TOWARD A THEORY-BASED NATURAL LANGUAGE CAPABILITY IN ROBOTS AND OTHER EMBODIED AGENTS: EVALUATING HAUSSER'S SLIM THEORY AND DATABASE SEMANTICS

by

Robin K. Burk

A Dissertation
Submitted to the University at Albany, State University of New York
in Partial Fulfillment of
the Requirements for the Degree of
Doctor of Philosophy

College of Computing and Information
Department of Informatics
2010

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Toward a Theory-Based Natural Language Capability
in Robots and Other Embodied Agents:
Evaluating Hausser's SLIM Theory and Database Semantics

by

Robin K. Burk

DEDICATION

To Roger Burk, whose encouragement and support never wavered
despite all the weekends we didn't spend relaxing together,

And to Charity Burk, whose keen literary critical perspective
challenged more than one assumption of mine over the years.

My love to you both.

ACKNOWLEDGEMENTS

This dissertation would be incomplete if it did not start by acknowledging the role of Andy Haas, whose intellectual curiosity first attracted me onto a path that sent me off programming Bayesian object recognition and a left associative grammar in a robot named Darth Brooks and later sparked an examination of the role of logic in computational grammar. Thank you!

I'm grateful to my other committee members, each of whom contributed to my doctoral program and directly or indirectly to this dissertation topic as well. Istvan Kecskes supported independent study that introduced me to the computational language challenge from a pragmatics and cognitive linguist's perspective. His discussion of the characteristics of English when used by non-native speakers as a *lingua franca* sent me off to learn more about the interaction of culture, cognition and language. Neil Murray first introduced me to the foundational question of what problems are and are not computable. And Kevin Knuth's focus on Bayesian data analysis and his background in intelligent systems for space exploration raised questions about how best to bridge the gap between the probabilistic layers in many robot architectures and higher level functions such as planning and natural language.

Two other Albany faculty deserve special acknowledgement. Jagdish Gangolly supported independent study in description logic-based ontologies and formal concept analysis in parallel with a real-world project in a complex domain that illustrated both the usefulness and the limitations of current semantic technologies. And Ozlem Uzuner's insightful questions about method and metrics gave focus to my analysis. My thanks to you both.

CONTENTS

TABLE OF FIGURES

ABSTRACT

Computational natural language understanding and generation have been a goal of artificial intelligence since McCarthy, Minsky, Rochester and Shannon first proposed to spend the summer of 1956 studying this and related problems. Although statistical approaches dominate current natural language applications, two current research trends bring renewed focus on this goal. The nascent field of artificial general intelligence (AGI) seeks to evolve intelligent agents whose multi-subagent architectures are motivated by neuroscience insights into the modular functional structure of the brain and by cognitive science insights into human learning processes. Rapid advances in cognitive robotics also entail multi-agent software architectures that attempt to parallel in many ways the sensory and cognitive processes of humans. Natural language capability is a key objective for both types of software, whether embodied in a physical robot or in a virtual world that emulates features of the physical environment.

Hausser's SLIM theory of natural language communication and associated Database Semantics computational instantiation are an ambitious attempt to bridge the gap between formal theory approaches to computational natural language capability and an embodied approach to language and meaning which requires integration of language with sensory perception, planning and social interaction. This dissertation evaluates Hausser's approach to the development of human-level computational natural language capability in embodied and socially situated agents and argues that a theoretical basis for such capability is emerging as a result of recent evidence from linguistics, cognitive science and neuroscience.

1. Introduction

The use of language to communicate is a key human capability. Many organisms are able to communicate with others of their kind through sounds, visual displays, odors and chemical excretions. But only humans, so far as we can tell, have the ability to truly create and use language, with its large vocabularies, syntaxes that allow for combinations of words into vast numbers of complex utterances and, above all, indirect or symbolic references not only to concrete facts but also to abstract concepts and hypothetical or counterfactual ideas.

Since as early as the mid 1950s computational understanding and production of natural language has been a key goal of both computer scientists (Figure 1) and computational linguists. There are many reasons for attempting computational language capability, among them to facilitate information retrieval and machine translation, to provide a human-friendly user interface to equipment and software systems, as capabilities inherent in intelligent artificial agents and as a means to validate theories about the nature of language and of cognition itself.

However, the complexity and diversity of natural language present several daunting challenges both for formal linguistic theory and for computation. Languages differ significantly from one another in syntax, semantic categories and pragmatic use as well as in details of phonology, word morphology and specific vocabularies. Moreover, within a given language, syntactic irregularities are common, as are multiple meanings

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

August 21, 1955

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it for a summer.

John McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M Corporation
C. E. Shannon, Bell Telephone Laboratories

The following are some aspects of the artificial intelligence problem:

2. **How Can a Computer Be Programmed to Use a Language**

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning and rules of conjecture. From this point of view, forming a generalization consists of admitting a new word and some rules whereby sentences containing it imply and are implied by others. This idea has never been very precisely formulated nor have examples been worked out …

3. **Neuron Nets**

How can a set of (hypothetical) neurons be arranged so as to form concepts. Considerable theoretical and experimental work has been done on this problem by Uttley, Rashevsky and his group, Farley and Clark, Pitts and McCulloch, Minsky, Rochester and Holland, and others. Partial results have been obtained but the problem needs more theoretical work …

6. **Abstractions**

A number of types of “abstraction” can be distinctly defined and several others less distinctly. A direct attempt to classify these and to describe machine methods of forming abstractions from sensory and other data would seem worthwhile.

Figure 1. Excerpts from the 1955 Dartmouth Summer Proposal on Artificial Intelligence

(polysemy) associated with individual word forms. And beyond these structural characteristics, meaning is often conveyed through language- and culture-specific metaphors, ironic phrasing and paralinguistic elements that are difficult even for humans to master when acquiring second language competency.

These challenges, while significant, have not discouraged researchers from attempting computational approaches to understanding and producing natural language. Progress has been much slower than was hoped in the optimistic days of the Dartmouth Proposal, however.

A variety of approaches to computational understanding and production of natural language have been proposed only to fall short in significant ways, beginning with model theoretic representations of meaning and generative grammars, through connectionist and other probabilistic learning models to the statistical natural language approaches which dominate in current research. Although significant progress has been made in specific application areas such as speech recognition, information retrieval and statistically-based machine translation, integrated human-level computational language capability remains an elusive goal.

Two current research trends bring renewed focus on this goal. On the computational side, the nascent field of artificial general intelligence (AGI) seeks to evolve intelligent software agents whose multi-subagent architectures are motivated by neuroscience insights into the modular functional structure of the human brain and by cognitive science insights into human learning processes. Rapid advances in cognitive robotics also entail multi-subagent software architectures that attempt to parallel in many ways the sensory and cognitive processes of humans. Natural language capability is a key goal for such software, whether embodied in a physical robot or in a virtual world which emulates features of the physical environment, as in many AGI projects.

In parallel with these trends and with the neuroscience and cognitive science which inform them, a significant amount of recent linguistic theory and research has been focused on the pragmatic use of language, on language acquisition both by children and by adults learning additional languages, and on the interaction between language acquisition, language use and cognition.

This research activity raises a fundamental question. Is there a theory-based approach on which to base the development of human-level computational natural language capability? Specifically, on what theoretical basis (if any) can we construct *embodied* software agents that are or can become capable of human-level natural language use?

Such an agent would be capable of both understanding and generating language utterances within a social and physical context. Its embodiment (whether in a physical or virtual sense) implies that its language use occurs within a specific context of space and time and that it has the ability to sense its external (and potentially changing) environment and to act upon it – a significant difference in emphasis from earlier logic-based attempts at computational language capability of the sort envisioned in the

Dartmouth proposal and later proffered by *e.g.* Montague. Embodied language capability also implies a distinction between the agent's interior context and the external environment. Interior context of some sort is required for such language and cognition features as past or future reference in time, counterfactual or hypothetical statements and the agent's ability to link or contrast the content of a linguistic utterance and knowledge regarding the state of the external environment. Such an agent would be both embodied and cognitive in nature (unlike simple reactive robots and software agents), possessing some analogue to the conceptual and deductive capabilities that characterize human cognition.

A theoretical basis for computational natural language capability, if such a theory is achievable, should be independent of any specific language or culture and would offer at least two benefits. First, by asserting general principles regarding the nature of natural language and the mechanisms on which natural language communication is based, such a theory would offer a unified approach by means of which to manage the overwhelming complexity and diversity among languages and within specific natural languages as they are actually used by native speakers. And second, implementation of software agents based on the theory could in turn provide a means to test the linguistic and cognitive assumptions on which the theory rests.

On what basis should such a theory be developed? Since natural language is a complex human phenomenon that occurs within human social contexts and that makes reference to concepts and experiences shaped by human sensory perception and cognition, it is reasonable to assume that a theory capable of offering a useful basis for computational natural language would benefit from the insights of neuroscience, cognitive science and linguistics – in other words, from the study of human language users. It is not surprising, therefore, that current attempts towards such a theory occur primarily within AGI and robotics research, where the embodied nature of natural language in human use is being translated to embodied computational sensory perception and cognition. These efforts begin with computational techniques and apply them to natural language as just one of many functions required for the desired intelligent software agent.

An alternate approach is that of Roland Hausser, a computational linguist. Hausser begins not with computation but rather with natural language itself as a complex communication capability characteristic of humans. He asks: what principles underlie natural language communication? How does a speaker formulate an utterance to convey a desired message and how does a hearer decode that utterance into meaning? What role does the speaker's and hearer's external environment play in this process? What is the nature of the internal context within which utterances are formulated and decoded?

Hausser has proposed the SLIM (Surface, Linear, Internal, Matching) theory of natural language communication and an instantiation of the theory in computational mechanisms he calls Database Semantics. He proposes that the best evaluative test of this theory is in fact to instantiate (growing) fragments of various natural languages within software agents that would ideally be fully embodied but that in any case are intended for embodiment. Unlike the perception and learning-based approaches of most roboticists and AGI researchers, Hausser’s work has its roots in formal logic and model theoretic semantics. However, unlike many previous researchers who have attempted logic-based approaches to computational natural language capability in the past, Hausser also has evolved several deep critiques of the assumptions inherent in model theory for natural language. In particular, he focuses on a clear delineation between the semantics and pragmatics of natural language in use, proposing specific mechanisms by which these two dimensions interact with sensory perception on the one hand and cognition and higher logic-based functions such as planning on the other hand. In this sense his theory has both formal/logical and embodied characteristics.

Hausser’s work raises the important question of how any theoretical basis for computational natural language capability can be judged, short of a full implementation of intelligent embodied software agents which have fluency in a natural language and in physical interactions with the external world. Since such agents are not yet achieved, do we have any basis on which to judge either his computational natural language theory and implementation or the work of AGI and robotics researchers in this area?

This dissertation proposes that such a basis is in fact emerging as a result of the combined insights of recent neuroscience, cognitive science, linguistics and developmental psychology research. Chapter 2 examines the historical development and the motivation for Hausser’s approach. Chapter 3 examines Hausser’s SLIM theory of natural language and the Database Semantics computational mechanisms that he proposes. Chapter 4 analyzes Hausser’s model in light of recent work in linguistics theory. Chapter 5 considers the implications of recent work in cognitive and neuroscience with regard to the goal of natural language capability for physical robots and for AGI software agents. And Chapter 6 concludes by suggesting a way forward for future research towards fully fluent, embodied computational natural language capability.

2. Development and Motivation for Hausser's Approach

Hausser's work has from the beginning been theory-based, having its roots in formal logic. The nature of his theory has evolved in significant ways, however, over several decades. In order to evaluate Hausser's mature theory regarding natural language communication and his associated computational model it is useful to understand how and why he has adopted the positions he currently espouses. This chapter traces the evolution of Hausser's work, taking his own perspective as he seeks to formulate and justify a theoretical basis for computational linguistics. Chapter 3 will describe his mature theory and implementation approach.

2.1 The Development of Model Theory and Its Application to Natural Language

In his early papers Hausser assumes he is addressing an audience familiar with model theory and its use for natural language representation. We needn't review in detail the twentieth century debate regarding formalism; however a summary of developments in formalism provides a context for understanding Hausser's earliest work and his later critique and partial rejection of standard model theory as a representation of natural language.

The formalist program arose out of an attempt to provide a rigorous and unambiguous foundation for mathematics. Since the time of Euclid, Western mathematics had long been considered the epitome of rational thought, offering clear definitions of key concepts via a set of axioms and explicit, rigorous deduction in the form of theorem proofs. Philosophers and theologians openly sought to emulate these characteristics in their own fields. This view of mathematics was brought into question, however, by the development in the nineteenth century of non-Euclidean geometries which, while not consistent with our sensory intuitions about the physical world around us, appeared to be logically consistent and therefore undermined the assumption that mathematics is true because it accurately describes the physical world which, being in existence, is logically consistent within itself.

The development of non-Euclidean geometries along with Riemann's work on number theory, infinities and countability, pushed defenders of mathematics as a rigorous discipline to seek the rigor and meaning of a mathematical system in the logical consistency of formally axiomatic representations of the system rather than in its correspondence to the 'real world'. Such representations made use of the new disciplines of symbolic logic and set theory, which were intended to avoid the ambiguities inherent in natural language-based argumentation.

Early efforts in mathematical logic and formal set theory, such as the work of Cantor and Frege, were however unable to deal satisfactorily with some logical paradoxes that were known to the ancients. Russell provided a partial solution to this problem through his analysis of the Epimenides paradox. Statements such as **This statement is false** demonstrate the limits of naïve understandings of truth and falsity for mathematical reasoning. Russell's solution was to recast such statements in terms of set membership, asserting that a set cannot be a member of itself and giving rise a few years later to the hierarchical type theory described in his and Whitehead's *Principia Mathematica* . During the same few years Zermelo's axiom of choice and Fraenkel's axiom of replacement provided a cardinality-based approach to axiomatizing set theory.

If the axioms of a mathematical system could be stated in set theoretic terms, and if set theory could be reduced to some system of formal logic (as Russell and Whitehead attempted in their *Principia* using the first order predicate calculus) then, it was believed, the rigor and consistency of that system could be demonstrated objectively. Hilbert in particular championed a program to axiomatize all areas of mathematics for this purpose and considerable attention was given in the second decade of the twentieth century to axiomatic definitions and proofs of consistency for key areas of mathematics.

In order to formalize an axiomatic representation one needs a non-ambiguous, symbolic language. Early attempts centered on representing propositions and combining then into valid chains of inference. However, propositional logic is inadequate for many forms of deduction because it does not easily represent quantifiers. What was wanted was a way to formally represent syllogisms such as:

All men are mammals.
John is a man.
Therefore, John is a mammal.

The first order predicate calculus (first order logic) was developed for this purpose. First order logic is to propositional logic as algebra is to arithmetic: it provides a way to define and reason about classes of propositions just as an algebraic function represents a class of possible calculations. In both cases this is accomplished through the use of variables whose values may be drawn from a specific domain along with operators on those variables. The predicate calculus describes a system of interest in terms of:

- an infinite set of variables
- quantifiers (such as $\forall$ 'for all' and $\exists$ 'there exists at least one')
- logical connectives (such as $\wedge$ for conjunction 'and', $\vee$ for disjunction 'or', $\rightarrow$ for implication, $\leftrightarrow$ for biconditional implication 'if and only if' and $\neg$ for negation 'not')
- punctuation marks such as parentheses to control the order in which connectives are applied
- an identity or equality symbol (=)
- grammatical rules that govern the syntax of a formal language defined using these elements.

In first order logic, formalizations often also include the constants $\top$ for 'true' and $\perp$ for 'false' which are either provided explicitly or defined in terms of the quantifiers in order to provide a basis for deduction.

The intent in axiomatic representation is to express the content of a particular mathematical system by assigning mathematical objects/meanings to variables and by defining the base axioms of the system in the resulting language. A key objective of formalization is to show the validity of this system, *i.e.* its logical consistency. Given the usually infinite number of variables involved, or of referents for those variables, such consistency generally cannot be established by examination. An alternate approach is to show that the set of logical formulae is satisfiable, *i.e.* that there is at least one non-self-contradictory object for which all of the formulae are simultaneously true.

Gödel provided three key initial contributions to the formalist effort. In 1929 he published his completeness theorem which established the correspondence between syntactic provability and semantic truth (validity, satisfiability) for the first order predicate calculus. He shortly thereafter also published his compactness theorem which established the finite nature of logical consequence and of the semantic models that can be defined using the syntax of first order logic. Together these theorems provided encouragement that the attempt to represent and prove the validity of mathematics by representing it in a formal language based on logic was viable.

Gödel also, however, demonstrated the limitations of the axiomatic formalist program in his famous paper on formally undecidable propositions. This theorem proved that, even in the apparently simple case of integer arithmetic, an axiomatic representation of the system contains statements that we know are true but which cannot be formally proven to be logically entailed in an explicit axiom set. One way to restate the import of this proof is that it demonstrates that for any mathematical system of interest, its formal representation defined using first order logic will always contain semantic content which cannot be completely formalized within that representation. Wittgenstein made similar criticisms of Russell's type theory when he noted that despite the intent of reducing mathematical content to the syntax of formal logic, Russell makes reference in his symbolic rules to the meaning of their signs.

In response to these criticisms of the axiomatic program, Carnap, Weyl and Tarski adopted and expanded on the work of Löwenheim and Skolem in model theory. Model theory responds to the problem of unformalizable meaning in axiomatic systems by distinguishing between a set of axioms defined in a given logical language and the mathematical system which they are intended to represent.

> What is usually called the construction of a model for a postulate set is the construction of an interpretation for this syntactical part. (Carnap 1942)

An interpretation consists of assigning entities to variables and meaning to predicates in an otherwise abstract set of formulae.

Tarski further refined this understanding with a more rigorous set of definitions:

> A possible *realization* in which all valid sentences of a theory T are satisfied is called a model of T.
>
> Consistency and completeness can also be characterized in terms of models: a theory T is consistent if and only if it has at least one model; it is complete if and only if every sentence of T which is satisfied in one model is also satisfied in any other model of T. Two theories T_1 and T_2 are said to be *compatible* if they have a common consistent extension; this is equivalent to saying that the union of T_1 and T_2 is consistent. (Tarski 1953b)

A theory or set of logical formulae must be expressed in a formal language. Tarski took the important step of asserting that in fact we need at least two such languages: the object language we intend to use for the formulae plus a meta-language that governs the creation of object languages, to include a definition of ‘truth’, or what he called the ‘convention T’.

Thus the problem of logical consistency and completeness is addressed at the meta-language level, resulting in a formal object language in which to express a theory (set of axioms and theorems stated as logical formulae). When the variables of the theory are given semantic meaning through reference to a model, the result, one intends, will describe the mathematical system of interest while avoiding logical paradoxes entailed in self-reference.

Tarski’s model theory supports recursive definition of languages, in which one can develop meta-meta-languages of increasing abstraction. This is possible when elements of the meta-language are treated as logical constants to be interpreted by reference to a model structure, rather than as primitives. Rather than assuming an intuitively agreed-upon meaning for logical constants such as the connectives and quantifiers, a model structure allows various interpretations of them, requiring that a given interpretation (model) define these basic elements with regard to a given domain of objects and a given definition of logical relations. For instance, to interpret quantifiers such as $\forall$ ’for all’ and $\exists$ ‘there exists at least one’ requires a definition of what elements within a domain these quantifiers range over. In this way mathematical concepts of increasing abstraction and power can be formally represented and deductions drawn from them. Among the most abstract of these is type theory which offers a very high level logic with which to group logical expressions into types defined by the ways in which they can combine into well-formed formulae. Church's type theory and λ calculus in particular have been fruitfully applied for this purpose both by mathematicians and by computer scientists in programming language theory.

Tarski's approach to model theory was broadened by Kripke's work in intensional and specifically modal logic. Intensional logic distinguishes between the standard quantifiers that range over all of the objects in a universe (extensions) and those that range over the possible terms that can have those objects as their value (intensions). Hence the intensional logics are of higher order than the first order logic described earlier. These logics include functors, *i.e.* incomplete expressions with arguments that can be filled in to create more complete expressions. Intensional functors added to axioms about extensions, along with Church's λ calculus used to describe the properties of functions and functors, form the basis for many meta-languages which in turn are used to define and prove the characteristics of specific formal languages and their models of interest.

Intensional logic is useful in part because it provides a way to formalize Frege's distinction between *Sinn* (sense) and *Bedeutung* (reference). This distinction attempts to distinguish an object being referred to from the way in which that reference occurs in a particular expression. In particular, descriptions and names may not, in some cases, be automatically substituted for one another. Intensional reference, in which an expression is used in place of the name of an object, may occur as an intermediate step in a mathematical proof but it also occurs quite frequently in natural language, as the following example demonstrates:

(1) **Mary saw the Evening Star.**
(2) **Mary saw Venus.**

(1) and (2) do not have exactly the same meaning despite the fact that **Evening Star** generally is taken to refer to Venus when it appears just after sunset. Here the sense (form of expression) is different from the reference (object referred to).

Because intensional logic does not assume a fixed set of logical connectives with established meanings it can also be extended to include logical elements beyond those of first order logic. For instance, modal logics include elements of modality such as necessity and possibility and temporal logics can qualify propositions in terms of time and tense.

Frege explicitly intended his *Sinn* / *Bedeutung* distinction for mathematical statements and like Tarski after him was skeptical regarding the possibility of formalizing natural language. Montague, who took his Ph.D. under Tarski, had no such hesitation:

> There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed, I consider it possible to comprehend the syntax and semantics of both kinds of language within a single natural and mathematically precise theory. (Montague 1970a; see also Montague 1970b)

Montague proposed a syntax of expression categories and syntactic rules for their combination based on Church's type theory. He then defined a corresponding intensional logic in which every rule defines a type and the syntactic categories have a corresponding type in the intensional logic. Semantic interpretation of a sentence consists in first mapping it into this intensional logic and then interpreting the intensions of terms via a model which is defined with reference to the expression's context or index, *i.e.* the time, place and facts about the environment within which it occurs. This mapping provides the link between language and meaning.

Montague grammars, or more broadly any categorial grammars applied to artificial or natural languages, have the useful property of being weakly equivalent to the context free grammars of computer science. A context free language allows expressions to be nested but not to overlap, a characteristic which Chomsky originally posited as obtaining in the deep structure of natural language and which supports parsing of programming languages by compilers. Montague's work in this area was presented in a series of three publications beginning in 1970 and ending with his famous PTQ ("Proper Treatment of Quantification in Ordinary English") in 1973.

2.2 Hausser: Early Work

While Montague was proposing his model theoretic approach to natural language, Roland Hausser was pursuing his doctoral work in theoretical linguistics at the University of Texas at Austin, where he defended his dissertation in 1974. It was an exciting time to be in Austin, where the faculty included Emmon Bach, Stanley Peters, Robert Wall and Lauri Kartunnen and where Hausser's fellow graduate students included David Dowty, Per-Kristian Halvorsen and Hans Uszkoreit[1], each of whom has made significant contributions to language theory or computational natural language practice in the years since.

Although he contributed as a graduate student to a project in transformational grammar led by Bach and Peters (Hausser 1971), in his own research Hausser adopted a model theoretic framework with an initial paper that addressed existence presuppositions using Kripke's modal logic of intensions (Hausser 1973). In that same year Montague published his PTQ. Shortly afterwards Hausser's second and third papers and subsequently his dissertation addressed quantification in Montague grammars extended to clarify the existential status of referents. The treatment of distributive and collective plural nouns in (Hausser 1974b) in particular foreshadows Hausser's mature approach in its concern to assign features in semantic formalisms so as to mirror the intuitions of native speakers while also seeking representational simplicity where possible, which he accomplishes in this case by assigning the feature of number to quantifiers rather than to nouns in order to clarify co-reference between collective and distributive plural noun phrases.

As he continued to extend model theoretic formalisms for the semantics of natural language Hausser began also to explore the role of pragmatics, partly in response to speech-act theorists such as Austin and Searle. In a paper addressing the treatment of non-declaratives in a PTQ-style intensional logic (Hausser 1978), Hausser makes an initial attempt to articulate the distinction and the relationship between semantics and pragmatics. He criticizes the assertion of Austin and Lewis that there is a *semantic* equivalence between sentence pairs such as

(3) **I order you to leave.** (declarative)
(4) **Leave!** (imperative)

on the grounds that while (3) denotes a proposition, (4) does not. Their relationship, Hausser contends, is one of overlapping use-conditions rather than of equivalent semantics. He concludes that it is a mistake to attempt to encode speech act properties in the semantic representation of syntactic mood. Semantics for Hausser at this point, as for others working in the model theoretic framework, has to do with formalizable meaning from which logical implications can be drawn. And at this point he is still attempting to formalize syntactic mood in terms of possible forms of denotation. Nonetheless, Hausser has already begun to seriously consider how pragmatics and semantics inter-relate.

In the 1978 paper Hausser also emphasizes two elements that will persist in all of his subsequent attempts toward formalizing natural language: the necessity of defining a formal system that supports both generation and interpretation of natural language in a unified formalism and the linguistic standard of basing such a system on the principle of surface compositionality. He writes:

> One reason why I have chosen to present my analysis of syntactic mood in form of an extension to PTQ is that PTQ is a *complete* grammar in the sense that the *generation and interpretation* of a fragment of English is *coordinated* in a rigorously formal generative system…

> The confusion of semantic properties and speech act features shows, furthermore, that in addition to the *methodological standard* of completeness we need some kind of *linguistic standard* to guide our use of mathematical power to linguistically motivated analyses. But which standard of linguistic analysis should we adhere to?
>
> The assumption that the semantics of natural languages works like the semantics of formal languages in that the meaning of complex expressions is the systematic result of the meaning of the basic parts (and the mode of syntactic combination) suggests a principle which I would like to call the principle of *surface compositionality*. According to this principle the semantic representation of a linguistic expression should contain nothing that does not have concrete surface syntactic motivation. Furthermore, a surface compositional analysis must characterize explicitly how the meaning of a *complex surface* expression is composed from the meaning of its *basic surface* constituents. (pp. 77-79; emphasis in original)

Compositionality was not a new idea when Hausser wrote this passage. Frege had proposed a version of this and it is a standard characteristic of formal logic. However, here Hausser is specifically asserting *surface* compositionality for *natural* language. He therefore rejects the transformational/generative approach that asserts the existence of deep structure in language beyond the surface structure of an expression, *i.e.* beyond the language expression as it is presented by the speaker. While Hausser's semantics / pragmatics distinction is quite underdeveloped at this early stage, it already provides a rationale for adopting surface compositionality and rejecting Chomsky. This distinction, overshadowed in this paper by some specific model theoretic concerns, will eventually form a central element in Hausser's mature agent-centric theory and computational model for natural language communication and will form the basis for his later ambitious claims regarding natural language capability for autonomous robots.

2.3 Propositions or Speech Acts?

Following his discussion of non-performatives Hausser briefly considered how Montague grammar might be extended in regard to question and answer discourse (Hausser and Zaefferer 1980). He then returned to and more broadly addressed the place of pragmatics in model theory (Hausser 1980). The latter paper is seminal in the development of Hausser's overall approach to computational linguistics in its call for a clear differentiation of syntax, semantics, lexicon context and pragmatics, a differentiation much to be desired, he notes dryly, since "of the components of grammar actually proposed in the literature, each has been expanded to handle a lot more phenomena than is advisable for its own good." He reiterates his call for a theory of discourse that delineates the boundaries of these components and that describes "how the different components interact in the course of interpreting the *use of an expression by a speaker relative to a context*".

Here again Hausser begins with extensions to Montague grammar, although he now critiques model theory as a whole for abstracting away from the use aspects of language. In this paper he also introduces a distinction he will maintain consistently thereafter between what he calls meaning$_1$ (the literal meaning of the surface expression) and meaning$_2$ (the communication effect of meaning$_1$). Using the notional goal of a computational 'speaker simulation device', he concludes there is a need to formalize *two* types of models, one for language tokens and one for referent contexts, and to provide referential mechanisms between them. Thus in response to the question “model theory or speech act theory?” Hausser begins to answer “aspects of both”.

2.4 Critiquing Model Theory

Hausser’s initial exploration of pragmatics led him to develop a new treatment of context in model-theoretic semantics. In (Hausser 1981a) he proposes to replace the traditional model theoretic approach, in which ‘denotation’ and ‘reference’ are used synonymously and in which the formal model operates as a substitute for reality, with an approach that *relativizes semantics to the speaker.* Rather than consider denotation conditions as instructions for determining the truth value of a sentence with regard to a model and index, as in Montague, here Hausser proposes to consider them as *instructions for synthesizing the sentence’s denotation in a lexical space.* Reference is then treated as the pragmatic process of matching that denotation to a context, *i.e.* to what the speaker “perceives and remembers at a given moment”, formalized as a model-theoretic structure in the same lexical space as the denotation of the surface expression. Hauser works out a formal treatment of indexical and anaphoric pronouns to illustrate how direct *vs.* denotational reference could be mapped to such a context.

In (Hausser 1981b) Hausser more completely describes and justifies his emerging theoretical stance. Returning to the notional speaker simulation device, he observes that model theoretic and speech act approaches to natural language have mirror-image shortcomings. Traditional model theory fails to provide satisfactory descriptions of language in use, while speech act theory cannot give an adequate account of the way in which meaning is linked to the form and semantics of expressions. A more appropriate theory, Hausser asserts, would be one in which a well defined semantic concept space is distinguished from the nonlinguistic, sensory based internal experience and memory of the speaker, to which semantic concepts are (more or less completely) mapped in pragmatic language use.

The remainder of (Hausser 1981b) is devoted to a close examination of the model theoretic implications of distinguishing denotation from reference. In a closely-reasoned critique of Tarski's treatment of the Sorites paradox Hausser finds that the paradox dissolves with the introduction of this semantics / pragmatics distinction into a model structure's formalization of meta-language. This, in turn, has both theoretical and practical implications for the formalization of natural language in model theoretic terms and for computational manipulation of that formal system.

Referring to his notional speaker simulation device (SID) Hausser writes:

> (O)nce the surface interpretation system, comprising a formalized natural surface language (object language) and an *operationalized* meta-language has been implemented as part of the SID, the infinite recursion of meta-languages inherent in Tarski's system will be of no further consequence for the processing of meaning by the SID. That is, the SIDs may continue to communicate with each other even if the native meta-meta-language use to build them (and specifically to define their formal meta-language) is suddenly forgotten and extinct. (Hausser 1981b, 162; emphasis added)

Buried in this formulation is a central assertion on which Hausser will base all of his subsequent work – an assertion that bears directly on our larger question of computational natural language capability for autonomous robots and embodied intelligent agents – namely that semantic processing of language must ultimately be grounded outside of the syntactic and semantic formalism adopted for representation and computation. In the notional SID this grounding would consist of implementation code run on a digital computer. In humans this grounding consists of direct experience of the external world gathered through non-verbal sensory capabilities of the human body. Whether agents capable of natural language communication are natural or artificial, Hausser will argue, this pragmatic reference underlies and is distinct from semantic denotation.

But why would traditional model theory be inadequate to represent natural language? Note that Hausser argues the necessity, and not simply greater practicality or existence in nature, of both semantic and pragmatic representations and mechanisms for natural language understanding and generation.

2.5 Two Contexts

Hausser offers a tightly reasoned critique of Tarski in (Hausser 1981b). The liar's paradox is one that must be addressed by any attempt to use formal logic as the fundamental basis for capturing meaning symbolically and reasoning therefrom. This paradox arises from the difficulty of establishing the truth condition of statements such as **This statement is false**.

Tarski attempted to resolve this paradox by rigorously defining the requirements for a meta-language and truth condition such that truth in the meta-language is distinguished from truth in the object language. He noted that the meta-language used must be essentially stronger than the object language if it is to provide a sufficient basis for this, leading to the possibility of an infinite regression of meta-languages. However, this potentially infinite regression is resolved, Tarsi asserts, because ultimately we rely on our intuition regarding the meaning of truth conditions in a given language when applied to a given model structure.

How clear is that intuition? In (Hausser 1981b) Hausser challenges Tarski's treatment of the Sorites paradox. This paradox emerges when categories are vaguely defined. If sand is dribbled in a spot, grain by grain, when does it constitute a heap (*soros* in Greek)? One grain does not constitute a heap. Adding a single grain more doesn't seem to either, nor does a third or a fourth. As a result we might paradoxically conclude that thousands of grains don't constitute a heap either. And yet that violates our intuitive sense of the meaning of 'heap'. Russell dealt with this paradox by proclaiming that logic does not apply to this situation (Russell 1923) because the terms involved are vague and Quine made a similar argument in support of his assertion that natural language cannot be formally represented.

Hauser's analysis of Tarski's treatment of Sorites hinges on cardinality assumptions in the proof but his critique is a broader and ultimately ontological one, and is of special import for formal representations of natural language.

We can, in standard model theory, define a formal model to represent any state of affairs we like, but we are constrained by the meaning of the terms under interpretation (words if we are formalizing natural language). Montague attempted to deal with this difficulty by adding meta-linguistic postulates that have the effect of eliminating from consideration those models that violate our intuitive meaning of the terms (the 'meaning postulate' approach). However, Hausser notes

> While the method of meaning postulates permits to maintain that assumption of the standard approach according to which the model structure is viewed as a representation of reality and the denotation conditions are viewed as instructions to find out whether a sentence is 1 or 0 relative to an index, meaning postulates are an extremely cumbersome method for formally implementing lexical interdependencies.

In other words, Montague's attempt to use postulates to constrain the possible referents of an expression is at best inadequate to cope with the richness of semantic relationships that are present in natural language and human cognition. For instance, the semantic relationship between 'man' and 'human' must be captured by one or more meaning postulates and so too *every other* semantic relationship possible between the terms in the model. This may be practical for some mathematical systems or for small fragments of natural language but it is utterly impractical for representing full natural languages, which have much richer sets of lexical terms and referents.

Standard model theory defines meaning as a direct reference or denotation between expressions in the object language and formal model-theoretic semantics. This raises the question: must the objects in the model be real or are they simply constructs? If they must be real, then the ultimate model structure is the external world, with an infinity of (changing) facts and details that cannot be completely examined in order to assert validity of an expression. But if they aren't real, then it is unclear how model theory can prove the truth of statements about the real world, which make up a large part of natural language in use.

Standard model theory also suffers from its inability to deal with context-dependent expressions (indexicals) such as pronouns, and hence of linguistic phenomena such as anaphora. Montague offered a coordinates approach in which space, time and other parameters are defined for each indexical term. Again, Hausser argues, this introduces unhelpful complexity without satisfactorily capturing our intuition about meaning:

> The coordinates approach permits to retain the assumption according to which meaning is a direct relation between expressions and referents by defining a context of use as an extended point of reference.
>
> The intuitive interpretation of a model structure as a representation of reality, however, suffers under the coordinate approach. Since the model structure is assumed to specify a state of affairs at an index, one would expect that this state of affairs *is* the context. Instead, the coordinates approach introduces a second kind of reference mechanism: while the denotation of regular constants is specified over the denotation function, the denotation of indexicals is specified over numerous additional context-coordinates. Furthermore, to define the context as an arbitrary n-tuple of external coordinates fails to capture the highly specific interaction between context-dependent expressions and a coherent context (*i.e.* situation).

Similarly, standard model theory has difficulty dealing with non-literal references in natural language such as metaphor, being forced to consider them as conditions of ambiguity. Here standard model theory faces both a cardinality issue when applied to natural language (the number of possible metaphors, ironic uses of language, etc. is unlimited) and also a failure to recognize the role of use-conditions in establishing the meaning of non-literal expressions.

The difficulty of applying standard model theory to natural language arises in part because of phenomena in natural language such as indirect reference and non-literal expressions that can be avoided in artificial languages. Hausser rejects, however, the assertion by Quine and other logical positivists that these difficulties demonstrate that natural language is illogical. Rather, he asserts, the problem is with an overly simplistic notion of reference. Lexical expressions *denote* concepts *within a semantic space*. Those concepts, in turn, *refer* more or less exactly *to the real-world context* that is perceived by humans through their sensory apparatus or remembered (accurately or inaccurately, completely or incompletely). This distinction was not deemed necessary by Tarski for formal languages that refer to mathematical objects but it is required, Hausser asserts, for natural language.

Hausser proceeds in the remaining sections of (Hausser 1981b) to provide a modified formal definition for meta-languages that conforms to his proposed denotational function described above, *i.e.* one in which the truth condition of a surface linguistic expression is determined with respect to a model that represents the speaker's semantic space rather than physical reality. He does this by proposing that the semantic model be treated as a partial model structure in which some elements are left in variable form. Thus we have two models: one that represents the model theoretic synthesis of the speaker/hearer's intuition of the literal meaning of expressions and the tokens that make them up (the token model) and one for the context, *i.e.* what the speaker/hearer perceives and remembers in a given utterance situation, with meta-linguistic functor relationships defined to link them. In this approach reference consists in matching the synthesized literal meaning of the surface language with the context – and that reference process, Hausser asserts, is a part of pragmatics, not semantics.

2.6 Grammar and Computational Implementation

With the combination of a rigorously formal treatment of semantics and (the goal of) an operationalized pragmatic mapping of semantics to the speaker's perceived experience of his external world context, Hausser has set the stage for his mature model for computational natural language communication. He has not yet, however, settled on appropriate formalisms at the syntactic level.

Hausser's first attempt at a surface compositional syntax, presented in (Hausser 1982) proposed a pure categorial, *i.e.* strictly context-free, grammar. However, unlike either Montague's syntax or the transformational grammars of the time, Hausser's ORTAX attempts to be surface compositional, applying an unusual, orthogonal parse tree structure in order to accommodate discontinuous but related words in expresssions.

In keeping with his coalescing theoretic stance Hausser followed the presentation of ORTAX with two papers that specifically address the criticism of Russell *et al.* to the effect that natural language, being inherently vague, cannot be rigorously formalized. Hausser again responds, as he did in (Hausser 1981b), by locating vagueness in pragmatic use mapping rather than denotation (Hausser 1983a; 1983b).

During the period of these publications Hausser held a position as Privatdocent (lecturer) at the University of Munich. In 1983, however, he was awarded a five year Heisenberg grant by the Deutsche Forschungsgeneinschaft in then-West Germany. This freed him to concentrate on research and publication, which he pursued as a visiting scholar in Stanford University's philosophy department (1983-4) and Center for the Study of Language and Information (CSLI, 1984-6), followed by the computer science department (1986-7) and the Laboratory for Computational Linguistics (1987-8) at Carnegie Mellon University. During his term at Stanford's CSLI he participated in the Foundations of Grammar project led by Lauri Karttunen; other participants included Mark Johnson, Ron Kaplan, Martin Kay, Fernando Pereira, Carl Pollard, Ivan Sag, Stuart Shieber, Hans Uszkoreit, Tom Wasow and Dietmar Zaefferer.[2] In 1989 he assumed the newly-created professorship for computational linguistics at the Friedrich-Alexander-Universität Erlangen-Nürnberg (Germany) where he continues to teach and research.

During the period of his Heisenberg grant Hausser turned his attention to computational implementations of natural language models. His first step was a more complete treatment of his ORTAX grammar in (Hausser 1984), a book whose aim he later summarized as

> (to show that) the choice between phrase structure grammar and categorial grammar is not merely a matter of terminological habit or professional expedience but rather has far-reaching consequences on the resulting linguistic analyses. For instance*, the categories of categorial grammar are combinatorially and denotationally transparent,* while those of phrase structure grammar are opaque …. (Hausser 1986, pg. 7; emphasis added)

Hausser's insistence on combinatorial and denotational transparency may be assumed to derive from his initial choice of a model-theoretic stance basis for linguistics theory since a main motivation for the development of model theory was precisely the desire to formally and transparently capture the syntax and denotation of mathematical expressions and the rules that govern their combination in inferences. A reviewer of the book who found Hausser's very detailed development of ORTAX somewhat 'hermetic' nonetheless acknowledged that Hausser addressed linguistic problems of interest, such as ellipsis, which are required for adequate computational linguistics theory and natural language applications.[3]

Having defined syntactic and semantic formalisms he believed to be appropriate for computational understanding and generation of natural language, Hausser then took the step while at Stanford's CSLI of instantiating ORTAX as a software parser. It was a seminal event in the evolution towards Database Semantics, for

> It became apparent that, well-motivated as the grammatical system of (Surface Compositional Grammar) seemed from a linguistic point of view, it was not a very suitable basis for an efficient parsing program. The reason for this is not an inherent property of categorial grammar but the irregular nature of conventional constituent structure trees, which are common to both categorial grammar and phrase structure grammar.

> … In hindsight the formalism of SCG may be regarded as a last ditch attempt to save constituent structure analysis, albeit in the form of 'orthogonal trees', while the new research led us to the conclusion that ***constituent structure analysis should be abandoned completely***. (*ibid.*; emphasis in original)

In rejecting the search for constituent structure in language as it is presented by a speaker, Hausser breaks with others pursuing logic-based linguistics who, in one way or another, assumed the necessity of transforming surface language through a series of manipulations into standard constituent phrase elements of sentences. Such analysis was deemed necessary in order to apply categorial grammar to natural language since the former allows nested expressions but requires that they not overlap. Natural language presents many uses in which, for example, a pronoun is separated from its referent in ways that challenge the direct application of such a grammar. A variety of phrase-oriented grammars motivated by Chomsky and by successful use in parsing artificial programming languages have been proposed to transform the surface syntax of natural language into a representation of the assumed underlying logical structure. Hausser, however, here rejects this approach entirely as the basis for computational natural language understanding. He will later reject it on other grounds as well.

Having rejected grammars that rely on constituent structure analysis for natural language, Hausser immediately began work on an alternative approach and in (Hausser 1985) introduced Left Associative Grammars (LAGs). LAGs are described in Chapter 3; here it is only necessary to note that they provide a computationally tractable (Hausser 1987), strictly surface-compositional syntactic formalism based on time linear parsing of tokens as they are presented in an expression. Unlike phrase structure grammars, which are based on possible substitutions and recursively range over the full extent of the expression, LAGs are based on possible continuations from the portion of the expression previously examined (sentence start).

For the remainder of his Heisenberg grant and as he settled into his new professorship Hausser's publications focused on the foundational definitions (Hausser 1988c), complexity proofs (Hausser 1988a; 1992) and decidability (Hausser 1989b) of this new grammatical formalism. He then turned to computational implementation, choosing network database structures to encode context information. Hausser introduced Database Semantics in a flurry of papers beginning with (Hausser 1996) and culminating in two complementary books: (Hausser 1999; second edition 2001), which integrates and presents foundational theory for computational linguistics, and (Hausser 2006a) which gives a detailed treatment of the instantiation of that theory in Database Semantics.

3. The SLIM Theory of Natural Language and Database Semantics

Hausser is a theoretical linguist by training, so it is not surprising that he bases his computational approach to natural language on a theory of language in use. This chapter examines that theory and the resulting computational mechanisms he proposes for human-level computational natural language capability in embodied artificial agents.

3.1 The SLIM Theory of Natural Language Communication

Hausser articulates four basic principles on which to base computational natural language systems. Together they form what he calls "a certain intuitive conception of natural language communication, called the SLIM theory of natural language" (Hausser 2004). This theory is intended to describe human language use and thereby to serve as the most appropriate basis for computational natural language capability as well (see section 3.2 below). The SLIM principles are:

(1) The methodological principle of Surface compositionality

(2) The empirical principle of time-Linear processing

(3) The ontological principle that semantic and pragmatic processing are processes Internal to the speaking or hearing agent, and

(4) The functional principle of Matching as the link between semantic literal denotation and pragmatic reference.

Section 3.2 below discusses Hausser's proposal for validating these principles. This section examines each principle in turn.

3.1.1 Surface Compositionality and Time Linear Processing

Hausser adopted the methodological principle of surface compositionality in the earliest days of his scholarship, as noted above (Hausser 1978), and has affirmed this principle to the present day. In (Hausser 2006) he clarifies what he means by surface compositionality:

> A grammatical analysis is surface compositional if it uses only the concrete word forms as the building blocks of compositions, such that all syntactic and semantic properties of a complex expression derive systematically from the syntactic category and the *literal* meaning of the lexical items. (pg. 18, emphasis added)

In affirming surface compositionality Hausser is explicitly rejecting the position of speech act theorists such as Searle and Austin who locate the meaning of an utterance in the utterance as a whole. In this sense his approach resembles that of traditional model theory in which the surface forms of words are translated (synthesized) into corresponding logical representations of meaning.

Hausser illustrates what he means by surface compositionality by contrasting it with grammars based on substitution, such as variations of phrase structure grammars, which seek to map the surface tokens as presented into standardized constituent structure representations. In such grammars parse trees must postulate the implicit presence of syntactic elements when they are lacking in a surface expression – for example, when a noun phrase such as **water** lacks an explicit determiner (such as **some** or **the),** as in the sentence **Mary drinks water**. As noted in section 2.6 Hausser rejects the search for constituent syntactic structure that extends beyond the explicitly presented surface language tokens (Hausser 1984).

In (Hausser 1999) he defends the principle of surface compositionality at greater length, noting that adopting this principle yields syntactic analyses that are maximally concrete (no zero surface or underlying forms may be used). Surface compositionality also yields syntactic and semantic analyses that are of the lowest complexity and that provide for internal matching between semantic literal meaning and context which can be extended systematically and transparently from single words to combinations of words and expressions. Here 'transparently' means that combining words into expressions and expressions into sentences concatenates meaning without the introduction of additional or non-surface elements.

Hausser's definition of surface compositionality explicitly refers to the semantic representations in the agent's concept space as literal meanings. Literal meanings in the concept space are consonant with a model theoretic desire to translate natural language into formal logic. For the most part Hausser has abandoned translation into intensional logic as a matter of practice; however, the mechanisms of Database Semantics which he proposes as appropriate implementation of the SLIM theory of natural language communication are carefully chosen to map to the propositional and predicate calcului (Hausser 2003). Section 3.3 below describes these mechanisms and examines Hausser's claims regarding their suitability; here it is useful to note that the ability to represent an utterance in propositional or first order logic holds open the possibility of a direct link between computational language representation and logic-based activities such as planning in robots and other artificial agents (Hausser 2002).

Since surface compositionality as Hausser defines it links surface tokens to literal concept meanings, non-literal language features such as metaphors are treated as a matter of pragmatics and not semantics. This is a point Haussser made in his critique of traditional model theory (*c.f.* Hausser 1981b) and it affirms to some degree the insights of speech act theories that language has its meaning in use situations. Although he never explicitly addresses the distinction between live metaphor and metaphorical uses that have become frozen (*i.e.* conventionalized), it is clear from Hausser's denotation / reference distinction that both sorts of metaphorical language use must be addressed at the pragmatic and not the semantic level of language processing. Accordingly he provides a brief treatment of metaphor in (Hausser 1999d) which is discussed in Chapter 4.

Unlike formal linguists who seek to combine transformational grammars motivated by Chomsky with semantics motivated by Montague (*c.f.* Portner and Partee 2002), Hausser's concern for surface compositionality caused him to completely reject phrase structure grammars (Hausser 1986), originally in favor of categorial grammars similar to those proposed by Bar-Hillel and Montague:

> (T)he choice between phrase structure grammar and categorial grammar is not merely a matter of terminological habit or professional expedience but rather has far-reaching consequences on the resulting linguistic analyses. For instance, *the categories of categorial grammar are combinatorially and denotationally transparent,* while those of phrase structure grammar are opaque …. (Hausser 1986, 7; emphasis added)

Categorial grammar is transparent because it is based on formal type theory and therefore explicitly identifies the categories of expressions and their combinatorial relationships. It is denotationally transparent insofar as expressions can nest but not overlap. However, Hausser's attempt in ORTAX to produce a computationally efficient implementation of categorial grammar for a major fragment of English failed, by his analysis, because categorial grammars, like phrase structure grammars, attempt to identify constituent structure within sentences. Irregular and complex constituent structure trees associated with the wide variety of possible expressions in natural language use caused Hausser to question not only whether there might not be a computationally more viable way to parse natural language but more fundamentally whether constituent structure is in fact inherent in the processing of natural language by humans.

In response to this concern Hausser adopted his second principle, that of strict time-linear order in the interpretation and production of surface language. He calls this the empirical principle of natural language because, he asserts, that is what we observe about ourselves and others as we use natural language:

> The most elementary relation between the words in a sentence is their time-linear order. Time-linear means linear like time and in the direction of time.

> The time linear structure of natural language is so fundamental that a speaker cannot but utter a text sentence by sentence, and a sentence word form by word form. Thereby the time-linear principle suffuses the process of utterance to such a degree that the speaker may decide in the middle of a sentence on how to continue. Correspondingly, the hearer need not wait until the utterance of a text or sentence has been finished before his or her interpretation can begin. (*ibid.,* 19)

Instead of phrase structure or categorial grammars, which transform a sentence through a series of substitutions into formal representation of standard structural elements, Hausser proposes the use of a new formalism he calls Left-Associative Grammars (LAGs). LAGs parse sentences word by word in strict time linear sequence and combine a sentence start (words parsed so far) with the next word based on possible continuations as specified by rules that specify legal patterns that can follow. Unlike phrase structure and combinatorial grammars LAGs do not recurse across the sentence as a whole.

3.1.2 Cognitive Agents and the Internal Locus of Meaning

Although his use of Left Associative Grammar is an obvious factor distinguishing Hausser's work from that of other computational natural language approaches, the more fundamental principle of his SLIM theory concerns the locus and ontology of meaning. Although he uses the term 'ontology', Hausser does not explicitly define it. An examination of his use suggests that he intends the term to mean the assumptions about the nature and state of being that underlie a theory of natural language communication's treatment of meaning and reference.

Hausser identifies two binary (±) features which in combination determine four possible ontologies of meaning (Hausser 2001c). The first feature is whether the theory is or is not constructive. This feature describes the relationship between the surface language forms, the language-capable agent and the referents in the external world. A [-constructive] theory of semantics regards the language-capable cognitive agent as being on the same level and having the same status as the referents for language terms and expressions. Such an agent can observe the relationship between language and the world, which is external to itself. A [+constructive theory], on the other hand, locates the relationship between language surfaces and referents *within* the cognitive agent. Hence, in a [+constructive theory], what the cognitive agent has not perceived cannot be part of language reference, although the agent's internal processes such as wishes, plans etc. can be.

The [±constructive] feature of a language theory is important because it constrains the type of semantics that are possible in that theory. A [-constructive] theory such as traditional model theory sees meaning in terms of the inherent truth of conditions of statements, without reference to a speaker or hearer. Such a theory, Hausser notes, *must* have a meta-language semantics, for all of the reasons behind the development of model theory, *i.e.* because of the contradictions that arise if the object language is used to evaluate the truth of statements made in that language. On the other hand, a [+constructive] theory must have a *procedural* semantics that describes the process within the cognitive agent by which sign and meaning are linked.

The second ontological feature of natural language theories that Hausser identifies is [±sense]. A [-sense] theory identifies the surface language meaning directly with the referent objects of that surface. A [+sense] theory posits a level of meaning between the surface language and the referent objects, as in Frege's *Sinn (*sense*)*, distinct from *Bedeutung* (reference). Hausser's early exploration of [+sense] semantics occurs in his discussion of constructive intensional contexts (Hausser 1982a) and of vagueness and truth conditionality (Hausser 1983c). These papers suggest that Hausser uses the term 'constructive' analogously to its use in logic to denote that the theory asserts a semantics which specifies how a proposition can be constructed to reflect the surface language and not merely evaluated as to its truth condition.

Combining these the two features of [±constructive] and [±sense] yields four possible approaches to semantic interpretation, as illustrated in Figure 2, reproduced from (Hausser 2001c).

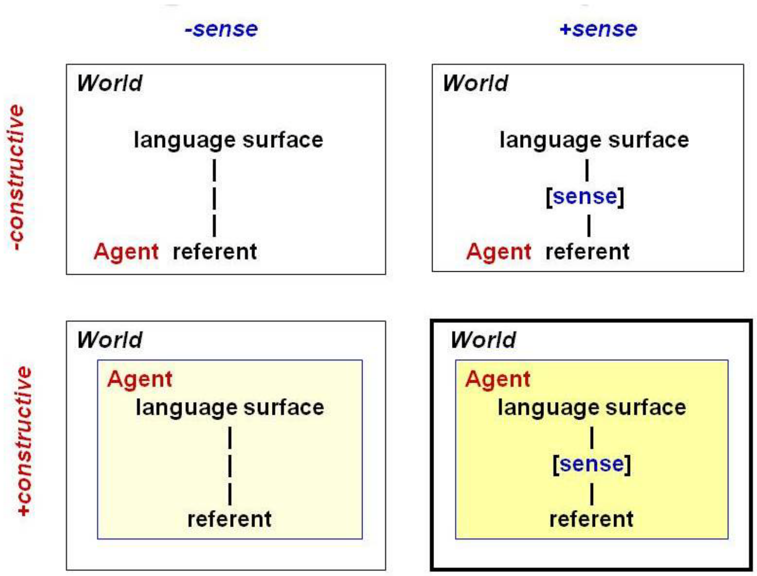

Figure 2. Four Ontologies of Semantic Interpretation

Within this typology, [-constructive, -sense] describes truth-conditional, logic-based approaches to semantics such as that of Russell, Carnap, Quine, and Montague. These approaches are concerned to limit the referents of language to objects that are ontologically real in the physical world or, on the part of mathematical realists, to abstract but well-defined objects such as sets and numbers. Their aim is to provide a rigorous basis for establishing the truth condition of a sentence.

The [-constructive, +sense] approach is typified by Frege's analysis of opaque intensional meanings in natural language which, while introducing the level of *Sinn*, explicitly was not intended to be 'psychologistic' or to imply cognitive states (Frege 1892).

The [+constructive, -sense] ontology of semantics is typified by computer programming languages, for which reference is inherently procedural – external signs in the form of source code are accepted and translated directly into procedures to be executed by the hardware. Hausser notes that, perhaps because they were implemented in such languages, many classical artificial intelligence approaches adopted a [+constructive, -sense] semantic as well, citing as an example Winograd's SHRDLU and the work of Newell and Simon, who were active at the Dartmouth Summer program in 1956. Hausser does not cite specific elements of Newell and Simon's work, being content to quote their explicit rejection of an internal level of sense when representing meaning. However, Newell's SOAR architecture for cognitive software agents and the Physical Symbols System hypothesis on which it is based, do reflect a [+constructive, -sense] semantics. Anderson and Bower's ACT-R architecture, on the other hand, reflects a [+constructive, +sense] semantic. Chapters 4 and 5 consider the issue of cognitive assumptions in more detail.

Hausser criticizes the [+constructive, -sense] semantics for being limited to closed, toy worlds and for providing no basis for "autonomous classification of new objects, in principle", noting:

> It is by no means accidental that these systems have no components of artificial perception: *because they lack the intermediate level of concepts* (sense) *they could not utilize perception* (e.g. artificial vision) *to classify and to automatically integrate new objects into their domain.* (Hausser 2001c, emphasis added)

Any AI approach that centers on production systems (*i.e.* systems that are able to derive propositions entailed in existing propositions in the knowledge base, but that require any new information to be translated into logical propositions before it can be integrated into the system) will be limited insofar as it fails to account in some way for perception as the basis for assignment of new objects to classes. Accomplishing such assignment using only formal logic is at best a difficult task according to Hausser since object recognition and classification in humans is based first on associative (connectionist) rather than formal (symbolic) mechanisms and only secondarily with some conceptual mapping that parallels formal predicate assertion. Here Hausser parallels his critique of direct reference in model theory with a corresponding critique of proposition-only 'old' AI.

As sections 3.3 and 3.4 below demonstrate, however, this critique does not mean that Hausser has abandoned his interest in propositional representation of knowledge or in logical inference. Instead, he turns to the semantics/pragmatics distinction to identify the interface between perception and concept-based cognition. This is the rationale for his choice of a [+constructive, +sense] semantics as the ontological principle underlying his SLIM theory of natural language.

In the SLIM theory, Hausser repeats his conclusions in (Hausser 1981a and 1981b) by positing a +sense level of conceptual meaning which, being not-vague and not-ambiguous is suitable for logical representation, and a separate context-level representation of what the agent perceives and remembers about the environment.

Hausser posits that natural language is first and foremost a means of communication between two agents. Given his distinction between the semantic model of concept meaning for linguistic terms and the pragmatic association of those concepts to what is experienced and remembered by the speaker or hearer, it follows that these agents must have a [+sense] level and are therefore cognitive in nature.

> Up to now, scientific analyses of natural language have been mostly limited to structural objects, fixed on paper or magnetic tape. Such objects are exemplified by a single word form, a sentence or a text. By concentrating on the structure of the signs, one has attempted to abstract away from the aspect of communication. The purpose of producing and interpreting language in the first place, however, is interaction between cognitive agents. Therefore, a scientific analysis of natural language cannot fail to be inadequate if it does not include the production and interpretation procedures inside the cognitive agents. The question should not be what a sign *is*, but rather what it *does*, and how it does it by virtue of what it is. (Hausser 2004a; emphasis in original)

Hausser has used the image in Figure 3 below to describe the relationship between an agent, its environment and linguistic signs (reproduced from Hausser 2004). Several elements of this figure are worth noting in particular. The primary element of the figure is the cognitive agent, with perceptual interfaces to the outside world. Those interfaces include sensory-based input and output of signs (sounds, images on a page) as well as non-linguistic perception. Here Hausser specifically refers to peripheral *cognition* associated with non-linguistic perception and specifies a non-language component to central cognition. This reflects his assertion that it is possible to have cognition without language, but not language without cognition. In a cognitive agent without language, cognition consists of direct conversion of percepts into representations of the context (environment) via pattern extraction / matching within peripheral cognition, followed by decisions in response to those representations within central cogition and and actions taken as a result, mediated through peripheral cognition again.

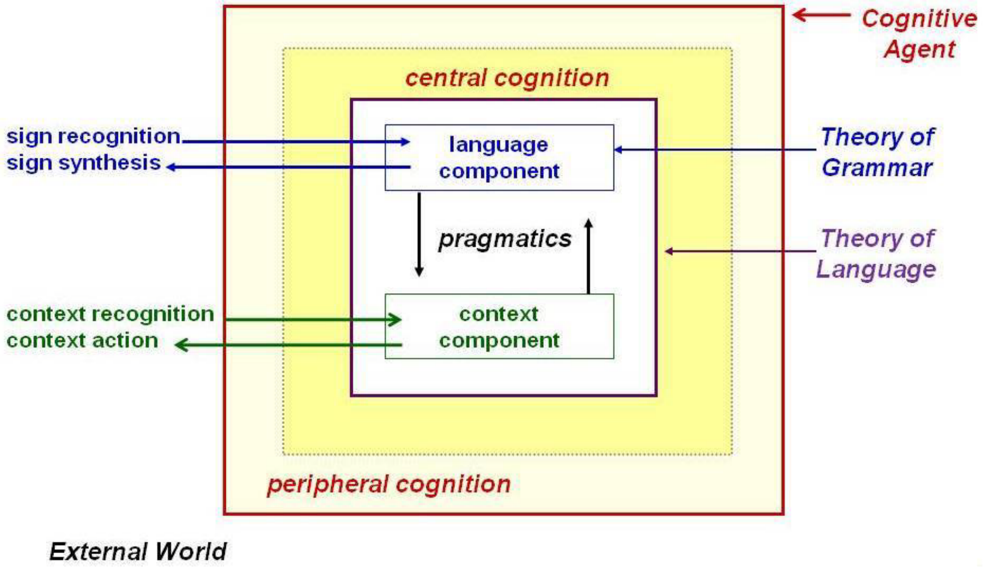

Figure 3. Hausser's Theory of Signs Relative to Agents

For instance, to add detail to Hausser's brief description in (Hausser 2006, pg. 21-22), a squirrel has a cognitive ability to identify nuts on the ground as a result of visual and other sensory inputs. Peripheral cognition would include the visual system's ability to distinguish color, texture and shape. The squirrel doesn't simply perceive, however: it also recognizes that those visual attributes signal the presence of a nut and that nuts are potential food. It then takes action to gather and store the nuts and to return to dig them up later when food is scarce. These are non-linguistic procedural capabilities in the squirrel as a cognitive agent, beginning with object recognition through planning and action steps based on memory.

The addition of language extends but does not change the more fundamental cognitive processes according to Hausser. This is true not only because sensory perception and peripheral cognition are invoked during natural language communication, as when we hear or read a sentence, but also because in accordance with the semantic/pragmatic distinction Hausser draws, language can denote concepts precisely because those concepts in turn refer, more or less accurately, to objects and activities in the environment which have already been or are being experienced or imagined. This process occurs within the cognitive, language-using agent. Meaning consists, therefore, not in an eternally valid truth condition external to the agent but rather with regard to an agent's pragmatic association of surface language with concepts and concepts with experience. And because this association is *procedural* rather than formally static, it requires the functional principle of the SLIM theory of natural language, namely that pragmatic association occurs through a matching process.

3.1.3 The Matching Principle of Pragmatics

Hausser's functional principle of language states that referring with language to past, current or future objects and events is best modeled in terms of pattern matching between language meanings (*i.e.* concepts at the semantic level) and context (*i.e.* the agent's internal representation of what it perceives, remembers, etc.). He adopts the metaphor (and, in Database Semantics, the implementation) of a database to represent both the linguistic/conceptual and the pragmatic context of a language-capable agent (Hausser 1996):

> In order to function objectively, a computerlinguistic model of natural language communication requires an explicit definition of the context of use. What, however, are the exact ontological, structural and computational properties of this component? As a first, most general answer to this question we have called the speaker/hearer-internal context a database in the widest sense of the word … For the operations of storage and retrieval in databases resemble the cognitive processes in speakers and hearers in important respects.

Hausser notes that one characteristic of databases is their ability to represent multiple relationships between various pieces of information. The choice of database design depends on a number of factors which must include both completeness (all valid associations must be representable) and efficiency. In standard software application databases, these characteristics are determined by table or record structure definitions, indexing schemas and the choice of primary keys for storage and retrieval of information, each of which depends on the nature of the information to be stored and retrieved.

On what basis can linguistic tokens, semantic concept representations and world knowledge be pattern matched to one another? More specifically, *what* is being matched and *how* is the matching accomplished?

Hausser presents seven pragmatics principles that govern this crucial step in linking language to denotation to reference (Hausser 2006a). The first three principles are:

1. The speaker's **meaning$_2$** is the use of the sign's literal **meaning$_1$** relative to an **internal context.** This internal context of interpretation must be determined and delimited correctly in order to ensure that the correct meaning$_2$ is assigned. Among the features

required for this determination are the Spatial location and Time at which the sign was produced, the Author of the sign and the intended Recipient (STAR).

2. A sign's STAR determines the entry context of **production and interpretation** in contextual databases of speaker and hearer, respectively.

3. Matching of signs with contexts is **incremental and time ordered**. In language production, signs are produced in the time order in which the underlying thought path proceeds across concepts. In language interpretation the thought path is constructed incrementally as linguistic signs are received.

Hausser does not provide an explicit rationale in (Hausser 2006a) for the emphasis on the STAR associated with a sign's production or interpretation, although it plays a key role in the ability of his Database Semantics mechanisms to handle statements about past, present and hypothetical events, distant objects and indirect references. However, we may assume that his motivation stems from his use and later critique of Montagovian model theory since the STAR for a sign plays a role very similar to the context coordinates proposed by Montague as a way to deal with indexical pronouns, anaphora and related linguistic phenomena (Section 2.1.4 above). This is underscored by Hausser's emphasis on indexicals and names as well as other linguistic symbols in the remaining four pragmatics principles he asserts:

4. The reference mechanism of a symbol is based on a meaning$_1$ which is defined as a concept type. Symbols refer from their place in a positioned sentence by matching their meaning$_1$ with corresponding concept tokens at the context level.

Hausser signals that he will introduce two kinds of representations in the agent's context database. Both types (general concepts) and tokens (instances of a type) are represented by feature structures called proplets (*i.e.* partial propositions) which are made up of attribute-value pairs. The difference between them is that some of the attributes in a type, namely those that are accidental, are populated by variables. Necessary attribute values for a concept (such as angles equal to 90 degrees in squares) are constants common to both a concept's type and the tokens that represent instances of that type. Accidental attributes values (such as the length of a square's sides) that are represented in the type structure as variables are given specific values in each associated token. Thus for each concept type there are potentially an infinite number of associated tokens possible. Which specific tokens are available in a given agent's context database depends on the agent's experiences up to the time a sign is produced or interpreted. As the agent receives sensory inputs through its peripheral cognition, the features of that experience are matched against concept types to produce concept tokens in the agent's context database. Hausser does not dwell on this step, but a close reading of his recent papers and an examination of his mechanisms suggest that this matching may be done probabilistically or associatively, so long as the requisite attributes can be extracted from the sensory data and a token created of the appropriate type.

A similar type/token relationship exists for sign recognition and for lexical lookup at the language level. For instance, when a surface sign such as **trees** is heard or read as part of an incoming sentence, the sensory patterns are matched to a surface for word recognition. The resulting surface token **trees** is then matched against a lexicon to produce a partially populated language token whose attribute of number is set to plural, part of speech set to noun, and so on and which points to the associated concept type which represents the literal value for that token. Although Hausser does not explicitly invoke in (Hausser 2006a) the denotation/reference distinction that is central to his critique of traditional model theory as applied to natural language (Hausser 1981b), the mapping of surface signs to surface tokens to concept types is the denotation step. During pragmatic matching, when the attributes of that token are further matched to the attributes of corresponding context tokens representing the agent's experience, the reference process has completed. Section 3.3 below describes the proplet data structure and the matching process in more detail.

In addition to symbols, whose meaning$_1$ is a concept type, there are two other basic kinds of signs for which Hausser articulates pragmatic principles. These are indexicals (**I**, **it**, **here**, **now** …) and names (**Mary**, **John**). Both indexicals and names have occasioned considerable analysis in logic-based approaches to natural language, as is evidenced in Montague's coordinates approach to indexical reference. Hausser's approach to them is defined in his remaining pragmatic principles:

5. The reference mechanism of an *indexical* is based on a meaning$_1$ which is defined as either a pointer into the agent's context or as a pointer to the agent.

6. The reference mechanism of a *name* is based on a private marker which matches a corresponding marker contained in the cognitive representation of the object referred to.

7. The *distinction between kinds of signs* (symbols, indexicals and names) is orthogonal to the *distinction between main parts of speech*. The part of speech controls the combinatorics in a sentence (horizontal relations) while the kind of sign determines the reference mechanism which relates the word meanings to the context (vertical relations) . Symbols occur as verbs, adjectives and nouns. Indexicals occur as adjectives and nouns. Names occur only as nouns.

Note that parts of speech combine into propositions regardless of the kind of sign used – the sign types, which have distinct pragmatic reference mechanisms, are distinguished from the main syntactic parts of speech which are the primary elements of semantic denotation. This is consonant with Hausser's call in (Hausser 1980) for a theory of discourse that delineates the boundaries between the components of syntax, semantics, lexicon context and pragmatics and that describes "how the components interact in the course of interpreting the use of an expression by a speaker relative to a context".

The seven pragmatics principles are intended to accomplish this overarching objective. Principle 1 separates the literal meaning$_1$ of the sign type (*i.e.* the concept it denotes) from the meaning$_2$ aspects of associated tokens (*i.e.* the specific context elements to which it might refer), which is a precondition for syntactic-semantic analysis with surface compositionality. Principle 2 provides a precondition for finding the context of use for a sign or an expression. Principle 3 establishes a uniform mechanism of interpretation for both speaker and hearer modes of communication. Principles 4-6 establish reference mechanisms for the different kinds of signs and Principle 7 determines which reference mechanisms must have procedural implementations for each part of speech.

3.2 Implementation as Validation

On what basis can Hausser's SLIM theory of natural language communication be validated (or invalidated)? Specifically, on what basis can we evaluate the adequacy of this theory for embodied computational agents as well as for humans?

Having stepped away from the criteria of satisfiability and validity posited in model theory, and having asserted the inherently procedural nature of pragmatic meaning reference, Hausser proposes that his or any other theory should be judged on the basis of its input/output equivalence to natural language processing by fluent human hearers and speakers.

Beginning with grammar, in (Hausser 1989a) he argues that grammars must be psychologically well-founded in order to be I/O equivalent to human language use. This requires that it meet two criteria, namely procedural adequacy and derivational order. Derivational order translates to time linearity and is the basis for the Left Associative Grammars that Hausser proposed in lieu of either phrase –based or categorial grammars. Procedural adequacy of a grammar for natural language requires that the grammar provides an explicit formal statement of grammatical rules, provides a bi-directional mapping between language surface and meaning, and is decideable. Hausser finds both phrase based and categorial grammars lacking in bi-directional mapping. As a result, he asserts, they fail to account for the ability of humans to move between speaker and hearer roles. He notes that Chomsky's claim that transformational grammar models human language capacity ("competence") rather than language use ("performance") leads Chomsky to assert that these two dimensions of language evolved separately from one another, a claim Hausser finds implausible.

Hausser goes farther, in (Hausser 2006a), by proposing that a functional model specified declaratively and implemented in a prototype should be the basis of verifying a theory of natural language communication. Such a model, implemented for an increasing portion of a specific natural language (or ideally multiple languages) should be based on what he calls the Equation Principle, namely that the more realistic the reconstruction of cognition, the better the model will function since natural language use is based on human cognition. Conversely, the better the model functions the more realistic the reconstruction of cognition may be assumed to be.

3.3 Database Semantics

Hausser proposes a computational approach which he calls Database Semantics (DBS) which is intended as a declarative functional model for the SLIM theory of agent-oriented natural language communication. Partial instantiations of DBS in the form of prototype implementation code for fragments of English, Russian and German are available for download[4]; Hausser and colleagues have also presented results for Korean (Lee 2002; Choe and Hausser 2006) .

The DBS model consists of data structures for language processing and context knowledge representation, along with three left-associative grammars: one grammar for parsing language when the agent is in hearer mode, a parallel grammar for generating language when the agent is in speaker mode and a third grammar that provides the mechanism for inferencing across formally represented context knowledge. The remainder of this chapter will briefly review those mechanisms.

3.3.1 Left-Associative Grammars

Left-associative grammars (LAGs) are intended to provide ‘simple syntax for complicated semantics’. Hausser initially introduced this formalism in (Hausser 1987a) during his research residency at Carnegie Mellon University and has since expounded on it in a series of papers and in his *Foundations of Computational Linguistics* book (Hausser 2001a). He provides a formal algebraic definition for LAGs as follows:

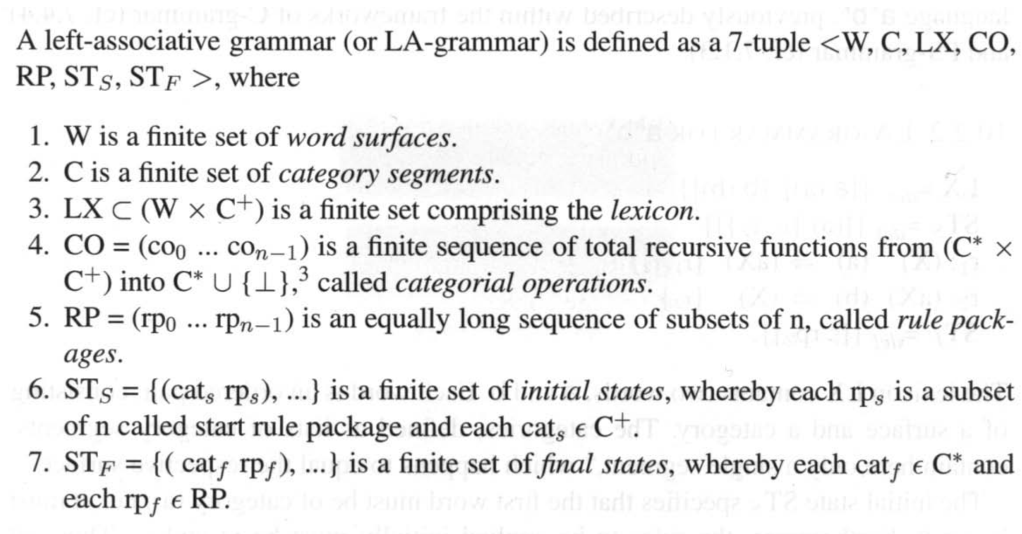

A left-associative grammar (or LA-grammar) is defined as a 7-tuple <W, C, LX, CO, RP, ST_S, ST_F >, where

1. W is a finite set of *word surfaces*.
2. C is a finite set of *category segments*.
3. $LX \subset (W \times C^+)$ is a finite set comprising the *lexicon*.
4. $CO = (co_0 \ldots co_{n-1})$ is a finite sequence of total recursive functions from $(C^* \times C^+)$ into $C^* \cup \{\perp\}$,[3] called *categorial operations*.
5. $RP = (rp_0 \ldots rp_{n-1})$ is an equally long sequence of subsets of n, called *rule packages*.
6. $ST_S = \{(cat_s\ rp_s), \ldots\}$ is a finite set of *initial states*, whereby each rp_s is a subset of n called start rule package and each $cat_s \in C^+$.
7. $ST_F = \{(\ cat_f\ rp_f), \ldots\}$ is a finite set of *final states*, whereby each $cat_f \in C^*$ and each $rp_f \in RP$.

In practice, the left-associative grammars in Database Semantics consist of a lexicon, a set of initial states (initial sentence starts), a set of rules defining the possible continuations from a given sentence start, and a set of final states. Rules are defined in terms of the categories of the next token that can continue a given portion of an expression as it is parsed from left to right, *i.e.* in time linear order. LAGs in Database

Semantics do not recurse to re-examine or substitute for tokens that have already been examined. Parallel interpretation paths are maintained and are pruned naturally as subsequent rule applications indicate a failure to match specific pattern sequences. It is possible for a sentence to be parsed without resulting in a single interpretation path, in which case the mechanisms of DBS provide for extra-sentential reference resolution.

Along with the algebraic definition of the LAG formalism, Hausser proposes a hierarchy of LAG types based on restrictions on the categorial operations and the degree of ambiguity associated with rule applications. This hierarchy, orthogonal to the Chomsky hierarchy familiar in computer science, is intended to aid in assessing the computational complexity of natural language approaches. However, Hausser provides a mapping from his hierarchy to that of Chomsky. A-LAGs are unrestricted and can accept and generate all recursive languages. B-LAGs bound the length of intermediate sentence start categories and can accept and generate all context-sensitive languages. C-LAGs, or constant LAGs, consist of grammars in which no categorial operation looks at more than k segments in the sentence start categories, for a finite constant k. This is the type of LAG that Hausser proposes for Database Semantics.

C-LAGs in turn can be divided into several sub-classes. C1-LAGs have ambiguities which are non-recursive, *i.e.* in which none of the branches produced as rules are applied return to the state that caused the ambiguity. C1-LAGs that meet this condition will parse an expression with a maximum of attempted rule applications $(n - (R - 2)) \cdot 2^{(R-2)}$ for $n > (R - 2)$, where n is the length of the input expression in tokens and R is the number of rules in the grammar; *i.e.* their computational complexity is linear. C2-LAGs allow exactly one of the parsing paths to return to the state generating a given ambiguity and parse in polynomial time. Other C-LAGs parse in exponential time. Hausser maps these classes into the categories more familiar to computer scientists and provides complexity proofs in a series of papers (Hausser 1988a, 1989b, 1992).

In comparing left-associative grammars with phrase structure grammars, Hausser notes that context-free forms of PS-grammar have been widely used for natural language because they provide the most generative capacity within the PS hierarchy while being computationally tractable. However, he asserts, context-free PS grammars do not fit the structures of natural language well and are at best a first order approximation even to most programming languages.

Within the LAG hierarchy, Hausser argues that natural languages are highly likely to be C-languages since the A- and B- types reflect category complexity that is not seen in natural language use. Within the C-LAGs, complexity of the natural language is equivalent to ambiguity of the language. This raises the question, how should ambiguity be treated?

Hausser argues that the ambiguity of a sentence like **The osprey is looking for a perch**, where **perch** can mean either a type of fish or a place to roost, is not a syntactic ambiguity but rather a semantic one. Parsing of such a sentence can be accomplished with a C1-LAG through *semantic doubling* in which parallel paths of semantic interpretation are maintained until they can be disambiguated pragmatically. Therefore, so long as there are no recursive ambiguities within a natural language, natural languages fall into Hausser's C1 class and parse in linear time. And that it is highly likely that there are no recursive ambiguities in natural language falls out of the semantic/pragmatic distinction and the fact that meaning reference in natural language consists of pragmatically linking literal semantic concepts to the agent's context database.

Hausser has spent considerable effort establishing formal foundational definitions (Hausser 1988c), complexity proofs (Hausser 1988a; 1992) and decidability proofs (Hausser 1989b) for Left-Associative Grammars (LAGs). These are reprised and extended in his 1999 *Foundations of Computational Linguistics* book, along with reconstructions of regular, context-free and context-sensitive languages in various flavors of left-associative grammar by simulating finite state automata, push down automata and linearly-bounded automata, respectively.

3.3.2 Proplets and Propositions

Implicit in the SLIM theory is the assumption that both language and central cognition are essentially propositional in nature, with associative sensory processing and basic feature extraction from sensory data assigned to peripheral cognition. In an invited lecture at the 1999 Natural Language Processing Pacific Rim Symposium Hausser identified three core attributes of basic propositions, whether they relate to logical formulae, the world context or natural language statements (Hausser 1999b):

	logic	*world*	*language*
1.	functor	relation	verb
2.	argument	object	noun
3.	modifier	property	adjective-adverbial

These core attributes form the basis for the main data structure in Database Semantics (DBS), namely the proplet (partial proposition), and correspond to content words in natural language. Function words such as articles, prepositions, conjunctions and auxiliary verbs are encoded within the proplets of associated content words.

A proplet is a flat structure (record) consisting of feature-attribute pairs. The specific features represented in the structure differ depending on whether the proposition is at the language or context level. Type proplets yield token proplets when attribute variables in the type proplet are assigned specific values (which must be atomic – no embedded structures permitted). Propositions are constructed by the concatenation of proplets at the appropriate level.

Proplets contain such features as the surface word form (null in the case of context proplets), the part of speech associated with that word form, and a proposition number that is common to all of the proplets in the proposition. Additional features link proplets semantically through core values. For example, the proplet for the subject of a sentence contains a functor feature which is mapped, during parsing, to the associated verb proplet. The verb proplet in turn is linked back to the noun and forward to the next proposition.

An intentional feature of the proplet structure is its ability to map concepts to word surfaces from multiple languages. Hausser assumes that concepts are independent of and prior to specific languages (an assumption revisited in Chapter 4), allowing word surfaces to be treated simply as another attribute of the underlying concept.

A basic example of this approach is illustrated in (Hausser 2006, 35-37). First, consider the propositional content of an agent perceiving a barking dog and taking the action of running away. These two context-level propositions are encoded in the following proplets, simplified to illustrate the basic data relationships:

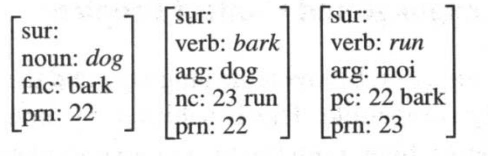

$$\begin{bmatrix}\text{sur:} \\ \text{noun: } \textit{dog} \\ \text{fnc: bark} \\ \text{prn: 22}\end{bmatrix} \begin{bmatrix}\text{sur:} \\ \text{verb: } \textit{bark} \\ \text{arg: dog} \\ \text{nc: 23 run} \\ \text{prn: 22}\end{bmatrix} \begin{bmatrix}\text{sur:} \\ \text{verb: } \textit{run} \\ \text{arg: moi} \\ \text{pc: 22 bark} \\ \text{prn: 23}\end{bmatrix}$$

The non-combinatorial semantic content of a proplet is coded as the value of its core attribute. For instance, the attribute **noun** in the first proplet refers to the concept *dog* and the attribute verb in the second proplet refers to the concept *bark*. (Figures in this section are reproduced from (Hausser 2006a).) The full context proplets would include STAR and other context information as well as these basic attributes.

In addition to concepts, core values may consist of pointers or markers in accordance with the pragmatic principle 7 outlined in section 3.1.3 above. In the dog barks – (I) run example, the **prn** value identifies the first and second proplets as constituting a single proposition.

Hausser asserts that concepts originate as 'patterns (within central cognition) for recognition and action in the agent's peripheral cognition' (*ibid.*). That is to say, features extracted from sensory inputs by the agent's peripheral cognition are matched against concept types to produce interpreted tokens of experience. Similarly actions planned within the agent's central cognition are matched against physical control mechanisms, resulting in their execution within the external environment. These cognitive functions do not depend on language and can be performed by a cognitive agent that lacks language capability. Hausser does not describe a non-cognitive agent, *i.e.* one that simply reacts to stimuli with pre-determined actions, but such agents are familiar from computer science (cellular automata) and robotics (behavioral/reactive robotic architectures). Being unsuited for natural language communication, non-cognitive agents play no part in the SLIM theory or DBS instantiation.

Concepts are encoded as token instances of type structures with features that define the concept in question. For instance, squares have four angle values equal to 90°. Individual square tokens, representing specific squares within the agent's experience or imagination, will differ with regard to the edge length and such attributes as color and location. Hausser's treatment of context knowledge representation is still somewhat sparse but may be applied to characterize the concept of dog to include features such as belonging to the class mammal, normally having 4 legs and a tail, being carnivorous predators, and so forth. Chapters 4 and 5 examine this assumption more closely.

Core values link proplets combinatorially, providing the DBS representation of semantic links between concepts within a proposition and across related propositions. The **fnc** (functor) attribute in the dog proplet above points to the concept *bark*, whose own proplet indicates that dog is the argument for this action in proposition 22. The **nc** (next conjunct) attribute in the second proplet points forward to the subsequent proposition, *i.e.* the next proposition in a related sequence, and the **pc** (previous conjunct) attribute in the third proplet points back to its predecessor. Since this is a sequence of context propositions, continuation implies contextual relatedness.

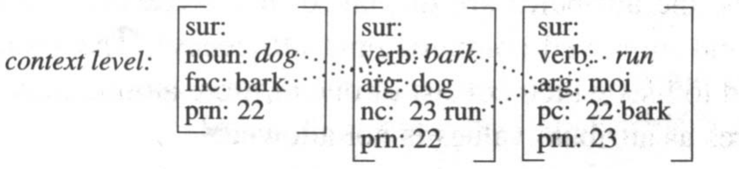

Figure 4. Partial Context-Level Propositions

In order to implement the functional interaction that he asserts between language and context levels, Hausser repeats the use of the proplet structure to represent language propositions:

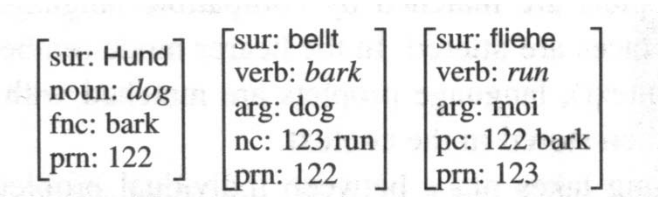

sur: Hund
noun: *dog*
fnc: bark
prn: 122

sur: bellt
verb: *bark*
arg: dog
nc: 123 run
prn: 122

sur: fliehe
verb: *run*
arg: moi
pc: 122 bark
prn: 123

Figure 5. Partial Language-Level Propositions

Here Hausser uses the German surfaces to illustrate that DBS is designed to support many languages against the same context database.

Proplets at the context level require a limited number of pre-defined attributes. The concepts *dog*, *bark*, *run* and *moi* link to context tokens in the artificial agent (or to neural representations in humans) which in turn link to procedural code that represents the pragmatic reference of those concepts. Proplets at the language level, however, require a number of specific additional attributes. These include:

mdd	modified	noun, verb or adjective which is modified
mdr	modifier	adjective or adverbial modifying a noun, verb or other adj-adv
idy	identity	used to establish whether two nouns are identical or not
sem	semantics	segments specifying non-combinatorial properties
cat	category	segments specifying combinatorial properties

Hausser also defines constant values which can be assigned to specific language level attributes to represent such features as number, person and case for nouns; number and tense for verbs; and semantics of determiners. The full list of proplet attributes and values is given in Appendix C of (Hausser 2006). In addition to generic values, Hausser includes specific values for language features such as the present, past, past participle etc. forms of the verb be and of the auxiliary verbs have and do in his English examples. These values augment the core values of the proplets, resulting in parallel and therefore transparently linked data representations, *i.e.* transparent between language and context levels. Section 3.3.3 illustrates the process of populating language proplets and describes the matching process by which an agent selects the appropriate context proplet (concept) during language interpretation.

Proplets are organized into token lines, *i.e.* sequences of proplets with the same core value, and stored persistently in what Hausser calls Word Banks, presumably on analogy to ‘database’. Language-level proplets with the same core value reflect the presence of different surface forms to indicate number, person, tense, aspect and other linguistic features. Language-level proplets with the same surface form but different core values reflect polysemy in natural languages. The potential for context proplets with the same core value but different pragmatic meanings is a key element in DBS and reflect the presence in the agent’s experience (context word bank) of multiple referents for the same core concept. When this occurs, the core value in the language-level proplet is appended with the appropriate context proposition number to indicate exactly which instance is being encoded. As a programming matter, the combination of core value and proposition number ensures unique database lookup.

Database Semantics organizes language and context proplets physically in token lines with the same core value, and through logical association by means of proposition numbers. This organization allows the agent to proceed from language to context in language interpretation and from context to language in language production. For persistent storage Hausser chooses the network (CODASYL) database model. Network databases can be described as graphs in which objects (records) serve as nodes with links to related objects. Network databases are seldom implemented today for information systems, having given way to relational and object-based data management systems in the last decade or more of programming practice, so at first glance this is an unusual decision. Hausser doesn't justify his choice but it is not difficult to do so on the basis of the strong flexibility network models offer. In particular, network databases permit such relationships as multiple parents of the same record type and asymmetry in the number or type of relationships a given record can have with records of the same or other types. The value of this flexibility will become apparent during the discussions below on left-associative grammars as applied to proplets. Object databases would also suffice as the storage and retrieval mechanism for Database Semantics.

3.3.3 LA-Hear, LA-Think and LA-Speak

Use of a single data structure for both language and context propositions allows the definition of aligned left-associative grammars for parsing and generating language and for traversing the context database of concepts. A full implementation of a language or language fragment in DBS requires the definition of three LAGs. The LA-Hear grammar parses language for interpretation. The LA-Speak grammar moves from context propositions to generate language statements. Hausser also proposes the use of a third left-associative grammar, LA-Think, for inference within the language-capable agent.

Language interpretation begins with mapping of surface tokens into partially-populated language-level proplets. This may be done with a simple lookup for small fragments of a language but requires morphological analysis for more complete systems. Hausser discusses morphological analysis in some detail in his 1999 *Foundations* book but otherwise focuses on language processing after the word form recognition step has occurred.

Whether specified fully for simple lookup or constructed through morphological analysis, the lexicon consists of stand-alone language proplets for which only the following attributes are populated:

attribute	*value*
sur	surface form
noun or verb	core value = matching key for context type proplets
sem	non-combinatorial properties
cat	combinatorial properties

The **cat** attribute encodes such features as number, person and case for nouns and pronouns, and number and mood for verbs. The **sem** attribute encodes such features as the gender of nouns, tense and negation of verbs. Within the lexicon other attributes are left unpopulated.

Mapping the unanalyzed surface token to a lexicon proplet is the first step accomplished in the LA-Hear process by the LA-Hear grammar. Parsing linearly, this grammar then constructs a sequence of completed language-level proplets that represent a sentence through a process of cross-linking, *e.g.* by populating the **fnc** attribute in a noun proplet with the core value of the verb for which it is a subject or object and similarly populating the **arg** attribute in the verb proplet with the core value of the noun along with **mdd** and **mdr** attributes when modifiers are present. Pronouns and determiners are also encoded into the proplets for relevant core values as they are encountered. The result of the LA-Hear process is a series of concatenated proplets that constitute a proposition representing the semantic meaning of the surface expression. For instance:

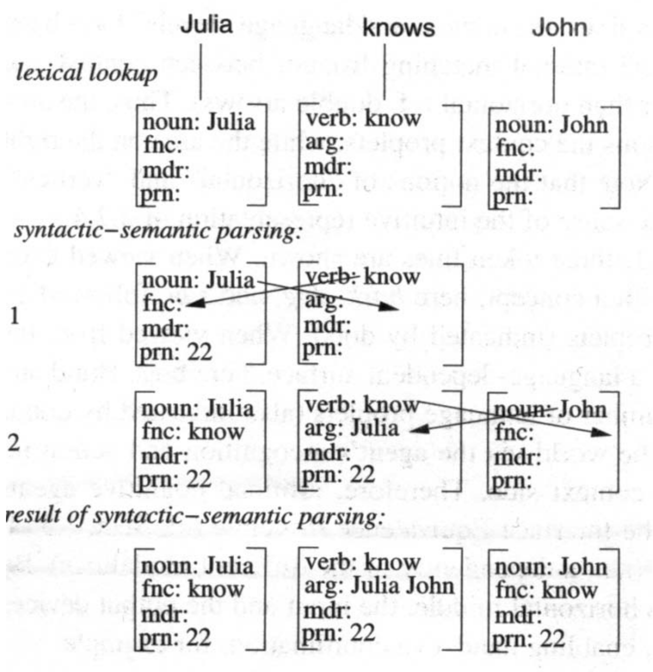

Figure 6. Populating Proplets During Syntactic-Semantic Parsing

As noted above, left-associative grammars are specified in terms of rules governing the possible syntactic continuation patterns, given the portion of a sentence already parsed (sentence start). Rules are specified in terms of binding and replacement variables, along with loading variables that represent inter- and extra-propositional coordination and reference and specify sequences of primitive operations that over write or append to the contents of arguments in a proplet in order to link it to other proplets.

As a sentence is parsed linearly, it is quite likely that there are multiple patterns which might come after the portion of the sentence parsed so far or that there are multiple context proplets associated with a language proplet's core value. In this case parallel interpretation paths are begun, only to be eliminated as soon as there is syntactic, or more commonly semantic, disambiguation. If the end of the sentence is reached without reducing the possible matches to a single representation then the following sentence, if any, is examined to see if there is extra-propositional coordination or reference which would disambiguate the reference syntactically. If not, Hausser proposes that the DBS implementation choose the *best match* by examining the pragmatic context. For instance, suppose the core value is **dog,** the modifier is **your** and I am in hearer mode. If I have had several dogs over my lifetime, a pragmatic examination would note that the associated verb is in the present tense, in which case the match would be made to the token for my current dog. If I own several dogs right now, it may be the case that other modifiers such as **old** or **red** would indicate unambiguous reference to a specific dog. Or, as is common in conversation, I might need to ask the speaker to which dog he is referring in order to complete the match.

The LA-Speak grammar in a Database Semantics implementation is a parallel grammar with production rules that proceed from context propositions to surface word forms to generate sentences of words in linear order.

Once the LA-Hear grammar has populated proplets and sorted them into the Word Bank, the LA-Think grammar supports inferencing in order to answer questions or generate thought patterns that form the basis for sentence production.

The agent in a mature DBS implementation is assumed to have both episodic and absolute propositions in its context knowledge base. Episodic propositions are linked to specific events with spatial and temporal indices, and may have author and intended recipient indices if the proposition encodes a linguistic expression. Absolute propositions represent content that holds independently of any STAR (such as mathematical or scientific assertions), as well as personal beliefs which the agent assumes to be true at a given time during language use. The latter are subject to change as the agent's experience changes, but are active in language interpretation at any given point in time.

Hausser does not provide a substantial fragment for an LA-Think grammar but does provide some basic mechanisms. For instance, he reconstructs the logical principle of *modus ponens* in both the propositional and the first order predicate calculus. Modus ponens consists of arguments of the form

If X is true then Y is true.
X is true.
Therefore Y is true.

Hausser begins the reconstruction by defining an attribute called **cnj** to encode conjunction relationships between context propositions. In truth-conditional semantics, conjunction is true if both propositions are true. From the perspective of natural language in common use, however, logical conjunction does not necessarily present a satisfying way to encode the pragmatic relationship between propositions since the conjunction of any two logically true but otherwise unrelated propositions yields a valid but vacuous expression value of **true** that fails to capture our intuition about the relatedness of the expressions. For instance the sentence **I slept well last night and am ready to face the day** conveys some degree of causality between the two clauses – I can face the day at least in part because I'm rested.

Hausser proposes the use of inference mechanisms that are capable, if need be, of non-monotonic reasoning. As an example he suggests that the **is-a** relationship be replaced within DBS by a new relationship he calls **instantiates**. He then reconstructs *modus ponens* at the first order logic level by representing the premise with a universal quantifier as an absolute proposition and the premise with an existential quantifier as an episodic proposition.

Hausser also proposes that indirect reference and such non-literal references as metaphors be handled at the context rather than the language level, through the pragmatic choice of appropriate instantiations of a concept. This is a departure from Gricean implicature as treated in other natural language approaches. For instance the utterance **Please sit in the apple juice seat**, which Lakoff and Johnson (1980) assert is not meaningful in its own right because it is not a conventional way of referring to anything, in Database Semantics has a clear meaning$_1$ since its surface forms individually denote well-defined literal concepts. The sentence can also have a meaning$_2$ if the agent's context database contains episodic propositions that refer to a seat at a table at whose place setting there is a glass of apple juice. Chapters 4 and 5 examine whether this approach is congruent with current linguistic theory and with research evidence regarding natural language processing in humans.

3.4 Summary of SLIM and Database Semantics

Hausser has proposed an extensively-articulated theory of natural language communication which he asserts is both descriptive of human language use and appropriate for computational natural language capability in embodied artificial agents such as robots. Proceeding from the principles of surface compositionality, time-linear processing of surface word forms and the ontological principle that meaning is assigned internal to a speaking or hearing agent based on context (world) knowledge and belief, he focuses attention on an expanded definition of pragmatics as a key element of language understanding and proposes left-associative grammars and a partial-proposition data structure as the basis for his computational model. Along the way he makes several key assumptions regarding the nature of natural language use and of cognition in humans. The next two chapters evaluate those assumptions in light of current evidence and theory from logistics, cognitive science and neuroscience.

4. Evaluating the SLIM Theory and Database Semantics

Hausser asserts that the basic scenario for natural language use is interactive communication between agents who alternately assume speaker and hearer roles. Implicit in this assertion is that static text corpora are limited and potentially misleading foci for developing human-equivalent computational natural language capability.

Hausser's ultimate goal is to provide a rigorous, formal theory and transparent computational mechanisms suitable for creating an autonomous agent capable of both understanding and generating natural language to the degree and in a way parallel to that of humans. Because human language use is based on cognition rooted in sensory perception and processing, such an agent must be, or simulate being, embodied. Embodied agents must be capable of interacting directly through sensory interfaces with a real or simulated external world that includes but is not limited to other natural language users.

Hausser proceeds from four core principles, the first two of which are that natural language is surface compositional in its syntax-semantics relationship and is processed time linearly without recursion over surface or assumed deep structure. His third principle is that meaning does not primarily reside in the external-to-the-agent truth conditionality of predicates asserted of entities in a given world, but rather is assigned procedurally within the language-using agent who is situated in time and space and whose past and present experiences and imagination provide the referents for semantic-level concepts. However, Hausser retains a propositional format for both semantic and context-level representations and proposes the use of a left-associative grammar as the mechanism to traverse context proplets in either drawing conclusions from statements that one hears or in assembling context propositions that will form the basis for statement production – a function that is only sparsely described in his publications to date.

Hausser's fourth principle is that language refers through procedural assignment at the pragmatic level. Surface language denotes unambiguous concepts defined in terms of feature-attribute structures. Pragmatic mapping of a semantic concept to a specific referent is made on a 'best fit' procedural basis from among the instances in context experience to which the concept might apply. Syntactic and semantic ambiguity during parsing is handled by semantic doubling, *i.e.* by maintaining parallel branches of interpretation that get pruned as successive words or subsequent sentences disambiguate the denotation.

Hausser asserts that, *contra* the logical positivists, vagueness in natural language is not inherent in the semantics of statements but rather occurs at the level of pragmatic interpretation. For instance, he cites the adjective **red** as an example of a concept whose denotation is intuitively clear. The phrase **red object** might present an unclear reference in a context where there are many objects in a wide variety of colors, but that same phrase might be easily assigned a single interpretation if the context has only one object whose color lies, say, somewhere in the pink-red-orange spectrum.

Language builds on cognition, which is not dependent on language to function. Central cognition is based on concepts and propositions and hence on inference. Propositions have three possible core attributes with equivalents in language, world context and logical formulae. Nouns, verbs and adjectives or adverbials in language are equivalent to objects, relations and properties in the world context and to arguments, functors and modifiers in formal logic and type theory.

Central cognition is surrounded by peripheral cognition, in which sensory inputs are received and patterns recognized, or conversely intent is translated into motor actions. A cognitive agent has both peripheral and central cognition – language is an extra that builds on these.

This chapter places Hausser's theory and computational approach in the context of linguistic theory and, where available, experimental evidence in order to highlight the important questions he raises regarding current work in the field and to discern future research questions raised by his approach. Chapter 5 examines evidence related to the specific challenges associated with natural language in an embodied and socially situated agent.

4.1 The Structure of Language: Syntax and Semantics

That natural language has structure follows from the fact that a finite number of elements (word forms and punctuation) can generate an infinite number of sentences in any given language and that language users can both understand and produce sentences they've never heard or produced before. The role of grammar is to characterize that structure in such way that the semantics (meaning) of the sentence is inferred as a result of that structure. A well-defined grammar is a critical element in computational linguistics because computation requires algorithmic description of processes. This is the case for statistical as well as logical inference.

Since the early 1970s a wide variety of grammatical approaches have been proposed, often in succession by the same researchers. These may be grouped into four families: categorial, constraint-based, functional and cognitive.

Categorial grammar is, as briefly described in chapter 3, an outgrowth of formal type theory in mathematics. A categorial grammar describes syntax primarily in terms of functor-argument relationships between words and expressions along with rules governing the application of the functors (Steedman 1996; 2000). Basic expression types might consist of Noun, Noun Phrase, Verb, and so on. When the grammar is restricted to these elements it is able to generate languages that are syntactically context-free, *i.e.* in which clauses may be nested within one another but may not overlap in scope. Despite early optimism regarding the possibility of representing natural language syntax and semantics using these elegant and computationally tractable formalisms, there is good reason to believe that many, if not most, natural languages are not context-free (*c.f.* Higginbotham 1984 regarding English).

However, categorial grammars are attractive because their formal rigor and the transparency of operations they generate allow them to be associated to semantic interpretation with both syntactic and semantic compositionality and a clear link between the two levels of structure. This is accomplished by assigning interpretation types to the basic phrase categories and then defining functions for each interpretation type. Sub-categories may be defined through the assignment of values to features such as person, tense and number. Intensionality and quantification may also be added, as in the example of Montague grammars. Such grammars are heavily lexicalized, *i.e.* they push most of the information required for parsing and semantic understanding into the lexicon rather than in category definitions (Szabolcsi 1992) with the goal being a 'variable-free grammar'. Bentham, Groendijk, de Johngh, Stukof and Verkuyl (working jointly under the pseudonym L. T. F. Gamut) present a clear exposition of categorial grammar from its foundations to early work in discourse representation in (Gamut 1991). Hausser presents a more formal description and critique in (Hausser 2001a).

Hausser worked with categorial grammars from the beginning of his academic career as a graduate student through his efforts while at Stanford to instantiate his ORTAX grammar in a computational implementation during the period 1983-1984. However, he diverged from the mainstream of categorial approaches beginning in 1978 with his insistence on surface compositionality and in 1980 with his analysis of the semantics / pragmatics boundary and his insistence that meaning assignment is not a matter of truth conditionality at the semantic level but rather of procedural reference to extra-linguistic experience at the pragmatic level. The break away from categorial grammar was complete by 1984 when he rejected the search for constituent structure in natural language in favor of time linear, surface compositional parsing using left-associative grammars. Nonetheless the feature-attribute structures that make up proplets in Database Semantics owe a great deal to the categorial approach to language structure.

Categorial grammars are no longer a dominant paradigm in computational linguistics but some researchers continue to propose extensions and modifications to the basic approach, including the use of combinatory logic rather than the λ calculus for combining expressions (Steedman 1987; 2000).

Hausser was not alone in finding categorial grammar inadequate for natural language theory and computational capability. Perhaps the largest and most diverse family of grammars are the constraint-based approaches. Sometimes called unification or information-based grammars, this family includes generalized phrase structure grammar (GPSG) and head-driven phrase structure grammar (HPSG). Phrase structure grammars are again built around feature-attribute structures, with the possibility that some attributes are themselves populated by complex data structures. These may take the form of matrices that nest sub-attributes and their values, of trees or of lists but in all cases constitute directed acyclic graphs – an important feature for computational search algorithms. (Note the difference from LAGs, in which proplets are flat. This difference in data structure is related to the difference in parsing technique: LAGs parse from linearly with no recursion whereas phrase structure grammars recurse.) A key operation on these feature structures is *unification*, a first order logic operation that combines the directed acyclic feature-attribute graphs according to the values of head nodes (hence 'Head-driven Phrase Structure Grammar') or of sub-graphs (Blackburn and Bos 2005). Unification considers logical expressions or feature structures to be bearers of partial information which can be merged, hence the occasional early use of 'information-based' for this family of grammars. Tree-adjourning grammars (Joshi and Rambow 2003), in which nodes of trees are re-written as other tree structures may also be considered as part of this family.

Pollard argues (Pollard 1996) for the name ‘constraint-based’ for this wide and somewhat diverse family because from a theoretical linguistics perspective, while they differ in the algorithms they bring to bear, what is characteristic and most important about each of them is the way in which they use features and attribute values to constrain the set of possible well-formed expression patterns and forbid other patterns as ill-formed and ungrammatical. He includes (although others might not) the lexical-functional grammars in this family. Lexical-functional grammars (LFGs) are defined in terms of both feature-attribute structures and syntactic structures that define legal constituents of well-formed sentences, along with a number of other structures that govern the arity (number of arguments) for various predicates, phonological and morphological constraints (Bresnan 2001). Lexical-functional grammars and Hausser’s Database Semantics have in common the use of the lexicon to relate *e.g.* active and passive forms of a verb to one another, although Hausser accomplishes this through a shared core value in the proplet while LFGs do so by positing an explicit relation formalism. Also compatible with this family is Jackendoff’s Representational Modularity theory which describes the language faculty in terms of mental modules and mutually co-constraining structures (Jackendoff 1996). Shieber provides an introduction to constraint-based grammar theories in (Shieber 1986) and formal treatment in (Shieber 1992).

In line with constraint-based grammars, Discourse Representation Theory provides constraint-based treatment of extra-sentential reference, anaphora and similar linguistic features through representation into first order logic (Kemp and Reyle 1993). Significant work has been done in DTR by a number of researchers since Kemp's original publications, which focused on unbound anaphora, *i.e.* pronouns whose reference crosses sentence boundaries. As with Hausser's model for extra-sentential reference (Hausser 2006a), this is accomplished through the use of partially populated formal representations whose attributes are mapped to values as the discourse proceeds.

Pollard identifies eight key properties of constraint-based grammatical theories:

1. **Generativity** – the theory must at a minimum tell us what the well-formed structures in a language are. This requires that the theory make determinate what formal objects are used for modeling linguistic entities, whether a string of symbols counts as one of the assertions or constraints in the grammar and, given a grammar and a specific object, whether the object satisfies the constraints imposed by the grammar
2. **Expressivity** – the language within which the grammatical theory is expressed should be maximally expressive, *i.e.* unconstrained. It is the theory and not the language in which it is expressed that should impose constraints on the grammars.
3. **Empirical adequacy** – since any theory can be re-axiomatized (expressed in new sets of axioms with the same entailments), there are no 'deep principles' from which to derive grammatical theory. Therefore constraint-based theories begin by identifying the constraints active in

languages and work back to sort those into axioms and dependent theorems.

4. **Psycholinguistic responsibility** – although constraint-based grammatical theories aim only at describing linguistic competence and not the actual means by which humans perform linguistic tasks, they must in principle be capable of being consulted or carried out by humans. In this sense a constraint-based grammar "is more like a data base or a knowledge representation system than it is like a collection of algorithms."
5. **Non-destructiveness** – a generalization of the property of monotonicity for unification, this means that the grammar should not irreducibly make reference to operations that destroy linguistic structure.
6. **Locality** – for any candidate structure, the question of whether or not that structure satisfies the grammatical constraints must be determined locally, solely on the basis of that structure, without reference to other 'competing' structures. This property excludes *c.f.* Optimality Theory (Prince and Smolensky 1993) from the constraint-based family.
7. **Parallelism** – given that any linguistic theory must make reference to multiple levels of representation, in a constraint-based grammatical theory no level can be created by transforming (operating destructively on) another level. This includes the restriction that no lexical items can be inserted into the lexicon during expression processing, *i.e.* any required recursion should be reflected in the lexicon entries from the start.

8. **Radical non-autonomy** – primarily of syntax. In constraint-based grammatical theories, constraints operate either within a specific level of representation or through mutual constraint at the interface between levels of representation,

In (Pollard 1996) Pollard is at pains to contrast the family of constraint-based grammatical theories with Chomsky's Minimalist Program, which fails to meet most these criteria. But his aim is broader, namely to lay the foundation on which to address the breadth of challenges beyond grammar that face computational linguistics. He writes:

> The last thing we want is an autonomous theory of syntax. Instead what we need are theories that deal simultaneously with all linguistically relevant factors, be they phonetic, morphological, syntactic, semantic or pragmatic. And once we get serious about interfacing the theory of competence with processing models, nonlinguistic factors such as world knowledge, frequency considerations, and the beliefs and goals of speakers must also be brought into the picture. It seems to me that, among the existing options, constraint-based grammar has the highest potential to rise to this challenge. (*ibid.*)

Do Hausser's SLIM theory and Database Semantics approach belong to the constraint-based grammar family? Taken together it seems that they do. LAGs impose some constraints not found in *e.g.* HPSG, such as the flat feature-attribute structure in proplets and the strict time-linear parse. Hausser is at pains to define formal foundational definitions for his grammatical model using a highly expressive language traceable to formal logic. He regards his SLIM theory to be empirically grounded and his Input/Output Equivalency criterion addresses Pollard's property of psycholinguistic responsibility, although Hausser arguably proceeds past language competency to tentatively propose (or more properly, assume) a model for language processing by humans. LAGs parse non-destructively and locally. Above all, Hausser's levels of representation from the lexicon through linguistic types and tokens to context types and tokens embody Pollard's criteria of parallelism and radical non-autonomy.

What is particularly striking is the degree to which Hausser's goals parallel Pollard's call for "interfacing the theory of competence with processing models", bringing into the picture "nonlinguistic factors such as world knowledge, frequency considerations, and the beliefs and goals of speakers". World (context) knowledge is a key element in Database Semantics, frequency considerations can (if one chooses) play a role in pragmatic disambiguation of reference and the beliefs and goals of speakers are expected, if under-described, features of the context knowledge base.

However, Pollard's enthusiasm for constraint-based grammatical theories notwithstanding, they do face specific challenges from the functional and cognitive grammar approaches.

Functional Grammar (FG) was first proposed by Dik and has since attracted significant research across many natural languages. (Note that Halliday's Systematic Functional Grammar is a different theory (Halliday 2004). SFG is a broad analysis of language in use.)

Functional grammar draws upon the basic ideas of situation semantics (*c.f.* Barwise *et al*. 1991) and is focused on function-based relations between elements within natural languages as expressions address states of affairs (Dik 1997a; 1997b):

- **Syntactic functions** such as Subject and Object define different perspectives regarding states of affairs as presented in linguistic expressions.

- **Semantic functions** such as Agent and Recipient define different perspectives through which states of affairs are presented in linguistic expressions.
- **Pragmatic functions** such as Theme, Tail, Topic and Focus define the informational status of constituents of linguistic expressions. They relate to the embedding of the expression in the ongoing discourse, that is, are determined by the status of the pragmatic information of Speaker and Addressee as it develops in verbal interaction.

A key goal of Functional Grammar is to minimize the number of rules, operations or procedures that link surface language and the structures which are postulated to underlie that surface. Correspondingly, structure-changing transformations are avoided, as are empty elements in the grammar. The grammar is organized around abstract underlying clause structures in which entities, properties or relations and states of affairs are represented by terms, predicates and predication, respectively. In addition, clause structures can represent interpersonal levels such as propositions (possible facts) and clauses proper (speech acts). The lexicon is organized primarily around predicate frames that specify fundamental semantic and syntactic properties, such as the syntactic category of the predicate, the number of arguments and the semantic functions of the arguments. The lexicon also contains entries for basic terms such as nouns and pronouns, along with composite terms in which modifiers are specified in a predicate form. States of Affairs are predicate/term pairs that specify something that might be true. SOAs are marked for a number of semantic parameters including Control, Dynamism, and Telicity, as well as satellite parameters of terms such as Time, Place, Manner and so on. Speech Act structures capture both linguistic and paralinguistic characteristics such as tonal emphasis. There are corresponding structures to indicate syntactic relations, as well.

As can be seen from the brief description, Functional Grammar attempts to account formally for the full, multi-dimensional structure of a language along with its use patterns. It is therefore much more ambitious in scope and representational richness than either categorial or constraint-based grammatical theories. Although Hausser's SLIM theory and Database Semantics mechanisms share with FG an interest in mapping syntax and semantic content to context knowledge, Hausser does not address semantic roles in his grammar and would reject the assumption that in fact those roles belong to semantics proper. He has not directly addressed FG, but a review of the development of his approach suggests he would assign these roles to pragmatic, procedural processing of language since they are not well-defined concepts (core meanings) but rather inferred use aspects of core meanings.

André Wlodarcyzk proposes a functionally-oriented formalism that also captures semantic roles but that is somewhat less broad in its intended scope than Dik's functional grammar. Like FG his Meta-Informative Centering Theory is based around situations within which infons, *i.e.* elements of semantic knowledge (declarative *know-what*), are derived from communicative interactions among language users (Wlodarczyk 2008a). Wlodarczyk posits that when infons are organized and complemented by other infons they are enriched with ontological knowledge in the form of procedural *know-how*. Taking his cue from robotics and other multi-agent intelligent systems, Wlodarczyk establishes a high level formal ontology of agent features and roles that agents play in situations. This ontology is augmented by domain-specific ontologies in the interpretation of language expressions. The result is an associative semantics that can be represented by a hypergraph which supports recursive transit. This formalism intends to bridge from semantics to pragmatics and permits pragmatic as well as semantic validation of utterances (Wlodarczyk and Wlodarczk 2008). Meta-Informative Centering Theory has been applied at the Sorbonne in the CASK project, where it informs knowledge discovery in databases that include freeform text (Wlodarczyk 2009).

MCT is an intentionally cognitive approach to natural language representation and analysis, as is Hausser's Database Semantics. Of the two approaches MCT is, however, much richer in its representation of the ontological and functional relationships between concepts and agents involved in a natural language exchange. While there is nothing in Database Semantics that prevents the use of formal ontologies for procedural assignment of meaning and pragmatic interpretation, there is nothing that recognizes or supports such use either, other than (as illustrated in Hausser 2005b) the use of the basic proplet and left associative grammar formalisms. Moreover, the notion of semantic role in the sense used in MCT and FG is alien to Hausser's focus of semantics on conceptual content elements alone.

Functional Grammar and especially MCT raise the question of the degree to which it is necessary to have a full model of cognition in order to create computational natural language capability at the human level. Chapter 5 describes one current project in natural language capability embodied in physical robots for which formal ontologies play a key role in pragmatic interpretation. Broadly speaking, however, the lack of semantic role representation or of formal ontologies that capture the relations between roles and between concepts means that the computational simplicity and flat data structures of Hausser's Database Semantics are achieved at the cost of pushing key elements of pragmatic interpretation down into unspecified procedural implementation where they receive no pre-determined formal representation or manipulation. This, of course, is equally true of categorial and other constraint-based grammatical theories. It does suggest, however, that such approaches are not sufficiently rich in formalisms to directly achieve Pollard's goal to interface the theory of competence with processing models that can address world knowledge and speaker/hearer attitudes and goals. This is a particular omission with regard to computational natural language capability in embodied agents, for which world knowledge and interaction are a key requirement.

4.2 Cognition and Content

The requisite role of models of cognition in theories of natural language and in computational natural language capability takes on another dimension in the fourth family of grammatical theories, namely Cognitive Grammar.

As proposed and developed by Langacker (Langacker 2007), CG is a strongly functional, as opposed to formal, approach to languages. Langacker places meaning primarily (but not solely) in the mind of the speaker/hearer. However, he associates it not with static concepts but with dynamic *conceptualization*, a term that includes new as well as established conceptions; sensory, motor and emotive experience; apprehension of the linguistic, social and cultural context; and the conceptions that emerge and develop over time as the result of a sequence of experiences. CG rejects a primary focus on propositional structure for cognition, stressing instead the imagistic character of basic concepts. Image schemas, *i.e.* "schematized patterns of activity abstracted to bodily experience, especially pertaining to vision, space, motion and force" (*ibid.*, 32) are seen as basic, pre-conceptual structures that give rise to more elaborate and abstract conceptions through combination and metaphorical projection. Langacker provides a simple example of an image schema for the notion **Enter** in Figure 7 (*ibid.*, 33):

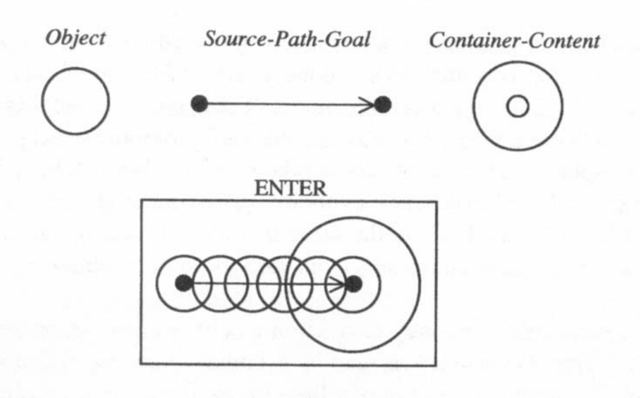

Figure 7. Image Schema for the Notion 'Enter'

Langacker distinguishes several kinds of fundamental notions, each of which he claims is fundamental in a different way. These include **minimal concepts** in particular domains of experience, such as line and curvature in the spatial domain, brightness and focal color in vision, precedence in time and so forth. **Configurational concepts** are independent of any particular domain and include such notions as boundary, contrast, separation, proximity, contact, inclusion, separation, and group. Langacker asserts that configuational concepts come closest to the spirit of image schemas. **Conceptual archetypes** are experientially grounded conceptual gestalts such as **object**, **seeing something**, **holding something**, **the human face**, **exerting pressure to move something** and so forth.

Langacker does not regard this three-way distinction as either definitive or clear-cut. This is because all conceptions are dynamic (*i.e.* residing in processing activity) and so there is no sharp boundary between simple concepts and certain basic cognitive abilities. From this flows the general Cognitive Grammar proposal about certain grammatical notions that are "fundamental and possibly universal", to include noun, verb, subject, object and possessive. This proposal has several parts (*ibid.*, 34):

> 1. Each such notion can be characterized semantically in terms of both a **prototype**, valid for central instances, and a **schema** instantiated by all instances.
> 2. The prototypical meaning consists of an experientially grounded conceptual archetype.
> 3. The schematic meaning resides in a domain-independent cognitive ability.
> 4. The basic abilities are initially manifested in the corresponding archetypes. Presumably innate, the abilities make it possible for structured experience, based on these archetypes, to occur in the first place.
> 5. At a later developmental stage, these same abilities are extended to other domains of experience, where their application is less automatic (hence more apparent.

Langacker works out as an example the case of nouns (*ibid.*):

> 1. The noun category has both a semantic prototype and a semantic schema.
> 2. The prototype is the conception of a physical object.
> 3. Providing the semantic characterization is our capacity for conceptual grouping.
> 4. Developmentally, conceptual grouping is first manifested in the apprehension of physical objects, the noun category prototype.
> 5. It figures subsequently in the apprehension of the many other kinds of entities also coded by nouns.

Thus we move from perceiving a discrete physical object such as a ball to a noun like **orchard** which for most linguistic purposes is also – but by extension, not by direct experience - treated as a singular, common noun referring to a countable 'object' (*i.e.* we can refer to **orchards** just as we refer to **balls** despite the fact that orchards are not discrete in the same way as balls – an issue that, as Chapter 5 discusses, arises with regard to representation of concepts in robots).

Because our species has evolved to cope with the world around us, we all perceive physical objects and employ a similar grouping capacity, *i.e.* our comprehension is largely commensurate because it is grounded in common bodily experience. Langacker writes (*ibid*., 35):

> Our apprehension of the world is thus active, dynamic, and constructive in nature. A fundamentally important consequence is that the conceptions evoked as linguistic meaning are nontransparent: they do not simply reflect or correspond to the world in a wholly straightforward manner, nor are they derivable in any direct or automatic way from objective circumstances. Instead, a conceptualist semantics must start by recognizing the prevalence – indeed, the utter pervasiveness – of imaginative devices and mental constructions … We further demonstrate our imaginative capacity in constructing and manipulating an extraordinary variety of mental spaces (Fauconnier 1985). Some types of mental spaces … are those representing a hypothetical situation, a particular person's beliefs, the situation obtaining at a certain time or place, and the content of reported speech.

The remainder of (Langacker 2008) provides an extensive introduction to grammatical analysis from a CG perspective, from simple expressions through sentences, discourse, speech acts and expressives. The theory is not presented in computational form.

Hausser's brief treatment of metaphor makes it plain that he rejects the approach of Lakoff and Johnson to metaphorical use of language on which Langacker builds his cognitive grammar theory. In (Hausser 1999d) he asserts the central role of literal meaning$_1$ denotation in metaphor meaning. As mentioned in Chapter 3, Hausser replies to their assertion that the sentence **Please sit in the apple juice seat** has no meaning at all when taken in isolation by asserting that the sentence has no meaning in a context either under their interpretation unless there is a seat with an apple juice setting to which to refer. Here Hausser emphasizes the way in which his approach to pragmatics extends beyond that of Grice to encompass the procedural process of finding the most likely referent for any expression.

Prototype theory plays a large role in cognitive science and was first proposed by Rausch, a cognitive psychologist interested in categorization (Rausch 1973). Hausser addresses categorization within Database Semantics in (Hausser 2005b) in a formalization that uses context proplets to encode **is-a** and **instantiates** relationships. In addition he borrows from a visual analysis technique called Recognition By Components (Biederman 1987) to illustrate how concept relationships can be encoded and processed using this proplet formalism to replicate the equivalent of a formal ontology. In this treatment, as elsewhere, Hausser's focus is on computational efficiency and transparency of derivation. He does not directly address contemporary ontological formulations, formal concept analysis or the description logics that currently form the basis of most concept space formalization in semantic applications. He does, however, attempt to integrate some aspects of visual pattern matching into the process of concept development.

In addition to Langacker's CG, Construction Grammar also falls into the cognitive family of grammatical theories. Developed by Langacker's student Goldberg and Lakoff (among others) Construction Grammar takes grammatical constructions rather than individual syntactic units and their combinatory rules as the primary unit of grammars, which then consist of taxonomies defining families of constructions. Some construction grammarians work only at the theory level while others adopt unification-based frameworks to describe syntax, as with Fillmore's extensive FrameNet project and the motivating theory behind it (*c.f.* Fillmore 1976). In particular Embodied Construction Grammar specifically focuses on embodiment and sensorimotor experience as foundational, positing that the content of all linguistic signs is based on mental simulation of such experience (Feldman 2006). ECG also adopts a unification-based framework for grammar representation.

Cognitive Grammar and related theories raise significant questions regarding the theoretical basis and computational approaches appropriate for embodied agents with human level natural language capability. Rejecting the propositional focus that is central not only to Hausser's approach but also to *e.g.* Anderson and Bower's ACT-R cognitive model (Anderson and Bower 1974), it more fundamentally rejects the idea that concepts or semantic content can be described in terms of clearly delineated feature-attribute pairs, suitable for manipulation using formal logic, proposing instead prototypical imagistic and sensory patterns more amenable, perhaps, to associative machine learning approaches.

Chapter 5 considers the challenge and implications of embodiment for cognition and natural language capability in greater detail.

4.3 Context, Salience and Time Linear Processing

Is Hausser correct to assert that humans process natural language linearly in time as individual words are perceived? The question of time linear processing bears on a number of issues in computational linguistics, including grammatical theory insofar as we adopt Pollard's goal of going beyond language competency models to address how humans process language and link it to world knowledge.

Hausser's Left Associative Grammar formalism was not the first left-to-right, non-recursive parsing approach for natural language. (Tomita 1987), for instance, applies Aho's LR parsing approach for context-free programming languages to natural languages represented as augmented context-free grammars. His algorithm addresses computational efficiency however and not the stronger assertion Hausser makes to the effect that human processing of natural language is time linear. Phillips on the other hand argues for linear order in establishing consitutency in a generative phrase-based grammar (Phillips 1998).

A stronger case for time linear language processing can be made on the basis of psycholinguistic and neuroscientific evidence advanced in support of the graded salience hypothesis proposed by Giora (Giora 2003). Salience refers to the relative importance or relatedness of a word or phrase to its context. In the case where a word of the appropriate form and syntactic category might have more than one meaning, how do language users decide which meaning was intended by the speaker?

The modular mind approach treats meanings as encapsulated and subject to rapid comparison and selection by the language faculty (Fodor 1983). This approach is closely linked to the representational or computational theory of mind which posits discrete mental tokens representing concepts which are manipulated according to a cognitive process unique to language. Modular approaches to salience assume rapid access to conceptual tokens (in some variants to all possibly relevant tokens simultaneously) without prior coordination to contextual information, leading to occasional misinterpretations that require revision for sentence comprehension. The interactionist or direct access view, on the other hand, asserts that context effects on meaning selection are primary. A highly biased linguistic and/or non-linguistic context strongly constrains meaning selection; a less constraining context requires additional processing in order to support comprehension (MacWhinney 1987; Bates and MacWhinney 1989; Bates 1999). Evidence for the interactionist view is presented by (*c.f.* Vu *et al*. 2000) who demonstrate that a sentence context that is very constraining may activate the relevant meaning neurally without activating alternative meanings, as in the sentence **The biologist wounded the bat** which constrains the meaning of **bat** compared to **He located the bat**. This is true even when, as in the example given, the activated meaning is less frequent across all uses of the word in the language.

The graded salience hypothesis blends elements of both approaches. Along with the representational/modular mind approach, graded salience accepts the existence of two distinct mechanisms involved in language interpretation, one that is sensitive only to linguistic information and functions ‘bottom up’ on a word by word or phrase by phrase basis and another that proceeds ‘top down’ from all contextual knowledge, linguistic and extralinguistic. However, the modular, lexical access mechanism is assumed to be ordered, resulting in access to meanings that are more conventional, more frequent, more familiar to the hearer or more prototypical much more rapidly than to others. As a result, the graded salience hypothesis predicts that coded meanings of high salience will be accessed as soon as the relevant neural structures are activated in the brain, regardless of contextual information or speaker intent. Contextual information may be sufficiently strong to arrive even faster than the salient lexical meaning, but even in this case the two retrieval paths remain separate and do not interact.

The assumption of two parallel meaning mechanisms, one encapsulated and linguistic and the other integrative and incorporating a wide range of sensory and non-linguistic contents, provides the basis for testing the graded salience hypothesis based on retrieval/response times for meanings under various contextual and salience conditions. Giora cites research results in her own lab and across the literature, especially with regard to processing of figurative language, and finds significant support both for graded salience at the linguistic level and for parallel processing mechanisms. Notably, rapid access to salient linguistic contents occurs as words and terms are encountered, *i.e.* thereby supporting the hypothesis of time linear processing.

To draw an analogy to computational reasoning, the modular approach is roughly equivalent to traditional AI based on symbolic propositional content, with the assumption that there is some mechanism for updating the knowledge base as a result of sensory inputs. The interactionist approach is strongly associative, with the possible assumption that there is an interface from which something like a symbolic representation of associative classifications can be drawn in some manner. The graded salience hypothesis is roughly the equivalent of positing both a semi-autonomous symbolic planning and reasoning layer and a separate, ongoing associative layer within a software agent. The former is invoked first during language interpretation while the latter eventually confirms or negates the choice of the most salient meaning for a word or phrase.

Giora's hypothesis is evocative of Paivio's dual coding approach to mental representation of cognitive contents (Paivio 1990) which posits both a relational and an associative model for meaning. However, Giora's work clarifies the parallel access and the conditions that establish precedence among the two mechanisms. Paivio's associative mechanism is strongly visual in nature and his explanation of semantic assignment has similarities to Langacker's notion of image schemas.

Not all context is internal to the experience of the natural language speaker/hearer. Language use is situated within social contexts as well as resulting from the embodiment of the user. This is true in several ways. First, the fact that word senses can be ranked in frequency of use across the population of speakers of that language at some given time means that one element of saliency is clearly established outside of the speaker/hearer him or herself. Such frequency becomes an element of the speaker/hearer's internal context through experience. Second, pragmatic elements of language such as the speaker's attitude as inferred from intonation or ironic word choice function within a specific social context and often have far less communicative effect when that context is removed.

Of equal importance, however, is the fact that language changes over time as new literal and figurative meanings for words and phrases are added to the language and gain currency. Support is lent to Langacker's focus on conceptualization vs. static concepts through an examination of the cognitive effects on people learning a second or later language. Thus Kecskes finds that bilingual and multi-lingual speakers evidence the development of synergic concepts that are not exactly the equivalent of the corresponding concept in either their native language or the newly-acquired language (Kecskes 2007). Kecskes posits two distinct levels of mental representation, one linguistic and one contextual. As in Hausser's theory, Kecskes distinguishes abstract concepts at the semantic level from the conceptual interpretations which an expression has in a certain context of use. However, Kecskes explicitly limits compositionality to the semantic level whereas Hausser attempts to maintain it through context token proplets and the LA-Think grammar.

Kecskes agrees with the focus in cognitive science that stresses the acquisition of new concepts incrementally over time through repeated exposure and refinement. He writes (*ibid.*):

> A word (label) is a symbol that pulls together all knowledge and information that has been connected with the use of that label. It encodes the history of the use of that label in various situational contexts. The amount of this knowledge and information with its fuzzy boundaries creates what we call a concept.
>
> A concept is a construct that blends knowledge gained from actual situational contexts in an individual-centered way. The reason why concepts convey relatively similar information for a particular speech community is that community members have had relatively similar experience with the given label in language use.

Keckes rejects the idea that word meanings are underspecified at the semantic level and get specified by the context. Instead he proposes that words to some degree provide their own context. Lexical units and situational context mutually define and depend on one another through a dynamical process that is particularly visible during new language acquisition. As a result Kecskes proposes a Dynamical Model of Meaning (Kecskes 2009) in which both lexical and situational context contribute to meaning attribution and creation. DMM asserts that natural language is never context-free since language encodes prior experience, including the experience of how a word or phrase has been used in the past. Moreover, since different cultures organize their background knowledge differently, no language is culture-free. Therefore both the meaning-construction and meaning-prompting systems of a language are culture-specific. Kecskes identifies two levels of meaning. Coresense is abstracted from prior uses of a word in a context and captures the interface between semantic and conceptual levels in terms of the most familiar, regular, typical or possibly most frequent uses. Coresense is "denotational, diachronic, relatively constant and objective" and reflects the use of the word in a community over time. Consense, on the other hand, is "active, subjective, referential and connotational and changed by actual situational context". Consense encodes *private* context and determines the position of a wordsense in the salience gradation. These are mutually influencing: coresense is learned through language acquisition and comprehension, but over time consenses that are repeated can lead to changes in the shared coresense.

Evidence for the dynamical nature of meaning within the ego-centric speaker/hearer (who nonetheless communicates within a community of shared language use) comes in part from examining how bilingual and multilingual speakers master the use of lexical idioms (**kick the bucket**, **eat one's words**) and pragmatic idioms (**I'll see you later**) as well as from evidence of shifts in consenses that underlie the use of both languages. Pragmatic idioms in particular encode significant shared cultural context which can be opaque to new language learners. Like Hausser, Kecskes extends our understanding of pragmatics beyond the Gricean sense of implicature and speaker intent by focusing on the assignment of meaning internal to the speaker/hearer. However, Kecskes' Dynamical Model of Meaning is significantly richer than Hausser's approach in its understanding of multiple contexts and in its description of the dialectical development of meaning over time both within and between language users.

4.4 Critiquing Hausser

A review of recent linguistic and cognitive theory and some supporting experimental evidence confirms the observation that it is not Hausser's grammatical formalisms but rather his cognitive assumptions that are most open to challenge. Recent theory and experimental evidence suggest that context knowledge is not organized in feature-attribute frames associated transparently with unambiguous, clearly delineated concepts. Instead, it appears that the conceptual framework that underlies language use is more heavily dominated in one way or another by associative processing and that concepts may function more as prototypical patterns than as feature templates. Hausser acknowledges this to some degree in (Hausser 2005b) but has provided very little in the way of examples or higher level processing mechanisms for encoding this context development. In addition, Kecskes and other linguists assert that semantic/pragmatic interaction is dynamical, bi-directional and has both an ego-centric and a social dimension for which Hausser does not explicitly account.

That said, Hausser's SLIM theory appears to be validated on several key points, including time linear processing, the internal locus of meaning assignment and the role of pragmatic procedures in assigning appropriate meaning within a context that includes internal representation of experiences from the external world. As with other approaches that proceed from constraint-based grammatical theories, his approach lacks formalization for semantic roles and has little to say about either the dynamical nature or the social context of meaning creation and meaning change over time.

Hausser stresses the role of context knowledge in language understanding and production but his treatment of context representation and reasoning is extremely under-developed, a weakness that matters because his intent is to lay the groundwork for human-level computational language capability in *embodied* agents whose use of language extends to interactive (and therefore, implicitly, socially *situated*) communication. Chapter 5 explores the implications of embodiment and social situatedness for such an agent.

5. The Challenge of Embodiment

The comparison of Hausser's approach with recent linguistic theory in Chapter 4 focuses attention on the nature of cognition, especially with regard to concepts and concept formation, and on the link between natural language use and cognition in humans. Cognitive linguistics asserts that the concepts which form semantic content in natural language emerge as abstractions from concrete physical experiences, but as prototypes with expanding sets of referents rather than in the feature-attribute sense proposed by Hausser and others whose work is grounded in formalist theory. However, the research undertaken and cited by Giora which establishes the key role of graded salience for a hearer's assignment of meaning to words and expressions makes it clear that, whatever the nature of concepts might be, they are sufficiently encapsulated to be accessible for very rapid retrieval during cognitive tasks. Hence she proposes two channels of meaning assignment process, one that functions rapidly through the application of encapsulated salient meaning to a linguistic expression and another that initially functions separately from the first and that applies context information in an integrative manner in order to validate or correct the selected meaning.

Kecskes' Dynamical Model of Meaning extends the locus of concept development and meaning to establish dueling contexts, one that encodes the relatively stable meaning of a word or expression which is shared within a community of language users and another that is ego-centric and includes the speaker/hearer's unique set of accumulated experience and associations with that linguistic unit. These two loci of meaning creation are dialectically interdependent, accounting both for the cognitive shifts that accompany second language acquisition and also for the change in common usage over time. Taken together Giora and Kecskes propose that natural language is both embodied and socially situated.

5.1 Implications of Embodiment in Situated Natural Agents

Embodied and situated cognition has received considerable theoretical and experimental attention recent years, with competing theoretical typologies pointing to a rapidly growing body of experimental and observational data. Shapiro refers to this as the embodied cognition research *programme* rather than *theory* to indicate that the commitments and subject matters of embodied cognition (EC) are somewhat nebulously defined at present (Shapiro, 2007). Contrasting EC to the cognition-as-computation approach which stresses the manipulation of symbolic representations internal to the cognitive agent, *i.e.* in which sensory inputs once received and translated play no continuing role, Shapiro identifies three new directions that EC research takes in comparison to earlier work in cognitive science and artificial intelligence. The first new direction places less emphasis on the role of semantically endowed symbols and more on properties that emerge from the physical attributes of the body. Here Shapiro notes extensive work on sensory perception that suggests that body structure itself creates constraints and opportunities for neural control. For instance the distance between one's ears and the density of one's cranial matter can be seen as directly involved in associative auditory processing without translation or representation into variables for algorithmic processing. Second, rather than conceiving of cognition as “the churnings of a brain isolated from the body and the environment in which it is situated” EC attempts to account for the content of cognition by appeal to the nature of the body containing the brain. Here Shapiro cites Lakoff and Johnson’s assertion that almost all of our concepts derive originally from the use of metaphorical reasoning extending from some basic concepts that stem directly from the body and its interaction with the environment (Lakoff and Johnson 1980). And third, instead of viewing cognition as “beginning with the stimulation of afferent nerves and ending with signals to efferent nerves”, *i.e.* in Hausser’s approach as beginning with inputs to peripheral cognition and ending with outputs in the form of action signals to the body, EC may see cognitive processes as extending into the environment in which the organism lives. On this point Clark notes that human beings go to great lengths, not only to create tools that are extensions of cognition (pencil and paper, PDAs and smart phones) but also to organize the environment to reduce the cognitive load required for daily activities such as placing keys by the door to ease the memory burden of remembering their location or cataloging files alphabetically to minimize searching demands. (Clark 1997).

Shapiro notes that traditionalists can respond to the claims and assumptions of the Embodied Cognition research programme with a number of pointed observations and questions. Is my PDA or smart phone really a part of my cognition? Are the tasks that have been accomplished for instance by Rodney Brook’s reactive robots, whose associative/connectionist processing algorithms enable them to navigate through rooms without bumping into things and to pick up objects for which they have been trained, sufficiently representative of interesting cognitive capabilities?

Markman and Dietrich go farther in their critique of some embodiment assertions, and in particular the rejection of cognitive representations for *e.g.* concepts. Analyzing research in cognitive psychology they assert that all theorists in fact accept or assume that cognitive processing is mediated by some sort of internal information-carrying states which (however) may or may not constitute representations (Markman and Dietrich 2000). Traditionally representations have been considered to be enduring over time, discrete, compositional, potentially abstract and to operate via rule-governed processes. While some radical anti-representationalists reject the presence of any such mediating states, Markman and Dietrich suggest that such critiques are more appropriately aimed at specific representational formalisms or models of thought than at the reality of some sort of representations in cognitive processing.

Wilson (Wilson 2002) identifies six distinct claims that are often advanced under the embodiment research program:

1. Cognition is situated in the context of the real world and inherently involves perception and action.
2. Cognition is time-pressured due to the requirements of interacting with the external environment.

3. Because of limits to our attention capacity, working memory etc. we off-load cognitive processing and information onto the environment and retrieve it on an as-needed basis.
4. The information flow between mind and the environment is so dense that the environment is a part of the cognitive system.
5. The function of the mind is to guide action and mechanisms such as memory and perception must be understood in terms of the primary function of situation-appropriate action.
6. Offline cognition continues to be body-based.

Wilson critiques each of these claims. In response to the claim that cognition is inherently situated, *i.e.* occurs in the context of task-relevant inputs and outputs, she notes that offline cognition (such as imagining counterfactuals and other introspection) is particularly characteristic of human cognition and need not occur with respect to task processing at all. It is important, she asserts, not to overstate the degree to which cognition is situated with respect to concrete goals in order to more carefully examine the nature and degree of tight coupling between environment and cognitive activity.

If claim 1 above asserts that cognition is situation bound, claim 2 focuses specifically on the time pressure associated with sensorimotor activity such as complex movement and problems that present a 'representational bottleneck' in which there is no time to build up a full representation of the environment or situation before action is required. Wilson notes that in fact humans often do poorly when challenged to perform an unfamiliar motor activity or cognitive task; much of our activity is familiar and when presented with novel challenges we often, if given the opportunity¸ choose an offline stance to step back, observe, assess, plan and only then take action. Certain sensorimotor activities such as locomotion and grasping of objects do seem to require rapid feedback/adjustment cycles to accomplish the fine muscle control required but this is not necessarily typical of cognition as a whole.

Claim 3, that we offload some cognition into the environment is often taken to mean that we archive information and processing tasks using technologies such as writing things down, organizing information in filing systems and the use of external tools to retrieve and manipulate information. A more interesting version of this claim is that we externalize cognition through action in the environment, as when a child counts using its fingers or when we utilize spatial arrangements such as Venn diagrams to determine the logical relationships among multiple categories, in what Wilson calls symbolic off-loading.

In response to claim 4, that cognition is spread across the environment as well as within ourselves, Wilson notes that any given system of such interaction is facultative, *i.e.* temporary for a local purpose, whereas the cognitive architecture of the individual mind or brain is an obligate, *i.e.* persistent and tightly coupled, system. Wilson finds claim 4 deeply problematic.

Claim 5, that cognition is for action, has its roots in research into perception and draws on observations suggesting that the same neural pathways invoked when we grasp an object are partially active when that object is simply viewed. Wilson raises the question of how far we can extend this model. She notes that many visual experiences do not involve such motor activity at all, as when we view a sunset or a very distant object. Moreover we seem to be able to store rich information about the properties of objects which we can later use in future unspecified and often novel ways, as when we draw on our understanding of a piano in order to consider its use to barricade a door against an intruder. In addition, the amount of information we store for unfamiliar objects appears to be less than we store for familiar ones, suggesting that our internal representations are additively enhanced over time with additional experience.

Claim 6 is, according to Wilson, the claim with the most evidence in research literature, although it is often under-emphasized. Many mental activities may in fact consist of covert, subtle simulation of sensorimotor activities, as when we mentally count on our fingers or mentally rehearse visual, auditory and even kinesthetic experience patterns. Such patterns appear to be associative in nature, not abstract feature-attribute sets and typically involve multiple sensory modes simultaneously.

Wilson cites substantial experimental evidence to demonstrate the direct involvement of the body in off-line cognition. Typical results include much poorer performance on tasks requiring short-term memory when repetitive sensorimotor actions are simultaneously required. On the other hand, tasks which are familiar place much lower demands on short-term memory and other cognitive processes than less familiar tasks, suggesting that tasks which become automatized bypass the representational bottleneck. Wilson suggests that learning a task well creates an internal representation of the regularities associated with the task, which thereby pre-emptively reduces the representational bottleneck when the task is actually undertaken. Similarly, Wilson cites strong evidence for the role of sensorimotor simulation in higher level tasks such as reasoning and problem solving, where again simulated visual images often play a key role.

Wilson notes the role of mental simulation in cognitive grammar in particular, which makes use of perceptual principles such as attention focus and figure/ground separation in establishing grammars based on such schemas. As briefly noted in Chapter 4, Embodied Constructive Grammar explicitly extends Langacker's cognitive approach in order to ground language in mental simulation. The Berkeley embodied cognitive grammar project is described later in this chapter. Other approaches include that of Barsalou, whose grounded cognition theory is described below, and the prototype theory of concepts proposed by Lakoff and his collaborators which forms the basis for most cognitive grammar approaches.

Finally, Wilson notes that motor simulation has been proposed as the basis for representing and understanding the behavior of conspecifics, *i.e.* other things which we can imitate, and in particular of other human beings. As such it would play an important role in social interactions and our partial understanding of the internal states of others.

Before examining some current evidence for embodiment with regard to language and social interaction and for simulation as a key process in cognition it is useful to consider Barsalou's version of embodiment and situated cognition theory, which he calls grounded cognition. Barsalou defines simulation as "the reenactment of perceptual, motor, and introspective states acquired during experience with the world, body and mind" (Barsalou 2008):

> As an experience occurs (*e.g.,* easing into a chair) the brain captures states across the modalities and integrates them with a *multimodal* representation stored in memory (*e.g.* how a chair looks and feels, the action of sitting, introspections of comfort and relaxation). Later, when knowledge is needed to represent a category (*e.g.* chair), multimodal representations captured during experiences with its instances are reactivated to simulate how the brain represented perception, action and introspection associated with it. According to this account, *a diverse collection of simulation mechanisms, sharing a common representational system*, supports the spectrum of cognitive activities. (emphasis added)

Barsalou notes that to the modalities identified above we might add a focus on situated action, social interaction and the environment, typically linked to goal achievement. He stresses the role of introspection in concept development, positing a theory of Perceptual Symbol Systems (Barsalou 1999) which accepts the importance of symbolic operations for interpreting experience but which implements them through simulator constructs that dynamically produce the equivalents of type-token binding, inference, productivity, recursions and propositions familiar from formalist theory.

Barsalou's PSS theory assumes a single, multimodal representational system in the brain that supports a diverse range of cognitive processings such as high level perception, conceptual knowledge and short-term/working, implicit and long-term memory. He notes (Barsalou 2008) that in humans the simulation system central to PSS is closely integrated with the linguistic system but it appears, *contra* Paivio, that deep conceptual processing does not occur in the linguistic system but rather in the simulation system. This is perhaps the biggest difference between Barsalou's PSS model and the formalist approach in that meaning according to Barsalou might be described as being not so much stored in the representational system as indexed by it. Section 5.3 describes recent robotics work that is built around PSS.

Barsalou also asserts (Barsalou 2005) that non-humans have a simulation system similar to that of humans but lack a linguistic system to control it and hence do not develop the powerful symbolic capabilities that characterize human cognition. Here Hausser's assertion that language rests on cognition that is independent of language finds partial validation but is subject to the significant caveat that language and mental simulation of bodily experience are seen by Barsalou as having a dialectical, interdependent relationship.

Barsalou's theory finds detailed and significant support from evolutionary and comparative neuroscience evidence, gathered painstakingly by biological anthropologist Terrence Deacon (Deacon 1997), who identifies a series of evolutionary changes in hominid brain structure specifically and uniquely resulting in the co-development of unique symbolic reasoning and language capabilities in humans. Among the neurological effects of those changes is a major increase both in prefrontal cortex structures as a proportion of the brain and in pre-frontal connectivity. Evidence from patients with damage to these areas has been clearly linked to specific language dysfunctions or to failures in sensory perception, depending on the specific localized area of damage – a finding that is highly suggestive of a tight link between perception and language use. However, this does not mean that the prefrontal cortex is a localized repository for symbolic representations in the brain or the locus of symbolic processing, His argument is complex but has significant implications for the nature of meaning and of language use and indirectly for the challenges of human-level computational capability for such use. He writes (*ibid.,* 265):

> We should not, however, make the mistake of thinking that prefrontal cortex is the place in the brain where symbols are processed. It is not. Massive damage to the prefrontal cortex does not eliminate one's ability to understand word or sentence meaning. The symbolic associations that underlie the web of word meanings are probably much more dependent on the mnemonic support of sensory-based "images". This is supported by the high incidence of semantic disturbances after posterior cortex damage, especially damage to areas surrounding the posterior temporal cortex. It is also supported by our intuitions of mental imagery generated as we read stories, or of sound motor "images" as we search our memories for the right words to say or the correct names to match with familiar faces.
>
> It would be misleading, however, to say that these images are all there is to symbols, any more than the words on this page suffice in themselves to convey their meanings. They are merely neurological tokens. Like buoys indicating an otherwise invisible best course, they mark a specific associative path, by following which we reconstruct the implicit symbolic reference. The symbolic reference emerges from a pattern of virtual links between such tokens, which constitute a sort of parallel realm of associations to those that link these tokens to real sensorimotor experiences and possibilities. Thus it does not make sense to think of the symbols as located anywhere within the brain, because they are relationships between tokens, not the tokens themselves; and even though specific neural connections may underlie these relationships, the symbolic function is not even constituted by a specific association or by the virtual set of

> associations that are partially sampled in any one instance. Widely distributed neural systems must contribute in a coordinated fashion to create and interpret symbolic relationships.
>
> The critical role of the prefrontal cortex is primarily in the construction of the distributed mnemonic architecture that supports symbolic reference, not in the storage or retrieval of symbols. This is not just a process confined to language learning. The construction of novel symbolic relationships fills everyday cognition. A considerable amount of everyday problem solving involves symbolic analysis or efforts to figure out some obscure symbolic association. As soon as language processing leaves the realm of relatively habitual phrases and uses, it, too, often involves some level of novel symbol construction … One of the most important uses of language is for inferential processes, for taking one piece of information and extrapolating it to consequences not obvious from the information given. This is essentially using symbols to elicit or construct new symbols.

Here Deacon lays out a neurological basis for the integration of sensory perception through associations that form distributed patterns which in turn form the symbols that constitute semantic meaning for language. The process is dynamic, ego-centric and yet situated and involves both an encapsulated representation in the form of established symbolic association patterns and the availability of overall neural associations capable of integrating all internal and environmental information in order to validate or modify the encapsulated representation, as Giora hypothesizes are active in the influence of saliency during meaning assignment.

Deacon's model resembles in several ways that of psychological anthropologist Christine Hardy (Hardy, 1998) who posits that humans develop *semantic fields* – "coherent organizations of meaning-clusters and related processes in a dynamical, evolving network" – and *semantic constellations* – "self-organized, coherent clusters within a person's semantic field". As with Deacon's account of brain and language co-evolution, Hardy's theory is sufficiently new to many in the computational sciences to warrant an extended passage in order to see the scope of her approach to meaning and cognition. She builds on a basic neural network model but posits nested levels of organization (*ibid.,* 4-5):

> Semantic fields theory adds two features to this basic network architecture. The first is the concept of the SeCo (Semantic Constellation), a specialized network clustering and organizing related experiences. As we shall be seeing, a SeCo is often part of a larger SeCo, and may include sub-SeCos; the semantic fields model thus posits a n.etworks-within-networks architecture.
>
> The second added feature is the premise that SeCos link all possible types of elements, not only linguistic items or propositions, but any psychological, physiological, or brain process (such as sensation, affect, procedure, gesture, behavior, and their related neurological processes).
>
> The introduction of this kind of architecture has some important implications. For one thing, it recasts the body-mind relation as a transversal network integration of mental and brain processes: major SeCos may reach from lower neuronal processes to higher rational ones. Another implication is the recognition that knowledge in the human mind is never strictly abstract; it is necessarily tied to numerous sensory-affective processes …
>
> We are coming to recognize that, while humans certainly engage in abstract reasoning, this is not the way our mind operates most of the time. Computational rule-bound processing, as expressed in logical or mathematical reasoning, must be seen as a high-level process – more akin to something we painfully learn and force our minds into, ra`ther than a basic, natural working of the mind.

The latter part of (Hardy 1998) asserts collective Semantic Constellations that are formed and evolve within social contexts. Deacon also devotes the last chapters of (Deacon 1999) to the impact of social context on the co-evolution of neural and language development. The details are too extensive to trace here other than to note that he asserts evolutionary advantage to the use of ritual and eventually of language in establishing social relations that do not obtain among other primates but that favor survival of the young and lessen the likelihood of physical aggression within and between small familial groupings. Such ritual has not disappeared from our language use today:

> Though speech is capable of conveying many forms of information independent of any objective supports, in practice there are often extensive physical and social contextual supports that affect what is communicated. Language acquisition still relies on an extensive gamelike ritualization and regimentation of the symbolic acquisition context, although the child's computational supports enable this process to take place without explicit *reductio ad absurdum* grounding of all symbols and possible combinations in the system …
>
> The evolution of symbolic communication … created a mode of extrabiological inheritance with a particularly powerful and complex character, and with a sort of autonomous life of its own … (T)he neuroanatomical evidence of massively altered brain proportions and the anthropological and clinical evidence for universality of symbol learning across a wide spectrum of circumstances indicate that the human brain has been significantly overbuilt for learning symbolic associations … Brain-language co-evolution has significantly restructured cognition from the top-down, so to speak, when compared to other species. (*ibid.,* 407-9)

A rapidly increasing body of experimental evidence from functional magnetic resonance imaging and other neuroscience techniques collected since this passage was written has begun to identify the neural patterns that underlie specific aspects of language and cognitive processing. The picture is by no means complete or entirely well focused at this time. However, as these passages from (Deacon 1999) suggest, it is a reasonable working hypothesis that language is indeed the result of embodied and situated cognition, that there is nonetheless some form of symbolic representation associated with that cognition – but not necessarily a representation that fits neatly into well-circumscribed feature-attribute sets – and that while social context plays a key role in language acquisition and meaning construction, it does so as a modification to and not a replacement for the dynamical internal processes by which a language user develops and modifies an internal context of meaning associations.

A somewhat different approach to the neurological basis of natural language and cognition has recently been offered by Feldman (Feldman 2008), who presents and defends the Embodied Constructive Grammar approach and who, unlike Deacon and Hardy, is working with his colleagues on a computational instantiation. Formerly Professor of Computer Science and head of the cognitive science program at the University of California, Berkeley, Feldman starts by noting a series of experiments that demonstrate unexpected connections between motor activity and language processing. Bargh, for instance, found that subjects who were asked to sort jumbled sentences that contained words associated with being elderly, such as **old** and **wrinkle**, walked significantly more slowly as they left his lab than subjects whose sentences contained no such words (Bargh 1996), a finding that has since been replicated by other researchers. Feldman interprets this as evidence that unconsciously activating a concept may influence motor activity (in this case walking) associated with that concept. Moreover the influence can flow the other way. Cacioppo reports that when subjects were asked to evaluate the likeability of abstract ideographs while simultaneously pushing something away or pulling it toward themselves, those ideographs paired with pulling were rated more favorably than those paired with pushing away (Cacioppo *et al.* 1993). Briñol and Petty (2008) review similar evidence from a wide variety of experiments.

Feldman goes on to construct a detailed description of neural connectivity that supports the idea that categories or concepts are image schemas linked at their base to the neurological processing of physical experience. A key aspect of his theory is that the most basic level of image schemas are characterized by their affordances, *i.e.* by the possibilities they present us. For instance, we generate an image schema for **chair** that doesn't necessarily specify the material it is made of, height, or other such parameters but rather by the fact that we can sit on one. Thus while the traditional approach, epitomized by model theory, assumes that the structure of the world determines what concepts are needed and that languages arbitrarily assign different words for these concepts, embodied cognitive grammar asserts that our concepts depend on how we interact with the world – and that that interaction in turn is culturally shaped. Within a given cultural context semantic spaces are constructed that describe colors, emotions, spatial relations or more obviously culture-dependent categories such as types of songs or dances. Languages may differ in how a given semantic space is organized but every language must have some words for the conceptual primitives that all people share. Thus various languages around the world have from two to twelve basic color terms but in all cases speakers of a given language choose as most representative of a given term the place in the color spectrum that is in the middle of the wavelength span to which the term can refer (Feldman *op. cit.,* 102-104).

Feldman posits that super-ordinate categories are built up over categories that are directly linked to sensory perception through metaphor. In particular, primary metaphors such as the association between affection and warmth, describe internal experience in terms of sensory-motor experiences which, being common to humans, are publicly accessible and therefore can serve as the basis for communication. More abstract metaphors are partial mappings onto more basic metaphors.

As an example Feldman describes the Event Situation Metaphor, which is based on the semantic frame of **journey**, itself entailing domain concepts such as **location**, **movement**, **paths**,and **forced movement**. From this a complex metaphor is built in which causes are forces, states are locations, changes are movements, actions are self-propelled movements, purposes are destinations, means are paths to destinations, difficulties are impediments to motion and so on. This metaphor underlies, according to Feldman, language uses such as **He's between a rock and a hard place** or **I'm getting nowhere on this project**.

Since Feldman sees concepts as patterns of neural connections across sensory processing and other brain domains, he explains more abstract concepts as building on the relative stability of those concepts that are already established and reinforced through language use. In this way he ties mental simulation to specific neural configurations associated with terms in a given language and cultural context. He does not have much to say (and asserts we don't know much yet) about subjective experience of the sort that informs the egocentric experience in Kecskes' model, but he does propose a partial example of how metaphor configurations, probabilistic learning and simulation might enable children to learn the grammar of the language used around them. This example assumes that the grammar of the language consists, not of the more traditional elements of syntax, but rather of situation frames.

Feldman and several colleagues at the International Computer Science Institute at Berkeley, led by Chang, have begun implementing a computational model for language acquisition structured around embodied cognitive grammar (Chang *et al.* 2002; Bergen *et al.* 2005). The grammar is built from four formalisms: constructions, schemas, semantic spaces and maps. A schema is a conceptual structure consisting of roles (parameters) and constraints on the roles and relations between roles. A construction represents a basic linguistic unit which pairs language form and meaning, representing entities or relations. A map identifies correspondences across conceptual domains between schemas or spaces. A space is a conceptual domain containing entities and relations between them. These primitives are asserted to be composable, which addresses a concern about the suitability of cognitive grammar for computational natural language processing (*c.f.* Kemp and Partee 1995). However, (Chang *et al.* 2002) notes that the scalability of the structured connectionist implementations of these primitives, and especially of the action that is connected to them, is currently unclear, as is implementation for such linguistically relevant phenomena as neural priming (related to salience assignment) and the role of context.

The ability of the Berkeley cognitive grammar project to provide a basis for embodied human-level computational natural language remains to be seen. This is true not only because it is relatively new but also because implementations are currently focused on extremely simple grammars and basic language acquisition mechanisms. In particular, it is not obvious how higher level cognitive functions such as inference with regard to abstract concepts, planning and representation of subjective state will be accomplished in this strongly connectionist approach. Feldman does give one hint in the last chapter of (Feldman 2008) when he suggests that such a robot should use simulations rather than rules to capture how humans experience, think and behave. This raises significant issues regarding the ability of connectionist simulation approaches to scale and to provide the equivalent of a logical framework for reasoning using higher level concepts. At a more practical level it also raises questions about the amount of hardware needed to support such computation, assuming the algorithmic approach were established.

Traditional grammatical syntax is relatively concrete and limited in the number of categories applied to phrases and sentences and appears to be more directly linked to the forms of language as they are presented, in contrast to a cognitive grammar linked to presumed underlying conceptual constellations which populate situational frames. Moreover it is highly unclear how many situation frame types are needed for a cognitive grammar; although presumably the exact number is culture- and language-specific, it appears to be quite high. In general it seems at least possible that the tight linkage Feldman and his colleagues draw between neural activity, concept creation and language learning in terms of situation frames unhelpfully blurs distinctions that would otherwise give insight into cognition on the one hand and language on the other hand and complicate rather than elucidate the basic categories of thought and speech.

That said, the insights from cognitive linguistics taken more broadly underscore the nature of the challenges associated with computational natural language capability for artificial agents. Section 5.2 examines one domain area of interest for embodied artificial agents, namely spatial cognition and spatial reference in natural language.

5.2 Spatial Cognition and Language

Spatial cognition and language offers an insightful example of the issues for a natural language-competant embodied and situated artificial agent. Movement of a robot through a physical environment or movement by the avatar of an intelligent agent in a virtual world is a basic function which is likely to be commanded or reported by such an agent using natural language. What evidence is there regarding the nature of spatial cognition in humans and of spatial references in natural language? That is to say, with what sorts of spatial language would the desired agent need to cope?

Spatial representation and reasoning have been extensively studied from a number of perspectives, among them basic concept acquisition, embodied cognition, mathematical learning and the generation and interpretation of route descriptions for autonomous robots or in GPS-based navigation systems. A substantial amount of research has been conducted into language use for navigation routes (*c.f.* Taylor and Tversky 1992), with a primary focus on ego-centric vs. allo-centric frames of reference and the use of landmarks as key organizing elements in direction giving.

A very simple navigation route description problem hints at the complexity of spatial cognition and hence the challenge of spatial language understanding and generation in artificial agents (Burk *et al.* 2007) . Experimental subjects were presented by Haas with a simple layout of a building interior, marked with icons indicating "the robot is here" and "destination". They were asked to generate written directions for the simulated robot to follow in order to reach the destination. No interaction was permitted: the subject was required to produce a complete set of navigation directions before the simulated robot began to move. The directions are typically organized as a set of steps which can be characterized in terms of:

- The type of destination for this step (doorway, side hall, end of hallway)
- Direction (left, right, forward)
- An ordinal characteristic for the destination (first, second, third, last)
- At-end (true if the destination for this step is the end of the hallway ahead of the agent as it begins the step)
- The action required for this step (advance, advance and turn, or do nothing until the next step)

Surprisingly, despite the simplicity of the layout subjects used a variety of conflicting descriptions with regard to ordinals, as illustrated in Figures 8 and 9 below:

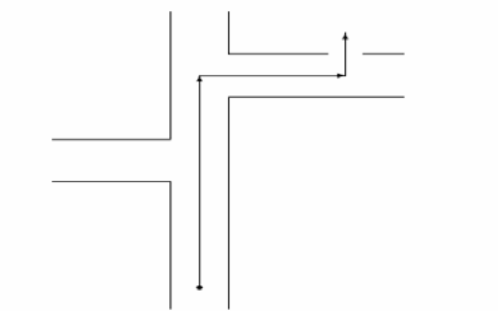

Figure 8. Turn right at the first (second) hallway …

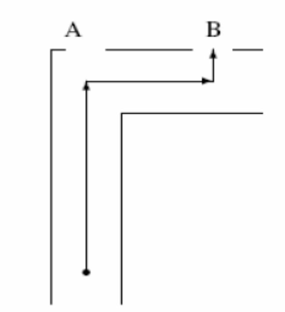

Figure 9. Go into the first (second) room …

The varied responses suggest that different subjects organized entities in space differently, at least with regard to ordinality. In Figure 8 some subjects anticipated the direction of the upcoming turn and provided a count of hallways on that side only, while other subjects took a more map-like view of the overall space through which the robot would transit. Similarly, in Figure 9 some subjects associated the room on the left with the current hallway while others associated it with the transverse hallway into which the robot would be sent.

What conclusions can be drawn regarding spatial cognition from this data? Not many, perhaps. It might be that ordinality is a convention established in one way or another, a somewhat abstract notion. Spatial orientation overall, however, is a very basic concept closely linked to physical experience. No matter what our culture or language, locomotion is a basic function of the bipedal human body guided by internal proprioceptic balance mechanisms and a forward-directed visual system sensitive to movement, shape, color, surface texture and spatial depth. It is perhaps surprising, therefore, that anthropological and cognitive linguists have documented substantial differences in spatial reference and reasoning across cultures (*c.f.* Shah and Miyake 2005; Carlson and Lee 2005; van der Zee and Slack 2003).

An ambitious research program conducted by the Max Plank Institute is illuminating (Levinson and Wilkins 2006). The Institute chose a dozen languages on several continents, mostly from unrelated language families, many of whose native speakers live in similar physical environments. In collecting information regarding spatial reference and reasoning in each culture, researchers identified which frames of reference were used, by whom, and in what context, as well as the spatial relationships distinguished in the languages. Frames of reference are coordinate systems used to designate the angle or direction in which a figure or object lies with respect a ground reference point when figure and ground are separated from one another. The three basic frames of reference (FoRs) are absolute, relative and intrinsic. Absolute FoRs are based on community-established bearing sets, such as the North-South-East-West bearings used to good effect by European sailors equipped with compass and sextant in the Age of Exploration. However, these compass directions are not the only possible absolute FoR. Many tribes in Alaska, for instance, orient on a major river in their geographic region, resulting in absolute bearings of Upstream, Downstream and Across. Relative FoRs map the bodily orientation of the viewer onto the scene, resulting in coordinates such as Right, Left, Ahead and Behind. Intrinsic FoRs designate coordinates projected from facets of the ground object: at the Back of the house, for instance.

Levinson asserts that intrinsic frames of reference are the most basic, evidenced by children two years old who when presented with objects as varied as toy trucks, buildings and chairs can line them up with their 'fronts' all facing the same way. Relative frames of reference appear to build on intrinsic FoRs by metaphor and can be quite complex in the varied extensions to the basic system. These frames start with the speaker's own front/back/right/left but then involve a secondary coordinate system mapped from the speaker onto the ground object. For instance, when I say **over there in front of the tree** I mean 'between me and the tree'; here the tree is given a human set of coordinates which are rotated so that my 'front' and the tree's 'front' are facing one another. At the same time, however, if I say that the ball is **to the right of the tree** I mean 'in the right side of my visual pattern as I look at the tree', not 'to the tree's own right side as it is conceived as facing me'. Thus the tree is rotated for one purpose but not for another.

At least it is in typical English usage. In Hausa the speaker's orientation is simply mapped onto the tree as well: it is viewed as facing away from the speaker and has the same right and left sides as well. Relative FoRs are not common in some languages while others such as Japanese and Dutch depend heavily on them (*ibid.,* 544). Moreover, while some languages permit all three FoRs, they are often used in different circumstances or by different groups of speakers within a given language. For instance, among the Yukatek Maya who participated in the study, men were most likely to use absolute FoRs based on cardinal directions, women almost never used them and relative and intrinsic FoRs were tied to social situations and tasks.

In addition to frames of reference, languages offer varying ways to describe the topology or degree of closeness or contact between items (including figure and ground) and the motion of objects across the reference ground. For instance, Yele Dnye (spoken on Rossel island in Papua New Guinea) is an isolate with no clearly established related languages. Its speakers use a very rich set of postpositional phrases to distinguish topological relations covered by the English notions 'in' or 'on'. Among these relations is the notion of attachment which is described by one of three postpostionals representing attachment 'on a vertical surface', 'by impalement on a hook, spike etc.' and 'stuck on, attached strongly'. Yele Dnye also distinguishes six different postpositionals for the notion designated in English by **on** or **above**. On the other hand Kilivia, spoken in the Trioband Islands, uses expression elements that distinguish movement towards or away from the speaker, without reference to the source or goal of the movement, but combines these elements with a wide variety of socially-oriented other elements to yield a set of very rich movement description possibilities (*ibid.*).

Levinson notes the strong cultural influence on spatial cognition as a result of day to day task functions whose characteristics are reinforced by the distinctions made in the language. This raises the question of whether language is here demonstrated to shape cognition. Levinson argues that it does, and that therefore so does culture in a dominant manner. However, differences between the genders in preferred forms of spatial reference and fluency on various spatial tasks have been demonstrated in Western cultures (Dabbs *et al.* 1998) as well as in the traditional cultures studied by the Planck Institute (*c.f.* Jones and Healy 2006) and there is evidence that testosterone levels in men are correlated with visiospatial cognitive function (Beauchet 2006), strongly suggesting that spatial cognition has a strong biological component.

What are the implications of diversity in spatial cognition and reference for human level computational natural language capability? At a minimum the evidence gathered by the Planck Institute and other studies (*c.f.* Munnich *et al.* 2001) shows that spatial cognition is not uniform across cultures and languages. This calls into question Hausser's assumption of a shared language-neutral semantic/conceptual space to which surface elements from a variety of languages can be transparently linked. That assumption is also challenged within a given language by neuroscience evidence for the involvement of different brain areas when subjects use different frames of reference (Committeri *et al.* 2004) and by evidence for gender differences in spatial cognition and preferential language use (as when women prefer to use landmarks rather than cardinal points when giving or receiving navigation directions), although both sexes are able to use various frames of reference with some degree of fluency which depends on the individual's innate ability and experience.

At the same time, the Institute evidence raises significant questions regarding the likely fidelity of machine translation across disparate languages and regarding the construction of language- and culture-independent formal ontologies for semantic interpretation of texts, as is illustrated in the Institute researchers' use of extensive semantic notation to capture the meaning elements inherent in various words and constructions across the languages studied. It would appear that language is indeed heavily situated in social contexts, while also being strongly embodied. Human-level computational natural language capability must in one way or another take this into account.

5.3 Implications of Embodiment in Situated Artificial Agents

The goal of an embodied artificial agent with human-level natural language capability introduces several significant challenges beyond language processing at the syntactic and semantic level. Natural language capability embodied in an autonomous mobile robot presents not only linguistic challenges but also the requirement that language be tied to understanding of and effective interaction with the physical world. Such an agent must have a number of complex capabilities, including:

- **Sensory perception** in an environment where input data is noisy, extremely diverse and unsorted or characterized.
- **Object recognition** – pattern extraction from sensory data and a subsequent analysis to identify those patterns which correspond to physical entities of interest or, in the case of spoken language, sound patterns that correspond to specific words. More broadly, to identify the type of a specific object whose exact pattern has not previously been

encountered and characterized according to some pre-determined or evolving typology.

- **Inference/reasoning** – the ability to maintain, update and draw conclusions from a base of knowledge about the physical context as well as the agent's internal state.
- **Planning** – identifying goals and actions to take to achieve those goals, monitoring success or failure of actions and adjusting goals and actions accordingly. This may involve coordination with other external agents, including humans.
- **Action in the external world** – control of hardware to effect actions that include directional movement of sensory elements, overall agent movement through an uneven or difficult environment, grasping and manipulation of objects and production of communicative sounds (speech) as well as control of other specialized communications capabilities such as digital communications equipment.

Natural language capability (as opposed to hard-coded recognition of a small set of verbal commands) implies that the agent is cognitive, *i.e.* that it can develop and use knowledge, beliefs, preference, goals and informational attitudes (collection, analysis, hypothesization) or some simulated equivalents to these. In accordance with Kecskes' Dynamical Model of Meaning such an agent must also be capable of dynamically acquiring and modifying concept structures and the semantic and pragmatic meanings associated with words and phrases in the language.

Thus robots (and virtually embodied agents) present a particularly direct instance of the more general symbol grounding problem in artificial intelligence. Natural language aside, such grounding has been implemented in many robots by linking probabilistic sensor data processing, often connectionist in design, with higher level symbolic reasoning through the use of feature extraction algorithms (*e.g.* from camera inputs) which provide posterior data for Bayesian reasoning (*c.f.* Knill and Richards 1996) about object identity, robot location etc.

One mature robot architecture is provided by the National Institute of Standards and Technology, which has sponsored the development of the Realtime Control Systems (RCS) reference architecture. Instantiations of RCS are in use in manufacturing robots and by NASA, among others. Based heavily on associative learning methods and fuzzy logic-based control theory, RCS has been extended to form a multi-resolutional intelligent agent architecture (Meystel and Albus 2002) but has not been linked with natural language capability, in part perhaps because the dynamical control systems approach does not yield semantic representations that facilitate such linkage. RCS does include a markup language that can support a message-based interface to client software or other systems. The Defense Advanced Research Projects Agency (DARPA) has developed the 4D-RCS version of the architecture to support autonomous ground navigation across complex terrains using machine learning techniques (Albus *et al.* 2007).

Natural language (beyond discrete verbalized commands) has been investigated as a robot command interface for some time. A current effort that combines traditional robot control mechanisms with limited domain-specific natural language understanding and interpretation is being jointly pursued by several universities in the Situation

Understanding Bot – Language Enabled (SUBTLE) project led by the GRASP lab at the University of Pennsylvania. Sponsored by the U.S. Army Research Office, the SUBTLE project has as its aim creation of a language-capable robot for urban search and rescue. This multi-faceted research effort primarily uses machine learning techniques for syntactic and semantic parsing, a formal ontology and parameterized action representation for mapping linguistic input to meaning to robot action plans. The pragmatics approach is oriented to Gricean implicature, *i.e.* to discerning and responding to intended meaning that may not be explicitly represented in the linguistic utterance. In one current research thread, a phrase structure grammar is being joined to statistical disambiguation for sentence parsing.[5]

A more explicitly cognitive and embodied approach to natural language capability in a physical robot has been pursued for some time by Deb Roy at the MIT Media Lab. Roy proposes the use of semiotic schemas as structured, mode-independent representations that are grounded in the agent's physical environment through a causal-predictive cycle of action and perception (Roy 2005). His Grounded Situation Model provides a unified set of representational primitives that span analog beliefs, categorial beliefs and a variety of projections between levels of representation. Natural language capability is built upon these representations and includes speech act analysis as well as syntactic/semantic analysis. The Grounded Situation Model objectives include question answering about the current physical context, temporal indexing of events and the ability to construct and discourse about imaginary (*i.e.* non-local, non-present-time) events.

The social situatedness of cognition and natural language is a primary focus of Breazeal's recent work in robotics, also at the MIT Media Lab. Breazeal's architecture is based on Barsalou's Perceptual Symbol System approach, extended to incorporate findings from developmental psychology (Breazeal *et al.* 2007), and has resulted in a robot that interacts collaboratively with humans and that uses internal simulation in planning game moves and adjusting beliefs.

In addition to the Barsalou/Breazeal approach, Sun has proposed his CLARION cognitive model for embodied and situated agents (Sun 2002). Adopting an explicitly dual-representational approach, CLARION replaces Lakoff's bodily schemata as the basis for building concepts with Heidiggerean comportments, *i.e.* direct, unmediated routine interactions with the external world. CLARION uses the temporal difference method of reinforcement learning to learn rules whose conditions are then represented at a higher level as symbolic concepts. These concepts are then linked associatively to the corresponding low-level conditions through a neural net structure. CLARION has been used to computationally instantiate a variety of lower and high-level cognitive tasks but has not been linked to robotic hardware at this point.

A somewhat more ambitious aim underlies the work of Cassimatis, namely the evolution of biologically-inspired human-level intelligence (Cassimatis 2006). With the goal of Artificial General Intelligence (AGI), Cassimatis and others are pursuing computational models of cognition that explicitly link heterogeneous reasoning techniques at various levels of conceptual representation and abstraction. Cassimatis posits a core set of common functions that underlie many AI algorithms and asserts that these functions can be implemented using multiple computation methods. The function set might include forward inference, subgoaling, simulating alternate worlds and object identity matching. He specifically rejects formalisms such as frames as being brittle to change under new knowledge and experiences. Instead, his Polyscheme architecture implements a variety of algorithms as sequences of attention fixation across external and internal data. Polyscheme has been implemented in physical robots (Cassimatis *et al.* 2004) and provides a common substrate for both language and world knowledge.

As this brief overview suggests, recent cognitive robotics approaches range from traditional control theory, to hierarchical architectures based on both symbolic and probabilistic sub-agents that include natural language parsing and the use of formal ontologies for pragmatics interpretation, to approaches explicitly based on embodied and situated cognition theory and instantiated in physical robots which already have some degree of language and social interaction capability. It remains to be seen how far the latter agents will develop and evolve over time in response to the learning challenges they are provided. Developing physical robots is expensive and complex, which is one reason some AGI researchers are turning to virtual worlds such as Second Life to provide their cognitive, intended-to-be-embodied and situated agents with social and other simulated interactions, as in the open source Open Cognition project[6].

6. Summary and Conclusions

Hausser's approach to computational linguistics is unusual for the combination of elements he puts forward and for his ambitious scope. Initially based on model theory, his work continues to be based on theory in the form of his SLIM principles. In left-associative grammars and Database Semantics he offers transparent, computationally tractable formalisms for language representation and processing.

Hausser is particularly noteworthy for his early focus on procedural, pragmatic assignment of meaning to language, motivated in part by speech act theory, and for his insistence that computational language processing should aim at human-level communication capability in an embodied and situated cognitive agent. In several papers he attempts to illustrate how peripheral, sensory inputs to such an agent might be pattern matched to create formal representations of episodic and absolute world knowledge in propositional form.

However, Hausser's overall approach is ultimately unconvincing at this point in its development due in large part to the sparsity of detail he has provided regarding pragmatic reasoning. This sparsity may not in fact be an accident but rather reflects the difficulty his approach faces when dealing with the full range of language beyond such referents as mathematical shapes and basic colors. Hausser's assumption that world context knowledge is inherently representable by predicates asserted against concepts defined as feature-attribute sets faces significant challenge from cognitive theory and cognitive linguistics. The nature of the relationship between language and culture is open to debate, but diversity across cultures and genders for such basic areas of cognition as spatial reasoning and language reference is now well established and contradicts a key assumption of Hausser, namely that we all share a common set of literal, well-defined concepts that are denoted transparently and interchangeably by any natural language.

Nonetheless, although his approach is still in many ways shaped by formalism, Hausser was prescient in his early and continued insistence on the bodily, non-verbal basis of language meaning. Moreover, it's not obvious that his continued use of propositional representation is wrong so much as it is incomplete. The anti-representationalists have yet to credibly demonstrate that their approach can achieve the levels of cognitive processing that humans demonstrate and that supports human-level natural language capability. Barselou's perceptual symbol system approach, which yields correlates to the basic categories and functions of formal logic, is more likely than the anti-representationalists to result in fruitful progress because while it acknowledges the embodied and situated context within which human cognition and language functions, it does not discard entirely the encapsulated representations and their manipulation that have formed the basis for Western thought since the classical Greeks. Notably, Barsalou's approach has already been implemented in a robot which functions interactively with humans in a limited domain.

A major theme in this dissertation is the increasing evidence that natural language and human cognition influence one another deeply and are based on multiple levels of neural structures and processes in the brain, some of which are multi-modal and sensory based and others of which are potentially amodal, encapsulated and similar to formal (symbolic) representations of knowledge. Support for the existence of multiple levels of representation in human cognition ranges from Giora's behavioral evidence for salient, encapsulated word and expression meanings validated by associative integration of neural contents, on the one hand, to detailed neuroscience discoveries linking specific areas of the brain and brain activity to separate frames of reference in spatial reasoning on the other hand. Deacon and others have argued for specific mechanisms leading to the co-evolution of human cognition and language, with an emphasis on the social context that provided both the necessity and the means for the development of ritual, symbolic thought and symbolic reference in language.

These insights are underscored by Kecskes' evidence for cognitive changes as a result of learning second and later languages and by his assertion of dualing contexts for meaning assignment. The assertion of a dynamical internal context is particularly interesting in light of evidence that mental simulation provides a key bridge between sensory experience, introspection and internal, offline cognition such as imagination.

This examination of Hausser's theory for computational linguistics and of relevant cognitive and neuroscience suggests several research questions going forward. One key question is whether it is in fact possible to bridge between symbolic and associative approaches to machine learning and reasoning, as Cassimatic and Barsalou assert. If so, which approach is better? Is Cassimatis right that there is a set of functions that underlie artificial intelligence and that can be instantiated using any of a variety of reasoning techniques? Is Barsalou correct in asserting the need for correlates to replace the specific categories and operations of standard formal logic? What is the best role for Bayesian reasoning and for *e.g.* multi-valued and fuzzy logics in embodied cognitive agents beyond their current use in object recognition and lower levels of sensory processing and hardware control algorithms, respectively?

A second question concerns functional grammar and related approaches to semantic functional ontologies and other category typologies. Do the categories of functional grammar better describe the basic nature of language than syntactically focused grammars? If so, why haven't they been proposed earlier in the centuries-long history of discussion regarding the nature of language and thought? The research of Levinson and his colleagues touched but did not focus on functional aspects of the spatial cognition and language reference they studied. However the initial data they did collect raise the question: is it the functions of language elements or of cognition that are of real value in the functional grammar approach? To what degree can those be distinguished from one another?

Finally, what are the implications of the embodied and situated nature of language for statistical natural language processing and the use of formal ontologies for semantic tagging? Given the increasing evidence for cultural and gender diversity in so basic a category as spatial cognition and reference, does statistically-based machine translation miss important levels of meaning? Do formal ontologies and semantic web technologies impose cognitive and language limitations in their applications and on users? Given an asserted tight coupling between language and cognition of some sort, will these technologies result in a bland common culture mediated by the Internet and the loss of richer forms of cognition and semantic content in various cultures?

Hausser was prescient in his focus on pragmatic, procedural meaning assignment based on non-verbal bodily based cognition. Based on that focus he called for the instantiation of language theory in computational, embodied artificial agents situated in communicative social contexts. Today that call is being answered by new approaches to cognitive and semantic processing in robots and in cognitive agents able to learn through interactions in virtual worlds. Along the way both the goal of theory-based computational natural language capability and the field of artificial intelligence are being fruitfully reinvigorated.

Notes

[1] Kartunnen, Lauri. 2007. Word Play. *Computational Linguistics* 33:447.

[2] *CLSI Monthly* **1**. Center for the Study of Language and Information, Stanford University.

[3] Polguere, Alain. 1986. In *Computational Linguistics* **12**.

[4] Partial Java implementations for fragments of English and other languages were available for download from http://www.linguistik.uni-erlangen.de/clue/en/research.html as of 02/21/2010.

[5] Current SUBTLE project information including a list of researchers and publications was available for review at http://www.seas.upenn.edu/~muri/index.html as of 03/27/10.

[6] See http://www.opencog.org/wiki/The_Open_Cognition_Project .

Hausser Publications

Hausser, Roland 1971. A Theory of Systematic Deviants. Bach, Emmon & Stanley Peters *ed. Report to the National Science Foundation* (GS-2468). The University of Texas at Austin.

——— 1973. Presuppositions and Quantifiers. *Papers from the Ninth CLS-Meeting.* Chicago:Chicago Linguistics Society.

——— 1974a. On Incorporating ANY into a Montague Grammar. *Kritikon Litterarum*.

——— 1974b. Syntax and Semantics of Plural. *Papers from the Tenth CLS-Meeting.* Chicago:Chicago Linguistics Society.

——— 1974c. *Quantification in an Extended Montague Grammar.* Ph.D. Dissertation, University of Texas at Austin .

——— 1976a. Presuppositions in Montague Grammar. *Theoretical Linguistics* **3**.

——— 1976b. Scope Ambiguity and Scope Restriction. Bartsch, Renate *et al. ed. Amsterdam Papers in Formal Grammar,* Vol. 1. Amsterdam: Universität.

——— 1977. Das Projektionsproblem bei Präsuppositionen Drachman, Gaberell ed. *Akten der 3. Salzburger Frühlingstagung für Linguistik.* Tübingen: Verlag Gunter Narr.

——— 1978. Surface Compositionality and the Semantics of Mood. Gruenendijk, Jeroen & Martin Stokhof ed. *Amsterdam Papers in Formal Grammar* (Vol. II). Universität Amsterdam. In: Searle, John R., Ferenc Kiefer & Manfred Bierwisch ed. *Speech Act Theory and Pragmatics.* Dordrecht, The Netherlands: Reidel.

——— 1979. How do Pronouns Denote? Heny, Frank & Helmut Schnelle ed. *Syntax and Semantics* (Vol. 10). New York: Academic Press.

. ——— 1980. The Place of Pragmatics in Model Theory. Groenendijk *et al.* ed. *Formal Methods in the Study of Language* (Mathematical Center Tracts 135). Amsterdam: Universität.

——— 1981a. A New Treatment of Context in Model Theory . *Sull' Anaphora, atti del Seminario Accademia della Crusca.* Florenz.

——— 1981b. The Speaker Simulation Device, and the Foundations of Semantics and Pragmatics. *Language and Linguistics* **7**, Seoul, Korea.

——— 1982a. A Constructive Approach to Intensional Contexts. *Language Research*, **18**. Seoul, Korea.

——— 1982b. A Surface Compositional Categorial Syntax. *Journal of Linguistics of Korea*. **7**, Seoul, Korea.

——— 1983a. On Vagueness. *Journal of Semantics* **II**(3/4).

——— 1983b. The Syntax and Semantics of English Mood. Kiefer, Ferenc *(ed.)* *Questions and Answers.* Dordrecht, The Netherlands: Reidel

——— 1983c. Vagueness and Truth. *Conceptus. Zeitschrift für Philosophie* **XVII**(40/41.

——— 1984. *Surface Compositional Grammar.* München: Wilhelm Fink.

——— 1985. Left-associative Grammar and the Parser NEWCAT. (Report IN-CSLI-85-5). Stanford CA: Stanford University, Center for the Study of Language and Information.

——— 1985. Zur Struktur des Lexikons in der generativen Grammatik. Posner, Roland & Thomas Ballmer ed. *Nach-Chomskysche Linguistik.* Berlin: Walter de Gruyter.

——— 1986. *NEWCAT: Natural Language Parsing Using Left-associative Grammar.* (Lecture Notes in Computer Science 231. Berlin: Springer.

——— 1987a. Left-associative Grammar, Theory and Implementation. *Report CMU-CMT-87-104. Center for Machine Translation.* Pittsburgh PA: Carnegie-Mellon University.

——— 1987b. Modelltheorie, künstliche Intelligenz und die Theorie der Wahrheit. Puntel, Lorenz ed. *Theorien der Wahrheit.* Darmstadt: Wissenschaftliche Buchgesellschaft.

——— 1988a. Computational Complexity of Left-Associative Grammar. *Proceedings of the Second International Conference on Theoretical and Methodological Issues in Machine Translation, June 12-14.* Pittsburgh PA: Carnegie-Mellon University.

——— 1988b. Left-Associative Grammar: An Informal Outline. *Computers and Translation* **3**(1), 23-67, Dordrecht, The Netherlands: Kluwer.

——— 1988c. Left-Associative Grammar: The Algebraic Definitions. *Computers and Translation* **3**(2), Dordrecht, The Netherlands: Kluwer.

——— 1989a. *Computation of Language. An Essay on Syntax, Semantics and Pragmatics in Natural Man-Machine Communication.* Berlin: Springer.

——— 1989b. Computational Complexity and Decidability in Left-Associative Grammar. *Proceedings of the International Workshop on Parsing Technologies, August 28-31.* Pittsburgh PA: Carnegie-Mellon University.

——— 1992. Complexity in Left-Associative Grammar. *Theoretical Computer Science* 106: 283-308, Dortrecht: Elsevier.

——— 1996. A Database Interpretation of Natural Language. *Korean Journal of Linguistics* 21: 29-55.

——— 1998. Häufigkeitsverteilung deutscher Morpheme. *LDV-Forum* 15:6-26.

——— 1999a. *Foundations of Computational Linguistics. Man-Machine Communication in Natural Language.* Berlin:Springer.

——— 1999b. Database Semantics and Temporal Inference. *Invited Lecture at NLPRS'99 conference, Beijing, China.*

——— 1999c. Database Semantics: Natural Language Interpretation, Inferencing, and Production. *Proceedings of the Second International Conference on Cognitive Science, Waseda University, Tokyo, Japan.*

——— 1999d. The Handling of Metaphor in Internal Matching Pragmatics. *Paper read at AISB '99 convention, Edinburgh.*

——— 1999e. Three Basic Approaches to Automatic Word Form Recognition. *Paper presented at VEXTAL conference, Venice, Italy.*

——— 2000. *Grundlagen der Computerlinguistik, Mensch-Maschine Kommunikation in natürlicher Sprache.* Berlin, New York, Springer.

——— 2001a. *Foundations of Computational Linguistics, Human-Computer Communication in Natural Language* 2nd Edition. Berlin: Springer.

——— 2001b. Spatio-Temporal Indexing in Database Semantics, in A. Gelbukh ed. *Computational Linguistics and Text Processing*, Lecture Notes in Computer Science Vol. 2004. Berlin: Springer.

——— 2001c. The Four Basic Ontologies of Semantic Interpretation. in H. Kangassalo *et al.* eds. *Information Modeling and Knowledge Bases XII.* Amsterdam: IOS Press.

——— 2001d. *Human-Computer Communication in Natural Language*, lecture presented at the UNESCO conference *Language in Cyberspace* held September 26/27 in Seoul, Korea

——— 2001e. Database Semantics for Natural Language. *Artificial Intelligence* **130**. Dordrech: Elsevier.

——— 2002a. Overcoming Language Barriers by Means of Computers, *Journal of Universal Languages*, Sejong University, Korea

——— 2002b. Replicating Quantified Noun Phrases in Database Semantics, paper presented at *7th International Workshop on Applications of Natural Language to Information Systems* (NLDB 2002), Royal Institute of Technology, Stockholm, Sweden, June 27-28, 2002.

——— 2002c. A Hypothesis on the Origin of the Sign Types, in A. Gelbukh ed. *Computational Linguistics and Intelligent Text Processing*, Lecture Notes in Computer Science, p. 16-31. Berlin: Springer-Verlag.

——— 2002d. Autonomous Control Structure for Artificial Cognitive Agents, in H. Kangassalo *et al.* eds. *Information Modeling and Knowledge Bases XIII*. Amsterdam: IOS Press.

——— 2003a. Comparing Music and Language: Data Structures, Algorithms, and Input-Output Conditions, in G. Wille, B. Schröder und H.-C. Schmitz ed., *Computational Linguistics, Achievements and Perspectives*. Sankt Augustin: Verlag.

——— 2003b. Reconstructing Propositional Calculus in Database Semantics, in H. Kangassalo et al. ed., *Information Modeling and Knowledge Bases XIV*. Amsterdam: IOS Press.

——— 2004a. A Theory of Signs for Database Semantics, paper presented at the EJC 13, Kyushu Institute of Technology, Kitakyushu, Japan, 2003. Published in H. Kangassalo et al. ed., *Information Modeling and Knowledge Bases XV*. Amsterdam: IOS Press.

——— 2004b. Applying Database Semantics to the WWW, paper presented at the *WISE 2004 International Workshops*, Brisbane, Australia 2004. Published in C. Bussler *et al.* ed., *Lecture Notes in Computer Science 3307*. Berlin: Springer-Verlag.

——— 2005a. Comparing Left-Associative Grammar with PSG, paper presented at the *9th Asian Logic Conference,* 16-19 August, 2005, Novosibirsk, Russia.

——— 2005b. Memory-based Pattern Completion in Database Semantics, *Language and Information*, 9:69--92, Korean Society for Language and Information.

——— 2005c. Turn Taking in Database Semantics. In H. Kangassalo *et al*. ed. *Information Modeling and Knowledge Bases XVI*. Amsterdam: IOS Press.

——— 2005d. What if Chomsky were right?, Comments on the paper Multiple Solutions to the Logical Problem of Language Acquisition by Brian MacWhinney. *Journal of Child Language* 31:919-922. Cambridge UK: Cambridge University Press.

——— 2006a. *A Computational Model of Natural Language Communication; Interpretation, Inference, and Production in Database Semantics.* Berlin:Springer.

——— 2006b. Reconstructing Geons in Database Semantics, in *Frontiers in Artificial Intelligence and Applications;*. 136:280-288 Heidelberg: Springer.

——— 2007a. Comparing the Use of Feature Structures in Nativism and in Database Semantics. In Hannu Jaakkola, Yasushi Kiyoki and Takehiro Tokuda eds. *Information Modelling and Knowledge Bases XIX*. Amsterdam: IOS Press.

——— 2007b. Handling Valency and Coordination in Database Semantics. In Thomas Herbst and Katrin Götz-Votteler ed. *Valency - Theoretical, Descriptive and Cognitive Issues, Trends in Linguistics - Studies and Monographs [TiLSM] 187.* Berlin: Mouton de Gruyter.

——— 2008. Center fragments for upscaling and verification in Database Semantics, *presented at the EJC XX, Tsukuba, Japan, June 2--6, 2008.*

——— 2009a. Erasmus Lectures, March 30th-April 3rd 2009, Univ. of Jyväskylä, Jyväskylä, Finland

——— 2009b. From Word Form Surfaces to Communication. In Takehiro Tokuda, Yasushi Kiyoki, Hannu Jaakkola, and Tatjana Welzer Družovec ed. *Information Modelling and Knowledge Bases XXI.* Amsterdam: IOS Press.

——— 2009c. Modeling Natural Language Communication in Database Semantics. In M. Kirchberg and S. Link ed. *Proceedings of the APCCM 2009, Australian Computer Science Inc., CIPRIT, Vol. 96*

Hausser, Roland & Dietmar Zaefferer 1980. Questions and Answers in a Context-Dependent Montague Grammar. Guenthner, Franz & S. J. Schmidt ed. *Formal Semantics and Pragmatics of Natural Languages.* Dordrecht, The Netherlands: Reidel

Other References

Anderson, John and Gordon Bower. 1974. A propositional theory of recognition memory. *Memory & Cognition* 2:406-412.

Albus, James, Roger Bostelman, Tommy Chang, Tsai Hong, Will Shackleford and Michael Shneier. 2007. Learning in a hierarchical control system: 4D/RCS in the DARPA LAGR program. *Journal of Field Robotics,* 23:975-1003.

Lawrence Barsalou. 1999. Perceptual symbol systems. *Behavioral Brain Science* 28:577-660.

——— 2005. Continuity of the conceptual system across species. *Trends in Cognitive Science* 9:309-311.

——— 2008. Grounded Cognition. *Annual Review of Psychology* 59:617-45.

Barwise, Jon, John Mark Gawron, Gordon Plotkin, and Syun Tutiya, eds. 1991. *Situation Theory and Its Applications Vol. 2*. Stanford University: Center for the Study of Language and Information.

Bates, Elizabeth. 1999. On the nature and nurture of language. In E. Bizzi *et al.* ed. *Frontiers of biology: The brain of homo sapiens*. Rome: Instituto della Enciclopedia Intalian fondata de Giovanni Trecanni S.p.A.

Bates, Elizabeth and Brian MacWhinney. 1989. Functionalism and the competition model. In Brian MacWhinney & Elizabeth Bates ed. *The crosslinguistic study of language processing.* Cambridge UK: Cambridge University Press.

Beauchet, Olivier. 2006. Testosterone and cognitive function: current clinical evidence of a relationship. *European Journal of Endocrinology* 155:773-781.

Biederman, Irving. 1987. Recognition-by-components: A theory of human image understanding. *Psychological Review* 94:2. American Psychological Association.

Blackburn, Patrick and Johan Bos. 2005. *Representation and Inference for Natural Language, A First Course in Computational Semantics*. Stanford CA: CSLI Publications.

Breazeal, Cynthia, Lawrence Barsalou, and Linda Smith. 2007. *An Integrated Architecture for Grounded Intelligence in Its Development, Experimental, Environmental, and Social Context*. DTIC ADA473768.

Bresnan, Joan 2001. *Lexical Functional Syntax*. Oxford UK: Blackwell.

Briñol, Pablo and Richard Petty. 2008. Embodied persuasion: Fundamental processes by which bodily responses can impact attitudes. In *Embodiment Grounding: Social, cognitive, affective, and neuroscientific approaches* ed G. R. Semin & E. R. Smith, 184-207. Cambridge UK: Cambridge University Press.

Carlson, Laura and Emile van der Zee. 2005. *Functional Features in Language and Space.* Oxford UK: Oxford University Press.

Carnap, Rudolph 1942. *Introduction to Semantics.* Cambridge MA*:* Havard University Press.

Cassimatis, Nicholas L., J. Gregory Trafton*,* Magdalena D. Bugajskab and Alan C. Schultz. 2004. Integrating cognition, perception and action through mental simulation in robots. *Robotics and Autonomous Systems* 49:13-23.

Cassimatis, Nicholas L. 2006. A Cognitive Substrate for Achieving Human-Level Intelligence. *AI Magazine* 27: 2.

Choe, Jae Woon, and Roland Hausser. 2006. Treating Quantifiers in Database Semantics. M. Duži *et al.* ed. *Information Modeling and Knowledge Bases XVIII*, ed. Amsterdam: IOS Press

Clark, Andy. 1997. *Being There: Putting Brain, Body and World Together Again.* Cambridge MA: The MIT Press.

Committeri, Georgia, Gaspari Galati, Anne-Lise Paradis, Luigi Pizzamiglio, Alain Berthoz and Denis LeBihan. 2004. Reference Frames for Spatial Cognition: Different Brain Areas are Involved in Viewer-, Object- and Landmark-Centered Judgements About Object Location. *Journal of Cognitive Neuroscience* 16:1517-1535.

Dabbs, James M., E-Lee Chang, Rebecca A. Strong and Rhonda Milun. 1998. Spatial ability, navigation strategy, and geographic knowledge among men and women. *Evolution and Human Behavior* 19:88-98.

Deacon, Terence. 1999. *The Symbolic Species.* New York: W.W. Norton & Co.

Dik, Simon C. 1997a. *The Theory of Functional Grammar. Part 1: The structure of the Clause.* Second, revised edition ed. Kees Hengeveld. Berlin: Mouton de Gruyter.

Dik, Simon C. 1997b. In Kees Hengeveld ed*. The Theory of Functional Grammar. Part 2: Complex and Derived Constructions*. Berlin: Mouton de Gruyter.

Fauconnier, Gilles. 1985. *Mental Spaces: Aspects of Meaning Construction in Natural Language.* Cambridge MA: The MIT Press.

Feldman, Jerome A. 2006. *From Molecule to Metaphor: A Neural Theory of Language.* Cambridge MA: The MIT Press.

Fillmore, Charles J. 1976. Frame semantics and the nature of language. In *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech* 280:20-32.

Fodor, Jerry. 1983. *The Modularity of Mind.* Cambridge MA: The MIT Press.

Frege, Gottlob. 1892. Über Sinn und Bedeutung (On Sense and Reference). Tr. Max Black in *Journal of Symbolic Logic*, 13:152-153.

Gamut, L. 1991. *Logic, Language and Meaning*, Vols. 1 and 2. Chicago IL: University of Chicago Press.

Giora, Rachel. 2003. *On Our Minds: Salience, Context and Figurative Language.* Oxford UK: Oxford University Press.

Hardy, Christine. 1998. *Networks of Meaning: A Bridge Between Mind and Matter.* Westpoint, CT: Praeger Publishers.

Higginbotham, James. 1984. English is Not a Context-Free Language. *Linguistic Inquiry* 15:225-234.

Hoffman, Guy and Cynthia Breazeal. 2008. Achieving fluency through perceptual-symbol practice in human-robot collaboration. *ACM/IEEE International Conference on Human-Robot Interaction*, Amsterdam, The Netherlands, March 12-15, 2008.

Jackendoff, Ray S. 1996. *The Architecture of the Language Faculty*. Linguistic Inquiry Monographs. Cambridge MA: The MIT Press.

Jones, Catherine M. and Susan D. Healy. 2006. Differences in cue use and spatial memory in men and women. *Proceedings of the Royal Society of London Series B* 273:2241-2247.

Joshi, Aravind and Owen Rambow 2003. A Formalism for Dependency Grammar Based on Tree Adjoining Grammar. *Proceedings of the Conference on Meaning-Text Theory*, Pari, 16-18 June.

Kecskes, Istvan. 2006. Synergic concepts in the bilingual mind. In Istvan Kecskes and Liliana Albertazzi ed. *Cognitive Aspects of Bilingualism.* Dordrecht, The Netherlands: Springer.

——— 2008. Duelling contexts: a dynamical model of meaning. *Journal of Pragmatics* 40:385-406.

Kamp, Hans and Barbara Partee. 1995. Prototype theory and compositionality. *Cognition* 57:129-191.

Kamp, Hans and Uwe Reyle. 1993. *From Discourse to Logic*. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Knill, David and Whitman Richards, ed. 1996. *Perception as Bayesian Inference.* Cambridge UK: Cambridge University Press.

Lakoff, Goerge and Mark Johnson. *Metaphors We Live By*. Chicago IL:Chicago University Press.

Lee, Kiyong 2002. A Simple Syntax for Complex Semantics. In Ik-Hwan Lee, Yong-Beom Kim, Key-Sun Choi and Minhaeng Lee ed. *Language, Information, and Computation: Proceedings of The 16th Pacific Asia Conference*, Keynote Speech 2-27, The Korean Society for Language and Information, Seoul.

Levinson, Stephen and David Wilkins, ed. 2006. *Grammars of Space: Explorations in Cognitive Diversity.* Cambridge UK: Cambridge University Press.

MacWhinney, Brian. 1987. The Competition Model. In Brian MacWhinney ed. *Mechanisms of language acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Markman, Arthur B. and Eric Dietrich. 2000. In defense of representation. *Cognitive Psychology* 40:128-171.

Meystel, Alexander M. and James S. Albus. 2002. *Intelligent Systems: Architecture, Design and Control.* New York NY:John Wiley and Sons.

Montague, Richard 1970a. Universal grammar. *Theoria* 36:373–398.

——— 1970b. English as a Formal Language. In ed. BrunoVisentini *Linguaggi nella società e nella tecnica* 1970:189–223.

Munnich, Edward, Barbara Landau and Barbara Anne Dosher. 2001. Spatial language and spatial representation: a cross-linguistic comparison. *Cognition* 81:171-207.

Paivio, Allan. 1990. *Mental Representations: A Dual Coding Approach*. Oxford UK: Oxford University Press.

Phillips, Colin. 2003. Linear Order and Constituency. *Linguistic Inquiry* 34:37-90.

Pollard, Carl. 1996. The Nature of Constraint-Based Grammar. *Pacific Asia Conference on Language, Information and Computation*, Seoul Korea December 20, 1996.

Portner, Paul and Barbara H. Partee 2002. *Formal Semantics: The Essential Readings.* Oxford UK: Blackwell Publishing.

Prince, Alan and Paul Smolensky. 1993. *Optimality Theory: Constraint Interaction in Generative Grammar.* Rutgers University Center for Cognitive Science Technical Report 2.

Rausch, Elizabeth. 1973. In T. Moore ed. On the Internal Structure of Perceptual and Semantic Categories. In *Cognitive Development and the Acquisition of Language.* New York: Academic Press.

Roy, Deb. 2007. Semiotic schemas: A framework for grounding language in action and perception. *Artificial Intelligence*167:170-205.

Russell, Bertrand. 1923. Vagueness. *The Australian Journal of Philosophy and Psychology* 1: 84-92.

Schabes, Yves and K. Vijay-Shanker. 1990. Deterministic left to right parsing of tree-adjoining languages. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics* 276-283, Pittsburgh PA.

Van der Zee, Emile and Jon Slack ed. 2003. *Representing Direction in Language and Space.* Oxford UK: Oxford University Press.

Shah, Priti and Akira Miyake, ed. 2005. *The Cambridge Handbook of Visuospatial Thinking.* Cambridge UK: Cambridge University Press.

Shapiro, Larry. 2007. The embodied cognition research programme. *Philosophy Compass* 2:338-346.

Shieber, Stuart M. 1986. An introduction to unification-based approaches to grammar. *CLSI Lecture Notes Number 4*. Stanford CA: CLSI Publications.

——— 1992. *Constraint-Based Grammar Formalisms.* Cambridge MA: The MIT Press.

Steedman, Mark. 1987. Combinatory grammars and parasitic gaps. *Natural Language and Linguistic Theory* 5:403–439

——— 1996. *Surface Structure and Interpretation.* Cambridge MA: The MIT Press.

——— 2000. *The Syntactic Process.* Cambridge MA: The MIT Press.

Sun, Ron. 2002. *Duality of the Mind: A Bottom Up Approach Toward Cognition.* Mahwah NJ: Lawrence Erlbaum Associates.

Szabolcsi, Anna. 1992. Combinatory grammar and projection from the lexicon. In Sag and Szabolcsi ed. *CSLI Lecture Notes 24*. Stanford CA: CLSI Publications.

Tarski, Alfred, Andrezj Mostowski and Rafeel M. Robinson 1953. *Undecideable Theories*. Amsterdam: North Holland Publishing Company.

Taylor, Holly and Barbara Tversky. 1992. Spatial mental models derived from survey and route descriptions. *Journal of Memory and Language* 31:261-292.

Tomita, M. 1987. An efficient augmented-context-free parsing algorithm. *Computational Linguistics* 13:31-46.

Wilson, Margaret. 2002. Six views of embodiment. *Psychonomic Bulletin & Review* 9:625-636.

Wlodarczyk, André. 2008. Roles and Anchors of Semantic Situations. *Etudes cognitive / Studia kognitywne 8*, SOW, PAN, Warsaw.

——— 2009. Interactive Discovery of Ontological Knowledge for Modelling Natural Language Resources. *International Scientific Conference COGNITIVE and CONTRASTIVE STUDIES* - MONDILEX Project Open Workshop "Representing Semantics in Digital Lexicography", Pałac Staszica PAN, Warszawa.

Wlodarczyk, André. and Hélenè. Wlodarczyk. 2008. The Pragmatic Validation of Utterances. *Etudes cognitive / Studia kognitywne 8*, SOW, PAN, Warsaw.