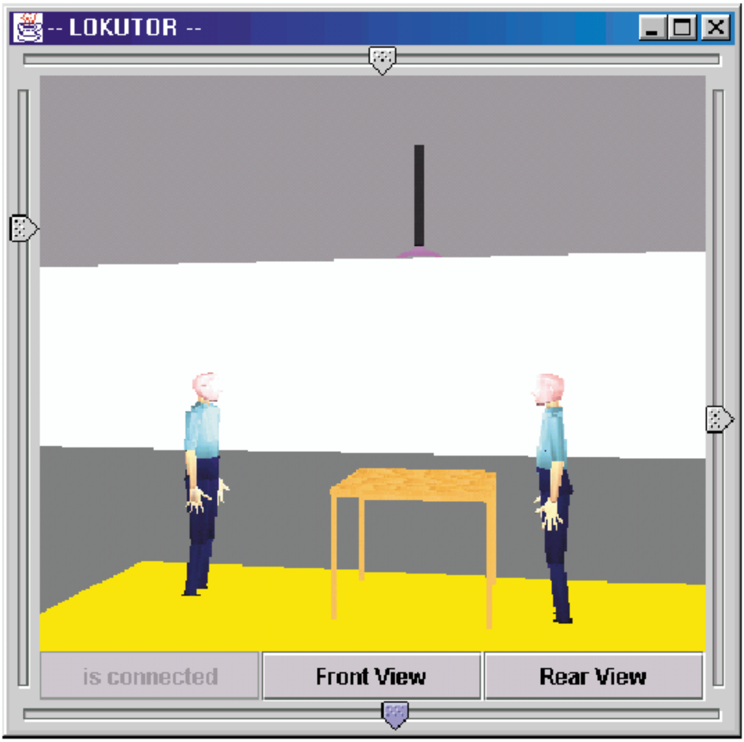

Figure 3: *Two Lokutors interacting.*

a certain threshold level in its emotional system. Otherwise the experiment starts all over with a different sensoric feature.

3) **Imitation game**: In the center of the table a bridge consisting of two blocks and a bar is placed. The Lokutors are trying to build a parallel bridge. First, each Loktur tries to identify one of the supporting blocks, then tries to find a similar block on the table and place it next to the supporting block. These actions have to be coordinated as each Loutor is choosing the position indepently. Once the placement of the objects is correct, the bar has to be moved on top of the blocks. One Lokutor chooses a bar and picks it up. The second Lokutor is instructing the carrying Lokutor as where to move the bar. Once the end position of the bar is correct for one of the blocks, the bar is

handed over to the other Lokutor and the roles change: now the instructor become the carrier of the bar and vice versa.

In future work we will try to refine the generic declarative representation format of the underlying domain knowledge to allow more complex discourse structures when communicating with the agent. We think that the presented XML-based approach is a step into the right direction when building communicative agents for intelligent virtual environments.

References

R. A. Brooks, Intelligence Without Representation, in *Artificial Intelligence*, volume 47, pp. 139–159, 1991.

S. Holzner, *XML complete*, McGraw-Hill, New York, 1998.

R. Jackendoff, *Semantic Structures*, Current studies in linguistics series, 18, MIT Press, Cambridge, MA, 1990.

S. S. Franklin & A. Graesser, Is It an Agent, or Just a Program: A Taxonomy for Autonomous Agents. In J. P. Müller, M. J. Woolridge, and N. R. Jennings, editors, *Intelligent Agents III. Agent Theories, Architectures and Languages*, 21–35. Springer, Berlin, 1996.

B. Jung & J.-T. Milde. An open virtual environment for autonomous agents using VRML and Java. In *Procedings VRML'99*, 1999.

H. Lobin, *Handlungsanweisungen. Sprachliche Spezifikation teilautonomer Aktivität*, Deutscher Universitäts Verlag, 1999.

J.-T. Milde & Tobias Ahlers, Mensch-Maschine-Kommunikation in IVUs: Der kommunikative Agent Lokutor, in *Procedings of the first German workshop on Intelligent Virtual Environments*, KI 1999, Bonn, 1999.

J.-T. Milde, Kornelia Peters, and Simone Strippgen, Situated communication with robots, in *Proceedings of the first international Workshop on Human–Computer Conversation*, Bellagio, Italy, 1997.

K. Peters, Natürlichsprachliche Steuerung eines behaviorbasierten Roboters, Report 94/8, Situierte Künstliche Kommunikatoren, SFB 360, Universität Bielefeld, 1994.

L. Steels, The artificial life roots of artificial intelligence, in *Artificial Life Journal*, volume 1, MIT Press, Cambridge, 1994.

NOTES

1 http://java.sun.com/products/java-media/3D/index.html

2 http://www.vrml.org

3 http://ece.uwaterloo.ca:80/~h-anim/

4 http://www.w3.org/TR/REC-xml.html

5 A detailed description of the characteristics of behavior-oriented architectures can be found in Brooks, 1999 and Steels, 1994.

6 wxTTS, http://web.ukonline.co.uk/julian.smart/

7 MBrola, http://tcts.fpms.ac.be/synthesis

ParParsing Preferences and Linguistic Strategies

Rodolfo Delmonte
Ca' Garzoni-Moro, San Marco 3417
Università "Ca Foscari"
30124 - VENEZIA
Tel. 39-41-2578464/52/19 - Fax. 39-41-5287683
E-mail: delmont@unive.it - Website: byron.cgm.unive.it

Abstract

We implemented in our parser four parsing strategies that obey LFG grammaticality conditions and follow the hypothesis that knowledge of language is used in a "modular" fashion. The parsing strategies are the following: Minimal Attachment (MA), Functional Preference (FP), Semantic Evaluation (SE), Referential Individuation (RI). From the way in which we experimented with them in our implementation it appears that they are strongly interwoven. In particular, MA is dependent upon FP to satisfy argument/function interpretation principles; with semantically biased sentences, MA, FP and SE apply in hierarchical order to license a phrase as argument or adjunct. RI is required and activated every time a singular definite NP has to be computed and is dependent upon the presence of a discourse model. The parser shows garden path effects and concurrently produces a processing breakdown which is linguistically motivated. Our parser is a DCG (Pereira & Warren, 1980) is implemented in Prolog and obeys a topdown depth-first deterministic parsing policy.

1 Introduction

In order for a parser to achieve psychological reality it should satisfy three different types of requirements: psycholinguistic plausibility, computational efficiency in implementation, coverage of grammatical principles and constraints. Principles underlying the parser architecture should not conform exclusively to one or the other area, disregarding issues which might explain the behaviour of the human processor. In accordance with this criterion, we assume that the implementation should closely mimick phenomena such as garden path effects, or an increase in computational time in presence of semantically vs syntactically biased ambiguous structures. We also assume that a failure should ensue from strong garden path effects and that this should be justified at a psycholinguistic interpretation level.

Since we base most of our grammatical principles on LFG we assume that lexical information is the most important knowledge source in the processing of natural language. However, we also assume that all semantic information should be made to bear on the processing and this is only partially coincident with lexical information as stored in lexical forms. In particular, subcategorization, semantic roles and all other semantic compatibility evaluating mechanisms should be active while parsing each word of the input string. In addition, the discourse model and external knowledge of the world should be tapped when needed to disambiguate ambiguous antecedents.

Differently from what is asserted by global or full paths approaches (see Schubert, 1985; Bear & Hobbs, 1988; Hobbs, Stickel, Appelt, Martin, 1993), we believe that decisions on structural ambiguity should be reached as soon as possible rather than deferred to a later level of representation. In particular, Schubert assumes "...a full paths approach in which not only complete phrases but also all incomplete phrases are fully integrated into (overlaid) parse trees dominating all of the text seen so far. Thus features and partial logical translations can be propagated and checked for consistency as early as possible, and alternatives chosen or discarded on the basis of all of the available information (ibid., 249)." And further on in the same paper, he proposes a system of numerical 'potentials' as a way of implementing preference trade-offs. "These potentials (or levels of activation) are assigned to nodes as a function of their syntactic/semantic/pragmatic structure and the preferred structures are those which lead to a globally high potential. The total potential of a node consists of a) a negative rule potential, b) a positive semantic potential, c) positive expectation potentials contributed by all daughters following the head (where these decay with distance from the head lexeme), and d) transmitted potentials passed on from the daughters to the mother (ibid., 249)." Whereas Hobbs et al. theorize in favour of a delayed ambiguity resolution in LF, they suggest heuristic strategies which can account for most of them, "...a very good heuristic is obtained by using the following three principles: 1) favor right association; 2) override right association if a. the PP is temporal and the second nearest attachment site is a verb or event nominalization, or b. if the preposition typically signals an argument of the second nearest attchment site (verb or relational noun) and not of the nearest attachment site (ibid., 241)".

Other important approaches are represented by Hindle et al., 1993, who attempt to solve the problem of attachment ambiguity in statistical terms. The important contribution they made, which was not possible in the '80s, is constituted by the data on attachment typologies derived from syntactically annotated corpora. We will report similar data on Italian, derived from our corpus of Italian, currently being annotated as a syntactic treebank, a portion of which will be used as sample with comparable length to the English one: the authors worked on a test sample made up of 1000 sentences in which their parser identified ambiguous attachment conditions. In order to simulate automatic disambiguating procedures we shall use data derived from our shallow parser, which has an output comparable to the Fidditch parser quoted by Hindle et al. and shown in the article. The final output as recorded in the Treebank is the result of the concurrent work done by manual annotators, automatic validation procedures and my final supervision of the resulting constituent structures, visualized by a tree-viewer.

1.1 The parser

The parser we work with is organized as shown in Fig.1 and can deal with a certain number of linguistic phenomena at sentence level, while leaving other problems to be solved at discourse level.

The parser we present was conceived in the middle '80s and started as a transfer module for a Machine Translation Expert system in a very restricted linguistic domain. Then it became a general parser for Italian and English, to be used with LFG students. German was added later on, in the beginning of '90s. Since the people working at it were interested in the semantics as much as in the syntax, it was soon enriched with a Quantifier Raising algorithm and an Anaphoric Binding Module. In 1994 the Discourse Model and the Inferential Processes algorithms were developed. Finally in 1996 work on a Situational Semantics interface and on the Discourse Structure was carried out. These experiments were finally enriched – two years ago – with a number of Parsing Strategies procedures like setting up a Lookahead mechanism, a Well-Formed Substring Table and a number of other semantically and/or lexically based triggering lookup procedures.

We worked from the very beginning within LFG framework, which already from the start allowed us to think in terms of a much richer representation, closer to semantics, than just a context-free syntactic constituency. In particular, all levels of control mechanisms which allow coindexing at different levels of parsing gave us a powerful insight into the way in which the parser should be organized.

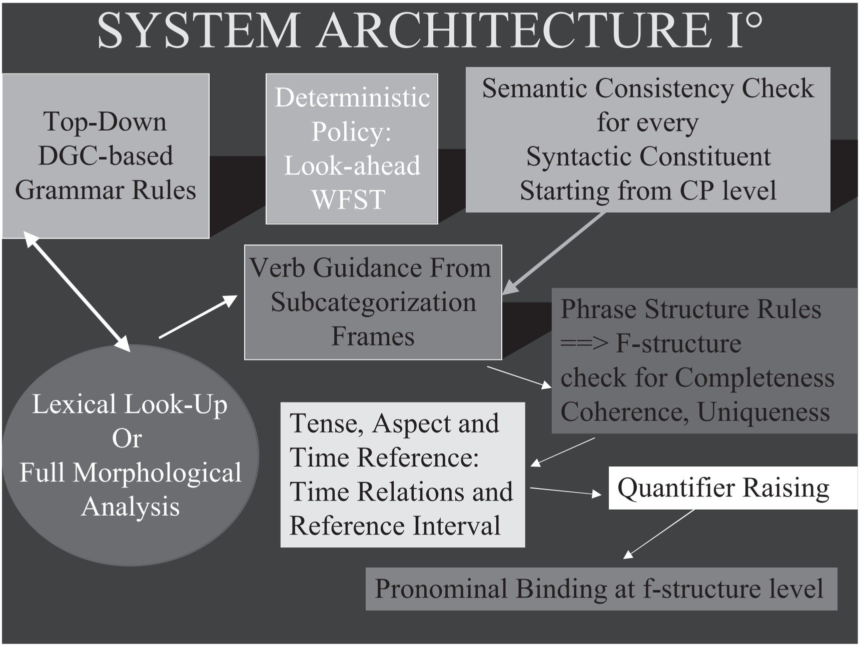

Figure 1: *Getaruns parser*

Yet the grammar formalism implemented in our system differs from the one suggested by the theory, in the sense that we do not use a specific Feature-Based Unification algorithm but a DCG-based parsing scheme. In order to follow LFG theory more closely, unification should have been implemented: but DCG gives us full control of a declarative rule-based system, where information is clearly spelled out and passed on and out to higher/lower levels of computation. The grammar is implemented in Prolog using XGs (extraposition grammars) introduced by Pereira (1981;1983). Prolog provides naturally for backtracking when allowed, i.e. no cut is present to prevent it. Furthermore, the instantiation of variables is a simple way for implementing the mechanism for feature percolation and/or for the creation of chains by means of index inheritance between a controller and a controllee, and in more complex cases, for instance in case of constituent ellipsis or deletion.

Apart from that, the grammar implemented is a surface grammar of the languages chosen. Also functional control mechanisms – both structural and lexical – have been implemented as close as possible to the original formulation, i.e. by binding an empty operator in the subject position of a propositional like open complement/predicative function, whose predicate is constituted by the lexical head.

Of course there are a number of marked differences in the treatment of specific issues, concerning Romance languages, which were not sufficiently documented in the linguistic literature at the time. In particular,

- we introduced an empty subject pronominal - little pro - for tensed propositions, which had different referential properties from big PRO; this had an adverse effect on the way in which c-structure should be organized. We soon realized that it was much more efficient and effective to have a single declarative utterance-clause level where the subject constituent could be either morphologically expressed or Morphologically Unexpressed. In turn MUS or little pros could be computed as variables in case the subject was realized in postverbal position. At the time LFG posited the existence of a rule for sentence structure which could be rewritten as VP in case there was no subject, MUS, or in case the subject was expressed in postverbal position, an approach that we did not implement;

- we also use functional constituents like CP and IP: CP typically contains Aux-to-Comp and other preposed constituents, adjuncts and others; IP contains negation, clitics, and tensed verbal forms, simple and complex, and expands VPs as complements and postverbal adjuncts;

- each constituent is semantically checked for consistency before continuing parsing; we also check for Uniqueness automatically by variable instantiation. But sometimes, in particular for subject-verb agreement we have to suspend this process to check for the presence of a postverbal NP constituent which might be the subject in place of the one already parsed in preverbal position!!;

- syntactic constituency is replicated by functional constituency: subject and object are computed as constituents of the annotated c-structure, which rewrite NP - the same for ncomp - this is essential to assign the appropriate annotated grammatical function; this does not apply to VP, a typical LFG functional non-substantial constituent;

- our lexical forms diverge from the ones used in the theoretical framework: we introduced aspectual categories, semantic categories and selectional restrictions in the main lexical entry itself;

- we also have semantic roles already specified in the lexical form and visible at the level of syntactic-semantic parsing;

- rather than generating a c-structure representation to be mapped onto the f-structure via an annotated c-structure intermediate level, we already generated a fully annotated c-structure representation which was then checked for Grammatical Principles Consistency at the level of number/type of arguments and of Adequacy for adjuncts, with a second pass on the output of the parser, on the basis of lexical form of each predicate and semantic consistency crossed checks for adjuncts.

All parser rules from lexicon to c-structure to f-structure amount to 1900 rules, thus subdivided:

1. Calls to lexical entries - morphology and lexical forms = 150 rules
2. Syntactic and semantic rules in the parser proper = 550 rules
3. Parsing strageties and other tools = 185 rules

    All syntactic/semantic rules = 850 rules

4. Semantic Rules for F-Structure

    Lexical Rules for Consistency and Control: - semantic rules 439

    F-structure building, F-command: - semantic rules 170

    Quantifier Raising and Anaphoric Control: - semantic rules 441

    All semantic f-structure building rules: 1050

2 Theories for Parsing Strategies

Parsing theories are of two kinds: the first garden-path theory (hence GPT) is syntactically biased, and the second incremental-interactive theory (hence IIT) is semantically biased. There is a crucial difference between the two theories: whereas GPT claims that a single analysis of a syntactic structure ambiguity is initially constructed and passed to the semantic module, IIT claims that the syntactic module offers all grammatical alternatives to the semantic one, in parallel, to be evaluated. There is evidence that is favorable to both sides on this issue (see G.Altman, 1989).

The basic claims of GPT are that the sentence processing mechanism (the parser) uses a portion of its grammatical knowledge, isolated from world knowledge and other information, in initially identifying the relationships among the phrases of a sentence. IIT permits a far more intimate and elaborate interaction between the syntax and the semantic/referential modules. Altmann (1989) offers a functional argument against a system in which choices are initially made by a syntactic processor, and later corrected by appeal to meaning and context. He says that if referential or discourse information is available, only a strange processor would make decisions without appealing to it. It is also our opinion that all lower level constraints should work concurrently with higher level ones: in our parser all strategies are nested one inside another, where MA occupies the most deeply nested level.

The list of examples here below includes sentences used in the literature to support one or the other of the two parsing theories, GPT and IIT (see Altman (ed), 1989).

1. Mary put the book on the table.
2. Mary put the book on the table in her bag.
3. Mary saw the cop with the binoculars.
4. Mary saw the cop with the revolver.
5. The thief stole the painting in the museum.
6. The thief stole the painting in the night.
7. John saw Mary in the kitchen.
8. John saw Mary from the bathroom.

Altmann and Steedman (1989) note that several of the ambiguities present in these sentences and resolved by Minimal Attachment involve contrasts between NP modification and other structures. However, examples 2, 3 and 5 require linguistic knowledge.

In example 1 we assume that the pp should be computed as an argument of the main predicate "put", thus following a MA strategy when parsing the np "the book". However, the same strategy would lead to a complete failure in example 2, where the pp should be taken as np modifier. If we look at subcategorization requirements of the verb, example 1 constitutes a clear case of Verb Guidance and of semantic role satisfaction requirements: the main predicate requires an argument which is a locative, so at every decision point in which a pp might be taken, argument requirements should be accessed and a MA strategy imposed locally by FP.

Example 3 contains an instrumental adjunct: when the head preposition “with“ is met, the parser will not close the np “the cop“ and will continue building an internal pp modifier since preposition “with“ heads a compatible np modifier, a comitative. In our dictionary the verb “see“ has one single lexical entry but a list containing two different lexical forms. The first form in the list has a higher number of arguments,

I. see <SUB/perceiv, OBJ/theme, PCOMP/locat>

where the Pcomp predicates a location of the Object. In the second form the Pcomp is absent. In order for a Location pp to be accepted, the head preposition should be adequate, and “with“ does not count as such. In examples 3 and 4, the first decision must be taken when computing np structure. In fact, a pp headed by preposition “with“ is a semantically compatible np modifier – a comitative – and the analysis should be allowed to continue until the pp is fully analysed. In other words SE should verify pp attachment consistency inside the NP constituent. However, this may only happen in case MA is deactivated by FP, after matching with lexical information has failed.

In the following examples (5, 6), argument structure plays no role whatsoever: instrumentals, comitatives, locatives with predicate “steal“ are all cases of sentential adjuncts, and only SE can apply. As a matter of fact, “in the museum“ might be freely attached lower at NP level as well as higher at Sentence level. This is due to the fact that head preposition “in“ constitutes a viable local NP modifier and there are no argument requirements from the main verb predicate. However, “in the night“ is not a possible NP modifier and example 6 is a clear case of minimal attachment sentence. On the contrary, in example 5 PP attachment is ambiguous between a np internal modifier and VP level attachment.

In example 7 we understand that the location at which Mary was when the seeing event took place is the kitchen: we also understand that John might have been in the same location or in a different one, already provided by the previous context, and this can be achieved by FP which activates MA and makes the locative PP available at VP level.

In example 8, on the contrary, we understand that the location from the which the seeing event took place is the bathroom and that John was certainly there; however we are given no information whatsoever about Mary’s location. This case is treated as the previous one, except that the PP is computed as sentence adjunct rather than as VP complement.

As for Referential Identification, the following examples are disambiguated at the level of pronominal binding:

9. The doctor called in the son of the pretty nurse who hurt herself.

10. The doctor called in the son of the pretty nurse who hurt himself.

Pronominal binding is a level of computation that takes place after f-structure has been completely checked and built in LFG –the same applies in GB framework, where S-structure gives way to L-structure and this is where binding takes place. In order to take pronominal binding into account we should be able to backtrack from one level of representation f-structure, to a lower level c-structure and to attach the predicative adjunct constituted by the relative clause at a higher level in case the reflexive pronoun is masculine. A very time-consuming and risky process.

We propose RI, instead, a strategy completely cast in IIT, which addresses a specific semantic level of representation: the Discourse Model, or History List of all entities appeared in the text under analysis with their properties. RI should be activated whenever a Definite NP is completely parsed and adjunct modifiers have to be attached locally, in other words while still in the process of c-structure building.

In examples 9 and 10, we assume that RI should be activated to ascertain whether in the DM there is an entity which has the property of being hurt, associated to the individual “son“ or to the individual “nurse“. In order to produce this kind of interpretation process, we also assume that relative clauses represent presupposed or known facts in case they are factual and these facts should be present in the DM. No syntactic and semantic constraints may be invoked in favour of one or the other interpretation: attachment of the relative adjunct is simply predictable from the DM. We query in the history list to recover the semantic identifier associated to the entity “book“. In turn this Id is used to recover all the properties associated with it and then to match them with the properties described in the relative clause which have to be organized in a semantically adequate representation. If the intersection is null the relative clause is attached locally. However the presence of a reflexive pronoun is the trigger for the local search of an adequate antecedent which in this case would push the system to apply MA and deposit the predicative adjunct in the WFST to be processed higher up.

2.1 Sorting out Language Dependent Differences

In our perspective we would like to take a very pragmatic and experimental stand on the problem of ambiguity. In the first place we want to look only at structural ambiguity and build up a comprehensive taxonomy from a syntactic point of view; secondly, we want to spot language dependent ambiguities in order to define the scope of our research appropriately.

A. Omissibility of Complementator

• NP vs S complement

• S complement vs relative clause

B. Different levels of attachment for Adjuncts

• VP vs NP attachment of pp

• Low vs high attachment of relative clause

C. Alternation of Lexical Forms

• NP complement vs main clause subject

D. Ambiguity at the level of lexical category

• Main clause vs reduced relative clause

• NP vs S conjunction

E. Ambiguities due to language specific structural proprieties

• Preposition stranding

• Double Object

• Prenominal Modifiers

• Demonstrative-Complementizer Ambiguity

• Personal vs Possessive Pronoun

3 Implementing Parsing Strategies

Among contemporary syntactic parsing theories, the garden-path theory of sentence comprehension proposed by Frazier (1987a, b), Clifton & Ferreira (1989) among others, is the one that most closely represents our point of view. It works on the basis of a serial syntactic analyser, which is top-down, depth-first –i.e. it works on a single analysis hypothesis, as opposed to other theories which take all possible syntactic analysis in parallel and feed them to the semantic processor.

From our perspective, it would seem that parsing strategies should be differentiated according to whether there are argument requirements or simply semantic compatibily evaluation for adjuncts. As soon as the main predicate or head is parsed, it makes available all lexical information in order to predict if possible the complement structure, or to guide the following analysis accordingly. As an additional remark, note that not all possible syntactic structures can lead to ambiguous interpretations: in other words, we need to consider only cases which are factually relevant also from the point of view of language dependent ambiguities.

We implemented two simple enough mechanisms in order to cope with the problem of nondeterminism and backtracking. At bootstrapping we have a preparsing phase where we do lexical lookup and we look for morphological information: at this level of analysis of all input tokenized words, we create a stack of pairs input wordform - set of preterminal categories, where preterminal categories are a proper subset of all lexical categories which are actually contained in our lexicon. The idea is simply to prevent attempting the construction of a major constituent unless the first entry symbol is well qualified. When consuming any input wordform, we remove the corresponding pair on top of stack.

In order to cope with the problem of recoverability of already built parses we built a more subtle mechanism that relies on Kay's basic ideas when conceiving his Chart (see Kay, 1980; Stock, 1989). Differently from Kay, however, we are only interested in a highly restricted topdown depthfirst parser which is optimized so as to incorporate all linguistically motivated predictable moves.

An already parsed PP is deposited in a table lookup accessible from higher levels of analysis and consumed if needed. To implement this mechanism in our DCG parser, we assert the content of the PP structure in a PP table lookup storage which is then accessed whenever there is an attempt on the part of the parser to build up a PP. In order to match the input string with the content of the store phrase, we implemented a WellFormed Substring Table (WFST) as suggested by Woods (1973). Now consider the way in which a WFST copes with the problem of parsing ambiguous structure in its chart. It builds up a table of well-formed substrings or terms which are partial constituents indexed by a locus, a number corresponding to their starting position in the sentence and a length, which corresponds to the number of terminal symbols represented in a term. For our purposes, two terms are equivalent in case they have the same locus and the same length.

In this way, the parser would consume each word in the input string against the stored term, rather than against a newly built constituent. In fact, this would fit and suit completely the requirement of the parsing process which rather than looking for lexical information associated to each word in the input string, only needs to consume the input words against a preparsed well-formed syntactic constituent.

To give a simple example, suppose we have taken the PP “in the night“ within the NP headed by the noun “painting“. At this point, the lookahead stack would be set to the position in the input string that follows the last word “night“. As a side-effect of failure in semantic compatibility evaluation within the NP, the PP “in the night“ would be deposited in the backtrack storage. The input string would be restored to the word “in“, and analysis would be restarted at the VP level. In case no PP rule is met, the parser would continue with the input string trying to terminate its process successfully. However, as soon as a PP constituent is tried, the storage is accessed first, and in case of non emptiness its content recovered. No structure building would take place, and semantic compatibility would take place later on at sentence level. The parser would only execute the following actions:

- match the first input word with the (preposition) head of the stored term
- accept new input words as long as the length of the stored term allows it by matching its length with the one computed on the basis of the input words.

3.1 Principles of Sound Parsing

• **Principle One:** Do not perform any unnecessary action that may overload the parsing process: follow the *Strategy of Minimal Attachment*;

• **Principle Two:** Consume input string in accordance with look-ahead suggestions and analyse incoming material obeying the *Strategy Argument Preference*;

• **Principle Three:** Before constructing a new constituent, check the storage of WellFormed Substring Table (WFST). Store constituents as soon as they are parsed on a stack organized as a WFST;

• **Principle Four:** Interpret each main constituent satisfying closer ties first - predicate-argument relations - and looser ties next - open/closed adjuncts as soon as possible, according to the Strategy of Functional Preference;

• **Principle Five:** Erase short-memory stack as soon as possible, i.e. whenever main constituents receive Full Interpretation.

•***Strategy Functional Preference*****:** whenever possible try to satisfy requirements posed by predicate-argument structure of the main governing predicate as embodied in the above Principles; then perform semantic compatibility checks for adjunct acceptability.

•***Strategy Minimal Attachment:*** whenever Functional Preference allows it apply a Minimal Attachment Strategy.

The results derived from the application of Principle Four are obviously strictly linked to the grammatical theory we adopt, but they are also the most natural ones: it appears very reasonable to assume that arguments must be interpreted before adjuncts and that in order to interpret major constituents as arguments of some predicate we need to have a completed clause level structure. In turn adjuncts need to be interpreted in relation both to clause level properties like negation, tense, aspect, mood, possible subordinators, and to arguments of the governing predicate in case they are to be interpreted as open adjuncts.

As a straightforward consequence, owing to Principle Five we have that reanalysis of a clause results in a Garden Path (GP) simply because nothing is available to recover a failure that encompasses clause level reconstruction: we take that GP obliges the human processor to dummify all naturally available parsing mechanisms, like for instance look-ahead, and to proceed by a process of trial-and-error to reconstruct the previously built structure in order not to fall into the same mistake.

4 Experimental Results

Here below we give parsing times on a SparcStation 5 for ex.1: `Mary put the book on the table.` With no parsing strategy activated, the PP is taken as adjunct in the NP constituent and the oblique argument is left indefinite.

Time: 1.3 sec.; with only Minimal Attachment activated, Time decreases: 0,87 sec.; with Functional Preference and Minimal Attachment activated, Time increases slightly: 1.05 sec. With Semantic Evaluation and no other strategy activated there is a big Time increase: 2.8 sec. The reason of this increase lies in the fact that SE for arguments of a given predicate is only activated at the end of the parsing process with Functional Unification: at this level the WFST is still filled with accepted constituents which will then be erased. Finally, for the strategy of

Referential Identification we have ex.9 which requires no strategy activation; however in ex.10 the relative clause should be attached adequately on a higher level. Time depends on the size of the DM: With a small one, time increases only a small fraction required to perform a query in the history list to recover the semantic identifier associated to the entity “book“. The same procedure would then apply for the higher NP “the son“. Overall time increases again and rises to 3.7 sec.

5 Treebank Derived Structural Relations

As noted above in the Introduction, an important contribution to the analysis of PP attachment ambiguity resolution procedures is constituted by the data made available in syntactic Treebanks. Work still underway on our Venice Italian Corpus of 1 million occurrences revealed a distribution of syntactic-semantic relations which is very similar to the one reported by Hindle et al. in their recent paper and shown in Table 1 below.

Argument noun	378	39.5%	
Argument verb	104	11.8%	
Light verb	19	2.1%	
Small clause	13	1.5%	
Idiom	19	2.1%	57%
Adjunct noun	**91**	**10.3%**	
Adjunct verb	**101**	**11.5%**	
Locative indeterminacy	**42**	**4.8%**	
Systematic indeterminacy	**35**	**4%**	
Other	**78**	**8.8%**	**39.4%**
	TOTAL	**880**	**100%**

Table 1: *Shallow Parsing & Statistical Approaches*
(Data from D.Hindle & M.Roth, “Structural Ambiguity and Lexical Relations”)

As the data reported above clearly show, most of the prepositional phrases are constituted by arguments of Noun, rather than of Verb. As the remaining data, adjuncts are represented approximately by the same amount of cases, 11% of the sample text.

Venice Italian Corpus
1 million tokens

- **All prepositions - 54 different types or wordforms: 170,000 occurrences**

Argument-like prepositions

DI/of and its amalgams	**78,077 --> 46%**	
A/to and its amalgams	**29,191 --> 17.2%**	
DA/by-from and its amalgams	**13,354 --> 7.9%**	**71.1 %**

Adjunct-like prepositions

IN and its amalgams -	***21,408 --> 12.6%***	
PER and its amagalms -	***12,140 --> 7.1%***	
CON and its amalgams -	***5,958 --> 3.5%***	***23.2%***

Table 2: *Shallow Parsing & Statistical Approaches*

We started looking at our data by collecting all information on prepositions as a whole and then we looked into our Treebank and looked for their relations as encoded in the syntactic constituent structure. Here below we report data related to prepositions for the whole corpus: Notice that in Italian as in English, the preposition *of / di* would be used mainly as a Noun argument/modifier PP.

In contrast to English, however, nominal premodifiers do not exist in Italian, and the corresponding Italian Noun-Noun modification or argument relation without preposition would be postnominal. Such cases are not very frequent and constitute less than 1% of Noun-Noun head relations. We then selected 2000 sentences and looked at all prepositional phrases in order to highlight their syntactic and semantic properties, and we found out the following: *(see page 72)*

PPs not headed by DA or DI	3977		51%		
Argument of verb		944		23.7%	
Argument of Noun		1300		32.7%	
Adjunct of Noun or Verb	1733		43.6%		
PPs headed by DA			*504*		*6.5%*
Argument of Verb		*164*		*32.5%*	
Argument of Noun		*114*		*22.6%*	
Adjunct of Noun or Verb	*226*		*44.9%*		
PPs headed by DI			*3314*		*42.5%*
Argument of Verb		*72*		*2.17%*	
Argument of Noun		*2733*		*82.5%*	
Adjunct of Noun or Verb	*509*		*15.4%*		
TOTAL				**7795**	**100%**
Arguments of Verb			1180		15%
Arguments of Noun			4147		53%
Ambiguous PPs			2468		32%

Table 3: *Quantitative Syntactic and Semantic Distribution of PPs in VIC*

- The number of prepositional phrases in Italian texts is four times bigger as the one reported for English Texts, and this might be due to the poor use of nominal modifiers which in Italian can only be post-modifiers, attested from an analysis of the sample text;
- PPs Arguments of Nouns are 53% in Italian and 39% in English, i.e. 14% more in Italian;
- PPs Arguments of Verbs are 15% in Italian and 17% in English – if we sum all argument types and idioms
- together -, i.e. 2% more in English;
- Adjuncts of Nouns and Verb are 31% in English and 32% in Italian.

Thus, the only real major difference between the two languages can be traced back to the behavior of PP noun arguments, which in turn can be traced back to a language specific difference: the existence of prenominal modifiers in English and not in Italian – not yet substituted by the use of postnominal modification.

References

Altman G.T.M. (ed.) (1989), Language and Cognitive Processes 4, 3/4, Special Issue – Parsing and Interpretation.

Bianchi D., R.Delmonte, E.Pianta (1992), GETA_RUN – A General Text Analyzer with Reference Understanding, in Proc. 3rd Conference on Applied Natural Language Processing, Systems Demonstrations, Trento, ACL 1992, 9–10.

Clifton C., & F. Ferreira (1989), Ambiguity in Context, in G.Altman (ed), Language and Cognitive Processes, op.cit., 77–104.

Delmonte R. (1984), La “syntactic closure“ nella Teoria della Performance, Quaderni Patavini di Linguistica, 4, Padova, 101–131.

Delmonte R. (1992), Linguistic and Inferential Processing in Text Analysis by Computer, Padova, Unipress.

Frazier L. (1987a), Sentence processing, in M.Coltheart (ed), Attention and Performance XII, Hillsdale, N.J., Lawrence Elbaum.

Frazier L. (1987b), Theories of sentence processing, in J.Garfield (ed), Modularity in knowledge representation and natural language understanding, Cambridge, Mass., MIT Press, 291–308.

Garrod S., Sanford T. (1988), Thematic Subjecthood and Cognitive Constraints on Discourse Structure, Journal of Pragmatics 12, 599–534.

D.Hindle & M.Roth (1993), Structural Ambiguity and Lexical Relations, Computational Linguistics 19, 1, 103–120.

Kay Martin (1980), Algorithm Schemata and Data Structures in Syntactic Processing, CSL-80–12, Xerox Corporation, Palo Alto Research Center.

Kennedy A., et al. (1989), Parsing Complements, in G.Altman (ed), Language and Cognitive Processes, op.cit., 51–76.

Pereira F. and Warren D. (1980), Definite Clause Grammars for Language Analysis. A Survey of the Formalism and a Comparison with Augmented Transition Networks. Artificial Intelligence, 13, 231–278.

Pritchett B.L. (1992), Grammatical Competence and Parsing Performance, The University of Chicago Press, Chicago.

Schubert L.K. (1984), On Parsing Preferences, Proc. of COLING, 247–250.

Steedman M.J. & Altmann G.T.M. (1989), Ambiguity in Context: a Reply, in Altman (ed), op.cit., 105–122.

Stock O. (1989), Parsing with flexibility, dynamic strategies and idioms in mind, in Computational Linguistics, 15, 1, 1–18.

Whitelock P., M.M.Woods, H.L.Somers, T.Johnson, P.Bennett (eds.) (1987), Linguistic Theory and Computer Applications, Academic Press, London.

Woods W.A. (1973), An Experimental Parsing System for Transition Network Grammars, in Rustin (ed) (1973), Natural Language Processing, Algorithmic Press, New York, pp.111–154.

Discourse Particles as Speech Act Markers

Henk Zeevat
Computational Linguistics/ILLC
Spuistraat 134,1012VB Amsterdam, NL
henk.zeevat@hum.uva.nl

Abstract
Eine Reihe von Diskurspartikeln werden mit dem Ziel analysiert, die Rolle, die sie bei der Markierung von Sprechakten spielen können, zu erklären. Die Analyse verwendet eine optimalitätstheoretische Rekonstruktion der Präsuppositionstheorie.

1 Introduction

When one tries to further develop Stalnaker's ideas (cf. Stalnaker (1978)) on the conditions for pragmatically correct assertion (informativity and consistency with respect to the common ground between speaker and hearer), it is natural to come up with conditions like the following[1].

(1) a. it is not common ground that the speaker believes A.
b. it is not common ground that the speaker believes that not A.
c. it is not common ground that the hearer believes A.
d. it is not common ground that the hearer believes that not A.

If one of these conditions is not fulfilled, the assertion is improper, or non-standard. In the first case there is little to no effect that the speaker can hope to gain by what she has said: it cannot be a proposal to eliminate possibilities from the common ground. In the second case the speaker is self-correcting, and so faces an inconsistency with her own beliefs as represented in the common ground. In the third case, the speaker is also doing something that is not an assertion in Stalnaker's sense: she is at best assenting to an assertion by the hearer. In the fourth case as well, the speaker is correcting the hearer rather than asserting something.

These theoretical speculations are confirmed by looking at Dutch or German sentences that realise such non-standard assertions: they invariably contain discourse particles, like toch (doch), inderdaad (tatsächlich), immers (ja), wel (doch).

The following examples bear this out. (a) can be a self-correction, (b) an assent to the hearer, (c) a reiteration, (d) a hearer correction.

(2) a. Peter is toch thuis.
a'. Peter ist doch zuhause.
a". Peter is at home (after all?).
b. Peter is inderdaad thuis.
b'. Peter ist tatsächlich zuhause.
b". Peter is indeed at home.
c. Peter is immers thuis.
c'. Peter ist ja zuhause.
c". As you know, Peter is at home.
d. Peter is wel thuis.
d'. Peter ist doch zuhause.
d". Peter IS at home.

It is important to make the following observations. In contexts for (2a) in which the common ground contains the speaker's opinion that Peter is not at home, omitting the *toch* makes the utterance infelicitous. Likewise (2b) without the *inderdaad* is infelicitous if, according to the common ground, it is already the hearer's opinion that Peter is at home. (2c) without immers is infelicitous if it is common ground that Peter is home and (2d) without *wel* is infelicitous if the hearer has just said that Peter is not at home. This is indeed just what follows from Stalnaker's views on assertion. The particles seem to have the power to make otherwise infelicitous assertions into specialised non-standard assertions that have other goals than standard assertions, like correcting opinions expressed earlier on or reconfirming established opinions.

An initial hypothesis might be that the particles are in the language just to mark the non-standard character of certain speech acts. But this hypothesis is easily refuted.

If this were so, it would not be possible to combine all four particles as in (3), which, though not easy to contextualise, is nevertheless perfectly acceptable Dutch.

(3) Peter is toch immers inderdaad wel thuis.

It follows minimally that the particles do not mark a particular combination of speaker and hearer commitments to the truth or falsity of the proposition, because that combination would be inconsistent. The hypothesis also has to go when one considers the full uses of the particles in question in Dutch or German, as we will later on. And finally, it turns out that although the particles may indicate a combination of speaker and hearer commitments, they also allow other interpretations.

This raises two questions. First, how is it possible that the particles can mark deviant speech acts, i.e. one would like to have an account of their use from which it follows that they can sometimes mark a hearer or speaker commitment? Second, can these insights be used to improve the recognition of the user intention in dialogue systems? In addition, the function of these particles is unclear and any elucidation is welcome.

This paper gives an experimental account of these four particles in terms of an extended presupposition theory and manages to explain the uses quoted in this introduction. It follows that there is a potential use of the particles in future dialogue systems, i.e. the ones that have a capacity for presupposition treatment. Section 2 introduces the presuppositional treatment of particles and sections 3, 4 and 5 apply the treatment to the four particles in question.

2 The proper treatment of the particle *too*

Kripke's notes on presupposition Kripke (s.d.) started a new period in the study of presupposition where the analogy with anaphora became more and more prominent. The two most successful accounts are Heim (1983) and Van der Sandt (1992). Yet, in terms of Kripke's original example these theories do not perform very well at all.

Kripke is puzzled by the example (4).

(4) John will have dinner in New York too.

The traditional theories predict that this sentence presupposes (5) which for (4) is a mere triviality.

(5) Someone other than John will have dinner in New York.

After all, New York is a vast city where millions have dinner every night. If this were the presupposition, the *too* would not give us extra information about the context. It would also be the case that we can always add a *too* to the sentence *John has dinner in New York*. Both of these predictions are wrong: *too* is infelicitous if the common ground does not entail that another person has dinner in New York and it gives us the information that the common ground has this property. Kripke's suggestion is that *too* tells us that the context and not the world contains another person who has dinner in New York and that the *too* is anaphoric to this part of the linguistic context.

Both Heim's and Van der Sandt's theories contain a resolution mechanism that can pick up the antecedent in the context (in that case the *too* does not give new information). But they also allow the presupposition to be accommodated. In that case, we get precisely the prediction that Kripke criticises, i.e. the requirement of an unidentified other person who has dinner in New York. The theories should rule out accommodation for *too*, but do not have the means to do that. In this way, the theories also predict that *too* can be freely added to our example, without truth-value change or infelicity.

There are some other aspects of *too* in which it is different from standard presupposition triggers, like factive verbs, definite descriptions and lexical preuppositions. The first is that *too* itself does not seem to give information. The following example of Heim brings this out. Two kids are secretly phoning each other after bedtime without the permission or knowledge of their parents.

(6) A: My parents think I am in bed.
B: My parents think I am in bed too.

In one of the interpretations of the utterance by *B*, the *too* belongs to the complement of the belief sentence. Yet, *B*'s parents know nothing about *A* being in bed or not. The example also illustrates another problem with *too*.

Too (and other particles) take antecedents that are not available according to Heim or Van der Sandt. The antecedent *A is in bed* in (6) is not entailed under the operator *B's parents think* and neither is it accessible according to the Discourse Representation Theory in which Van der Sandt's theory is couched. The last property of *too* that is unexplained by the two theories is that its occurrence is obligatory in the sense that in most of the utterances in which it occurs it cannot be omitted without resulting infelicity.

My proposal (Zeevat(2000)) is to (a) liberalise the set of allowed antecedents for presupposition triggers to the veridical contexts and to (b) assume a generation constraint. (c) Embedding the theory within a form of Bidirectional Optimality

Theory then allows an explanation of the absence of accommodation for *too* and other presupposition triggers. I will sketch the three steps.

Veridical contexts were proposed by Giannakidou (1998) as a characterisation of the contexts that do not license negative polarity items and include beliefs, dreams, suggestions, possibilities and iterations of these. Properly inaccessible antecedents (and negative polarity items) must be in the scope of at least one non-veridical operator. (7) shows some of the possibilities with *too*.

(7) A. Maybe John will go to Paris.
B. I will go there too.
John suggested that Mary left and Bill said Susan did too.

There are some limitations to the antecedents *too* can take, as illustrated by (8) which some people do not like.

(8) John dreamt that Bill is Paris and Tom will go there too.

The English *indeed* is more liberal and (9) illustrates the wider range of antecedents it can take. I do not know why *too* is less liberal than other particles in this respect.

(9) John dreamt that he passed the exam and indeed he passed.
John thinks that Mary hates him and Bill said that she does indeed.

Generation constraints are defeasible constraints that the human generator tries to optimally satisfy when generating a sentence from a characterisation of the semantics. The generation constraint needed for *too* is **ParseOther**, a principle that forces the marking of the presence of another entity of the same type in the context. *Too* marks the presence of another element of the same type, like *also, another or a different*. It is possible to defend the view that this is all that we have to say about the semantics of *too* and that its function provides the explanation of its lack of semantic content.

A similar principle is **ParseOld**, a principle that forces the marking of material already the context as old material. *Indeed* is one of the linguistic elements that carries out this job, other are pronouns and definite descriptions.

In a bidirectional optimality theoretic framework we can combine the above generation principles with Blutner (2000)'s reconstruction of Van der Sandt's presupposition theory by two interpretation principles: **DoNotAccommodate** and **Strength**. The first principle, ranked above the other, militates against accommodations, the second one selects the strongest reading from among the different readings that come out of the accommodation possibilities. In the resulting system, the following principle (Blutner's Law) can be derived.

(10) If a presupposing expression has simple non-presupposing alternatives, it does not accommodate.

The motivation is simple: with a common ground that requires accommodation, a speaker will always select the non-accommodating alternative because it does not lead to a violation of **DoNotAccommodate**. (In the particular version of bidirectional optimality theory advocated by Blutner interpretation constraints are scored together with generation constraints in both directions.)

The predictions that our theory makes for *too* are non-accommodation (this does not rule out a fair amount of partial resolution), the availability of all veridical antecedents, and obligatory occurrence when the veridical context contains another element of the same type. Non-accommodation is a consequence of existence of the simple expression alternative where *too* is omitted. The lack of semantic content is responsible for the possibility of veridical antecedents: it does not matter where the antecedent comes from because it does not need to exist locally.

In these respects, *too* contrasts sharply with a trigger like *regret*. First of all *regret* does not have simple expression alternatives, which means that it allows accommodation. Second, its presupposition makes a strong semantic contribution: it identifies the fact to which the subject has her emotional reaction. This fact must at least be a belief of the subject for the subject to have an emotional reaction to it. Therefore, only real facts and beliefs of the subject can be antecedents and other veridical antecedents are ruled out. The strongest requirement arises when the antecedent identifies a participant, a cause or a precondition of the event described by the clause that contains the trigger (pronouns or definite descriptions). Here the only antecedents are proper constituents of the context of the trigger.

The specification of a trigger is exhausted by a statement of its presupposition and its semantic contribution. The overlap between presupposition and semantics filters away unwanted veridical antecedents. Accommodation or not is controlled by the inventory of the language.

For further details I refer to Zeevat(2000).

3 *Inderdaad* and *Immers*

My hypothesis about *inderdaad* (*tatsächlich*, *indeed*) is that it is just a presupposition inducer, in this case presupposing the positive version of the sentence to which it attaches. As such, it is an old marker and the generation constraint **ParseOld** is responsible for its obligatory occurrence. It takes veridical antecedents, because it does not contribute to the semantics of the clause. It does not accommodate, because as a particle it has a simple expression alternative: the sentence without the particle.

What does this predict about the speech acts in which it occurs? Basically, it says that the hearer, or the speaker or both can have an old opinion that the sentence is true. But it is not necessarily the opinion of one or both of the conversational partners, since the antecedent can also be the opinion of a third party or even weaker, the content of a dream, a suggestion etc. A dialogue system can conclude from an occurrence of *inderdaad* that what is said is already present and it is only the presupposition resolution itself that forces the selection of a speech act of reconfirmation, when resolution is to the speaker or the common ground. It can be the speech act of assenting if the resolution is to a hearer opinion that is not shared. Absence of *inderdaad* when no other old-marker is present, can lead to the conclusion that we have a proper assertion and not a reconfirmation or assent.

The same holds for an occurrence in a question.

(11) Is Harry inderdaad thuis?
Is Harry indeed at home.

(11) presupposes that Harry is at home. In imperatives, it can only presuppose the imperative itself (or the desirability of the course of action).

If we look at a sample of actual uses the hypothesis is largely confirmed, except for an antiquated use as a synonym for *feitelijk* (in fact). This older use is important, because *inderdaad* seems to imply that the new information is better than what we had before. This is either because *inderdaad* retains some properties of *feitelijk* or it is a pragmatic implication of *reconfirmation* or *assent* as such. If *inderdaad* does not add semantical content, the purpose of reconfirmation or assent can only be that new evidence has been found. There is also a subtle distinction between an assent with an isolated *inderdaad* and one with *ja* (yes) or a nod of the head. If *inderdaad* is used, the speaker claims to have better information than the other speaker whose assertion she assents to. We could capture the distinction by claiming that a sentence with *inderdaad* must still be informational in the sense of Stalnaker, in indicating that the speaker believed it not as a result of what the interlocutor asserted, but already before that. If we supply our reconfirmation or assent with an assertion containing *inderdaad*, the new information can only be the elimination of an existing uncertainty.

Immers is like *inderdaad* in presupposing the truth of the clause to which it attaches, but it is quite different at the same time. *Immers* makes a quite clear semantic contribution. It turns the clause into a reason for accepting what was said just before. Now reasons why something is the case must be the case as well in order to qualify as proper reasons. That John dreamt he was in Spain, or that Charles has suggested so are not reasons why John is away from home. That is why *immers* in simple clauses only takes proper antecedents and no non-entailed veridical contexts. It also does not bring the effect of the new and better view that we noticed with *inderdaad* and we would not expect that, since *immers* contributes to the semantic content of the clause.

Like *inderdaad*, it is obligatory. If the statement is already common ground, *immers* is needed to mark the fact that we are dealing with old information. This leads to the following curious fact. *Omdat* like its English counterpart *because* is a presupposition trigger. This gives Dutch two ways of expressing the sentence (12a).

(12) a. He did not come because he is in Paris.
b. Hij kwam niet omdat hij immers in Parijs is.
c. Hij kwam niet omdat hij in Parijs is.

(12b) is obligatory resolved to the common ground. (12c) is obligatory accommodated, because, if it were old information, *immers* would have to appear. *Omdat* without *immers* is a presupposition trigger that is marked for obligatory accommodation, comparable to a complement of *regret* that has new intonation, or perhaps also indefinite NPs.

Formally, *immers* A has two presuppositions, the one we discussed and the current last sentence. It asserts that the first presupposition is a reason why the second one holds.

Looking at our data, one finds complete confirmation, although there cases where the causal connection is not very clear. *Immers* is not a high frequency item unlike its german approximate equivalent *ja* which has quite a number of other uses next to the one discussed here. Questions and imperatives with *immers* are not possible and the analysis given here explains why.

The occurrence of $immers$ in a user utterance is a reliable indication for assuming that the user is not making a normal assertion, but assumes both that the material is already established and relevant at the current point in the dialogue.

4 *Wel*

The marker *wel* in the uses we are focussing on is the typical marker of a correction to a negative utterance made by the other party. It is accented in that case and the most likely explanation is that *wel* is entering in a contrast relation with the negation in the corrected sentence.

(13) A: Jan is niet thuis. (Jan is not at home)
B: Jan is WEL thuis. (Jan IS at home)

In corrections to non-negated sentences, accented *niet* takes over this role.

(14) A: Jan is thuis.
B: Jan is NIET thuis.

But it is not clear there is an element here with which *niet* contrasts. Nevertheless, the relation of contrast with the corrected sentence is so strong that the correct explanation is probably that the whole sentence bears contrast, with everything except *niet* deaccented as old material.

There are many other uses of *wel*. Typical is the use in a concession:

(15) A: Jan kwam het boek toch gisteren terugbrengen.
A: John was going to return the book yesterday, wasn't he?
B: Jan kwam WEL, maar hij had het boek niet bij zich.
B:John came allright but he did not have the book.
B1: Jan kwam niet, maar hij heeft het boek WEL teruggegeven.

B1: Jan did not come, but he gave the book back allright.

Here the *wel*-clause marks the part where the speaker agrees with the other speaker. But this can be reversed, as in B1. The A. sentence invokes a context in which the plan that Jan was bringing the book yesterday is assumed and evidence is available that the plan has not been carried out. Another case is (16).

(16) A: So they came?
B: Jan WEL, maar Marie NIET.
(Jan did, but Marie did not)

(17) So they did not come?
Jan WEL maar Marie NIET.

Almost idiomatic are the unaccented combinations with modal verbs.

(18) Het moest wel.
It had to be.
implies: I/we did not want to but I/we had no choice.
presupposes: opposite inclination (?)

(19) Het zal wel beter gaan in het voorjaar.
It will probably be better in spring
In context: denies that the current bad situation will continue

(20) Het zal wel.
Ironical way of expressing disbelief.

(21) Het lijkt wel of je nooit meer thuis bent.
It would appear that you are never at home anymore.
presupposes falsity of what appears to be the case (?)

(22) John shows Mary his new dog.
M: Het lijkt wel een varken.
M: It looks like a pig.
presupposes it is not one (?)

(23) Kom je WEL? (presupposes the opposite)
Kom je wel? (expresses doubt)

(24) Wil je wel?
DO you want?
expresses doubt

Quite generally, we seem to be able to say that *wel p* presupposes $\neg p$. In concessive phrases, the presupposition can disappear and the main function is the contrast with the negation in the other half of the pair. The presupposition is perhaps still around in the unacccented cases, but it may be that a case for case analysis like the one I will supply later on for *toch* is in order.

The accented uses require overt negations to contrast with, either within a concessive pair or outside one. In the last case, the negated clause coincides with *wel*'s antecedent.

The explanation of *wel*'s appearance in a sentence must be two-fold. We need a principle that inserts it in a concessive pair, if the concession is built around a positive and negative element, but the generation of concessive constructions does not concern us in this paper. The other occurrences are due to the **ParseOld** principle we discussed before.

Wel takes veridical antecedents, as shown in (25).

(25) Karel droomde dat hij niet voor zijn examen zou slagen,
maar hij haalde het WEL.
Karel dreamt he would not pass his exam, but he
passed it allright.
Piet zei dat Marie niet zou komen, maar ze kwam
WEL.
Piet said that Marie would not come, but she did.

The use of *wel* can help in identifying the dialogue move the speaker is making. It is helpful in identifying corrections, though it must be distinguished from concessive uses and from other presupposing uses.

5 *Toch*

This is by far the most complicated of the four particles that are the protagonists of this paper. Compare the examples in (26), based on clauses meaning: *he is in Amsterdam or come to Amsterdam*

(26) a. Laten we hem vrijdag opzoeken. Hij is dan toch in Amsterdam.
a'. Let us visit him on Friday. He is then in Amsterdam anyway.
b. Hij is toch in AmStErDaM?
b'. He is in Amsterdam, isn't he.
c. Hij is TOCH in Amsterdam.
c'. He is in Amsterdam after all.
d. Is hij TOCH in Amsterdam?
d'. Is he in Amsterdam after all? (We thought he would not be)
e. Kom toch naar Amsterdam. (exhortation)
e'. Come to Amsterdam. (you know you'll like it).
f. Kom TOCH naar Amsterdam.
f'. Come to Amsterdam, (although I see why you do not want to).

The emphatic uses of *TOCH* are pretty straightforward. They indicate that the speaker presupposes the negation of the statement or question she is making. In the case of the imperatives, it is the opposite plan or the desire not to that is presupposed. But the non-emphatic uses are difficult to accommodate in this scheme.

Example (26b). is the most involved. Often it is treated as a question (a confirmation question) but the form is of an assertion and the intonation is not that of a normal question. Also the facial expression appropriate to its utterance indicates that it is really an assertion uttered expressing surprise at the content, like the assertion in (27).

(27) Hij is in AmStErDaM?

The surprise indicates that the speaker believes to know that what he says is false, in (27). It is a reaction to information that „he" would be in Amsterdam. What the *toch* does in (26b) is to invert these speaker assumptions: the speaker now believes that „he" is in Amsterdam and reacts to information to the contrary.

We could perhaps say that *toch* resolves to the negation of the statement made by the interlocutor. But then after resolution we have assertion with the expression of surprise, which is quite different: the speaker is not surprised that „he" is in Amsterdam, she is surprised that „he" is not. It would seem that this indicates that the *toch* here resolves to the positive information that „he" would be in Amsterdam and – because that rules out surprise at the positive information – the surprise is caused by something else, nl. the information supplied by the interlocutor. If we look at (26a) this confirms that pattern.

The *toch* here is a device of reminding the interlocutor of some old information and it is functioning not unlike *immers* which could take the place of *toch* in this context. In fact, there are dutch speakers who never use *immers* and always use unaccented *toch* instead. In my dialect, it normally just means that the fact he is in Amsterdam is independent of the current issue, more or less like the english *anyway*. It can be common ground that he is, but it can also just be unrecorded speaker information that he is.

Uses of unaccented *toch* in questions seem to be impossible. In imperatives, it softens the appeal made on the interlocutor. It does not seem to be impossible to understand this as presupposing a similar desire in the interlocutor. Again the opposite of the accented *TOCH* which presupposes a contrary attitude to the action ordered in the imperative.

In my corpus, by far most uses of *toch* are pro-concessives, i.e. single word concessives (like isolated *though* in English) that can be paraphrased by full although-sentences whose content is given by the context. This is a weakening of what we find in (26c) which seems naturally characterised by presupposing the negation of the clause. Though concessive sentences provide reasons for thinking that the main-clause is false, they do not (cannot) provide the information that the negation is true. It is possible to bring them closer by the notion of a suggestion. The contextually given concessive material can be taken as a suggestion that the clause is false and this would be an appropriate veridical antecedent. Alternatively, we should start from the notion of a reason to be false and let (accented) *TOCH* presuppose a reason for the clause to be false. I prefer the first alternative, since the second alternative makes the integration of the unaccented uses even more problematic than they are already.

What can we make of *toch* in our presuppositional theory? I am not very sure. I would like to say the following. *Toch* is just an old-marker without a preference for positive or negative antecedents. If the antecedent of *toch* has the same polarity as the current clause, no accent is provided by the speaker because there is no contrast between the clause and the recovered presupposition. If the antecedent has opposite polarity, accent results from the recovery of the antecedent. The accent would just be the result of the existence of an alternative in the speaker's mind, here created by the speaker's awareness that she is old-marking a clause for the prior occurrence of a negated version of the clause. I do not have a fully worked out accenting theory from which this accenting pattern would follow, but such a theory is needed. The alternative is that we have a tonal distinction between two lexical items *toch* and *TOCH* with different semantic properties. But this runs against the following argument that I owe to Manfred Bierwisch (p.c.). It would then be completely incomprehensible how it can be that Dutch and German have almost exactly the same *toch*/*doch* and *TOCH*/*DOCH* and the same for other accented and deaccented particles. In addition, it would make Dutch and German into tonal languages, a claim for which we do not have independent evidence.

Let us go through the examples. In (26a), the lack of accent indicates that the presupposed material is of the same polarity as the current clause: the second sentence presupposes that the unknown *he* will be in Amsterdam on the Friday. A problem is that it seems to be possible as well that the antecedent is not – even veridically – common ground. The speaker may merely indicate that the *he* is Amsterdam for reasons unrelated to the current purposes of the conversation.

In (26b) we meet a pure confirmation of the hypothesis. The presupposition is that *he* is in Amsterdam and this is expressed by the unaccented *toch*. The sentence itself is triggered by surprise over contrary information supplied by the other speaker.

In (26c), the accent indicates that the old material is of contrary polarity. The sentence corrects the old material. The same in (26d) which must be prompted by a suggestion that contrary to what we appeared to know he is in Amsterdam.

(26e) is more problematic. The insertion of the *toch* tones down the imperative to an exhortatation, and this can perhaps be explained by assuming that a positive inclination on the part of the hearer to do just that is assumed by the speaker. But like in (26a) it is not strictly required that the positive inclination is registered in the common ground. It may be, but it need not.

(26f) finally confirms our hypothesis in much the same way as (26c) and (26d). It would be appear that we find full confirmation for the hypothesis looking at the accented *tochs* and rather shaky confirmation for the unaccented uses. Here it seems that its function of marking a specific speech act has partly usurped its semantic contribution. But there is another way to look at this. We started by assuming that the function of the *toch* is being an old-marker and we derived from that accommodation is impossible: there is a simple expression alternative which does not do old-marking, nl. the sentence without the particle. Now in the problematic cases, it is questionable whether the sentence without the *toch* means the same. Leaving out the *toch* in (a.) fails to express that *he* will be in Amsterdam on Friday for independent reasons, (e.) becomes a full imperative. Now if the meaning changes when the *toch* is omitted, accommodation is possible, and that, I would suggest, is what happens. The story about *toch* can be as I indicated but it must take account of idiomatic further meanings acquired in discourse. Here I would suggest that an unaccented *toch* in imperatives, presupposing the action ordered, naturally changes the imperative into an exhortation to follow one's inclination to do as ordered. And, unaccented *toch* in an assertion, presupposing the truth of what is asserted, makes the truth of the assertion independent of the current discussion: it becomes a reiteration. In the absence of other means of expressing exhortations or reiterations, unaccented *toch* will also become a marker of these special speech acts.

(26b) is as we noticed also a case where a simple expression alternative is lacking: the surprising fact changes polarity if we take out the *toch*. This would predict that, also in this use, accommodation is possible a prediction that seems to be borne out.

In other respects, *toch* seems to follow the pattern of the other particles discussed in this paper. It takes veridical antecedents as in (28), it makes no contribution to the content of the clause and it cannot be omitted (but sometimes replaced) where it occurs.

(28) Jan droomde dat hij was gezakt voor het examen,
maar hij had het TOCH gehaald.
Jan dreamt he failed the exam, but he passed.

Toch is useful for future dialogue systems as an indicator of corrections when it is accented and when the corrected element can be found in the common ground.

6 Conclusion and Further Research

My first encounter with particles occurred half-way the eighties when I was working on pronoun resolution. Hypotheses about discourse and dialogue structure can have dramatic consequences for correct resolutions. It was then —as it is only marginally less now— difficult to recognise discourse and dialogue structure and in our system we did not even have the resources to reconstruct speaker plans. Particles seemed a way out: in German they are extremely frequent and together with tense shifts and topic they seemed to offer a heuristics that would make our recognition of the discourse and dialogue structure better.

This did not work because particles are not very well understood: many meanings are normally distinguished and few of the meanings seem to be very relevant for the discourse grammarian. The *anyway*, the „pop-marker“ of classical discourse grammar is almost an isolated case. And *anyway* is not a pop-marker at all. It marks that what is said in the current clause does not depend on the issue of the last clause or paragraph. The discourse function of closing of a topic is derived from this more primary function.

It is much the same I believe with the particles I have focussed on in this paper. Their function can be clarified to a large extent by analysing them as as presupposition triggers with a number of special properties. It follows that they have certain discourse functions, but those functions are not their primary function. As I hope to have shown in this paper, a reduction to presupposition makes it feasible to use certain particles for the recognition of the speech act the user is making.

There is a considerable class of particles that can be analysed as presupposition triggers. For *again*, I refer to c. Next to *again* we find *still*, *yet*, *already* and *notanymore*. Our four old-markers should also include *instead* and perhaps *dan*. As presupposition triggers, they have overwhelming similarities, like the avoiding of accommodation and a strong preference for partial resolution. The dividing line is the question of semantic contribution.

The temporal particles clearly sit with *immers* in requiring proper antecedents, because next to marking old material, they also make a contribution to the temporal semantics.

The implementation of the current approach to particles is not much more involved than the general approach to presupposition and anaphora resolution implemented in e.g. Johan Bos's DORIS system, an approach that could clearly be integrated in logic based dialogue systems. The main but unimportant difference is that a larger class of contexts needs to be searched to take care of veridical antecedents as well. A difference —really an advantage— is that the generation constraints also allow inferences about the absence of certain antecedents. The most serious obstacle to a full implementation is the difficulty of doing partial resolution, but this is a difficulty shared with any computational treatment of presupposition. A good discussion of the task for German *wieder* can be found in [Kamp/Rossdeutscher (1994)].

Future research will have to determine what other discourse particles can be captured in the presuppositional analysis proposed here.

References

[Blutner2000] R. Blutner. 2000. Some Aspects of Optimality Theory in Interpretation. *Journal of Semantics* (to appear).

[Giannakidou1998] A. Giannakidou. 1998. *Polarity Sensitivity as (Non) veridical Dependency*.John Benjamins.

[Heim1983] I. Heim. 1983. On the projection problem for presuppositions. In: WCCFL, 2: p. 114-26.S. Kripke. s.d. Presupposition. (MS)

[Kamp/Rossdeutscher1994] H. Kamp & A. Rossdeutscher. 1994. Remarks on Lexical Structure and DRS Construction. *Theoretical Linguistics*, 20: p. 97-164.

[Van der Sandt1992] R. van der Sandt. 1992. Presupposition projection as Anaphora Resolution. In: *Journal of Semantics,* 9: p. 333-77.

[Stalnaker1978] R.C. Stalnaker. 1978. Assertion. In: P. Cole (ed.). *Syntax and Semantics 9: Pragmatics*. New York, p. 315-32.

[Zeevat1997] H. Zeevat. 1997. The Common Ground as a Dialogue Parameter. In G. Jäger and A. Benz (eds.) Proceedings Mundial 1997, CIS, Universitaet Muenchen.

[Zeevat2000] H. Zeevat. 2000. Explaining presupposition triggers. (to appear).

NOTES

1 For a full discussion of these conditions see [Zeevat1997].

2 Giannakidou's notion is more restricted and omits suggestions and *maybe*-environments that in some languages allow certain negative polarity items.

3 I used a net-version of Multatuli's Max Havelaar, a classic dutch novel.

4 *Feitelijk* is not a presupposition trigger, though it can indicate another point of view on the issue at hand. Its analysis is not straightforward.

Imperatives, Commitment and Action: Towards a Constraint-based Model*

Paul Piwek
ITRI — University of Brighton
Watts Building, Moulsecoomb, Brighton BN2 4GJ, UK
Email: Paul.Piwek@itri.brighton.ac.uk
Web: http://www.itri.brighton.ac.uk/-Paul.Piwek

Abstract

The aim of this paper is to provide an analysis of the function of imperatives in dialogue. In particular, the focus is on the use of logically complex imperatives (e.g., 'Say hello to John, if you meet him') and temporal reference in imperatives (e.g., 'Open the bottle of champagne at midnight'). Two specific problems involving logically complex and temporally referring imperatives are introduced. These problems are addressed within a framework for communicating agents. It is argued that such a framework needs to allow for partial models of communicating agents and their environment and for the intentionality of agent states.

1 Introduction

This paper reports on a study into the function of imperatives in dialogue. Imperative clauses are traditionally characterized by the lack of a subject, the use of a base form of the verb, and the absence of modals as well as tense and aspect markers. Examples of naturally occurring imperatives are: 'Get off the table', 'Don't forget about the deposit' and 'Hold on, are we late?'. (The definition and the examples are taken from Biber et al., 1999:219). The function of imperatives (i.e., the task which they perform in discourse)[1] is typically to get the addressee to do something. The aim of this paper is to provide a formal analysis of this function. Eventually, this analysis may prove to be useful beyond the study of imperatives, since the same function can be expressed by different means (E.g., by a question such as 'Will you be quiet for goodness sake!'; from Biber et al., 1999:220). Let me now introduce two specific problems concerning the function of imperatives which I address in this paper.

(P.1) Imperative clauses, like declarative ones, can be used in combination with connectives such as 'or' 'and' and 'if'. For instance, 'if' can be used to construct rule-like imperatives such as: 'Say hello to John, if you meet him' (cf. Hamblin, 1987:84). A problem emerges when we consider the translation of the truth conditional content of this imperative into first-order logic (FOL). Suppose we choose the following translation: *meet(addressee,john)* $\rightarrow$ *say_hello_to(addressee,john)*, i.e., if the addressee meets John, then he says hello to John[2]. Furthermore, the example has been formalized without any explicit reference to time. Temporal aspects of imperatives are, however, dealt with later on in this paper. Let us assume that such a formalization of the content of imperatives is a legitimate move. Then we can use that formalization as a basis for formalizing the function of imperatives. It can, however, be shown quite easily that one very tempting analysis of imperatives cannot be correct. According to this analysis the function of an imperative is to instruct the addressee to make its content true. In FOL a conditional is true if its antecedent is false or its consequent is true. This analysis predicts that an addressee of a conditional imperative can comply with the imperative by simply making its antecedent false. For the aforementioned example this means that if the addressee were to avoid meeting John, this would count as acting in accordance with the imperative. This counterintuitive prediction casts doubt on the correctness of the aforementioned analysis.

(P.2) A further problem concerning logically complex imperatives has been attributed to Ross (1941). Ross points out that the relation of logical consequence between imperative clauses appears to be different from that between declarative clauses: in FOL, $\phi \vee \psi$ is a logical consequence of ϕ, whereas the imperative 'Post the letter or burn it' does not seem to follow from 'Post the letter'.

Generally speaking, I want to provide an alternative to what Hamblin (1987) has termed 'Giddap' — 'Whoa' theories of imperatives. According to such theories, imperatives in dialogue are immediately followed by an action, as in 'Instructor: Change to second gear. Pupil: (Changes to second gear)'. Such theories fail to extend to logically complex imperatives and imperatives which contain explicit temporal references. In order to explicate the function of imperatives I construct a framework of communicating agents in which a formal characterization of this function can be specified. More specifically, I use this framework to provide an outline of a succession of increasingly more elaborate and, hopefully, more realistic theories. I acknowledge that the end result is still only a rough approximation of the use of imperatives in real human-human conversations.

2 Framework: Models for Communicating Agents

In order to study the function of imperatives a model is constructed of the situations in which imperatives are being used. This model includes the agents who issue and receive imperatives and also the environment in which these agents operate. The idea is that the aforementioned agents are approximations of human beings, although only very rough ones[3].

The aforementioned approach is not very different from that in the other sciences such as biology or physics. However, for a functional analysis of imperatives it will not do to construct a purely physical model of communicating agents and their environment. Such a description would gloss over the fact that communicating agents carry and use *information* to guide their actions. For instance, when a person is told to 'Open a bottle of Champagne at midnight', s/he will have to store this information and act upon it at the right time. The fact that an agent carries some piece of information with it can be analysed as a state of that agent. Such a state, which an imperative can give rise to, is traditionally known as an Intentional state. Intentional states are states which are directed at or about objects or states of affairs in the world[4]. The crucial characteristic of Intentional states is that we can speak about whether they have been satisfied or not (cf. Searle, 1983). The analysis which is proposed here starts from the assumption that the satisfaction of an imperative is grounded in the satisfaction of the Intentional state (of the addressee) which the imperative gave rise to[5]. It will, however, transpire that quite a few technical difficulties arise when we try to formalize this idea. Primarily, the difficulty is that the information conveyed by an imperative has to be integrated with the information which an agent already avails over. For instance, if I know that it rains and I am told that I should bring my umbrella if it rains, then the combined information should incite my to take my umbrella with me. Thus, the relation between an action of an agent and an imperative can be influenced by other information available to the agent.

In summary, I am going to construct a surrogate world, populate it with a couple of agents and then examine which constraints this world should obey if it were to mirror the real world; in particular, situations in the real world where one individual issues an imperative to another individual. The world together with the constraints to which this world is subjected will be called a model (of reality). As external observers of this model, we could choose to view such a model from an ‘eternal point of view’. In other words, all of the past and future in the model would lay bare open for us to inspect. This approach is, however, not very fruitful. In practice, we eventually want to compare our constructed model with reality and determine whether it is a satisfactory picture of it. In reality we have, however, no direct access to the future and even our information about the past may be incomplete. Thus if we want our model to be comparable with reality, it should also allow for partial information.

Two desiderata for the framework have emerged: it should enable us to model both partiality (of information) and Intentionality. In order satisfy the first requirement, the framework is based on a variant of partial predicate logic (cf. Muskens, 1989). Intentionality is dealt with by extending the standard notion of a model for a language. I propose to add a function to my models which maps specific objects of the model (representing information which an agent carries with it) to (sets of) formulae of the language. These formulae can then again be evaluated in the model as usual. Thus, certain objects in the model become Intentional: they are representing information about the model itself (this information can be true or false depending on how the translations into the language evaluate in the model).

The foundations of the framework are build by means of techniques from formal logic (specifically, model theory). In formal logic, models are specified to provide a semantics for a formal language. I also specify a formal language and its models. The models are used as a representation of a collection of communicating agents and their environment. More specifically, the models include a world and a set of constraints on this world. These constraints will be formulated in terms of the aforementioned formal language. In this respect, the formal language is used in a non-standard way[6].

So let us start by introducing the formal language (henceforth *L*). *L* consists of operators, individual constants (i.e., names), individual variables and predicate constants. We assume that there are various operators of the form $\simeq_C$, where the *C* is a meta-variable which ranges over sets of constraints (these are discussed in detail shortly). The operators will be interpreted as various forms of possibility. An operator $\Diamond_C$ takes a formula ϕ as its argument. $\Diamond_C \phi$ is read as ‘Relative to the constraints *C* it is possible that ϕ’.

We use a multi-sorted logic. This means that our models contain several domains. There are (individual) constants to name the individuals in most of these different domains. The constants are grouped according to the domains over which they range. For each of the domains, there is furthermore a sufficiently large stock of (individual) variables ranging over the domain:

(1) (Sorted) Individual Constants and Variables

1. AGENTS: We have logical constants for two agents. They are **A** and **B**. Furthermore there are variables a_1, a_2,... ranging over the domain of agents. Note that we use boldface for individual constants and italics for variables.

2. OBJECTS: We assume a stock of constants for objects such as **door1**, **window1, window2**, **car1**, etc. Furthermore there are the variables o_1,o_2,... ranging over objects.

3. EVENTUALITIES: There are no specifically named eventualities. We do, however, have the variables e_1,e_2,... ranging over eventualities. Our treatment of eventualities is along the lines of Parsons (1990); we assume that an eventuality is either a state or an event and some events can be seen as the actions of particular agents.

4. INSTANTS OF TIME: We have both constants and variables for instants of time. Examples of constants are **at_5**, **at_ midnight**, etc. Additionally, we have the variables t_1,t_2,... ranging over instants of time.

5. MOODS: In this domain there are two sentence moods: imperative and declarative. They are named by the constants **imp** and **decl**.

6. COMMITMENT PEGS: A commitment peg is basically a commitment of a particular agent which this agent acquired at a particular time. We call it a commitment peg because the content and the time of introduction of the commitment are attached to it as if it were a peg by means of the two-place predicates c_ time (the first argument is a commitment peg and the second one an instant of time) and c_ content (the first argument is a commitment peg and the second one a singleton set consisting of a formula representing the content of the commitment). For this domain we have only variables.

7. SETS OF FORMULAE: There are constants and variables available which range of sets of formula. The constants are given mnenomic names from

which the formula which they represent can be read off directly. In general, a set of formula {ϕ} is referred to by the constant »{ϕ}«. The available variables are: s_1,s_2,...

8. COMMITMENT CONTEXTS: Agents can be committed to sets of formulae. In a moment, we will explain in more detail what it means for an agent to be committed to some piece of information. For now, we want to discern between different types of commitments. In particular, an agent, say **A**, can be committed to a particular set of formulae {ϕ}, but s/he can also be committed to, for instance, {ϕ} being a joint or mutual commitment between **A** and **B**. Whereas in the former instance, we will say that {ϕ} inhabits **A**'s nil context (represented with **-**), in the latter instance, we speak of **A**'s **AB** context. We have a domain of such different commitment contexts and a name for each of them: **-**, **AB**, **AB-A**, **AB-B**, **AB-G(A)** and **AB-G(B)**. That is, we have, respectively, the nil context, the mutual commitment context, the context in which **A** and **B** are mutually committed that **A** has a particular commitment, a similar context for commitments of **B** and mutual commitments that **A** or **B** has a particular goal.

In addition to this stock of variables and individual constants we have a collection of predicate constants. For each predicate constant we indicate which types of variables and constants it can take as its arguments. We use the notation: predicate name (arg1:type of argument, arg2:type of argument, ...). The set of predicate constants is presented in groups of constants which more or less belong together.

(2) PREDICATE CONSTANTS

1. (TIME)< (arg1: time, arg2: time). This predicate is to be interpreted as temporal precedence/identity of instants of time. succ(arg1: time, arg2: time). succ(*t*,*t'*) means that the instant *t′ immediately* succeeds the instant *t*.

2. (EVENTUALITIES) agent (arg1: agent, arg2: eventuality). addressee (arg1: agent, arg2: eventuality). patient (arg1: agent/object, arg2: eventuality). occur_at (arg1: eventuality, arg2: time). hold (arg1: eventuality, arg2: time). say (arg1: eventuality). open (arg1: eventuality). be_open (arg1: eventuality). u_content (arg1: eventuality, arg2: set of formulae). u_type (arg1: even-

tuality, arg2: mood). state (arg1:eventuality). event (arg1: eventuality). The names of most constants should betray their intended interpretation. u_content and u_type stand for utterance content and utterance type, respectively.

3. (INTENTIONALITY) true (arg1: set of formulae). false (arg1: set of formulae). With these predicates we can express that a particular object from the domain which contains sets of formulae is true or false (i.e., the conjunction of the formulae is true or false). Earlier on we have called such objects intentional objects. The truth conditions of the predicates true and false make use of an (Intentionality) function from objects in the domain to propositions of the formal language *L*. The details are spelled out below.

4. (COMMITMENT) commit (arg1: agent, arg2: context, arg3: set of formulae, arg4: time). With this predicate we can express that a certain set of formulae can be derived from a particular subcontext of an agent's commitments. For instance, commit(**A**, **AB** , $\{\phi\}$, **at_5**) can be paraphrased as: 'from the subcontext which represents the commitments which **A** thinks to be joint commitments with **B** all the formulae in $\{\phi\}$ can be derived at the instant of time denoted by the constant **at_5**'.

   base_commit (arg1: agent, arg2: context, arg3: commitment peg, arg4:time). These are base commitments (as oppposed to derived or inferred commitments) which are tagged with the time at which they were introduced into the agent's base of commitments.

   c_content (arg1: commitment peg, arg2: set of formulae). c_time (arg1: commitment peg, arg2: time). With these predicates we can record the time and the content of a commitment. The latter predicate will not be used in this paper. We have include it to illustrate how we can model that the agent stores information about a commitment which is different from its content. We could introduce further predicates to store, for instance, details of the surface realization of the utterance with which the commitment was introduced.

5. (FORMULAE) $\subseteq$ (arg1:set of formulae,arg2:set of formulae). $\vdash$ (arg1: set of formulae, arg2: set of formulae). This is a derivation relation between sets

of formulae. singleton (arg1: set of formulae). agent_action (arg1: agent, arg2: set of formulae, arg3: time). This should be read as the set of formulae (which is stipulated to be a singleton) denotes an action of which the individual denoted by the first argument is the agent at a particular time (the third argument).

We define terms to be either individual constants or variables (we use the meta variables $T_1, ..., T_n$ for terms). The set of formulae is defined as follows:

(3) FORMULAE

1. If P is an n-ary predicate constant and $T_1, ..., T_n$ are terms of the correct sort, then $P(T_1, ..., T_n)$ is an atomic formula;

2. If T_1 and T_2 are terms of the same sort, then $T_1 = T_2$ is an atomic formula;

3. If ϕ and ψ are formulae, then $\neg\phi$, $\phi \wedge \psi$, $\phi \vee \psi$, and $\phi \rightarrow \psi$ are formulae;

4. If ϕ is a formula and X is an individual variable, then $\forall X.\phi$ is a formula;

5. If S and S' are variables ranging over the sort 'set of formulae' and ϕ is a formula, then $S = \bigcup S' : \phi$ is a formula;

6. If ϕ is a formula and C_i denotes a set of constraints, then $\Diamond_{Ci}\phi$ is a formula.

The truth conditions of these formula are defined with respect to their intended models. Formally, a model is defined as follows:

(4) CONSTRAINT-BASED PARTIAL INTENTIONAL MODEL

A model is a $\langle W, C \rangle$, consisting of a (possibly) partial world W and a list of sets of constraints $C = C_1, ..., C_n$ on W (constraints are expressed with formulae of L). A world W is a quadruple $\langle D, I^+, I^-, Int \rangle$. It consists of a list of domains $D = D_1, ..., D_n$, a positive and a negative interpretation function I^+ and I^-, and an intentionality function Int.

Standard models for predicate logic are tuples $\langle D, I \rangle$, consisting of a domain and an interpretation function. The interpretation function maps individual constants to members of D and n-ary predicate constants to subsets of D^n. For instance, a constant **john** is mapped to the corresponding individual in the domain and a two-place predicate constant such as *love* is mapped to a subset of $D^2 = D \times D$. In other words, *love* is mapped to a set consisting of tuples of individuals, such that the first individual of the tuple stands in the relation *love* to the second one. In our models, there are two interpretation functions (I^+ and I^-) instead of one. We stipulate that for individual constants I^+ and I^- are identical. In order to bring partiality into our models, the interpretation functions differ for predicate constants. Consider again the predicate constant *love*. In ordinary models, the interpretation function gives us the pairs of individuals which stand in this relation. Any pair of individuals which is not returned by the interpretation function is automatically classified as not standing in the relation. It is impossible to express lack of information with regard to whether a particular pair of individuals stands in the relation or not. By introducing two interpretation functions, partiality of information becomes expressible: a pair of individuals can be standing in the relation (i.e., it is a member of I^+(*love*)), not stand in the relation (i.e., it is a member of I^-(*love*)) or there is no information about whether the relation holds between the individuals or not (i.e., it is neither a member of I^+(*love*) nor of I^-(*love*))[7].

Another non-standard ingredient of our models is the (Intentionality) function *Int*. It is a function from a subdomain D_k of D to sets of formulae of the language L. Note that within this framework we can express an analogue of the Liar sentence, i.e., 'This sentence is false'. We can form a formula which says that the singleton set consisting of that formula (denoted by a constant from D_k) is false: false(**c**) (where **c** denotes {false(**c**)}). In our framework, it is not possible to show that this formula is true or false with respect to a model (it can be shown that both attempts to construct a proof of truth and falsity lead to an infinite recursion). Since we use a partial logic, this does, however, not lead to a paradox: rather the formula comes out as undefined. Intuititively, the introduction of *Int* means that the members of D_k can carry (true or false) information about the model itself.

We have given a general definition of partial Intentional models. Let me now define a particular subclass of these models, that is, communicating agents models:

(5) COMMUNICATING AGENTS MODEL

A communicating agents model *M* is a Constraint-based Partial Intentional Model $\langle\langle D, I^+, I^-, Int\rangle, C\rangle$ such that:

1. $D = D_a, D_o, D_e, D_t, D_m, D_{cp}, D_{sf}, D_{cc}$. We have domains for, respectively, agents, objects, eventualities, instants of time, moods, commitment pegs, sets of formulae and commitment contexts;

2. $Int: D_{sf} \mapsto$ Formulae of L;
3. $C = C_{utt}, C_{comm}, C_{env}, C_{time}, C_{action}$. We have sets of constraints pertaining to, respectively, utterances, the environment, the temporal structure of reality and actions by the agents.

Before we can finally give our semantics for *L*, we need to introduce some further notions. Our partial models come with a natural structure, that of *informational subsumption*; we write $\sqsubseteq$. Given two models M_1 and M_2 such that $C_1 = C_2$, we say that $M_1 \sqsubseteq M_2$ (read: M_2 informationally subsumes M_1) if and only if for both $p = +$ and $p = -$: $I_1^p(c) = I_2^p(c)$ for all individual constants *c* and $I_1^p(c) \subseteq I_2^p(c)$ for all predicate constants *c*. A model is called *total* if there is no other model which informationally subsumes it. Furthermore, we write $M_{<t}$ for models which are undefined for precisely those n-tuples such that the n-th member is an instant of time and this instant is bigger than *t*. In other words, these are models which are fully defined up till time *t*.

We use the usual notion of an assignment *a* given a model *M*. Such an assignment maps variables of a particular sort *z* to members of the corresponding domain D_z of *D*. An assignment $a[d/x]$ is defined as being identical to the assignment *a* except for assigning *d* to *x*. The value of a term *T* with respect a model *M* and an assignment *a* (written $\|T\|^{M,a}$, and abbreviated as ||T||) is $I^+(T)$ if T is a constant and $a(T)$ if *T* is a variable.

(6) SEMANTICS: TARSKI TRUTH DEFINITION

1. $M \models P(T_1,...T_n)\,[a]$ iff $<\|T_1\|,...\|T_n\|> \in I^+(P)$
$M \dashv P(T_1,...T_n)\,[a]$ iff $<\|T_1\|,...\|T_n\|> \in I^-(P)$

(where *P* is not equal to true or false)

2. $M \models T_1 = T_2\,[a]$ iff $\|T_1\| = \|T_2\|$
$M \Dashv T_1 = T_2\,[a]$ iff $\|T_1\| \neq \|T_2\|$

3. $M \models \neg\phi[a]$ iff $M \Dashv \phi[a]$
$M \Dashv \neg\phi[a]$ iff $M \models \phi[a]$;

$M \models \phi \wedge \psi[a]$ iff $M \models \phi[a]$ and $M \models \psi[a]$
$M \Dashv \phi \wedge \psi[a]$ iff $M \Dashv \phi[a]$ or $M \Dashv \psi[a]$;

$M \models \phi \vee \psi[a]$ iff $M \models \phi[a]$ or $M \models \psi[a]$
$M \Dashv \phi \vee \psi[a]$ iff $M \Dashv \phi[a]$ and $M \Dashv \psi[a]$;

$M \models \phi \rightarrow \psi[a]$ iff $M \Dashv \phi[a]$ or $M \models \psi[a]$
$M \Dashv \phi \rightarrow \psi[a]$ iff $M \Dashv \phi[a]$ and $M \Dashv \psi[a]$;

4. $M \models \forall X.\phi[a]$ iff $M \models \phi[a[d/x]]$ for all $d \in D_{sort(X)}$
$M \Dashv \forall X.\phi[a]$ iff $M \Dashv \phi[a[d/x]]$ for some $d \in D_{sort(X)}$

5. $M \models \mathit{true}\,(S)[a]$ iff $M \models \bigwedge(\mathit{Int}\,(S))[a]$
$M \Dashv \mathit{true}\,(S)[a]$ iff $M \Dashv \bigwedge(\mathit{Int}\,(S))[a]$;

$M \models \mathit{false}\,(S)[a]$ iff $M \Dashv \bigwedge(\mathit{Int}\,(S))[a]$
$M \Dashv \mathit{false}(S)[a]$ iff $M \models \bigwedge(\mathit{Int}\,(S))[a]$;

6. $M \models S_1 = \bigcup S_2{:}\ \phi[a]$ iff $\|S_1\| = \bigcup \|S_2\|{:}\ M \models \phi[a]$
$M \models S_1 = \cup S_2{:}\ \phi[a]$ iff $\|S_1\| = \cup \|S_2\|{:}\ M \models \phi[a]$

7. $M \models \Diamond_{Ci}\phi\,[a]$ iff there is an M' such that $M \sqsubset M'$, M is total, $M \models \bigwedge C_i[a]$ and $M' \Dashv \phi\,[a]$

$M \Dashv \Diamond_{Ci}\phi\,[a]$ iff there is no M' such that $M \sqsubset M'$, M is total, $M \models \bigwedge C_i[a]$ and $M' \Dashv \phi\,[a]$

We read '$M \models / \Dashv \phi\,[a]$' as ϕ is true/false in model M under assignment a. The clauses 1., 2., 3. and 4. are along the lines of those of Muskens' (1989:49) partial predicate logic. Clause 5. makes essential use of our Intentionality function. For instance, the first item of this clause says that the formula *true* (S) (where S is a term denoting a set of formulae) is true iff the conjunction of the members of Int(S) is true. Int(S) stands for the set of formulae which is denoted by the term S. Clause 6. allows us to construct a set which is the union of a set of formulae which have a particular property ϕ. Finally, according to the first item of clause 7., ϕ is possible given a model M and a set of constraints C_i (and an assignment) if and only if there is model M' which properly subsumes M ($M \sqsubset M'$), is total, makes the conjunction of the constraints which are a member of C_i true and, finally, makes ϕ itself true.

Our semantics for possibility is similar to the semantics which Landman (1986:53) assigns to the word 'may'. Landman's definition is, however, given for propositional logic and does not relativize possibility to a set of constraints. In particular, our proposal to make the constraints part of the model is different from Landman's treatment. Furthermore, Landman argues that the definition requires some refinements to deal with all the correct inferential patterns in which the word 'may' can occur. For our purposes, such refinements would, however, unnecessarily complicate our treatment of possibility.

3 Theory

In this section, I specify the sets of constraints C_{utt}, C_{comm}, C_{env}, C_{time} and C_{action} which feature in our models. For reasons of space it will be impossible to provide a complete formal version of each and every constraint which I use. The ones which are most pertinent to the problems (P.1) and (P.2) are, however, spelled out in detail. I hope that this section provides the reader with an idea of how to use the framework which has been introduced in the previous section to formulate concrete theories of communicating agents.

Time and Eventualities Let me start with C_{time}. I assume that the time line is a discrete linear order which is infinite in both directions. We stipulate that events are seen as *transitions* between instants of time. Given an event *e* and an instant of time *t*, we write *occur_at*(*e*,*t*) to say that the event *e* took place between *t* and the instant of time which immediately succeeds *t*. Furthermore, we assume that each event is associated with at most one transition between two instants of time:

(7) $$\forall e.((\exists t.occur\_at(e,t) \wedge event(e)) \\ \rightarrow (\forall t'.occur\_at(e,t') \rightarrow t = t'))$$

Whereas we have the predicate *occur_at* for events, we have the predicate *hold* for states. A state can hold at several instants of time. However, these instants have to be a connected series. In other words, if a state occurs at two times, then there exists no time in between those times at which it does not occur:

(8) $$\forall e.\forall t.\forall t'.((hold(e,t) \wedge hold(e,t') \wedge state(e)) \\ \rightarrow \forall t''.(t \leq t'' \leq t' \rightarrow hold(e,t'')))$$

Commitment Let us proceed with C_{comm}. At each instant of time an agent has for each of its commitment contexts a (possibly empty) set of base commitments. For each base commitment, the agent also has information about the time at which that commitment was taken on. These base commitments give rise to a set of derived commitments. Basically, a set of formulae *s* is a derived commitment for an agent *a* at some given instant of time *t* iff *s* can be derived from the union of the base commitments which the agent maintains at that moment of time[8].

(9) $$\forall a.\forall c.\forall s.\forall t.commit(a,c,s,t) \rightarrow \\ \exists s'.(s' = \bigcup s'' : \exists p.base\_commit(a,c,p,t) \wedge c\_content(p,s''))) \wedge \vdash(s',s))$$

It is beyond the scope of this paper to provide any constraints for the relation '$\vdash$' which relates direct commitments to indirect (inferrable) ones. It is assumed that it corresponds to some (computable) relation of derivation between sets of formulae. I already pointed out that an agent has several different commitment contexts. There are, however, constraints which relate the content of these contexts to each other. Again, I will only provide an example of such a constraint (cf. Zeevat, 1997):

(10) $\forall a.\forall s.\forall t.(((a{=}\mathbf{A} \vee a{=}\mathbf{B}) \wedge commit(a, \mathbf{AB\text{-}A}, s, t)$
$\wedge\ commit(a, \mathbf{AB\text{-}B}, s, t)) \rightarrow commit(a, \mathbf{AB}, s, t))$

This constraint says that if both the contexts **AB-A** and **AB-B** of an agent *a* (**A** or **B**) contain some commitment *s*, then the context **AB** contains this commitment as well. In other words, if an agent thinks that it is a mutual commitment that **A** is committed to *s* and it is a mutual commitment that **B** is committed to *s*, then this agent (assuming that the agent is **A** or **B**) thinks that *s* itself is a mutual commitment.

I have used the notion of commitment without first providing a definition. I have relied on the reader's intuitive understanding of this notion. The definition is given implicitly in the course of this paper by means of the constraints which are imposed on commitments. We have seen some constraints which relate different types of commitments with each other. In a moment, we provide further constraints which explicate the role of commitments in the behaviour of an agent.

Utterances An agent can acquire new commitments through observation and through communication with other agents[9].In the real world, an agent can also change his or her mind and retract a commitment. For our purposes, this would, however, introduce complications which detract from the main issues which this paper addresses. Therefore, I assume that once an agent has acquired a commitment, that commitment will persist through time.

New commitments can be introduced through communication. The constraints in C_{utt} regulate the relation between utterance events and the introduction of new commitments. We assume that the agents cannot simultaneously produce an utterance. For that purpose we have a constraint (which we will not spell out in formal detail) which says that only one utterance event can occur at an instance of time. This is, of course, a simplification, and at a later time we might want to relax this constraint. The relation between the utterance of an imperative and the commitments of the addressee of an imperative are spelt out by the following constraint:

(11) $\forall e.\forall t.\forall t'.\forall s.(agent(\mathbf{A},e) \wedge say(e) \wedge addressee(\mathbf{B},e) \wedge$
$occur\_at(e,t) \wedge u\_content(e,s) \wedge u\_type(e,\mathbf{imp}) \wedge succ(t,t') \rightarrow$
$\exists p.(base\_commit(\mathbf{B},\mathbf{AB\text{-}A},p,t') \wedge c\_time(p,t)))$

Roughly speaking, according to this constraint if agent **A** utters an imperative to **B** with content *s* between *t* and *t* +1, then **B** updates his base commitments in the context **AB-G(A)** with *s* at time *t* +1. In other words, an imperative with content ϕ causes the speaker to think that it is now a mutual commitment that ϕ is a goal of the speaker. There is a further constraint which says the same for the **B**'s context **AB-G(A)** and two further constraints which apply when **B** is the speaker instead of **A**. Furthermore, we have a constraint which deals with declaratives. It says that if the content of the declarative is *s* and the speaker is **A**, then both speaker and addressee (**B**) add *s* to their **AB-A** context. In other words, they now both think that it is a mutual commitment that **A** is committed to *s*.

Of course, these constraints are idealizations. We have not taken into account situations which involve miscommunication. Furthermore, the relation which we have posited between sentence mood and the update of the speaker's and addressee's commitments does not take into account indirect speech acts (Searle, 1975). We trust that such more elaborate theories can be formulated within the framework which is presented here. However, such a theory is at the moment not our main concern (but see, e.g., Beun (1994) for a more elaborate account of the relation between utterances and intentional agent states).

Typically, a commitment is taken on through a declarative utterance. For instance, in 'A: Open the door. B: Ok (I will)', B's utterance is taken to be such a declarative. According to an earlier mentioned constraint, after the utterance, A and B each think that it is a mutual commitment that B is committed to opening the door. Let us assume that, as described, the addressee of an imperative signals that s/he will comply with the imperative. Thereby the agent agrees that it is a mutual commitment (amongst speaker and addressee) that the agent is committed to content of the imperative. So according to our analysis so far the function of an imperative is to induce such a commitment. The next question then is what such a commitment amounts to. I want to argue that a proper answer to this question requires us to look from two different perspectives at such a commitment. (**I**) On the one hand, we can ask whether the content of the commitment is true, or can still become true in the world. (**II**) On the other hand, we have an agent who's actions are influenced by the commitment, ideally, in such a way that the content of the commitment does become true. In other words, agents execute a certain policy in order to make sure that their commitments satisfied in the truth conditional sense. However, my claim is that satisfaction of a commitment is not judged purely in terms of truth conditional satisfaction but also in terms further constraints on the policy which led to that satisfaction.

I want to argue that the two problems which are discussed in the introduction of this paper arise from not properly taking into account the second perspective on the satisfaction of imperatives. Consider again the problem (P.1). Suppose an imperative with the content $\phi \rightarrow \psi$ is issued to an agent. If the agent subsequently and as a *result* of this imperative goes about making sure that *i* is false, then that agent does not act in the spirit of the imperative. And yet, according to perspective (I) there is nothing wrong with such behaviour: we can speak of the satisfaction of an imperative if its content is true in the world. In order to explain the infelicity of the agent's actions, we need to bring perspective (II) into play. I want to argue that agents are expected to go about satisfying imperatives only in line with policies for action of a certain type. In particular, an adequate policy should exploit the fact that conditionals of the form 'If such and such a state holds or event occurs, then do such and such' can be seen a *rules* for guiding the behaviour of an agent in a given situation. The picture emerges of an agent moves about in the world, maintaining a clock which indicates the time and checking whether at the current time there are any actions which s/he needs to execute in order to satisfy his or her commitments.

Actions, Commitment and Environment In our framework, the triggering of a rule relative to the other commitments which an agent maintains can be formalized in terms of the (logical) derivability of the consequent of the rule. A rule, i.e., conditional commitment $\phi \rightarrow \psi$, is triggered if the agent's commitments, which include the conditional commitment, allow the agent to derive the consequent ψ of the conditional commitment. More generally, we have the following policy which relates an agent's actions to his or her commitments: At every instant of time *t*, the agent checks which actions (of him or herself) at time *t* can be derived from the commitments. Precisely these actions, s/he then carries out. We can formulate this as a constraint by saying that for all times if an agent is committed (directly or through a derived commitment) to some action of him- or herself at that specific instance of time, then this action is carried out by her or him:

(12) $\forall a.\forall s.\forall t.((commit(a,\text{-},s,t) \wedge agent\_action(a,s,t)) \rightarrow true(s))$

Now compare this with the following constraint which is in the spirit of perspective (I) on imperatives and commitment:

(13) $\forall a.\forall s.\forall t.((commit(a,\text{-},s,t) \rightarrow true(s))$

In words, all commitments of an agent should be true in the world. Notice that constraint (13) is stronger than -i.e., entails- (12). The former requires all commitments to come true, not just those pertaining to actions of the agent[10].

Worlds which satisfy constraint (13) can be seen as ideal worlds, whereas the constraint (12) can be seen as a basis of a policy for an agent's behaviour to end up in such a world. In order for an agent to approximate (13) through such a policy we need to invoke additional assumptions about the behaviour of that agent. Let us start with a very simple model which requires almost no further assumptions about the agent's behaviour. As we progress to more complex models, the number of assumptions will increase. We assume that in this world C_{env} is empty: there are no interactions between states and events of the environment (i.e., no pre- and postconditions on events and no situations in which one event is part of another event). Furthermore, we assume that the agent's commitments are conjunctions of atomic formulae. These formulae contain no variables. All reference to objects, agents, times, eventualities, etc. are achieved by means of individual constants. Such a set of commitments will contain a subset which denote actions by the agent. The remainder will be actions by other agents, events and states. Let us assume that the latter come true by definition[11]. Commitments pertaining to the agent's own actions will come true if s/he simply carries out each and every one of them. Let us add negation to this model. Now, commitments are conjunctions of possibly negated atomic formulae. In this new setup, we have to extend the agent's policy with the following clause: s/he refrains from an action if its negation follows from his or her commitments.

Let us now move to a slightly more complicated model. Assume that commitments can also be conditionals (although quantification is still not permitted). Consider a conditional commitment of the form $\phi \rightarrow \psi$, where ψ denotes an action by the agent. Assume that this is the only commitment of the agent. It is evident that the policy of simply carrying out all actions which directly follow the commitments will no longer guarantee that is (13) satisfied. For instance, there could be a situation where ϕ is true, without the agent being committed to ϕ. In that situation, the agent ought to carry out ψ in order to make $\phi \rightarrow \psi$ true. But since the agent does not avail over the information p, s/he will do no such thing. In this new set up, an agent will have to actively be on the look-out for information which can trigger conditional commitments. A possibility might be that the agent tries to check all the atomic subformulae of his or her commitments (which can not be derived from his or her commitments) for their truth or falsity. The thus obtained information could then be added to the agent's commitments. Given such a set of commitments, the agent's policy would again be complete with respect to constraint (13).

Finally, let us assume that commitments can also be quantified formulae. Now, checking for the truth of subformulae will no longer do, since they can contain unbound variables. A complete policy could be attained by checking for the truth of such formulae under all possible substitions for their variables. For any realistic domain (i.e., with a sufficiently large domain of individuals), this would, however, not be a feasible policy. We might raise the question how human agents solve this problem. Clearly they have no fool proof policy. It would be interesting to investigate whether their actual policy approximates the aforementioned one in any way.

(14) (EXAMPLE) The content of the imperative 'Say hello to John if you see him' (issued to **A**) is formalized as:

$$(\forall t.\forall e.(see(e) \wedge agent(\mathbf{A},e) \wedge patient(\mathbf{j},e) \wedge occur\_at(e,t) \wedge \leq(\mathbf{utterance\_time},t)) \rightarrow(\exists e'.say\_hello(e) \wedge agent(a,e') \wedge patient(\mathbf{j},e') \wedge occur\_at(e',t)))$$

I now want to argue that the approach I have sketched also provides a solution to the problem (P.2). (P.2) arises from Ross' (1941) observation that an imperative with the content ϕ does not seem to entail an imperative with the content $\phi \vee \psi$. Hence, imperatives do not appear to conform to the standard laws for logical inference (e.g., in FOL). I agree with Ross' intuitions that an addressee who has been told to 'burn a letter', should not, subsequently, infer that s/he should 'burn or send the letter'. Let me first describe an approach (within my framework) which would lead to such a paradoxical conclusion, and then indicate how the approach I actually advocate solves the problem. According to perspective (I), an agent should make sure that all his commitments come true. How an agent actually goes about making sure that this really happens belongs to the realm of perspective (II). One possible policy is that an agent simply takes a commitment which follows from his commitments and tries to make it true, and then proceeds to make the next commitment which follows true, until all her or his commitments are true.

Now consider an agent which is only committed to ϕ (where ϕ describes an action by the agent and might have been acquired as a result of the result of imperative). The agent might first try to make a commitment true which follows from ϕ, e.g., $\phi \vee \psi$. In particular, the agent might choose to make ψ true. That would, however, be a bad policy, since there might be a logical connection between ϕ and ψ, e.g., it might be the case that ϕ and ψ cannot be true at the same time (one cannot burn a letter and send it at the same time). In that situation, making ψ true prevents the agent from making ϕ true. The policy I propose cannot lead to such an impasse since it requires an agent to only make true those commitments which describe primitive actions by the agent and which can be derived from his or her base commitments. Now, clearly we cannot derive the action description ψ from ϕ. Thus, our solution to the problem is to maintain the logical connection between ϕ and $\phi \rightarrow \psi$, but provide a new account of the role of commitments in the behaviour of an agent.

Up till this point I have assumed that the environment in which the agents operate is not subject to any interactions. For instance, there are no pre- and postconditions on actions. Adding these introduces a further complication. In particular, if we want to evaluate an agent's behaviour at a specific point in time *t* with respect to the commitment s/he has taken on. We then have to evaluate this agent's behaviour with respect to a world $W_{\leq t}$ which is specified only up till *t*. Since some of the commitments might be about future actions of the agent, we can no longer simply demand that all his commitment should be true in $W_{\leq t}$. Here, we need to take recourse to the concept of possibility:

(15) $\forall a.\forall s.\forall t.(commit(a,\text{-},s,t) \rightarrow \Diamond_{Cenv} true(s))$

Whether an agent behaviour satisfies his commitments at a given time *t* now depends on whether it is still *possible* (with respect to the constraints to which the environment is subjected) to make all his or her commitments come true.

4 Conclusions

A framework for communicating agents has been introduced. This framework combines and extends various techniques from formal logic. In particular, the framework allows for both Intentional states of agents and Partial models of reality. Furthermore, constraints are taken to be explicit ingredients of our models. They regulate possible extensions of partial descriptions of reality. Within the framework, we can model agents who carry commitments around which they took on as a result of imperatives. I provide an outline of several, increasingly more complex, models of how these commitments should influence the behaviour of an agent (so-called *policies*) for him/her to be said to comply/satisfy the imperative which gave rise to the commitment. The policy which I propose provides a solution to the problems (P.1) and (P.2) which were presented in the introduction of this paper.

Although I have attempted to achieve a certain level of formal rigour in my analysis, there are still many loose ends which require further development. For instance, the proposed model does not explicitly deal with the interpretation process; the question of how an agent arrives at a particular commitment given the surface form of an (imperative) utterance. The context dependence of this process will have to be incorporated into the model for phenomena such as indirect speech acts (cf. Power, 1979; Beun, 1994) and anaphora resolution (e.g., Ahn, 2000; Kamp & Reyle, 1993; Krahmer, 1998; Piwek, 1998; Van Deemter, 1991)[12]. Furthermore, the proposed model of time and eventualities makes some strong assumptions (e.g., the assumption that events are instantaneous) which cannot be maintained in the long run. A further issue which requires more discussion is the relation of the proposed model to planning theories of discourse (e.g., Lochbaum, 1994). Lochbaum describes how discourse can lead to complex plans (although, not logically complex, as the commitments described in this paper, but rather complex in the sense of involving hierarchical structures of subplans and actions). The focus of this paper is complementary to that work: I have focussed on how an agent's behaviour is influenced given that s/he has adopted a set of (logically) complex commitments (such a set of commitments can be seen as a partial plan). Despite all these shortcomings, I hope that the work provides a formal basis for an agent-based analysis of linguistic phenomena[13], such as, imperatives.

References

Ahn, R. (2000). *Agents, Objects and Information: Modelling the Dynamics of Interaction* (working title), PhD dissertation, Tilburg University (To appear).

Beun, R.J. (1994). 'Mental state recognition and communicative effects', *Journal of Pragmatics*, 21, 191—214.

Biber, D. S. Johansson, G. Leech, S. Conrad and E. Finegan (1999). *Longman Grammar of Spoken and Written English*, Harlow: Longman.

Hamblin, C. (1987). *Imperatives*, Oxford: Blackwell.

Hausser, R. (1999). *Foundations of Computational Linguistics*, Berlin: Springer Verlag.

Kamp, H. & Reyle, U. (1993). *From Discourse to Logic*, Dordrecht: Kluwer Academic Publishers.
Krahmer, E. (1998), *Presupposition and Anaphora*, Stanford: CSLI Publications and Cambridge University Press, CSLI Lecture Notes Series 89.

Landman, F. (1986). *Towards a Theory of Information*, Dordrecht: Foris Publications.

Lochbaum, K. (1994) *Using Collaborative Plans to Model the Intentional Structure of Discourse*, PhD. Dissertation, Harvard University

Muskens, R. (1989). *Meaning and Partiality*, PhD. Dissertation, University of Amsterdam.

Muskens, R. (1995). 'Order-Independence and Underspecification', In: J. Groenendijk (ed)., *Ellipsis, Underspecification, Events and More in Dynamic Semantics*, DYANA Deliverable R.2.2.C.

Parsons, T. (1990). *Events in the Semantics of English*, Cambridge, Massachusetts: MIT Press.

Piwek, P. (1998). *Logic, Information and Conversation*, PhD. Dissertation, Eindhoven University of Technology.

Power, R. (1979). 'The Organisation of Purposeful Dialogues'. *Linguistics*, 17, 107-152.

Ross, A. (1941). 'Imperatives and Logic', *Theoria*, 7, 53-71.

Searle, J. (1975), 'Indirect Speech Acts', In: Cole, P. & J. Morgan (eds), *Syntax and Semantics 3 (Speech Acts)*. New York: Academic Press, 59-82.

Searle, J. (1983). *Intentionality*, Cambridge: Cambridge University Press.

Van Deemter, C. (1991). *On the Composition of Meaning*, PhD. Dissertation, University of Amsterdam.

Zeevat, H. (1997). 'The Common Ground as a Dialogue Parameter'. In: Benz, A. & G. Jäger (eds), *Proceedings of MunDial '97*, CIS-Bericht 97-106, Department of Computational Linguistics, University of Munich, 195-214.

ENDNOTES

* I could not have written this paper without first having been introduced to and taught about natural language pragmatics by Robbert-Jan Beun and Harry Bunt. I am particulary indebted to Robbert-Jan Beun who took care of the day to day supervision when I was writing my PhD dissertation. I am grateful that never ceased to urge me to start my studies into natural language semantics and pragmatics from an overall view on communicating agents. I have tried to write this paper in that spirit.

1 I use the term *discourse* to refer to both written and spoken discourse. This paper focuses on spoken discourse, i.e., dialogue. Biber et al. (1999:221) found that imperatives are most frequent in spoken discourse.

2 This translation avoids several issues, such as the interpretation of pronouns, which are beyond the scope of this paper. Furthermore, the example has been formalized without any explicit reference to time. Temporal aspects of imperatives are, however, dealt with later on this paper.

3 By describing the agents in sufficient formal detail, they might also be realized as software agents. This could facilitate the evaluation of the model. However, the currently described model has not yet been implemented.

4 Following Searle (1983), I assume that such states need not be conscious and that they are not tied to the verb 'intend'. When I intend to do something, I am in an intentional state, but there are many other types of intentional states such as belief, desire, love and hate.

5 I am referring to the state which ensues if the addressee *accepts* the imperative. For the moment, let us forget about situations where the addressee refuses to comply with an imperative.

6 Normally, constraints on models are formulated as axioms which are independent of the models. For our purposes, it is essential that the constraints are part of the models. This

allows us to capture in a model the possible extensions of a partial world (roughly speaking, the total extensions of the world in which all the constraints are true). We can then provide an alternative definition of 'possibility', given a model, as as truth in at least one of the possible extensions of the partial world. Possibility is a concept which will be needed in the analysis of imperatives.

7 We exclude overdefinedness, that is, situations where a (sequence) of individuals is both a member of $I^+(R)$ and $I^-(R)$ for some relation R. Our intention is to use our models as representations of reality. Since reality is supposed to be consistent, we require our models to be so as well.

8 For the sake of uniformity it would have been desirable to represent (base) commitment as any of the other states which can hold at a given instant of time. For reasons of conciseness, I have, however, chosen for a more compact notation.

9 Inferred or derived commitments are not taken to be new commitments.

10 The latter is expressed by the subformula *agent_action*(*a*,*s*,*t*). We can see this subformula as a partial description of the object *s*, which itself again stands for a set of formula. Currently, we assume that in our models elements of the domain D_{sf} (of sets of formulae) either satisfy such a predicate or not and that this corresponds with the syntactic structure of the formula. This idea could be worked out more explicitly by using a tree description logic, where formulae are explicitly modelled as trees instead of primitive objects of the domain D_{sf} (see, e.g., Muskens, 1995).

11 Roughly speaking, these commitments correspond to the agent's beliefs. The agent does not actively make sure that they come true, although s/he should not perform actions which make them come out false.

12 Extending the current framework in that direction could be realized a move from Predicate Logic to Discourse Representation Theory (Kamp & Reyle, 1993). We would need to change the domain D_{sf} to contain objects representing Discourse Representation Structures, instead of sets of Predicate Logical Formulae. The Intentionality function *Int* would then be defined as a truth preserving mapping of these structures into sets of Predicate Logical formulae of language *L* of the framework.

13 Such an approach is argued for at length in, for instance, Hausser (1999).

Perspectives and Derived Extensions of Dialogue Acts

Anton Benz

Abstract

Wir betrachten die Rolle, die die Perspektiven der Dialogteilnehmer bei der Erweiterung des Gebrauchs von Dialog-Akten spielen. Wir konzentrieren uns auf Beispiele des referentiellen Gebrauchs von Kennzeichnungen und von Deklarativsätzen. Wenn wir eine Klasse von Dialogsituationen bestimmt haben, die für den jeweiligen Gebrauch prototypisch ist, dann erlaubt es die beschränkte Information — oder Desinformation — der Teilnehmer den Gebrauch systematisch auf neue Situationen zu übertragen, wobei der erweiterte Gebrauch allerdings zu anderen Interpretationen führen kann.

We consider the role of perspectives of dialogue participants for the extension of dialogue acts to new situations. We concentrate on the examples of the referential use of definite descriptions, and on the utterance of sentences in declarative mood. If we have characterised a class of prototypical situations where such a specific act can be performed, then the limited information — or misinformation — of the participants allow to derive systematically new situations where the same act can be performed, but might be interpreted differently.

1 Introduction

We investigate how the different perspectives of dialogue participants give rise to derived uses of already given dialogue acts. We understand by *perspective of a participant* the information the participants has about the dialogue situation, including e.g. the situation talked about and the beliefs of other participants. Let's consider the following prototypical example of an assertion *I hold an ace*.

> **(1)** *S* and *H* sit at a table playing with cards. *A* notices that he has an ace, and says: „I hold an ace.“

The assertion is, of course, a true assertion. Now, consider the slightly different situation:

> **(2)** Assume now that S has no ace but mistakes a joker for an ace. Again he says: „I hold an ace.“

This is still an assertion but no more a true assertion. S is justified to make this assertion because *out of his perspective* it is still a true assertion.

(3) Now assume *S* knows that he does not hold an ace but knows that *H* can't know this. He says: „I hold an ace."

Now, it is not only a false assertion but a lie. *S* must believe that *out of H's perspective* it is still a true assertion.

(4) Assume that in situations **(2)** and **(3)** *H* can see through a mirror the playing cards of *S*.

In this case, *H* can recognise that *S* was lying to her. It is now important for us to notice that she can only understand *S's* utterance as a lie because she knows that he must believe that she can't know that it is one. In the following example it is common knowledge that *I hold an ace* is obviously not true:

(5) Both players have only one card. *A* puts his card face up on the table. It is a king. Hence, both can clearly recognise that they mutually must know that it is a king. Then *A* says: „I hold an ace."

Without special context, the sense of the utterance of *S* must be unintelligible for *H*. It can't be a false assertion, nor a lie. This shows that the knowledge of the participants about their beliefs and their mutual knowledge plays a central role in the interpretation of sentences.

A similar phenomenon can be found in the case of the referential use of definite descriptions. In (Benz, 1999) we could show how the systematic exploitation of perspectives provides extended referential uses and interpretations of definites for defective states, i.e. states where prima facie the conditions for a successful use are violated. We started with a felicity condition for *basic* dialogue situations, i.e. situations where both participants have only true beliefs, and where this is mutual knowledge:

(6) There are two playing cards c_1 and c_2 lying face down down, side by side on a table. *A* and *B* can see them both and that they mutually see them. Then a supervisor turns the first card, c_1, around, so that both can see that it is an ace. And this will be, of course, common knowledge.
Now, says to *B*: „Please, point to *the ace*."

Then we extended this use and the interpretation to defective situations by use of four principles connected to the following types of situations:

- The speaker can believe that the felicity condition holds.
- The hearer can believe that the felicity condition might hold.
- The speaker can believe that the hearer believes that the felicity condition might hold.
- The hearer can believe that the speaker might believe that the felicity condition holds.

We could explain why in the following example the use of the definite description *the man with the red walking stick* is successful:

(7) There is a couple, *A* and *B*, sitting in the park. In some distance there are two men walking. One of them has a red umbrella. *A* thinks that he can see that it is a red walking stick. He believes that *B* would not be able to say what exactly the man carries with him, because she is somewhat short-sighted. Then he remembers that he knows the man.

A: Look there, *the man with the red walking stick*. Yesterday I had a game of chess with him.
B: Oh, really. I know him too. We talked together just before you met me. I saw that he does not have a walking stick but a very slim red umbrella.

Now, we can show that the same mechanism can be applied to assertions. This gives us an outline of a theory for how in general perspectives give rise to derived extensions of already given dialogue acts. This mechanism is especially interesting in dialogue theory as it allows us to explain the felicity and effects of uses in defective situations.

In Section 2 we will consider related approaches which try to explain the relation between the different uses of declarative sentences as in the examples **(1)-(4)** within a theory of rational behaviour. In Section 4 we outline the basic ideas of our approach, and define four operations which allow us to derive extended uses of dialogue acts. Then, in Section 5, we apply our theory to examples. In an appendix we provide a short overview of the basics of the so-called possibility approach, which we adopt for our representations of dialogue situations.

2 Rationally Based Speech Act Theories with Default Rules

The problems around the dependency of closely related speech acts, e.g. *assertions, lies, unsuccessful lies etc*., have been discussed e.g. in the papers (Cohen & Levesque 1985,1990a, b; Perrault, 1990). Their principal idea is as follows: They postulate a correlation between the uttering of a sentence with a certain syntactic feature (the locutionary act, e.g. imperative, declarative) in a certain context (specified by *gating conditions*, and a complex propositional attitude expressing the speaker's mental state. The speaker's uttering of a sentence under those gating conditions results in the hearer's beliefs that the speaker has the corresponding attitudes. Then, general principles governing mental states allow to derive other consequences of the speaker's having the expressed state.

> If *gating-condition* holds and locutionary act *X* is performed, then *consequence—condition* holds.

The general axioms which allow to derive further consequences describe general properties of co-operative agents (like *sincerity*, *helpfulness*), and of some propositional attitudes (*believe, intend*, *mutual belief*).

C. Raymond Perrault (Perrault, 1990) presents as a point of departure for his theory a possible version of such an approach[1]. It characterises speech acts by axioms of the form described above. E.g. for sentences uttered in *declarative* mood with propositional content p there is an axiom which postulates that the following consequence condition will hold:

$$BMB_{H,S}\, G_S\, B_H\, G_S\, B_H\, B_S p,$$

i.e. it will be the case that it is mutual belief ($BMB_{H,S}$) between hearer H and speaker S that the speaker has the goal (G_S) that the hearer believes (B_H) that the speaker wants that the hearer believes that the speaker believes that p. All stronger consequences, e.g. the standard case that the hearer believes after he has heard the utterance that the speaker believes p, or that he believes that p really holds, have to be derived by additional axioms characterising special circumstances. Of course, this condition is weak enough to cover all cases where a declarative sentence is used in a proper assertion, a lie, or the case where the hearer recognises the lie. As Perrault points out, such a theory can't handle *ironic* uses of declaratives, e.g. in case somebody says *This is the best meal I ever had* where it is obvious for speaker and hearer that the meal tasted quite bad. We will see that we can find even non-ironic uses of declaratives where this axiom is not weak enough.

Hence, the overall strategy of the approach is to formulate an axiom with very weak consequence condition, the strongest which capture *all* uses of a sentence with a certain syntactic feature. This is reflected in the fact that the *gating* condition for such axioms are very weak. In the above example they state no more than that S is the speaker in an utterance-event with hearer H and a sentence with content p. Special constraints, like *sincerity*, *competence*, *helpfulness*, have to be added to derive stronger consequences. But such a strategy meets the problem that it is difficult to find a consequent condition which is really weak enough to capture *all* uses of a certain class of sentences. We will adopt the converse strategy, i.e. we start with very strong conditions on the utterance situations with strong consequence conditions. They are intended to describe the most prototypical, or *basic*, cases for the use of a sentence with the relevant syntactic feature. Then we apply operations which reflect the influence of perspectives of participants on the dialogue situation to derive extensions. We can iterate this process, and in this way the conditions on the utterance situation and consequent situation would become weaker and weaker.

The approaches by Cohen & Levesque and Perrault handle the problem by *default mechanisms*. Perrault explicitly develops his theory as an application of default logic (Perault, 1990). Cohen & Levesque use a classical monotonic framework but add an axiom Cohen & Levesque, 1990a, Def. 5, p. 236) (Cohen & Levesque, 1990b, Def. 6) which works as a kind of default axiom. Hence, they too have a mechanism which allows to start with strong conditions. E.g. in Perrault's system we can prove the following default rules.

$$(\mathrm{DR1})\ \frac{B_{H,t}\,B_{s,t}\,p : B_{H,t}\,p}{B_{H,t}\,p} \qquad (\mathrm{DR2})\ \frac{B_{H,t+1}\,DO_{S,t}(p.) : B_{H,t+1}\,B_{S,t}\,p}{B_{H,t+1}\,B_{S,t}\,p}$$

(DR1) says: If the hearer believes at time t that the speaker believes at time t that p, and if it is not inconsistent to assume that H believes p, then he will believe that p. (DR2) says: If the hearer believes at time t+1 that the speaker uttered a sentence with propositional content p at time t, and if it is not inconsistent that he believes at this time that the speaker believed at time t that p, then H will believe at t+1 that the speaker really believed p. If the hearer knows that the speaker does not believe in p, then the default rules don't apply, and we arrive only on weaker consequence conditions. There are also default rules for intentions.

We can't go into a detailed discussion of the approaches. Cohen & Levesque and Perrault place their theories into a general theory of rational interaction. Especially, they explicitly describe the role of *intentions*. Of course, a full theory needs to deal with intentions and updates, but this needs more space than is available in this paper. So we can treat them only on an informal level. We claim that the fact that actions are performed relative to the perspectives of participants offer a real explanation for why extended uses of dialogue acts are possible. Hence, the theory of perspectives which we present in this paper can be seen as an empirical justification for such systems of default mechanisms.

3 The General Framework

We think of a dialogue situation as an example for what is known as a *multi-agent system*. In the following we take a dialogue situation to contain at a certain point of time a number of *dialogue participants* and an *outer situation*. We confine our considerations to the case where there are only two participants S and H. The outer situation may contain information about the immediate environment but most importantly it provides information about a situation talked about. The performance of a dialogue act like an assertion, question, use of a definite description etc. leads to changes in the state which describes the situation. An assertion e.g. aims at a change of the information state of the hearer. A question ultimately aims at a change of the information state of the speaker. We introduce this framework because it provides us with a conceptual basis for our theory of perspectives and dialogue acts.

In fact, we will not be interested in the actual representations an agent uses in his local state but only in the information which this representations contain about the state of the environment. Hence, we may identify the local states of the participants with the set of all states of the environment which are correlated to his actual representations, i.e. with the set of all *outer* situations which belong to a global dialogue situation where his local state is identical with the actual one. This set represents the *knowledge* of X about the world relative the overall system. This construction allows us to use a classical *possible worlds* approach to model the *information* of participants. Of course, this model would only contain the information the agent has about the environment. It is clear, that we also need to know what a dialogue participant knows about the knowledge of the other participant. Therefore, we represent dialogue situations as *possibilities* $w=\langle s_w,w(S),w(H)\rangle$, where s_w is the outer situation, and $w(S)$ and $w(H)$ are again sets of possibilities, so-called *information states*. This construction seems at first sight not to be well defined. In fact, it can't be defined in classical set theory. But it can be developed in (AFA) set theory (Aczel, 1988). The approach is known as *possibility* approach (Gebrandy & Groeneveld, 1997). We will present our results in this framework. We introduce it in more detail in an appendix.

We assume that there are some minimal restrictions on the class of possibilities. First, that all information states of the participants are non-contradictory, that *full introspection* holds, i.e. that all participants know what they believe, and that this is common knowledge. We denote the class of all these possibilities by $\dot{\mathcal{I}}$. Hence, if $w=\langle s_w, w(S), w(H)\rangle \in \dot{\mathcal{I}}$, then for all participants X

$$(1)\ w(X) \neq \emptyset,\ (2)\ \forall v \in w(X),\ v(X) = w(X), \text{ and } (3)\ w(X) \subseteq \dot{\mathcal{I}}.$$

Then, our basic dialogue situations will always be situations where both dialogue participants have only *true* beliefs, and where this is common knowledge. We denote the class of all such situations by $\mathcal{T}$. Hence, if $w \in \mathcal{T} \subseteq \dot{\mathcal{I}}$, then

$$(1)\ w \in w(S) \cap w(H), \text{ and } (2)\ w(S) \cup w(H) \subseteq \mathcal{T}.$$

The intended applications include speech acts like assertions, lies and questions, but also examples like the (referential) use of a definite description. Therefore, the notion of *dialogue act* has a wider meaning to us than the concept of *speech act*. Roughly, we can characterise our use of *dialogue act* as applying to any utterance of certain type of phrase which aims at a change in the information states of the participants. All considered examples of dialogue acts have in common that they depend on and affect only the information states of speaker and hearer. If the question whether a dialogue act can be performed or not is dependent on facts related to the outer situation, then it is clear that we can not expect that an extended use should be possible just because speaker and hearer *believe* it to be possible.

4 Dialogue Acts and Perspectives

4.1 The Basic Consideration

Let us look again to the example of the assertion *I hold an ace*.We abbreviate it by ψ . As a statement it would be felicitous just in case the sentence ψ is true. But, of course, the speaker *S* can relay for what he says only on what he believes to be true.

More generally, let $M \subseteq \dot{\mathcal{I}}$ be any class which represents a property of possibilities, e.g. a class where a dialogue act can be performed successfully. We explicate the fact that this property obtains *under the perspective* of a dialogue participant *X* in world *w* as $w(X) \subseteq M$. This means that *all* of *X*'s epistemic possibilities are elements of *M*.

For assertions this means that $w(S) \models \psi$ must hold, if the speaker should be justified to make them. If *M* specifies a class where some definite description d=def $x.\varphi(x)$ can be used felicitously, then it means that speaker *S* seems to be justified to utter *d* if and only if all his epistemic possibilities are elements of *M*.

Now, we turn attention to the hearer *H*. If he hears an utterance, is it necessary for him that *all* his epistemic alternatives are elements of *M* in order to make sense of what he has heared?

> **(8)** There are two playing cards on a table. The left one is an ace the right one a queen. *S* and *H* can see it and each other see it. Then they leave the room. An hour later they come back, the cards still there but face down. *H* has forgotten whether the first card is an ace or a king. He still knows that the second is a queen. *S*, who hasn't noticed this, says to *H*: „Give me the ace, then we leave and I invite you for a coffee."

Here, *H* can guess that *S* means the first card. The possibility that there is a king and a queen is ruled out by *S*'s use of „the ace." Therefore, we arrive at a weaker condition on the perspective of *H*, namely that there should be at least one possibility *v* in his set of epistemic alternatives *w*(*H*) that is in *M*.

We can find the same phenomenon in the case of assertions. If *H* is convinced that the speaker does not hold an ace, then he would not accept *S*' assertion *I hold an ace*. Again, we have as a minimal requirement for success that there is at least one epistemic alternative in *H*'s set of possibilities where the uttered sentence is true. Moreover, it is essential for assertions that the addressee does not know that the uttered sentence is true, i.e. that there is at least one epistemic alternative where ψ is false.

If M specifies a property of dialogue situations, then the actual situation w *may* have this property *under the perspective of* H iff $w(H) \cap M \neq \varnothing$ And, if M specifies the felicity conditions of some dialogue act, then the respective act can be understood as *reasonable* by H just in case he thinks that the actual situation might be an element of M.

We define two operators on subclasses of possibilities for each dialogue participant X. They are closely related to the modal operators $\Box_X$ and $\Diamond_X$, so that we denote them by the same symbols:

$$\Box_X M := \{w \in \dot{\mathcal{I}} \mid w(X) \subseteq M\}$$
$$\Diamond_X M := \{w \in \dot{\mathcal{I}} \mid w(X) \cap M \neq \varnothing\}$$

With these operators at hand we can reformulate our observations as: S is convinced that the actual world w belongs to M iff $w \in \Box_S M$; H can accept that w belongs to M iff $w \in \Diamond_H M$.

4.2 The Derivation of Dialogue Acts

We first show that there are four operations due to the perspective of one participant which extend the classes where some dialogue act can be performed.
As an example we use again the assertion

(9) S: I hold an ace. (ψ)

Assume that S is convinced of the truth of ψ. Even if it is in fact false, he will think to be justified to make the assertion. If, furthermore, the hearer H trusts into the beliefs of S, they should both accept this assertion as if S had told the truth. Then, S can use the assertion ψ to mislead the hearer H, if H is in a situation where he accepts the utterance. This is, of course, no more an assertion but a lie. We can see here, how the limited perspective of one dialogue participant can give rise to an extension of a dialogue act. To explain the precise connection of lies and assertions it is, of course, necessary to introduce *goals* of participants into the model. In this paper, we concentrate only on the restrictions for extensions of dialogue acts which are due to epistemic perspectives. The acceptability of the utterance for *S* and *H* forms a necessary condition for a derived use. Hence, the class of situations where this condition holds forms a class of *candidates* for a possible extension. Additional constraints may enter to determine the real extensions.

Hence, if M is a class of possibilities where some dialogue act can be performed, then this act might be extended to the class where *S* is convinced that *H* accepts the act. It is the class $\Box_S(M \cup \Diamond_H M)$. If $M \subseteq \Diamond_H M$, the definition of the extension can be simplified to $\Box_S \Diamond_H M$.

Now we can again turn to the perspective of *H*. Assume that the speaker *S* is convinced of the truth of ψ but *H* knows it to be false. *H* can make sense of the utterance, if it is possible for him that *S* might believe ψ. *Make sense* means here: he can understand it as an attempt of an assertion. Let us look again to a situation where this is *not* possible.

(10) Both players have only one card. *A* puts his card face up on the table. It is a king. Hence, both can clearly recognise that they mutually must know that it is a king. Then *A* says: „I hold an ace.“

Probably, *H* will be quite puzzled about this utterance. It seems to be impossible to make sense out of it. In contrast to the following scenario:

(11) Now assume *A* knows that he does not hold an ace but believes that *B* can't know this. He says: „I hold an ace.“ Assume that in this situation *B* could see through a mirror the playing cards of *A*.

Of course, *A*'s utterance was a lie, and *H* can not accept it as a true assertion. But he can make sense of it as he can recognise it as a lie. He can then react with a rejection, or he might *accept* it as assertion, thereby misleading the speaker.

If we look to referential uses of definite descriptions, we can find quite clear examples for this reasoning of the hearer.

(12) There are two playing cards on a table. They lay face down side by side. *S* gets told that the left one is an ace and the right one a queen. In fact, they are both jokers. Then *H* is brought to the table. He has seen the cards before, so he knows that they are jokers. A supervisor tells them that he has just informed *S* that the left card is an ace. Then *S* says to *H*: „Give me the ace!“

Of course, *H* should be able to identify the card, which *S* wants to refer to with *the* *ace*, although he knows that there are no aces. But he knows that out of the perspective of S there is one card in the common ground which is an ace, and that there is only one.

So if M is a class of possibilities where some dialogue act can be performed, then this act might be extended to the class where H thinks that it is possible that S is convinced that he can perform this act. We get as extension the class $\Diamond_H \Box_S M$.

We find in this way four operations which give us new classes where some dialogue act can be performed due to the perspective of one participant. Let M be given. Then we arrive at the following classification:

	direct	indirect
speaker	$\Box_S M$	$\Box_S \Diamond_H M$
hearer	$\Diamond_H M$	$\Diamond_H \Box_S M$

We have to mention that the simple form $\Box_S \Diamond_H M$ for the indirect operation for the speaker is sufficient only, if we can show that $M \subseteq \Diamond_H M$. Else, it should have the form $\Box_S (M \cup \Diamond_H M)$! If M characterises a dialogue act where $M \subseteq \mathcal{T}$, then we trivially have $M \subseteq \Diamond_H M$. In general, we can't expect this.

To get a real possible extension it is necessary that S is convinced that he can perform the act, and H must be able to make sense of this. It is not enough, if only one participant thinks that the act can be performed. Therefore, we have to build intersections of the derived classes. We get the following four groups:

	direct H	indirect H
direct S	$\Box_S M \cap \Diamond_H M$	$\Box_S M \cap \Diamond_H \Box_S M$
indirect S	$\Box_S \Diamond_H M \cap \Diamond_H M$	$\Box_S \Diamond_H M \cap \Diamond_H \Box_S M$

5 Applications

In this section we apply our theory to two examples: assertions and the referential use of definite descriptions. In the latter case we can refer to (Benz, 1999) for a more extended treatment. According to our strategy outlined in the last sections, we start in each case with a class of dialogue situations which characterises the most basic, or prototypical use. Then we show whether the examples belong to the basic or the derived new situations.

For the basic cases we assume that the participants mutually know that they have only *true* beliefs, i.e. for such a situation $w=\langle s_w, w(S), w(H)\rangle$ we assume that $w \in w(S) \cap w(H)$, and that the same holds for all $v \in w(S) \cup w(H)$. We introduced the class of all such situations in Section 3 as $\mathcal{T}$.

We assume that in the basic case where a participant S can utter a declarative sentence with propositional content ψ felicitously (1) ψ should in fact hold, (2) the speaker should be convinced of ψ, and (3) the hearer should not know whether ψ. If $[\psi] := \{w \in \mathcal{T} \mid w \models \psi\}$ denotes the class of possibilities in $\mathcal{T}$ where ψ holds, then the basic cases form the class $M := [\psi] \cap \Box_S[\psi] \cap \Diamond_H[\psi] \cap \Diamond_H[\neg\psi]$. In addition, we make an informal assumption about the intentions of the speaker. In the basic case he should always be sincere, i.e. if he says that ψ, he should really believe that ψ, and this should be mutually known. As already mentioned, we don't include this in the formal representation of a dialogue situation.

If we call an utterance of a sentence with content ψ *possible*, then we mean in the following that the utterance is *reasonable* for the speaker, and that the hearer can *make sense* of it. This was part of the general ideas motivating the operators introduced in the last section. Hence, we will call an utterance of ψ which is a lie, and which the hearer can recognise as a lie *possible* although it is unsuccessful, if we consider it as *speech act*.

We reconsider the examples given in the introduction. Hence, let ψ be the proposition *"S* holds an ace". As we are only interested in this proposition, we allow only for *two* outer situations, one where ψ holds, and one where $\neg\psi$ holds. We can therefore denote dialogue situations by triples $\langle \psi / \neg\psi, w(S), w(H)\rangle$.

Example **(1)** describes the following situation w_1:

$$w_1 = \langle \psi, \{w_1\}, \{w_1, v\}\rangle$$
$$v = \langle \neg\psi, \{v\}, \{w_1, v\}\rangle.$$

It is the situation where ψ holds, where the speaker knows all about the real situation, i.e. the only epistemically possible situation for him is w_1 itself, and where the hearer does not know whether ψ, but he knows that the speaker knows it, i.e. there are two possibilities for him, w_1 and another one which is identical to w_1 but where we replaced $\neg\psi$ for ψ. Of course, this interpretation is not fully justified by the way how the example was stated:

(13) *A* and *B* sit at a table playing with cards. *A* notices that he has an ace and says: "I hold an ace."

There are a lot of dialogue situations where this can be part of the description. It seems to be quite natural to give it a very strong interpretation. But we don't want to explain how we arrive at such a strong reading, but only, given the reading, why the utterance of „I hold an ace" is felicitous.

w_1 is clearly an element of $[\psi]$, i.e. ψ holds and $w_1 \in \mathcal{T}$. As $w_1 \models \psi$ it follows that $w_1 \in \Box_S[\psi]$), and as $v \models \neg\psi$ it follows that $w_1 \in \Diamond_H[\neg\psi]$. This proves that $w_1 \in M = [\psi] \cap \Box_S[\psi] \cap \Diamond_H[\psi] \cap \Diamond_H[\neg\psi]$, and that therefore the utterance of „I hold an ace" is felicitous.

The situation in Example **(2)** is described by

$$w_2 = \langle \neg\psi, \{w_1\}, \{w_1, v\}\rangle,$$

i.e. both participants have the same convictions as in the first example but this time ψ does not hold. Hence, $w_1 \neq w_2 \notin w_2(S)$, therefore $w_2 \notin \mathcal{T}$. But $w_2 \in \Box_S M \cap \Diamond_H M =: M_1$. Hence, the utterance is possible although it is not true this time. It is not uninteresting to write out the definition of M_1. After some simplifications we arrive at $\Box_S([\psi] \cap \Diamond_H[\neg\psi]) \cap \Diamond_H([\psi] \cap \Box_S[\psi]) \cap \Diamond_H[\psi]$. We can see in the second conjunct that it is necessary that the hearer believes it to be possible that the speaker can know that ψ. It would not be enough, if there are just any two basic possibilities in his belief-state, one where ψ holds and one where $\neg\psi$ holds.

In Example **(3)** the beliefs of the hearer are the same as in Example **(2)**, but the speaker knows now that $\neg\psi$ holds.

$$w_3 = \langle \neg\psi, \{w_3\}, \{w_1, v\}\rangle.$$

This is an element of $\Box_S \Diamond_H M \cap \Diamond_H M =: M_2$. Hence, it can be derive from the basic case, and our theory predicts that it can be reasonable for the speaker to say that

ψ. Of course, he must have the intention to mislead the hearer, i.e. his utterance must be a lie. More informally, his reasoning can be like this: *I know that I don't hold an ace. H does not know whether I hold it or not, but he knows that I know it, and believes that I am trustworthy. Hence, if I say to him that* ψ*, then he will believe me, and I want him to believe it.* But this reasoning makes essential use of the different perspectives.

In Example **(4)** the hearer recognised that the speaker believes to be in situation **(3)**. The new situation is an element of $M_3 := \Box_S \Diamond_H M \cap \Diamond_H \Box_S M_2 = \Box_S \Diamond_H M \cap \Diamond_H \Box_S (\Box_S \Diamond_H M \cap \Diamond_H M) = \Box_S \Diamond_H M \cap \Diamond_H \Box_S \Diamond_H M$.

$$w_4 = \langle \neg \psi, \{w_3\}, \{w_4\} \rangle.$$

The hearer can guess that the speaker reasoned in the way stated above, and can interpret his utterance as a lie, therefore, it makes sense for him. Again this is only possible because he can recognise that the utterance is reasonable out of the perspective of *S*.

We can see now how to construct further, more derived examples. Assume that the hearer believes himself to be in a situation as in Example **(4)** but the speaker knows this and has the intention to make the hearer believe that he is lying to him. Then it is reasonable for him to utter ψ. The underlying belief—states can be derived by a further application of the operator $\Box_S \Diamond_H$. In principle, we can repeat these constructions again and again, and get always new situations where an utterance of ψ is possible. We can find especially a situation where it is no more the case that it is mutually known that *the speaker has the goal (G_S) that the hearer believes that the speaker has the goal that the hearer believes that the speaker believes that* ψ, i.e. where it is no more mutually known that

$$(\Theta) \qquad G_S \; \Box_H G_S \Box_H \Box_S \psi .$$

In this case it would be necessary that the speaker himself believes that Θ. But this means that he wants that $\Box_H G_S \Box_H \Box_S \psi$. If we iterate our construction three times, then we get

$$w_5 = \langle \neg \psi, \{w_5\}, \{w_4\} \rangle$$
$$w_6 = \langle \neg \psi, \{w_5\}, \{w_6\} \rangle$$
$$w_7 = \langle \neg \psi, \{w_7\}, \{w_6\} \rangle,$$

which can represent the information of *S* and *H* in a situation where the speaker wants that the hearer believes that *S* wants him to believe that he is lying to him. In this case, after the utterance, the hearer will of course not believe that the speaker wants him to believe that he believes ψ. Hence, this is a non-ironic use which is a counterexample to the claim that Θ is mutually known in *all* possible uses of a declarative sentence. Of course, this example is already quite artificial, but this has to be expected from a counterexample.

Until now we have only considered examples where the use of a declarative sentence was possible. We now want to examine some examples where the dialogue situation does not belong to one of the derived classes.

The most simple case is represented in Example **{5}**. It is represented by

$$v_1 = \langle \neg\psi, \{v_1\}, \{v_1\} \rangle.$$

This is, of course, an element of $\mathcal{T}$ but not of [ψ], hence, not of *M*. It is provable that (v_1) is not an element of any of the derived classes[2].

(14) *H* is brought into a room where a supervisor puts some cards on a table. He can see all of them before they are turned around. The first card is an ace. Then *S*, who waited outside where he couldn't see the cards, enters and says: „The first card is an ace."

If ψ denotes the proposition *the first card is an ace*, then we can represent the utterance situation by v_2:

$$v_2 = \langle \psi, \{v_2, v_3, v_4, v_5\}, \{v_2\} \rangle,$$
$$v_3 = \langle \neg\psi, \{v_2, v_3, v_4, v_5\}, \{v_3\} \rangle,$$
$$v_4 = \langle \psi, \{v_2, v_3, v_4, v_5\}, \{v_4, v_5\} \rangle,$$
$$v_5 = \langle \neg\psi, \{v_2, v_3, v_4, v_5\}, \{v_4, v_5\} \rangle,$$

None of the possibilities v_2 to v_5 belongs to *M* because *S* is never convinced that ψ holds. Hence, his utterance can't be a basic use, and it is again provable that the situation v_2 does not belong to any of the derived classes. Intuitively, it is not a reasonable sincere assertion as the speaker lacks the necessary knowledge, and H knows this, and it is even common knowledge that he lacks this information. Therefore, *H* may not only reject this utterance, it should also be difficult for him to make any sense of it.

Finally, we want to consider the referential use of a definite description. For the basic case we assume that a speaker can refer with def $x.\varphi(x)$ to an object a iff it is common knowledge that $\varphi(a)$[3], and that a is the only object where this is commonly known. An example is **(6)**. We now reconsider Example **(8)**. There are two cards, c_1 and c_2. c_1 is an ace and lies to the left of c_2, a queen. The speaker wants to pick out c_1 by use of *the ace*. The utterance situation introduced in **(8)** can be described by w:

$$
\begin{aligned}
w &= \langle [A(c_1),Q(c_2)],\{u\},\{w,v\}\rangle \\
v &= \langle [K(c_1),Q(c_2)],\{u'\},\{w,v\}\rangle \\
u &= \langle [A(c_1),Q(c_2)],\{u\},\{u\}\rangle \\
u' &= \langle [K(c_1),Q(c_2)],\{u'\},\{u'\}\rangle
\end{aligned}
$$

We denote the class of basic cases again by M. We see that only u is an element of M, and therefore we can show that $w \in \Box_S M \cap \Diamond_H M$. The theory predicts that the use of „the ace“ should be successful.

6 Conclusion

I have argued that perspectives play an essential role in the derivation of extended uses of dialogue acts. The essential idea was to start with a class of basic, or class by application of four operators which reflect the four ways how the partiality of information can give rise to uses in situations which do not belong to the basic class.

- The speaker can believe that the performance is possible.
- The speaker can believe that the performance might be possible.
- The speaker can believe that the hearer believes that the performance might be possible.
- The hearer can believe that the speaker might believe that the performance is possible.

We concentrated on the referential use of a definite description, and especially the use of a sentence in declarative mood. In the last case, our theory could show how a speech act like lying can be derived from the act of asserting.

References

P.Aczel (1988): *Non-Well-Founded Sets*; CSLI-Lecture Note 14, Stanford.

A. Baltag (1999): *A Logic of Epistemic Actions*; Proceedings of the ESSLLI99 Workshop on Foundations and Applications of Collective Agent Based Systems, Utrecht.

A. Baltag, L. Moss, S. Solecki (1998): *The Logic of Public Announcements*, *Common Knowledge and Private Suspicion;* to be published.

J. Barwise, L. Moss (1996): *Vicious Circles*; Stanford, CSLI.

A. Benz (1999): *Perspectives and the Referential Use of Definite Descriptions in Dialogue*; ms, Berlin.

P.R. Cohen, H.J. Levesque (1985): *Speech Acts and Rationality*, in: *Proceedings of the Twenty-third Annual Meeting*, Association for Computational Linguistics, Chicago, IL.

P.R. Cohen, H.J. Levesque (1990): *Rational Interaction as the Basis for Communication*; in: P. Cohen, J. Morgan, M. Pollack (ed.): *Intentions in Communication*, Cambridge, Massachusetts.

P.R. Cohen, H.J. Levesque (1990):*Performatives in a Rationally Based Speech Act Theory*; *In Proceedings of the Twenty-eighth Annual Meeting*, Association for Computational Linguistics, Pittsburgh, Pennsylvania.

R. Fagin, J.Y. Halpern, Y. Moses, M.Y. Vardi (1995): *Reasoning About Knowledge;* Cambridge, Massachusetts.

J. Gerbrandy (1997): *Changing the Common Ground*; in: A.Benz & G. Jäger (ed.), *Proceedings of the Munich Workshop on Formal Semantics and Pragmatics of Dialogue MunDial'97*, pp.40-58 München.

J. Gerbrandy (1998): *Bisimulations on Planet Kripke*; ILLC Dissertation Series, Institute for Logic, Language and Computation, Universiteit van Amsterdam.

J. Gerbrandy, W. Groeneveld (1997): *Reasoning about information change*; Journal of Logic, Language and Information.

W. Groeneveld (1995): *Logical Investigations into Dynamic Semantics;* PhD thesis, ILLC dissertation series 1995-18, University of Amsterdam.

J. Jasparsm (1994): *Calculi for Constructive Communication. A Study of the Dynamics of Partial Sets*; PhD thesis, ITK, Katholieke Universiteit Brabant.

C.R. Perrault (1990): *An Application of Default Logic to Speech Act Theory* in: P. Cohen, J. Morgan, M. Pollack (ed.): *Intentions in Communication*, Cambridge, Massachusetts.

ENDNOTEN

* This work is part of the DFG—Project LA 633/5-1 on *Dialogue Semantics* located at the Humboldt-Universität Berlin, http://www2..rz.hu-berlin.de/asg/blutner/dialog/index.html. I want to thank the organisers of the *Communicating Agents* workshop for the opportunity to discuss and publish these results

1 He attributes this version to Cohen & Levesque. But there are significant simplifications in this picture, at least according to our reading of (Cohen & Levesque.1990 a,b).

2 If we denote by $T(w)$ the *transitive hull* of $\{w\}$, i.e. the smallest set σ, such that $w \in \sigma$ and $\forall v \in \sigma v(S) \cup v(H) \subseteq \sigma$ then it holds that: If w is in one of the derived classes, then $T(w) \cap \neq \emptyset$). We can see that $T(v_1)=\{v_1\}$. The same criterion helps also in the next examples.

3 We provide a precise definition of *common knowledge* in the appendix.

A The Possibility Approach and the Representation of Beliefs of Dialogue Participants

The possibility approach is essentially a possible worlds approach, i.e. it identifies the beliefs of an individual with the set of all worlds which are *possible* according to those beliefs. We denote the set of participants by DP $=\{S,H\}$, where S will denote the *speaker*, H the *hearer*. A *possibility* consists of a model for the *outer* situation, and *information states* for each participant, where those states are again sets of possibilities. The *outer* situation describes the *non—modal* part of the dialogue context. In case of e.g. assertions, this outer situation will be identified with the situation *talked about*.

The possibility approach was first developed by *J. Gerbrandy* and *W. Groeneveld* in (Gerbrandy and W. Groeneveld 1997). It makes use of an extension of classical set theory, the theory of *Non-Well-Founded Sets* developed by *P. Aczel* (P. Aczel, 1988)[4]. The original problem motivating the development of the possibility approach was to define suitable *update operations* for dialogue. The approach proved here to be especially useful[5]. For a proper understanding of the details the reader may need to have some familiarity with the underlying set theory. We hope that he can get an intuitive understanding without it.

Let S be a class of models for the possible outer situations. For simplicity we assume that all models in S have the same set of individuals. We define *possibilities* and *information states* in the following way:

- A *possibility* w is a triple $\langle \neg s_w, w(S), w(H)\rangle$ where $s_w \in \boldsymbol{\mathcal{S}}$ and $w(S)$ and $w(H)$ are information states.
- An *information state* σ is a set of possibilities.

s_w describes an outer situation, $w(S)$ and $w(H)$ the set of worlds S and H believe to be possible. We denote the class of all possibilities with $\boldsymbol{\mathcal{W}}$. The theory of non-well-founded sets allows for sets containing themselves, so it is possible that there exist possibilities w with $w \in w(X), X \in$ DP.

The definition may seem to be circular, and therefore ill-defined. It is, of course, not a recursive definition. For an explanation we need some machinery from (AFA)-set theory. Generally, we can define the class $\mathcal{W}$ as the *largest fixed point* of a set-continuous operator Φ with $\Phi(x) := \{\langle s,i,j\rangle \mid s \in \mathcal{S}, i,j \subseteq x\}$. It is provable that this fixed point is also the class of all *solutions* for a certain class of *systems of equations*. The definition above directly translates into a definition of this *certain* class: We assume that there are two classes of *urelements*, P and I. A system of equations over P and I belongs to the desired class if it contains only equations of the form:

- $w = \langle s_w, i, j \rangle$ for $w \in P$, where s_w belongs to $\mathcal{S}$ and i,j to I.

- $i = \sigma$ for i in $\in I$, where σ is a subset of P.

In this way, it is really a proper definition[6].

We introduce a formal Language $\mathcal{L}^M$. Let $\mathcal{L}$ be a language of predicate logic for the class $\mathcal{S}$. We assume that $\mathcal{L}$ contains all the predicates the dialogue participants can use to talk about an outer situation. Then $\mathcal{L}^M$ should be the smallest language containing $\mathcal{L}$ and the following sentences for $\varphi, \psi \in \mathcal{L}^M$ and $X \in DP$:

$$\neg\varphi, \varphi \wedge \psi, \Box_{X\varphi}, \Diamond_{X\varphi}, E\varphi, C\varphi$$

Let $w = \langle s_w, w(S), w(H)\rangle$ be a possibility. We define truth conditions for $\varphi, \psi \in \mathcal{L}^M$:

1. $w \models \varphi$ *iff* $s_w \models \varphi$, φ a sentence in $\mathcal{L}$.
2. $w \models \neg\varphi$ *iff* $w \not\models \varphi$.
3. $w \models \varphi \wedge \psi$ *iff* $w \models \varphi \;\&\; w \models \psi$.
4. $w \models \Box_X\varphi$ *iff* $\forall w' \in w(X)\,(w' \models \varphi)$.
5. $w \models \Diamond_X\varphi$ *iff* $\exists w' \in w(X)\,(w' \models \varphi)$.
6. $w \models E\varphi$ *iff* $w \models \Box_S\varphi \wedge \Box_H\varphi$.
   Let $E^0\varphi := E\varphi$, $E^{n+1}\varphi := E(E^n\varphi)$.
7. $w \models C\varphi$ *iff* $\forall n \in \mathbf{N}\; w \models E^n\varphi$.

For a dialogue participant X a possibility w is *epistemically possible* in v iff $w \in v(X)$. X *believes* that φ in w iff φ holds in all this epistemic alternatives in w, i.e. iff

$w \models \Box_X\varphi$. $w \models E\varphi$ means that *everybody* believes φ in w. φ is *common belief* in w iff $w \models C\varphi$. For information states we define

$$\sigma \models \varphi \textit{ iff } \forall w \in \sigma\ w \models \varphi.$$

Until now, we did not restrict the properties of possibilities. A subclass $M \subseteq \mathcal{W}$ is called *transitive*, iff

$$\forall w \in M \forall X \in DP\ w(X) \subseteq M.$$

Let $\mathcal{I} \subseteq \mathcal{W}$ be the largest transitive subclass with

$$\forall w \in \mathcal{I}\ \forall X \in DP\ \forall v \in w(X) : w(X)=v(X).$$

This property is called *introspectivity*. It means: (1) If a dialogue participant believes φ, then he knows that he believes it; (2) if he does not believe that φ, then he knows that he does not believe φ; and (3) it means that (1) and (2) are common knowledge. We will always assume that introspectivity holds.

Let $\mathcal{T} \subseteq \mathcal{I}$ be the largest transitive subclass with

$$\forall w \in \mathcal{T}\ \forall X \in \mathrm{DP}\ w \in w(\mathrm{X}).$$

If $w \in \mathcal{T}$, then w is for both participants an element of their sets of epistemic alternatives. Hence, if a participant believes that φ, then φ must in fact hold. Therefore, $\mathcal{T}$ denotes the class of possibilities where (1) the dialogue participants can only have *true beliefs*, i.e. *knowledge*, and (2) where this fact is common knowledge.
The *Anti-Foundation-Axiom* (AFA) of the underlying set theory guaranties that $\mathcal{T}$ is not empty. We can easily see that the $S5$-Axioms hold for $\mathcal{T}$:
If $w \in \mathcal{T}$, $X \in$ DP, we have

$$(1)\ w \models \Box_X\varphi \Rightarrow w \models \varphi,$$
$$(2)\ w \models \Box_X\varphi \Rightarrow w \models \Box_X\Box_X\varphi \text{, and}$$
$$(3)\ w \models \varphi \Rightarrow w \models \Box_X\Diamond_X\varphi\,.$$

We are only interested in non-contradicting information states of participants. This means that the set containing all their epistemic possibilities should contain at least one element. Let $\dot{\mathcal{W}}$ denote the largest transitive subclass of $\mathcal{W}$ with

$$w \in \dot{\mathcal{W}} \Rightarrow w(S) \neq \phi \neq w(\mathrm{H})$$

If M is any class of possibilities, then we denote by M the intersection of M and $\dot{\mathcal{W}}$. Note that $\dot{\mathcal{T}}=\mathcal{T}$.

ENDNOTES OF APPENDIX A

4 For more information about (AFA) Set Theory we can also refer to *Barwise & Moss* (Barwise&Moss 1996). For a more thorough discussion of the possibility approach we can refer to the thesis of Gerbrandy (Gerbrandy1998).

5 There is some literature concerning the proper definition of updates in dialogue: (Jaspars, 1994), (Groenefeld, 1995), (Gerbrandy & Groenefeld, 1997), (Gerbrandy, 1998), (Balltag, Moss, Solecki, 1998), (Baltag 1999).

6 For the set-theoretic machinery we must refer to (Barwise & Moss 1996). A very readable account can be found in (Gerbrandy 1998).

Roland Hausser

Foundations of Computational Linguistics. Man-Machine Communication in Natural Language

Berlin Heidelberg: Springer, 1999. ISBN 3-540-66015-1

Rezensiert von Winfried Lenders, Bonn

Das vorliegende Werk von Roland Hausser, „Foundations of Computational Linguistics“, ist, auch wenn es sich streckenweise einen solchen Anstrich gibt, weit mehr als ein bloßes Lehrbuch der Computerlinguistik. Das Buch erhebt vielmehr den Anspruch, Sprachtheorie zu bieten, ja darüber hinaus Kommunikationstheorie im Sinne einer Darstellung der theoretischen Grundlagen von Mensch-Maschine-Kommunikation in natürlicher Sprache. Kommunikationstheorie auf der Grundlage von Prinzipien der Computerlinguistik darzustellen, dahinter steht ein besonderes wissenschaftstheoretisches Verständnis von Computerlinguistik: Denn dieser noch relativ jungen Disziplin wird die Aufgabe zugeschrieben, Modelle sprachlicher Kommunikation, Modelle kommunikativen Verhaltens zu entwickeln, und zwar auf dem Wege der Algorithmisierung und Formalisierung. Es geht darum, wie man auch mit Einleitung S. 4 sagen könnten, einen Computer als Kommunikator auszubilden, also ein funktionales Modell der natürlich-sprachlichen Kommunikation aufzustellen. Hausser hat für eine Theorie dieser Art eine besondere Bezeichnung geprägt. Er nennt sie SLIM-Sprachtheorie, eine Abkürzung, hinter der sich folgendes verbirgt:

1. S = *Surface compositional* (Methodologisches Prinzip)

   SLIM verwendet eine oberflächenkompositionale Syntax und Semantik. in der nur konkrete Oberflächen zusammengebaut werden dürfen, unter Verzicht auf Null-Elemente, Identitätsabbildungen oder Transformationen.

2. L = *Linear* (=Empirisches Prinzip)

   SLIM verwendet eine zeitlineare Grammatik, die die empirische Tatsache formalisiert, dass in den natürlichen Sprachen immer ein Wort nach dem anderen geäußert bzw. wahrgenommen wird.

3. I = *Internal* (Ontologisches Prinzip)

SLIM behandelt die sprachliche Interpretation als einen Sprecher-Hörer-internen kognitiven Prozess.

4. M = *Matching* (Funktionales Prinzip)

SLIM behandelt Referenz auf der Grundlage eines Abpassens (*matching*) zwischen wörtlicher Sprachbedeutung und Verwendungskontext.

(Zitiert nach der Rohfassung der noch im Erscheinen begriffenen dt. Ausgabe des Buches)

Durch diese Markierungspunkte ist der Rahmen abgesteckt für ein umfassendes Theoriegebäude, das Hausser in 24 Kapiteln, gleichmäßig verteilt auf 4 große Teile präsentiert. Das Buch als ganzes beruht auf der Idee, Sprache nicht als abstraktes Konstrukt zu behandeln, sondern als Medium, das zusammen mit den Prozessen des Wahrnehmens, des Denken und des Verstehens der Kommunikation zwischen Individuen dient. Zu Beginn eines jeden Teils und eines jeden Kapitels zeigt der Autor selbst dem Leser den roten Faden, der das Buch durchzieht: Der Weg führt von den einfachen zu den komplexen Hilfsmitteln der sprachlichen Kommunikation und von den einfachen zu den komplizierteren Methoden der formalen Beschreibung. Gleichzeitig werden zahllose in der Sprach- und Kommunikationstheorie, der Mathematik und Logik, der Informatik und Physik schon vorliegende Entwürfe zur Sprache und ihrer Beschreibung vorgestellt und besprochen.

Der erste große Teil (Kapitel 1–6) trägt die Überschrift „Theory of Language". Ausgangspunkt sind hier allgemeine Grundlagen der Kommunikation und der Darstellung sprachlicher Zeichen im und für den Computer. Es werden die grundlegenden Aspekte maschineller Sprachanalyse im Rahmen der umfassenderen Aufgabe der Mensch-Maschine-Kommunikation, die Möglichkeiten einer technologischen Umsetzung dieser Forschungen, z.B. in Information Retrieval Systemen und maschinellen Übersetzungssystemen, beschrieben. Ab Kapitel drei wird Sprachtheorie als Kommunikationstheorie dargelegt. In diesen Kapiteln wird durch die „Konstruktion" einer Roboters namens Curious (S. 51: „in terms of constructing a robot named curious") die Komponenten der sprachlichen Kommunikation in einem prototypischen Modell dargestellt, das es – ähnlich wie Winograds SHRDLU-System – mit einer vereinfachten Welt zu tun hat. Im Unterschied zu und in Erweiterung von SHRDLU geht es Curious jedoch auch um das

Wahrnehmen und Erkennen der Welt bzw. um die Art und Weise, wie die „Welt“ als „Umwelt“ oder „Kontext“ des Roboters von diesem wahrgenommen und erkannt werden kann. Der Übergang von Wahrnehmen zum Erkennen wird dabei als eine Art Mustererkennung aufgefasst, durch welche die Parameter eines in der Umgebung wahrgenommene „Objekts“ gleichsam in das Innere des „kognitiven Agenten“ transportiert werden und dort ein ‚Instantiierungs-Konzept‘, I-Konzept genannt, bilden. Dieses wird dadurch „erkannt“, dass es „matcht“ mit einem M-Konzept, das sich als „type“, vielleicht auch „prototype“, innerhalb des kognitiven Agenten befindet. Es wäre hier – wie auch später – zu fragen, wie sich die von Hausser hier eingeführten Begriffe wie I-Konzept, M-Konzept etc. zu verschiedenen älteren kognitiven Modellen verhalten, etwa zu der Theorie der Mentalen Modelle von Johnson-Laird, der ältern Frame-, Script- und Schema-Theorien etc. Abgesehen von dieser Frage gelingt es Hausser jedoch sehr gut, mit Hilfe der beinahe genialen Annahme eines einfachen Roboters die grundlegenden Handlungen beim Wahrnehmen und Erkennen, also bei kognitiven Operationen darzustellen. In den Kapiteln 4 und 5 wird der zunächst sprachlose Curious mit Sprache ausgestattet, die ihn dann – über seine Fähigkeiten als wahrnehmenden, erkennenden und handelnden Agenten hinaus – auch zu sprachlicher Kommunikation befähigt. Hierzu wird – auch unter Hinweis auf Sprachtheorien wie die von Austin, Frege, Grice und Wittgenstein – im Gesamtrahmen der SLIM-Theorie eine Referenztheorie entwickelt. Gemäß dieser Theorie wird der Referenzprozess als eine interne ‚matching procedure‘ zwischen der literalen Bedeutung, d.i. dem M-Konzept, und einem entsprechenden kontextualen Referenten, dem I-Konzept verstanden. Die ‚literale Bedeutung‘ (*literal meaning*) kommt durch Konvention zustande. Hausser unterscheidet also zwei Arten von Bedeutung; er nennt sie meaning_1 und meaning_2: meaning_1 ist die literale Bedeutung, meaning_2 die der Äußerung eines Sprechers zugeordnete Bedeutung, also wohl diejenige, die sich aus dem I-Konzept ergibt. Auch hier stellt sich die Frage, wie diese Referenztheorie, die sich leicht zu einer Verstehenstheorie erweitern ließe, im Vergleich zu anderen Konzeptionen einzuordnen ist (z.B. in Relation zu Winograds Konzept eines operationalen Verstehens und zu den entsprechenden Vorstellungen der Kognitionspsychologie). Für Hausser ist der Unterschied klar: es handelt sich um eine ‚interne‘ matching procedure, während andere Theorien den Vorgang des Referenz (und auch des Verstehens) gleichsam von aussen betrachten, als etwas, das sich zwischen Sprecher und Hörer abspielt, und nicht in ihnen (intern). Genau hier aber wäre wiederum zu prüfen, wie sich Haussers Vorstellungen zu anderen Theorien verhalten, z.B. zu Winograd, der einerseits eine ‚interne‘

Repräsentation seiner (Modell-)welt konstruiert, wenn diese auch gesetzt und nicht wahrgenommen wird, und andererseits eine interne Repräsentation der Äußerungen des Benutzers, die beide aufeinander abgebildet werden müssen, damit es zu Referenz und Verstehen kommen kann.

Die letzten beiden Kapitel dieses Teils sind im Großen und Ganzen dem Zeichencharakter der Sprache, den Zeichentypen und ihrer Interpretation im Rahmen der SLIM-Theorie und den innnertextualen Referenzen und Koreferenzen gewidmet.

Im zweiten Teil „Theory of Grammar“ (Kap. 7–12) steht die Sprache als wichtigstes Kommunikationsmittel im Mittelpunkt. Es geht nicht mehr um das einzelne Zeichen, mit dem zu kommunizierende Entitäten bezeichnet werden, oder um das singuläre Ereignis selbst, das kognitiv und sprachlich verarbeitet wird, sondern um Zeichenkomplexe, um die Art und Weise, wie Zeichen zu komplexen Gebilden verknüpft werden und um die wissenschaftliche Beschreibung dieser Objekte und Vorgänge. Es geht – kurz gesagt – um die grundlegenden Methoden der formalen Beschreibung komplexer Sprachgebilde, also um Grammatiktheorie. Was aber heißt Grammatiktheorie? Um dies zu klären, geht Hausser - auf der Basis des immer noch die Diskussion beherrschenden Paradigmas der Generativen Grammatik - auf die Ursprünge des generativen Konkatenierens von Symbolen zurück und unterscheidet unter Gesichtspunkten einer formalen Sprachtheorie drei elementare Formalismen, die Kategoriale Grammatik *(C-Grammar*), die Phrasenstrukturgrammatik (*PS-Grammar*) und die Links-Assosiative Grammatik (*LA-Grammar*). Zunächst werden in den Kapiteln 7 und 8 die beiden ersten Grammatiktypen ausführlich bezüglich ihrer grundsätzlichen Eigenschaften und ihrer Mächtigkeit erörtert. Aufgrund seines sehr systematischen Ansatzes gelingt es Hausser dabei, die wichtigsten Unterschiede und Besonderheiten dieser Grammatiktypen klar zu machen. Ganz nebenbei oder vielmehr folgerichtig werden in diesem Kontext auch die für ein Lehrbuch unverzichtbaren Grundlagen des Arbeitens mit formalen Grammatiken und des Parsing (deklarativ vs. prozedural, kontextfrei vs. kontextsensitiv, top-down vs. bottom-up etc.) erklärt. Die Darstellung mündet in einer Liste von Desiderata besonders der PS-Grammatiken, ja eigentlich kommt Hausser zu einer vernichtenden Beurteilung dieses Grammatikformalismus, und zwar in allen zur Zeit aktuellen Versionen (GPSG, HPSG, LFG), die er allesamt als dem *nativism* verpflichtet sieht, also der These von der Existenz allgemeiner angeborener sprachlicher Grundmechanismen. Bei aller Prominenz, zahlreichen Revisionen, vielen Anhängern etc. sei die Entwicklung der PSG dennoch ein *textbook example* für mangelnde Konvergenz. Ein wichtiges Argument für diese

Sichtweise sei, dass in der Tat die praktischen Systeme der maschinellen Sprachverarbeitung sich entweder zu einer Theorie nur im Sinne eines Lippenbekenntnisses äußern oder auf Theorien gänzlich verzichten. Gründe dafür sind nach Hausser, dass sich auf der Basis des *nativism* keine funktionale Theorie der Kommunikation aufbauen lasse und dass PS-Grammar-Formalismen inkompatibel seien mit „the input-output conditions of the speaker-hearer." (179). Diese Desiderate würden, wie Hausser meint, durch den dritten Grammatik-Formalismus, die LA-Grammatik, aufgehoben, wie in den folgenden Kapiteln 10–12 zunächst allgemein, und dann im dritten Teil des Buches praktisch dargelegt wird.

Der Formalismus der LA-Grammatik (LAG) wurde von Hausser entwickelt und wird hier erstmals in umfassender Form vorgestellt und mit anderen Formalismen (C-Struktur, PS-Struktur) konfrontiert. Er sei „input-output equivalent with the speaker-hearer" (S. 183). Es wird damit ein Grammatiktyp vorgeschlagen, der sprachliche Äußerungen, die wir als Mensschen ‚zeit-linear' produzieren und auch ‚zeit-linear' wahrnehmen, in gleicher Weise ‚zeit-linear' abarbeitet bzw. ableitet. Dabei bleibt die Interpretation eines erkannten Elements, z.B. einer Präpositionalphrase, deren Zuordnung zu einem Verb oder zu einem Nomen, dem Referenzprozess, der in der Kommunikation zwischen Sprecher und Hörer abläuft, überlassen. Auf diese Weise erhält Hausser weit weniger komplexe Strukturen, denn es wird darauf verzichtet, alle möglichen syntaktischen Beziehungen zwischen Konstituenten und semantischen Lesarten darzustellen. Die LAG vermeidet also die in anderen Formalismen immer anzunehmenden vielfältigen Interpretationen von Sätzen; dies sei der natürlichen Kommunikation adäquater. Hauser wird damit zwar der Tatsache gerecht, dass wir uns ‚zeit-linear' sprachlich äußern und sprachliche Äußerungen ‚zeit-linear' wahrnehmen; es sei hier aber offen gelassen, ob damit auch die planerischen oder strategischen Akte, die beim Entwurf und bei der Strukturierung einer Äußerung in ihrer Gesamtheit, noch ehe sie hervorgebracht ist, am Werke sind, gerecht wird.

Der dritte Teil (Kap. 13–18), der mit „Morphology and Syntax" überschrieben ist, konkretisiert dieses methodologische Konzept der LAG in Bezug auf die grundlegenden sprachlichen Phänomene der Formenbildung und Syntax. In schöner Klarheit werden hier sowohl die grundlegenden Begriffe der Morphologie (Kap. 13) und Syntax (Kap. 16) dargelegt, als auch – unter Anwendung der LA-Grammatik – konkret Formalismen zur Erkennung von Wortformen (Kap. 14) und syntaktischen Strukturen (Kap. 17 und 18) vorgestellt. Auf dem Gebiet der Morphologie steht dabei das von Hausser entwickelte System La-MORPH im Zentrum, auf das hier nicht näher eingegangen wird, da es schon früher beschrieben und seine

Funktionstüchtigkeit in praktischen Tests nachgewiesen wurde (vgl. Hausser [Hg.]: Linguistische Verifikation. Dokumentation zur Ersten Morpholympics 1994. Tübingen: Niemeyer, 1996). In der Syntax strebt Hausser eine reine Oberflächensyntax an, in der es nicht um syntaktische Funktionen oder Rollen geht, sondern nur um die Wohlgeformtheit der Sätze, die sich aus den morphosyntaktischen Merkmalen ergibt, die in den Wortformen angelegt sind. Sätze sind Konstruktionen, zwischen deren Teilen syntagmatische Beziehungen bestehen, und zwar solche der Valenz, der Kongruenz und der Stellung. Nur um diese drei „Grundprinzipien der natürlichsprachlichen Syntax" geht es. Bei Ableitung eines Satzes ist zu garantieren, dass die Valenzbedingungen der Wörter erfüllt sind, Kongruenz zwischen den Wortformen erreicht ist und nur zulässige Wortstellungen erzeugt werden. Jede Wortform ist Valenzträger und hat ‚Valenzstellen', gleichsam Eigenschaften, die auf die nachfolgende Wortform schließen lassen. Diese Eigenschaften müssen in dem Lexikon, auf das die syntaktische Beschreibung zugreift, kodiert sein oder sich aus der vorausgehenden morphologischen Analyse der Wortformen ergeben. Zu diesen ‚Eigenschaften' gehört z.B. die Angabe der zweiten Person bei „du" und „liest", die – wenn diese Wortformen im Text zeitlinear auftreten – zur Überprüfung ihrer Kongruenz dienen. Die Valenzstellen fungieren gleichsam als *slots*, die im Verlauf der zeitlinearen Analyse durch geeignete *filler* ausgefüllt werden müssen. Wie dies im LA-Formalismus umgesetzt wird, zeigt Hausser unter Verwendung algebraisch formulierter Satzmuster. Die Anwendung dieser Formalismen wird schließlich an Fragmenten der deutschen und englischen Syntax veranschaulicht. Dabei beschränkt sich Hausser auf den deklarativen Hauptsatz, und hier vor allem auf die Nominalgruppen als nominale *filler* des verbalen Valenzrahmens, jedoch werden auch diskontinuierliche Sequenzen und Distanzstellungen, z.B. bei Auxiliarkonstruktionen, betrachtet.

Wenn Hausser sich in der Syntax ganz auf Oberflächenkonkatenationen beschränkt, so tut er das aus guten Gründen, die in seiner umfassenderen SLIM-Theorie motiviert sind. Denn die Auswahl der Satzkonstruktion aus einer Vielzahl von Möglichkeiten, die Wahl eines Tempus zum Ausdruck einer bestimmten Befindlichkeit in der Zeit, die Wahl eines Modus, die Wahl gerade ‚dieser' Wortstellung und nicht einer anderen usw., dies alles sind Prozesse, die im Sprachbenutzer ablaufen. Sie finden außerhalb der Syntax in den Teilen des ‚kognitiven Agenten' statt, die Hausser Semantik und Pragmatik nennt.

Diesen Teilen widmet er sich nun im vierten Teil (Kap. 19-24) des Buches, der mit „Semantics and Pragmatics" überschrieben ist. Man könnte meinen, nun folge er der alten semiotischen Trias Syntax, Semantik, Pragmatik, nach der sich so viele

Lehrbücher der Linguistik in ihrem Aufbau richten. Weit gefehlt! Hausser geht vielmehr ganz folgerichtig seinen Weg, wie er in den ersten theoretischen Teilen angelegt ist, weiter, ein Weg, der jetzt von den komplexen sprachlichen Zeichen, den Wortformen und Sätzen, zum Bezeichneten führt, oder, um Haussers Begriffe zu benutzen, der von der Ebene I, der sprachlichen Oberfläche (*language surface*) zur Ebene II, dem semantischen Inhalt (*semantic content*) führt. In konkreten natürlichsprachlichen Systemen wird der Weg von Ebene I nach Ebene II gemeinhin als semantische Interpretation bezeichnet. In diesem Zusammenhang tritt ein Kerngedanke des Buches und seiner leitenden Theorie besonders deutlich hervor: Hausser erläutert in knapper und präziser Form, was den drei Typen semantischer Systeme, dem semantischen System der logischen Sprachen, dem der Programmiersprachen und dem der natürlichen Sprache gemeinsam ist und was sie trennt. Ohne hier zuviel zu verraten: Es wird in einfacher, doch überzeugender Form dargestellt, was semantische Interpretation in künstlichen Sprachen und in der natürlichen Sprache heißt. Dabei werden die Grundkonzepte der modelltheoretischen Semantik von Tarski verständlich erörtert und in ihrer Anwendbarkeit auf natürliche Sprachen geprüft. Indem er weiterhin erläutert, was vom Standpunkt der modelltheoretischen Semantik aus Wahrheit ist, was man in dieser Hinsicht unter Bedeutung eines Ausdrucks zu verstehen hat und in welchem Sinn in diesem Zusammenhang von Ontologie gesprochen werden kann, schlägt Hausser die Brücke zwischen logischer und linguistischer Semantik. Es versteht sich, dass dabei auch die sprachtheoretischen Konzeptionen von Frege, Carnap und anderer Autoren der analytischen Sprachforschung zur Sprache kommen. Es geht nicht um bloße Deskription der Gegebenheiten, also nicht z.B. um die vielfältigen Möglichkeiten der Interpretation eines Satzes, oder um die Möglichkeiten, ein Wort von einem anderen hinsichtlich seiner ‚Bedeutung' zu unterscheiden, sondern um den Prozess der Benutzung von Sprache in kommunikativen Zusammenhängen. Indem Hausser dies auseinanderlegt, wird auch deutlich, auf welche Weise die Computerlinguistik Begriffe wie Bedeutung, Interpretation und Verstehen klären kann.

Eine Konsequenz des modelltheoretischen Ansatzes liegt darin, dass das Wissen von Sprecher und Hörer in der Kommunikation intern repräsentiert sein muss. Hausser schlägt hierfür eine *database* vor, die gleichsam das interne Weltmodell des Sprechers und des Hörers enthält. In dieser sind propositionale Ausdrücke enthalten, von Hausser *proplets* genannte, sowie *woplets*, wortbezogene Ausdrücke, und Begriffe, die sog. *coplets*. Zwischen diesen drei ‚Entitäten' der systeminternen Datenbank und den früher schon definierten Konzepttypen M-concept und I-concept bestehen Beziehungen, die ausführlich erläutert werden. Ebenso wird in einer ganzen Reihe von Modellen vorgeführt, wie sich der Autor den Ablauf der syntaktisch-semantischen Interpretation vorstellt. Auch hier wird, wie schon in anderem Zusammenhang bemerkt, leider kein Bezug zu einschlägigen Modellen der kognitiven Psychologie der 80er und 90er Jahre hergestellt, zu denen ähnliche Überlegungen führten.

Der Einbau in den kommunikativen Zusammenhang soll schließlich in Haussers Gesamtgebäude die Pragmatik leisten. Hausser nennt dies ‚pragmatische Interpretation' (S. 467). Allerdings bestehen hier noch begriffliche Unsicherheiten, denn offenbar bedeutet ‚pragmatische Interpretation' Verschiedenes, je nachdem, ob man sich im ‚hearer mode' oder im ‚speaker mode' befindet (S. 484 ff.). Es hat den Anschein, als ob Pragmatik hier ganz anders verstanden wird, als es seit Morris, Austin und Searle in der Linguistik üblich ist. Auch hier fehlen leider klärende Worte.

Was mit Haussers Buch insgesamt für die Computerlinguistik gewonnen ist, kann man vielleicht am besten verdeutlichen, wenn man einige Leitideen, die das Buch durchziehen, nochmals zusammenfassend beurteilt.

Eine erste Leitidee besteht darin, dass das Buch ein Gesamtkonzept der sprachlichen Kommunikation und ihrer Beschreibung mit algorithmischen Mitteln präsentieren will. Gegenüber anderen Ansätzen ist dabei ungewöhnlich, dass maschinelle Sprachanalyse nicht als „Spiel" angesehen wird, in welchem das analysierende System alle strukturellen und interpretatorischen Möglichkeiten einer Äußerung abzuleiten und Mehrdeutigkeiten aufzulösen hat. Vielmehr findet Sprachanalyse zwischen Kommunikationspartnern statt, von denen einer ein Mensch, der andere auch ein Computer sein mag. Hausser konstatiert, dass wir erfolgreich kommunizieren, mit Menschen und mit Computern. Daraus ist die Aufgabe der Computerlinguistik abzuleiten, nämlich zu klären, wie Sprache dabei funktioniert. Die Auflösung von Problemen, auch solchen der Sprachanalyse selbst, findet im übergeordneten – pragmatischen – Problemlösungsprozess durch den Nutzer (also den Interaktionspartner Mensch) statt, für den die Sprachanalyse die notwendigen Strukturen ermittelt. Ein solches Gesamtmodell ist begrüßenswert. Es liegt damit im Grunde die Ausformulierung einer Idee vor, die schon Ende der 70-er Jahre in Gerold Ungeheuers M-C-Modell der Computerlinguistik entwickelt wurde. Es handelte sich dabei um den Vorschlag, sprachverarbeitende künstliche Systeme nicht als autonome Systeme zu konzipieren, sondern als Systemverbünde mit den Partnern Mensch (M) und Computer (C). Natürlichsprachliche Prozesse sollten in Form problemlösender Interaktionen zwischen Computer und Mensch stattfinden, also in Form pragmatischer Prozessen, die den Benutzer einbeziehen und für die durch die Maschine ,nur' die strukturell aufbereiteten Daten bereitgestellt werden. Computerlinguistische Lösungen haben in diesem Modell nichts mit der Beschreibung der angeborenen Sprecher-Hörer-Mechanismen und Kompetenzen zu tun, eine These, die mit Haussers vehement vorgetragenen Argumenten gegen den von ihm so genannten Nativismus übereinstimmt.

Als weitere Leitidee des Buches ist auszumachen, dass gleichsam als Manifestation des vorgetragenen Konzepts ein ,kognitiver Agent' modelliert wird, der mit Sprache begabt ist und sich in einer Umwelt bewegt. Die ,Sprachkompetenz' dieses kognitiven Agenten folgt aber nicht den Ideen der ,nativistischen' Theorien, sondern den Prinzipien der Linksassoziativen Grammatik, seine kognitive Kompetenz folgt denen der modelltheoretischen Semantik. Wenn man auch über die Wertungen streiten kann, die das Buch in Bezug auf andere Grammatiktheorien enthält, so sind doch die vorgetragenen Grundgedanken plausibel, und die Konsequenz, mit der dieses Konzept durchgehalten wird, bewundernswert.

Eine letzte Leitidee, die hier zu würdigen bleibt, wird durch den Anspruch des Buches markiert, neben der Präsentation einer computerorientierten Kommunikations- und Sprachtheorie sich auch als Lehrbuch zu eignen. Dieser Anspruch wird erfüllt. Denn es finden sich, wie schon dargestellt werden konnte, neben den theoretischen ,Hintergründen'auch viele einfache praktische Beispiele und Erklärungen auch kompliziert scheinender Sachverhalte (z.B. die Ableitung der drei elementaren Grammatikformalismen). Der Lehrbuchcharakter des Buches wird auch durch die Übungen herausgestellt, die sich an jedem Ende der 24 Kapitel finden und die weniger dem schulischen Abprüfen des ,gelernten' Stoffes dienen, als vielmehr zu weiterem Nachdenken anregen sollen. Allerdings ist Haussers Buch eher in die Kategorie der anspruchsvollen Lehrbücher einzuordnen. Denn es enthält sehr viele Querbezüge (z.B. zur Semiotik, zu Karl Bühler, Tarski, Frege etc.) und Hinweise auf Hintergründe, die der Anfänger erst nach ergänzender Lektüre verstehen und einordnen kann. Dies betrifft z.B. die Ursprünge der PSG in der Automatentheorie und in der Theorie formaler Sprachen bei E. Post 1936, die Hintergründe der Theorie Chomskys (vgl. S. 142) und den oben schon genannten Komplex der modelltheoretischen Semiotik mit all ihren logik- und philosophiegeschichtlichen Facetten. Der Autor spannt damit einen ,Background' auf, der dem wirklich Interessierten hervorragende Anregungen zu weiterem Nachdenken und Nachlesen gibt. Wenn das Buch also in vielerlei Hinsicht auch einfachere Grundlagen der Computerlinguistik verständlich darstellt, so liegt doch vor allem eine anspruchvolle Gesamtdarstellung vor, die zu Recht die Bezeichnung *foundations* führt. Es wäre freilich zu wünschen, dass der kognitive Agent, den dieses Gesamtkonzept modelliert, auch in seinen komplexeren Teilen (der semantischen und pragmatischen Interpretation) als Pilotsystem realisiert werden könnte.

Bericht über die 11. Jahrestagung der GLDV

Multilinguale Corpora: Codierung, Strukturierung, Analyse

vom 8.–10. Juli 1999 in Frankfurt

Nachdem vor zwei Jahren, vom 21.-24. Oktober 1997, mit der „6th International Conference about the Use of Computers in Historical and Comparative Linguistics“ (6. Internationale Konferenz über den Gebrauch von Computern in der Historischen und Vergleichenden Sprachwissenschaft) bereits eine arbeitskreisinterne Tagung an der Universität Frankfurt stattgefunden hatte, war es eine besondere Ehre für unser Fach, nunmehr die 11. Jahrestagung in Frankfurt zu organisieren.

Es hatten sich etwa 90 Teilnehmer, darunter 45 Vortragende, zur Tagung angemeldet. Viele davon waren Mitglieder der GLDV, viele Studierende und Hochschulangehörige, auch aus dem Ausland (den weitesten Weg hatte Tom Lai aus Hong Kong). Es kamen aber auch Vertreter aus dem kommerziellen Bereich, wo an linguistischen, übersetzungstechnischen oder Kodierungsproblemen usw. gearbeitet wird.

Der vielbeachtete Festvortrag „The Encoding of Language Corpora: The TEI Recommendations in Principle and Practice“ wurde am Abend des Donnerstags (8.7.) von Lou Burnard vom Oxford Text Archive gehalten. Die Vorträge in den Arbeitskreisen verteilten sich während der drei Tage auf 2 große und 3 kleinere Hörsäle und gliederten sich in folgende Sektionen:

Codierung und Anordnung (3 Vorträge), Statistische Corpusanalyse (4), Texttechnologie (6), Lexikographie (6), Sprachspezifische Vorträge (5), Analyse und Generierung (5), Multi-linguale Corpora (7 Vorträge) sowie 4 Präsentationen. In einer eigens organisierten Sektion „Translation Memory Systems“ wurden in 9 Vorträgen die neuesten Entwicklungen auf dem Gebiet der automatischen Sprachübersetzung vorgestellt.

Leider hatten wir auch 4 kurzfristige Absagen, dies gab den Teilnehmern jedoch Zeit für Diskussionen und eingehendere Beschäftigung mit den Präsentationen.

Es versteht sich von selbst, daß eine solche Tagung mit den vergleichsweise geringen Mitteln, die uns zur Verfügung standen, nur schwer zu realisieren war; herzlich gedankt sei deshalb unseren Sponsoren, die uns unterstützt haben: die BHF-Bank mit einem großzügigen finanziellen Zuschuß, die Frankfurter Sparkasse, die Hessische Landesbank, und die Firmen Compaq und Misco sowie der Verlag Mouton deGruyter mit Sachspenden, wie Tagungsmappen, Blocks, Bleistiften, Kugelschreibern usw.

Dank der Kooperationsbereitschaft der Autorinnen und Autoren konnte in erstaunlich kurzer Zeit – nicht einmal ein halbes Jahr später – der Tagungsband[1] erscheinen, indem Beiträger und Herausgeber in ständigem e-mail-Kontakt standen; so konnten auch ursprünglich nicht vorgesehene Korrekturen eingearbeitet werden. Bei einem Band dieses Umfangs (knapp 400 Seiten) ist natürlich eine hundertprozentige Fehlerfreiheit unmöglich; es mag als Tribut an die erforderliche Aktualität der wissenschaftlichen Beiträge gelten.

Peter Olivier
Vergleichende Sprachwissenschaft der
Johann Wolfgang Goethe-Universität Frankfurt am Main
e-mail: olivier@em.uni-frankfurt.de

1 Multilinguale Corpora. Codierung, Strukturierung, Analyse. 11. Jahrestagung der Gesellschaft für Linguistische Datenverarbeitung. Hrsg. von Jost Gippert in Verbindung mit Peter Olivier. Praha: Enigma Corp., 1999. ISBN 80-86126-04-8. Preis: US$ 48,-. Siehe: http://titus.uni-frankfurt.de/curric/gldv99/papers.htm (Tagungsband) sowie http://titus.uni-frankfurt.de/curric/gldv99.htm (Index).