December 1995

Defining Cognitive Science at IRCS: Proceedings from the April 1995 Postdoc Workshop

Anne Vainikka
University of Pennsylvania

University of Pennsylvania Institute for Research in Cognitive Science Technical Report No. IRCS-96-07.

Defining Cognitive Science at IRCS: Proceedings from the April 1995 Postdoc Workshop

Abstract
The nature of language is the enigma linguists try to solve. What makes their task difficult is that in a number of cases, the logical links between the different components of a linguistic unit are difficult to establish. For example, in tone languages, the basic units that are the words are composed of two kinds of minimal, significant sound units, the phonemes on the segmental tier and the tonemes on a distinct tier. The word consists then of these two types of linguistic elements, related by phonological principles.

In Mawukakan, a Manding language spoken in Western Ivory Coast, the word resembles an architectural construction where the function of some structures only becomes evident when you consider the building with its surroundings. This is illustrated below by the tonal assignment of toneless clitics and the lengthening of the vowel of the focus marker *le* in Mawu.

Comments
University of Pennsylvania Institute for Research in Cognitive Science Technical Report No. IRCS-96-07.

Institute for Research in Cognitive Science

Defining Cognitive Science at IRCS:
Proceedings from the April 1995
Postdoc Workshop

Edited by Anne Vainikka

University of Pennsylvania
3401 Walnut Street, Suite 400C
Philadelphia, PA 19104-6228

December 1995

Site of the NSF Science and Technology Center for
Research in Cognitive Science

IRCS Report 96--07

Defining Cognitive Science at IRCS: Proceedings from the April 1995 Postdoc Workshop

Edited by Anne Vainikka

December 1995

Table of Contents:

Some Linguistic Speculations About Broca's Aphasia

Sergey Avrutin

Psycholinguistics as a subfield of Cognitive Science aims at understanding the relationship between the brain, the mind and language. The three major fields of research in Psycholinguistics - language acquisition, real time processing and language impairment - attempt to explain what is so special about human brain that enables us humans to acquire and understand language. In this sense, studies of brain-damaged patients provide a unique opportunity to investigate the relationship between a specific cognitive capacity (e.g. language) and the brain.

I begin my presentation with a brief overview of the development of aphasia research. Focusing on two types of aphasia (Broca's and Wernicke's), researchers have traditionally associated the former with a production deficit and the latter with a comprehension deficit. The picture thus appeared to be very simple: a certain area in the brain comprehends language and another area produces it. Such a simplified picture was challenged and has been proven wrong with the development of linguistic theory. Specifically, researchers have demonstrated that Broca's aphasics exhibit very *specific* problems with language comprehension, problems that can be characterized only with the help of a linguistic theory.

After discussing the results of previously reported comprehension experiments carried out with Broca's aphasics, I turn to to an experiment that I carried out in collaboration with Gregory Hickok. The goal of our experiment was to investigate to what degree two Broca's aphasics understand Wh-questions. Our interest was stimulated by the proposals in the literature that made certain predictions (although not overtly) about these constructions. In particular, the existing explanations do not differentiate between varios kinds of Wh-questions: *which*-questions and *who*-questions. In both cases, a poor performance is predicted.

Our approach, however, predicted a different pattern of results for these two kinds of questions. Based on data from real time processing experiments, we argued that only *which*-questions should present a problem for the subjects because these (so-called *D(iscourse)-linked constituents*) require more processing capacity. This prediction was borne out: both patients demonstrated above-chance performance on *who*-questions, and around chance performance on *which*-questions.

In addition, I hope to discuss results of an experiment with Italian speaking aphasics on adverb placement (reported in the literature) and show how our approach accounts for these results, as well as present two experiments that I have designed recently that will hopefully (!) support my proposal. In conclusion I will present some central ideas of the model that I hope to develop which would explain the abnormal pattern of comprehension not in terms of loss of some linguistic knowledge but rather in terms of a limited processing capacity.

Some Aspects of Mawukakan Prosody

Moussa Bamba

1. Introduction

The nature of language is the enigma linguists try to solve. What makes their task difficult is that in a number of cases, the logical links between the different components of a linguistic unit are difficult to establish. For example, in tone languages, the basic units that are the words are composed of two kinds of minimal, significant sound units: the phonemes on the segmental tier and the tonemes on a distinct tier. The word consists then of these two types of linguistic elements, related by phonological principles.

In Mawukakan, a Manding language spoken in Western Ivory Coast, the word resembles an architectural construction where the function of some structures only becomes evident when you consider the building with its surroundings. This is illustrated below by the tonal assignment of toneless clitics and the lengthening of the vowel of the focus marker *le* in Mawu.

2. The Prosodic Structure of the Mawukakan Noun Phrase

2.1 The tone assignment of the toneless determiner *-lu*

In Mawu, the noun phrase (NP) can be composed of a single noun root or a noun root plus one or many determiners. Any noun root, independently of its syllable structure, is assigned one of the four tonal patterns H (high), L (low), $H(L)$ (high-low) or $LH(L)$ (low-high-low). The final (L) of the two last patterns is never associated with the syllables of the root, as shown in (1).

(1)a.	*sisi*	-->	[*sisi*]	"*chest*"
	H		H	
b.	*sisi*	-->	[*sisi*]	"*smoke*"
	L		L	
c.	*gbusu*	-->	[*gbusu*]	"*gizzard*"
	$H(L)$		H	
d.	*wotolo*	-->	[*wotolo*]	"*handcart*"
	$LH(L)$		LH	

In (1a) where the underlying tone of the noun is H, each syllable has a high tone on the surface, and in (1b) with an underlying L, each syllable is realized low. In (1c), the noun is realized entirely high – as in (1a) – and the (L) of the underlying pattern is missing. In (1d), the first syllable of the noun is associated with the first tone of the underlying $LH(L)$ tonal pattern, and the two others syllables are specified by the H. Again, the final (L) does not show up.

However, when a toneless determiner is suffixed to the noun roots in (2), it is assigned an H after the patterns H and L, but bears an L after $H(L)$ and $LH(L)$, as illustrated in (2).

(2)a.	*sisi + lu*	-->	[*sisi*	*−lu*]	"*chests*"
	H		!*H*	*H*	
b.	*sisi + lu*	-->	[*sisi*	*−lu*]	"*smokes*"
	L		*L*	*H*	
c.	*gbusu + lu*	-->	[*gbusu*	*−lu*]	"*gizzards*"
	H(L)		*H*	*L*	
d.	*wotolo + lu*	-->	[*wotolo*	*−lu*]	"*handcarts*"
	LH(L)		*LH*	*L*	

Given that *-lu* is realized with an H after both H and L, its tone cannot result from the application of a normal tone rule in (2a,b). On the other hand, the L it bears in (2c,d) reveals an association of the suffix with the floating (L) of the noun root, in this context.

In addition to the appearance of H in (2a,b), there is a further reason to consider the H on *-lu* as the default tone at the tonal level in this derivation. The reason is that these two types of noun roots behave differently in relation to an auxiliary, as can be seen in (3).

(3)

a.	*ku(H)*	*wEE(H)*	*ye(H)*	-->	[*ku*	*wEE*	*ye*]
	yam	*Aux.perf*	*see*		!*H*	*H*	!*H*
	"A yam has been seen"						
b.	*saa(L)*	*wEE(H)*	*ye(H)*	-->	[*saa*	*wEE*	*ye*]
	sheep	*Aux.perf*	*see*		*L*	*H*	!*H*
	"A sheep has been seen"						
c.	*gbusu(H(L))*	*wEE(H)*	*ye(H)*	-->	[*gbusu*	*wEE*	*ye*]
	gizzard	*Aux.perf*	*see*		*H*	!*H*	*H*
	"A gizzard has been seen"						
d.	*wotolo(LH(L))*	*wEE(H)*	*ye(H)*	-->	[*wotolo*	*wEE*	*ye*]
	handcart	*Aux.perf*	*see*		*LH*	!*H*	*H*
	"A handcart has been seen"						

In (3a), the H of the auxiliary is realized higher than either the noun's or the verb's H tones. In (3b), where the noun has an L tone, the H of the auxiliary is higher than that of the verb. However, in (3c,d), the noun's final syllable is realized as a higher tone than what the auxiliary bears.

To account for the behavior of the nouns with a simple tonal pattern (H, L) and those with a complex pattern ($H(L)$, $LH(L)$), it is sufficient to assume that (i) nouns with a complex pattern bear underlying final accent, and (ii) in a sentence, the noun root and the morpheme adjacent on the right are gathered in an *iambic* foot if the noun is not underlyingly accented, but in a *trochaic* foot if the noun has a lexical accent.

2.2 The plural form of *-le*

In Mawu, the focus marker is a morpheme *-le* which is suffixed to the constituent. Like the plural morpheme *-lu*, *-le* is a determiner when it appears in the NP, where it can appear with all the other determiners. In (4), *-le* is realized as *-le* when it appears in final position or before the question marker *-a*. But when followed by the plural marker *-lu*, it becomes *-lee*, with its vowel lengthened.

(4)

a.	*saa+*	*le*	$-->$	[*saa−*	*le*]	
	H	*L*		*H*	*HL*	

"It's the squash seed."

b.	*saa+*	*le+*	*a*	$-->$	[*saa−*	*le−*	*a*?]	
	H	*L*			*H*	*H*	*L*	

"It's the squash seed?"

c.	*saa+*	*o+*	*le+*	*lu*	$-->$	[*saa−*	*o−*	*lee−*	*lu*]
	H	0	*L*	0		*H*	*H*	*H*	*L*

"It's the squash seeds."

In final position in (4a) and before *-a* in (4b), the variation affecting *-le* is only tonal. In (4c) where it is followed by the plural morpheme *-lu*, *-le* has its vowel lengthened and becomes *-lee*.

This lengthening is surprising given that *-le* and *-lu* are completely independent in the language. For example, *saa(H) + lu* realized as *[saa-luHH]* 'squash seeds' is attested. To account for this unpredictable lengthening, I suggest the underlying representation in (5) for *-le*, where **S** = syllable, **O** = onset, **N** = nucleus and **0** is an empty segmental position.

```
(5)     S
      / \
     O   N
     |  / \
     l e   O
```

The 0 at the segmental level in (5) will be filled only if the following morpheme is the plural morpheme *-lu*.

One may ask why such a rare and exceptional structure as in (5) would become necessary. The answer to this question is not purely phonological, but also semantico-pragmatic. In fact, the strongest argument in favor of (5) is that in Mawu, *-lee* is attested in an other contexts where the plurality of the element in focus is inherent to this element. This context is the emphatic plural pronouns which are composed of the non-emphatic forms plus the focus marker, as in (6).

(6)*a.*	*ã+*	*le*	$-->$	[*ã−*	*nee*]	"*we(emphatic)*"
	H(L)	*L*		*H*	*L*	
b.	*a+*	*le*	$-->$	[*a−*	*lee*]	"*you − PL(emphatic)*"
	H	*L*		*H*	*H*	
c.	*i+*	*le*	$-->$	[*i−*	*lee*]	"*they(emphatic)*"
	L	*L*		*L*	*L*	

In (6), it is the long form *-lee* which is suffixed to the plural emphatic pronouns, while the singular emphatic pronouns appear with the short form *-le*.

These two peculiarities of Mawu Phonology come to the rescue of certain abstract structures in the search for a logical link between what is supposed to be found in the input and what is obtained in the output structure.

3. Phonology is a cognitive science

The development of Phonological Theory is characterized by the discovery of the hidden connections between things which are apparently not related. For example, the *syllable* today represents one of the most important discoveries, even if the notion refers to an abstract phonological unit; this is also the case with the notion of *foot* which represents a relationship between syllables. Therefore, Phonology is a branch of Cognitive Science.

Discourse Structure, Inference, and Marked Syntactic Constructions

Betty J. Birner

Comprehension of any coherent discourse involves the tracking of information as it is introduced and developed. This tracking is facilitated in part by the use of marked (i.e., non-canonical) syntactic constructions such as those in (1):

(1) a. There is a rabbit in the garden.
  b. He's a genius, that guy.
  c. Sitting at the table was Aunt Mary.

The felicitous use of these constructions depends on a variety of factors relating to the 'information status' of their constituents – for example, whether the information they represent is already present in the interlocutors' shared model of the discourse or is assumed to be shared background knowledge. The use of existential *there*, as in (1a), requires that the postverbal noun phrase (here, *a rabbit*) represent information that is new to the hearer, whereas the use of right-dislocation, as in (1b), requires that the sentence-final noun phrase represent previously evoked information (Birner & Ward 1996). The use of inversion, as in (1c), on the other hand, requires that the sentence-initial constituent represent more familiar information within the discourse than does the sentence-final constituent (Birner 1994). Thus, these constructions can be used to indicate the relationship that obtains between information being presented in the current utterance and information already evoked in the discourse, or in the store of knowledge assumed to be shared by the speaker and hearer.

One such relationship is that of identity; another is inference. While much work has been done on the tracking of identity relationships in discourse (through, e.g., the use of pronouns), much less has focused on information that is rendered 'inferrable' (Prince 1981) on the basis of other information presented in the discourse, as exemplified by the italicized phrases in (2):

(2) a. I'll be up late writing this term paper; I haven't even finished *the introduction*.
  b. That's a great mystery novel; you'll never guess who *the murderer* is.
  c. I wanted to make some homemade bread last night, but *the recipe* was illegible.

In each of these examples, the entity represented by the italicized NP has not been previously evoked, but its likely existence is inferrable from information presented in the prior discourse.

An examination of natural discourse shows that inferential links can serve the same connective function in a marked syntactic construction as do links of identity. In English inversion, for example, inferrable information appears in exactly the same range of positions and contexts as does information that has been explicitly evoked in the prior discourse. Topicalization in English and inversion in Farsi show a similar patterning with respect to word order; that is, inferrable information shows the same distribution as does explicitly evoked information. An examination of such constructions suggests that wherever the felicitous use of a construction depends upon the information status of its constituents, inferrable information may appear in the same range of positions as does explicitly evoked information.

References:

Birner, B. 1994. "Information status and word order: An analysis of English inversion." *Language* 70:233-259.

Birner, B., and G. Ward. 1996. "A crosslinguistic study of postposing in discourse." *Language and Speech: Special Issue on Discourse, Syntax, and Information.* To appear.

Prince, E.F. 1981. Toward a taxonomy of given/new information. In P. Cole, ed., *Radical Pragmatics.* New York: Academic Press. 223-54.

Genetics and Cognition

Juergen Haas

In recent years remarkable similarities between Genetics and Cognitive Science have been uncovered, notably in the linguistic nature of DNA and other genetic material. The source of such similarities may in part be due to the fact that evolution and cognition require solutions to similar problems. It is interesting that the human genome has characteristics that are usually only associated with the human mind. The 23 human chromosomes are made of about 3 billion base pairs which corresponds to roughly 1.5 gigabytes of information. While it is not clear what percentage thereof is actually used, there appears to be a large amount of "knowledge" encoded in the human genome. Furthermore, the genome appears to have developed sophisticated techniques to acquire new knowledge and to improve itself and the organism it builds, i.e., it can "learn". The most striking cognitive aspect of the genome is that is uses a type of "language" to organize its knowledge.

One of the problems faced by both genetics and cognition is the importance of rapid faithful replication of knowledge. In addition, it must be possible to "reason" about that knowledge, that is, to make "statements" about pieces of information or to add new knowledge. In genetics this amounts to manipulating small regulatory elements on the DNA which determine whether a gene will be expressed and to what extent it is expressed. This is achieved mainly through genetic recombination and other operations on DNA molecules, which is analogous to forming quantitative statements in natural language. The data structure best suited for these tasks appears to be the one-dimensional digital array. Since most copies are made from other copies, the digital nature of the data structure facilitates faithful reproductions over many generations without distortions. In Genetics this array is essentially the DNA using an "alphabet" of four nucleotides: A, C, G and T. In Cognitive Science the corresponding construct is language with letters or phonemes as the elements of the arrays.

In both domains the task is to encode information to build complex multi-dimensional structures in one-dimensional strings, which leads to a structuring of the string that can be captured by grammars. We can say that a sentence in natural language roughly corresponds to a gene in genetics. The functional units of a gene, such as exons, coding regions and regulatory elements, are hierarchically organized into "phrases" to form the complete structure of the gene. We have developed grammars and parsers for a large class of genes and applied it to detect inconsistencies in genetic databases by checking the "grammaticality" of the given information. An extension of the parser can even correct certain types of such errors and infer additional information, thus making it possible to automatically construct databases with uniform and consistent gene structure information. Such databases have been successfully used for applying machine learning techniques, such as various forms of Case-Based Reasoning, to infer additional information. Further development of such grammars is necessary to capture the structure and arrangement of regulatory regions that surround the core sequence of a gene.

The function of most genes is to construct an amino acid sequence that folds into a particular functional protein (the "semantics" of a gene). However, certain classes of genes result in functional molecules by constructing RNA sequences that fold into characteristic configurations. The folding pattern (secondary structure) is determined in large part by pairwise interactions between the elements of the sequence. A class of grammars has been developed that can detect patterns in DNA or amino acid sequences that correspond to secondary structure, thus facilitating prediction of secondary structure. Typically, such grammars are stochastic in nature and are trained on a set of sequences whose secondary structure is known.

While the linguistic aspect of genes plays a crucial role in genetics and evolution, the three-dimensional physical properties of the proteins constructed from the DNA are what keeps an organism functioning and alive. In particular a large number of proteins interact with each other and DNA molecules to regulate gene expression (the rate at which proteins are generated). Such interactions are responsible, for example, for cell differentiation. In order to understand these mechanisms researchers have developed a number of software packages for modeling molecules (such as proteins) and simulating the forces within and between them. However, in order to discover molecular mechanisms involved in gene regulation we are developing a reasoning/learning component in addition to a specialized modeling/simulation component that can hypothesize molecular structures and mechanisms and verify or reject them through simulation.

Literature:

J.S. Aaronson, J. Haas, and G.C. Overton (1993) "Knowledge Discovery in GenBank", Proceedings of the First International Conference on Intelligent Systems for Molecular Biology, AAAI/MIT Press, 3-11.

S. Dong and D.B. Searls (1994) "Gene Structure Prediction by Linguistic Methods", Genomics, Vol 23, 540–551.

J. Haas, J.S. Aaronson, and G.C. Overton (1993) "Using Analogical Reasoning for Knowledge Discovery in a Molecular Biology Database", Proceedings of the Second International Conference on Information and Knowledge Management, ACM Press, New York, NY, 554-564.

A. Krogh, M. Brown, I.S. Mian, K. Sjolander and D. Haussler (1994) "Hidden Markov Models in Computational Biology: Applications to Protein Modeling", J. Mol. Biol., Vol 235 No.5, 1501-1531.

G.C. Overton, J.S. Aaronson, J. Haas, and J. Adams (1994) "QGB: A System for Querying Sequence Database Fields and Features", Journal of Computational Biology, Vol 1 no.1, 3-14.

Y. Sakakibara, M. Brown, R. Hughey, I.S. Mian, K. Sjolander, R.C. Underwood and D. Haussler (1994) "Stochastic Context-Free Grammars for tRNA modeling", Nucleic Acids Res., Vol. 22 no. 23, 5112-5120.

D.B. Searls (1992) "The Linguistics of DNA", American Scientist, Vol. 80, no.6, 579-591.

D.B. Searls (1995) "Formal Grammars for Intermolecular Structure", First International IEEE Symposium on Intelligence in Neural and Biological Systems, 30-37.

An Adapting Brain

Lance Hahn

In order to survive it is important for a human, or any animal, to attend to what is important (i.e. a predator or food) and to ignore what is irrelevant. By attending we are simply devoting available mental resources to the important target. This is done at many different processing levels. At one level, we choose to pay attention to a speaker and ignore the person we are sitting next to, at another level we pay attention the changes in light levels within a room and ignore the absolute light level. I am interested in how the human visual system modifies itself in order to pay attention to relative light levels within one context. In particular, I am interested in light adaptation mechanisms and neuron-based descriptions of these mechanisms.

One of my projects this year has been to resolve a debate between anatomists and visual psychophysicists. Anatomists have shown that each cone photoreceptor in the retina is connected to its neighbors via electrical synapses. This implies that adjacent photoreceptors interact. Visual psychophysicists, on the other hand, use behavioral measurements to show that the response of a receptor in the high acuity region of the retina (the fovea) is independent of its neighbor's response. One hypothesis which can resolve this dispute, is that the electrical synapses are only functional under certain restricted circumstances. At low luminance levels, where previous behavioral measurements are ambiguous, electrical synapses could improve the representation of light levels at the cost of spatial resolution. I will discuss the recent measurements which I have made and relate them to this hypothesis.

The new psychophysical measurements, and a collaborative work on which the potential photoreceptor interaction was modeled, suggest that the electrical synapses may not change as a function of light level, but could be useful in reducing noise at all light levels. At low light levels the signal-to-noise ratio is reduced by photon noise. At high light levels photon noise does not reduce the signal-to-noise ratio as much as at low light levels, but intrinsic cellular noise can play a larger role in reducing the signal-to-noise ratio. Thus, noise reduction via photoreceptor coupling may be desirable across light levels.

In order to answer the large questions of how we perceive, think and act we must understand what neural tools are available to the brain. By establishing links between known neural structures and known behavior, we begin to understand neural tools which shape the way we perceive, think and act at many neural loci.

Language Complexity and Theories of Syntax

James Rogers

Early theories of syntax within Generative Linguistics were based primarily in the realm of rewriting systems—Phrase Structure Grammars of varying types and Transformational Grammars. This approach was formally very fruitful, yielding, in particular, hierarchies of language complexity classes. While it is questionable how appropriate specific results in this realm are for linguistic theory, the underlying form of the results corresponds reasonably closely to issues that are linguistically relevant. Moreover, these classes have dual characterizations in terms of the structure of the languages in the class and in terms of the types of resources that are required to process these languages. Thus hypotheses about where the class of human languages falls with respect to these hierarchies make predictions about the nature of the human language faculty which correspond precisely to predictions about the structure of natural languages. Since the structure of language can be studied more or less directly, there is reasonable expectation of finding empirical evidence falsifying such hypotheses, should it exist.

More recently there has been a de-emphasis of the role of the mechanism deriving strings in the language in favor of well-formedness conditions on the structures that analyze their syntax. Government and Binding Theory is an example of such an approach. While theories in this framework are usually modeled as a specific range of Transformational Grammars, the connection between the underlying mechanism and the language the theory licenses is quite weak. In an extreme view, one can take the mechanism to generate the set of all finite trees (labeled with some alphabet) while the linguistic theory is actually embodied in a set of principles that filter out the ill-formed analyses. At least some of the motivation for taking this approach is an awareness that the structural characteristics that distinguish natural languages as a class may not be those that can be distinguished by existing language complexity classes. Thus the insights to be gained from formal language theory may actually be misleading in trying to develop a theory of human language. Nonetheless, as we have just argued, there are good reasons for establishing language complexity results for the class of languages licensed by such a theory. Berwick has suggested that one might consider the GB approach as aiming to "discover the properties of natural languages first, and then characterize them formally."

The difficulty of the first step of this program is well attested. Unfortunately, because GB theories are based more in the notion of licensing than in the notion of rewriting, the second step has proved quite difficult as well. One of the things I have been studying is a flexible and quite powerful method of establishing such results. This starts with a typical approach to formalizing GB theories—defining the set of structures they license with a set of logical axioms—but requires that the axioms be stated within a particular restricted logical language, termed $L^2_{K,P}$. On one hand, the restriction is reasonably weak; it is not difficult to capture most of the principles employed in GB theories reasonably transparently. On the other hand, it turns out to be a quite powerful restriction formally; one can show that a set of finite trees can be defined in $L^2_{K,P}$ if and only if it is (modulo a projection) a local set—the set of derivation trees licensed by a Context-Free Grammar. Thus definability in $L^2_{K,P}$ characterizes the strongly Context-Free Languages.

Using this result, one can explore the relative complexity of various constellations of principles within the GB framework. One can show, for instance, that free-indexation, at least in its most general form, cannot be enforced by a Context-Free Grammar. Somewhat more surprising, perhaps, is the fact that a reasonably complete GB account of English can be captured within $L^2_{K,P}$, thus establishing that the languages licensed by a specific GB theory are strongly Context-Free. While I will discuss these results briefly, the focus of this talk will be the potential implications such results may have with respect to the nature of the human language faculty. As it turns out, refinements of GB theories have tended to make them more amenable to definition in $L^2_{K,P}$. This suggests that the regularities in natural languages that motivate the linguistic theories might, in fact, be reflections of properties of the human language faculty that can be characterized, at least to some extent, by language complexity classes.

The Morpho-Phonological Basis of Syntactic Parameters

Bernhard Rohrbacher

Current linguistic theory assumes that the human brain is genetically endowed with a language faculty referred to as Universal Grammar (UG) which contains at least two types of elements: principles and parameters.

UG *principles* are fixed at birth, or mature independently of the language specific input; they are responsible for the immutable facts of human language, such as the fact that no human language allows extraction from one of two conjuncts, as shown by the ungrammatical English example **Who did you see the boy and?*

UG *parameters* are fixed by the child on the basis of the language specific input; they are responsible for the variation among human languages, such as the fact that some but not all human languages allow missing subjects in constituent questions, as shown by the grammatical Italian example *Quante pietre hai preso?* (How-many stones have-2SG taken) and its English translation *How many stones did you take?*

It has recently been suggested that UG parameters are fixed exclusively on the basis of language specific properties of *functional categories* such as tense, negation and subject-verb agreement (cf. Chomsky 1989). The parametrically relevant properties of functional categories are often assumed to be abstract, but this assumption is problematic since it often results in a circular theory which 'explains' a syntactic phenomenon by appealing to abstract features for which the syntactic phenomenon in question constitutes the only evidence. To avoid this problem, I propose that the parametrically relevant properties of functional categories are *concrete* in the sense that they are reflected in the overt morpho-phonology of these categories.

In this talk, I discuss the feasability of this theory with respect to the missing subject parameter illustrated above (cf. Rohrbacher 1994, Speas 1994). The data come from young children acquiring English, German or Swedish who frequently omit subjects although the target gramar does in general not allow this. I argue that such early subject omissions follow from the (non-adult like) morpho-phonological tense and agreement systems of the children and in particular from their tendency to utter unembedded non-finite clauses (cf. Roeper & Rohrbacher 1995, Rohrbacher & Vainikka 1995).

If parameterization involves only the overt morpho-phonological properties of functional categories, then children have to learn only the bare essentials, i.e. the idiosyncratic and irreducible shapes of the words of their target language. As far as I can see, this represents the simplest conceivable theory of language acquisition. This theory may in fact very well turn out to be overly simplistic, but that seems to me to be worth finding out.

References:

Chomsky, Noam (1989). "Some Notes on the Economy of Derivation and Representation." MIT Working Papers in Linguistics 10:43-74.

Roeper, Tom and Bernhard Rohrbacher (1995). "Null Subjects in early Child English and the Theory of Economy of Projection," to appear in University of Pennsylvania Working Papers in Linguistics 2.

Rohrbacher, Bernhard (1994). "Explaining the Syntactic Consequences of 'Rich' Agreement Morphology: On the Licensing of V-to-AgrS Raising and pro," to appear in Proceedings of the 13th West Coast Conference on Formal Linguistics.

Rohrbacher, Bernhard and Anne Vainikka (1995). "Verbs and Subjects before Age 2: The Earliest Stages in Germanic L1 Acquisition," to appear in Proceeding of the 25th Meeting of the North East Linguistic Society.

Speas, Margaret (1994). "Null Arguments in A Theory of Economy of Projection." University of Massachusetts Occasional Papers in Linguistics 17:179-208.

Motion Perception and the Spherical Eye

Inigo Thomas

Mobility is a crucial ability which divides the universe of living things into two categories: animals that move and plants that do not. Needless to say, it is of great importance for animals to sense self-motion as well as to perceive the world while moving. Designing motion sensors has played a prominent role in building artificial animals (robots), especially using the modality of vision.

As an observer moves, the images on his or her retina continuosly change. These changes are systematically related to the observer's motion and the location of objects in the world. For example, while looking out a train window, objects close to you move faster on your retina than objects further away. Although the *accuracy* with which humans can tell their direction of motion is still unknown, it seems quite obvious that humans *can* tell which way they are moving with respect to the environment purely from visual information. Such perception of motion – purely from visual input – is very striking when watching a movie in an Omni-Max (large, dome-screened) Theater.

Introspection reveals that when a human observer moves, the eye continually fixates on (or tracks) targets in the world. Further, the spherical human retina allows humans to see a very large field of view, almost 180 degrees along most directions. Although both fixation and a very large field–of–view are involved in human vision, the computational advantage of either (if any) for the perception of self–motion has not been established. The main contributions of this work include (i) formalizing changes occuring in the retina for a moving, fixating observer, (ii) using a hemispherical retina/imager to simplify the patterns of image changes, and (iii) thereby isolating very simple patterns at the retinal periphery that could be useful in determining an observer's direction of heading.

We will focus on very simple image changes that occur when an observer moves. More specifically, we will focus only on whether objects move upwards or downwards in one's retina during one's motion. Both mathematics and intuition indicate that if an observer moves direct ahead, i.e. in the direction he or she is looking, the retina is divided precisely into two halves: one region where all objects move upwards and the other where the objects move downwards. However, if instead the observer heads in any other direction than directly ahead, the retina is divided into four regions! Furthermore, the ratio of the areas of these regions is precisely and directly related to the direction of one's heading.

Since we do not have any experiments where human observers were shown these patterns and asked about the perceived direction of heading (owing primarily to the difficulty in generating large field–of–view movies), we are unable to make any claims regarding whether these patterns are actually employed by humans. However, since these patterns are very simple, it is possible that human visual perception uses them. If I had a psychologist-genie I would ask it to test humans with images in Omni-Max Theaters with these simple patterns. If I had a second genie, I would like it to find the portion(s) of the brain that is (are) involved in the perception of one's motion. If I had a third genie, I would like it to philosophize on movement: What is movement? Is the notion of movement an inescapable consequence of a changing, material universe? Or, quite the contrary, what would be the consequence of the lack of motion? That is, would a universe where matter remained stationary become timeless?

Instead of testing the predictions about direction of heading perception on humans, I have tested them on robots using the first actual hemispherical lens-camera system in motion research, involving a 180 degree field-of-view lens. I will conclude my talk with a video clip that depicts a mobile robot – equipped with the new lens – continuously monitoring its direction of motion, and thereby servoing to a visually identified target.

The Emergence of Sentence Structure

Anne Vainikka

A profound discovery of modern syntax is that sentences consist of three basic components: VP, IP and CP. The core of a sentence is the VP component, which contains the verb and its arguments. The IP component contains inflectional morphology which typically anchors the VP in real time (tense) or in possible worlds (modality). The third component, CP, contains elements which connect the sentence to other sentences or to discourse.

By the age of about three, children have clearly mastered these three components of sentences. For example, Fortescue (1984) reports that a child learning Greenlandic Eskimo used a set of about 60 inflectional affixes at the age of 2 years and 3 months; the adult language has over 700 such affixes. In this language, sentences often consist of a single word, a verb with a number of affixes attached to it. Around the age of two, however, verbs can have at most two affixes at the same time (e.g. subject-verb agreement and aspectual marking, related to the IP component). By the age of about three, Greenlandic Eskimo speakers produce verbs with four affixes, including subordinating enclitics which are related to the CP component (Fortescue & Lennert Olsen 1992).

What happens before the age of three? Are all three components present from birth? (This is a common view; cf. Hyams 1992, Poeppel & Wexler 1993.) Is just the core component (VP) present from birth, while the other two develop due to biological maturation? (This is the view developed in Radford (1990).) Are all three learned based on the input? This is the view that I have put forth (Vainikka 1993/4), with the proviso that such 'learning' is firmly guided by an innate language acquisition mechanism.

Under this view, it appears that the VP-component is acquired prior to the age of 2, while the IP-component appears to come in around the age of 2, and the CP-component around 2 1/2 to 3 years of age. The hypothesis of gradual acquisition of sentence structure has been tested for English (Vainikka 1993/4 and Vainikka & Roeper 1995), German (Rohrbacher & Vainikka 1994), and Tamil (a Dravidian language; Vainikka & Thomas 1994). Work on the acquisition of Swedish, Greek and Hebrew is in progress.

Although some of the predictions made by my approach are similar to those made by Radford's biological maturation approach, I am convinced that maturation is not involved in the acquisition of sentence structure. This conviction stems from our work on the naturalistic (non-classroom) acquisition of foreign languages by adults, who also appear to first posit the VP component in the new language, followed by the IP component, followed by the CP component (Vainikka & Young-Scholten 1994). Biological maturation could not be responsible for such a pattern of development in adult language acquisition; assuming that the two similar patterns should be accounted for by the same mechanism, the child pattern would not involve maturation either.

References:

Fortescue, M. (1984) West Greenlandic. Beckenham, Kent: Croom Helm.

Fortescue, M. & L. Lennert Olsen (1992) 'The Acquisition of West Greenlandic', in D.Slobin (ed.) pp.111-219. The Crosslinguistic Study of Language Acquisition, vol.3 Hillsdale, NJ: Erlbaum.

Hyams, N. 1992. 'The genesis of clausal structure'. In Meisel, J. (ed.), The acquisition of verb placement. Dordrecht: Kluwer.

Poeppel, D. & K. Wexler (1993) 'The Full Competence Hypothesis of Clause Structure in Early German'. Language 69:1-33.

Rohrbacher, B. & A. Vainikka (1994) On German Syntax under Age 2. IRCS Technical Report 94-24, University of Pennsylvania.

Vainikka, A. (1993/4) 'Case in the Development of English Syntax'. Language Acquisition 3:257-325.

Vainikka, A. & T. Roeper (1995) 'Abstract Operators in Early Acquisition'. The Linguistic Review 12:275-310.

Vainikka, A. & I. Thomas (1994) 'Development of Functional Projections in Tamil', Penn Review of Linguistics v.18.

Vainikka, A. & M. Young-Scholten (1994) 'Direct Access to X'-Theory: Evidence from Korean and Turkish Adults Learning German', in T.Hoekstra & B.D.Schwartz (eds.) pp.265-316. Language Acquisition Studies in Generative Grammar. Amsterdam: John Benjamins.

The Representation of Syntactic Knowledge

Spyridoula Varlokosta

The word *cognitive* refers to knowledge. Cognitive Science is a science that seeks to understand various mental processes such as perceiving, remembering, understanding, or learning. Hence, what constitutes the domain of inquiry of Cognitive Science is essentially the architecture of the human mind which is viewed as a complex system that receives, stores, and processes various kinds of information.

One of the disciplines that comprise Cognitive Science is Linguistics, which concerns the scientific inquiry to the human language. The goal of Linguistics is to understand how language is represented in the human mind, how it is acquired, how it is perceived and used, and how it relates to other components of cognition. Language has been traditionally viewed as a relationship between a physical signal and a meaning. A formal account of such a relationship is called a Grammar. A Grammar must be able to characterize both what we know about the physical signal and the meaning, as well as the mapping between the two. The latter is what the area of Syntax deals with.

The specific syntactic topic discussed here is *reference of empty subjects* and in particular the relationship between the Tense properties of embedded clauses and the reference of their empty subjects. The linguistic evidence is drawn primarily from Modern Greek.

Unlike English or the various Romance languages, Modern Greek does not have infinitives. Instead, Modern Greek makes use of a finite construction (known as the subjunctive construction) in the corresponding English or Romance infinitival contexts. These constructions are introduced by the particle *na* and take a verb which is always inflected for Agreement and Aspect.

The possibility of obligatory coreference (as in (1)) or its absence (as in (2)) is shown to be associated with the tense and aspectual properties of the subordinate clauses in (1) and (2), respectively, namely to the ability of the embedded clause to form an aspectually and temporally independent domain:

(1)a.	*O*	*Yanis*	*arxise*	*na*	*kolimbai*
		John – Nom	*started*	*PRT*	*swims – 3Sg/Present*
b.	**O*	*Yanis*	*arxise*	*na*	*kolimbuse*
		John – Nom	*started*	*PRT*	*swam – 3Sg/Past*

(2)a.	*I*	*Maria*	*elpizi*	*na*	*pai*	*sto*	*Parisi*
		Mary – Nom	*hopes*	*PRT*	*goes – 3Sg/Pres*	*to*	*Paris*
b.	*I*	*Maria*	*elpizi*	*na*	*pige*	*sto*	*Parisi*
		Mary – Nom	*hopes*	*PRT*	*went – 3Sg/Past*	*to*	*Paris*

As cognitive scientists our goal is to provide a formal characterization of the form in which human beings represent knowledge and process information. As linguists, our task is to uncover general characteristics, properties, or principles of human languages that reflect aspects of the architecture of the human mind. If in the course of investigating a particular language we uncover something we think is a significant general principle of language, we are led to hypothesize such a principle as part of the biological architecture of the linguistic processing system. This is a very powerful tool because it directs us to look for evidence which our hypotheses are consistent with in data from other human languages. Hence, when a certain claim about a universal principle makes interesting predictions, it motivates further research concerning its role in other languages and ultimately concerning its significance for the structure of the human mind.

Commentary: What Is Cognitive Science?

Gary Hatfield
Dept. of Philosophy

These ten interesting and diverse presentations are part of a workshop in which Penn IRCS postdoctoral fellows present their work while speaking to, or perhaps answering through example, the question: What is cognitive science? As a framework for my commentary, I first present two contrasting global answers to this question. I associate the first answer with the "symbolist view," the second with the "cognitive capacities view" of cognitive science.

¿From ca. 1975-1985 it was widely accepted that cognitive science is or should be defined as a discipline through its commitment to a particular theoretical vision of the structure of human cognition, according to which the human mind is a computational device with a basic architecture like that of a von Neumann machine (Kosslyn & Hatfield 1984; Pylyshyn 1989). The human brain is attributed an "internal symbol system" or "language of thought" (Fodor 1975) in which all cognitive operations are performed. The brain is conceived as a device that is built to use a specific "machine language"; this built-in symbol system has been dubbed the "classical architecture" (Fodor & Pylyshyn 1988). Commonly, adherents of the classical architecture envision three levels of analysis: a *semantic level*, which pertains to referential semantics or to representation/world relations; a *syntactic level*, pertaining to the formal properties of representations and the rules defined over those representations solely in virtue of their formal properties (rules that define the basic computational architecture of a cognitive system); and a *neural level*, pertaining to the hardware realization of the symbol-processing mechanisms. Under this conception, cognitive science is unified into a single discipline by its commitment to the symbol-system architecture for human cognition.

The "cognitive capacities" view defines cognitive science by its subject matter, rather than by its commitment to a particular theoretical paradigm. "Cognitive science" is a cross- disciplinary, cooperative program for studying the cognitive capacities of human beings and related species. These capacities include perception, language processing, memory, learning, problem solving, and reasoning. They all involve representing the world, or at least ostensibly doing so. The disciplines of cognitive science study the structure of these representations and their relations to the world, and they seek to analyze the processes by which these representations are formed and used in various cognitive tasks. Three levels of analysis may be described here too: the *task analysis*, which describes the function of a given system, or its task in relation the larger cognitive system (e.g., in mammals an important task of the visual system is to form a representation of the ambient spatial layout); a second level often described as the "algorithmic" (Marr 1982) or "representational" (Hatfield 1991) level, but which might better be termed the level of *functional analysis*, since it involves an analysis of the task into subtasks, which in the case of representational tasks focuses on the processes by which representations are formed; and finally a level of *neural implementation*. Under this conception of cognitive science the taxonomy of cognitive research is set by the task analyses. The functional architecture of cognitive subsystems and the relations between such systems can be analyzed prior to any particular commitments regarding the formal properties of elementary representations. In particular, the question is left open of whether symbolic representations are themselves primitive to the human cognitive system or are constructed from still more primitive representational elements.

These two answers to the title question yield two distinct conceptions of cognitive science as a discipline. The symbolist view would attempt to assimilate all cognitive processing to AI-inspired computational models. If it is the correct view, linguists, AI researchers, psychologists, and neuroscientists will eventually all be in the same building on the same grant (Fodor, in conversation), for they will all be studying the structure and implementation of a single "language of thought." This vision has the advantage of providing a strong identity to cognitive science as a discipline; it has the disadvantage of yielding the implication that if the computer-metaphor is wrong, then cognitive science is ill-founded as an area of study. The cognitive capacities view fits easily with the conception of cognitive science as a federation of independent disciplines that enter into cooperative relations for investigating and analyzing various cognitive tasks. The specific degree of theoretical unification among the disciplines is left an open question, to be decided empirically in the long run.

After hearing the presentations, I've tried to discern whether each speaker's research practice corresponds to the "symbolist" or "cognitive capacities" view. Haas' talk concerns the linguistic metaphor as applied to genetics, and so may contain an implicit commitment to symbolism. Each of the other nine presentations was consistent with the cognitive capacities approach, and none showed evidence of an explicit commitment to symbolism. Viewed from the cognitive capacities perspective, all ten presentations involve a task analysis. Six of them describe or functionally characterize subprocesses within that analysis: Avrutin, Haan, Haas, Rogers, Rohrbacher, and Vainikka. One, Avrutin, examines the relation between task analysis and disruption of a cognitive system caused by neural damage.

In sum, nine of the ten presentations show no practical commitment to the symbolist paradigm. This small sample supports the conclusion that commitment to the "classical architecture" is not a standard component of the research practice of cognitive science today.

References:

Fodor, Jerry A. (1975). *The Language of Thought.* New York: Crowell.

Fodor, Jerry A. & Pylyshyn, Zenon W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition* 28: 3- 71.

Hatfield, Gary (1991). Representation in perception and cognition: Connectionist affordances. In *Philosophy and Connectionist Theory*, ed. by W. Ramsey, D. Rumelhart, and S. Stich. Hillsdale, NJ: Lawrence Erlbaum.

Kosslyn, Stephen M. & Hatfield, Gary (1984). Representation without symbol systems. *Social Research* 51: 1019-1045.

Marr, D. (1982). *Vision.* San Francisco: Freeman.

Pylyshyn, Z. (1989). Computing in cognitive science. In *Foundations of Cognitive Science*, ed. by Michael I. Posner. Cambridge, Mass.: MIT Press.