

# Why grammar needs geometry more than lambda-terms

Jens Erik Fenstad, University of Oslo

May 25, 1999

## **Abstract**

To understand how language and cognition are rooted in the behavior of large and complex assemblies of nerve cells in the brain is a major challenge to science today, and it is a challenge that cuts across a large number of disciplines.

## **Contents**

<b>1 Grammar, mind and brain</b>	<b>2</b>
<b>2 Computers and cognition: some remarks</b>	<b>5</b>
<b>3 Reductionism or emergent structures</b>	<b>6</b>
<b>4 Programmes for teaching and research training</b>	<b>7</b>

To understand how language and cognition are rooted in the behavior of large and complex assemblies of nerve cells in the brain is a major challenge to science today, and it is a challenge that cuts across a large number of disciplines.

The first task is to understand the basic structure of the brain itself, and the neurosciences have in recent times made significant progress, using a vast array of advanced techniques, including new methods of brain imaging. This has led to an increased understanding of a number of cognitive functions, including, in particular, vision and the memory system.

But cognition and understanding are more than the anatomy and the physiology of the brain. Various scanning techniques have given us precise information about the location in the brain of specific cognitive functions, and we have observed (e.g. through the study of brain damages) how changes in such areas have altered cognitive behavior. But this is often a kind of *black box* understanding; the grand challenge is to understand the precise mechanisms involved, i.e. precisely how mind emerges out of brain.

To meet this challenge the neurosciences must join forces with a large number of other disciplines, such as linguistics, cognitive psychology, logic and the computational sciences, including the new field of neuroinformatics. This is truly an interdisciplinary challenge. Thus the first task of this paper is to survey some examples of the result of this activity; the next is to see how these developments have been - or should be - translated into programmes for teaching and research training inside Academia.

## 1 Grammar, mind and brain

It is not always easy to decide who are the *true* or *correct* actors in the study of the brain and cognition. Some people talk about the neurosciences and the cognitive sciences as being separate fields of study and, hence, with different requirements as to scholarship and competence; others proclaim the advent of a new - and presumably integrated - field of cognitive neuroscience. Be this as it may, the central themes are language, vision, memory and movement. And current results seem to justify both the excitement of today and the hopes for the future.

In this report I shall restrict myself to the interplay between grammar, mind and the brain. Broader surveys can be found in e.g. Kosslyn and Koenig (1992), Churchland and Sejnowski (1992) and Crick (1994). The latter book is focused on the study of the brain and of visual awareness as a point of entry to a general understanding of consciousness. Crick's book has also a useful annotated reading list intended for the general reader.

In the study of the complex interplay between grammar, mind and the brain much of current linguistic theory postulate two main modules; one being a conceptual module which encodes the meaning content and the world knowledge of the speaker/listener; the other being a computational module which encodes the syntactical and morphological structure of utterances. Let us refer to the latter as grammatical space and the former as semantical space, or - to use somewhat more colorful language - as mind.

MIND

LOGIC

GEOMETRY

GRAMMAR

BRAIN

Let us first comment on the left line of the diagram. A major part of contemporary linguistics has been focused on the investigation of grammatical space. There is no need to review this work here; see e.g. the Handbook of Logic and Language (1997). Grammatical structures are typically represented either in tree form, as in the theory of syntactic structures developed by Chomsky (1965), or in an attribute-value format, as in the theory of lexical-functional grammar proposed by Kaplan and Bresnan (1982). But this is only part of the theory, the ultimate goal of linguistic theory is to give an account of the link between linguistic structure and meaning. This means that one must move from syntax alone to a study of the syntax/semantics interface. The link or *connecting sign* between syntax and semantics has been either a formula or term in some logical formalism, or a symbolic representation derived from some attribute-value matrix. But whatever the nature of the connecting sign the final representation has been some model structure derived from the semantics of formal logic; see Fenstad (1997). This explains why we have used the word *logic* in the above diagram to symbolize the link.

Logic is the study of the general and the abstract. A model in the sense of formal logic is a set-theoretic structure consisting of a non-empty set, the domain  $A$ , and a collection of relations,  $R_1, R_2, \dots$  defined on the domain  $A$ , where an  $n$ -ary relation  $R$  over  $A$  is nothing but a collection of  $n$ -tuples  $(a_1, \dots, a_n)$  formed from elements  $a_1, \dots, a_n$  belonging to  $A$ . Beyond this there is no further structure. In a certain sense an  $n$ -ary relation can be specified by two lists, one list of *positive* facts, viz. those  $n$ -tuples which belong to the relation, and one list of *negative* facts, viz. those  $n$ -tuples which do not belong to the relation. Anyone acquainted with computers will recognize that a finite model in the above sense is nothing but a database. And, indeed, technological applications of current linguistic theory always represent the conceptual module - i.e. semantical space - as a database. Linguistic engineering is today a significant component of many computer applications and has been well supported in the RTD Programmes of the European Union.

But natural language understanding and mind is more than database theory, one missing component is geometric structure. A database or, equivalently, a model for first order logic, is in essence an algebraic structure, and there has indeed been a fruitful partnership between algebra, arithmetic and logic. The link to geometry survived within other fields of study, in particular in measurement theory. The study of multi-dimensional scales, needed in the understanding of perception, leads necessarily to geometry. One example of this is the development of the notion of a perceptual space - of which a prime example is color space.

From the geometry of perception it is but a short step to a similar ge-

ometrization of the more general conceptual component involved in language understanding. One early example of this geometrical trend is represented by the class of mental models introduced by P. J. Johnson-Laird (1983). He demonstrated clearly how geometrical representation of knowledge combined with the use of symmetry and invariance properties lead to psychologically more plausible models of reasoning than the traditional AI approach using formal proofs from first order logic. A more recent trend is the *hyperproof programme* of Barwise and Etchemendy (1991), which is a combined system for visual and logical reasoning.

Perceptual spaces and mental models can both be subsumed under the notion of a conceptual space as introduced by P. Gardenfors (1991 and 1996). A conceptual space  $S$  is given by a number of qualitative dimensions,  $D_1, \dots, D_n$ . A point in  $S$  is a vector  $v = (v_1, \dots, v_n)$ , where each  $v_i$  is an element of  $S$ . As remarked, perceptual spaces are examples of conceptual spaces, e.g. colorspace as a cognitive construct is a three-dimensional space, where the dimensions are determined by hue, saturation and brightness. Another example of a conceptual space is the two-dimensional space generated by the two first formants of vowels frequencies.

The dimensions determining a conceptual space may either be inborn or culturally acquired. In either case the crucial fact is that each dimension comes equipped with a metric or, more generally, a topological structure. This has important applications to language understanding and allows us to transcend the formula *mind = database*, which is the working hypothesis of linguistic engineering. In formal logic a property  $P$  in a domain  $A$  is any subset of the domain. But the property RED relative to the color circle has a definite geometrical structure: it is a convex subset of the color circle, i.e. any point on the circle between two RED points is itself a RED point. On the basis of this and many other examples Gardenfors (1991) makes the proposal to identify *natural properties* with convex subsets of suitable conceptual spaces. This suggestion has an interesting application to prototype theory, which was developed as an alternative to the standard logical approach based on lists of necessary and sufficient conditions; see Rosch (1978). Natural properties can be interpreted as convex subsets of some conceptual space; some exemplars of these properties are seen to be more central to the concept and can therefore be taken as *prototypes* of the property. The extent of the property is then a convex neighborhood of the prototype. Conversely, given a set of prototypes, the properties defined by these prototypes can be obtained by a convex partition of the space; for a sample of the mathematics and the algorithms involved see Okabe et al (1992).

We have now completed the left part of the diagram, which has traced a link from grammar via formal semantics to geometric structure. We believe that this account has justified the claim made in the title of this essay. Turning next to the right part of the same diagram, the basic observation from our point of view is that neuronal activity generates a complex brain dynamics and that the associated processes lead to certain geometric structure spaces, the *phase space* or *energy surface* of the process. But left and right need to be linked, our working hypothesis is to identify the geometric structure space derived from the brain dynamics with the geometric model theory derived from the semantics of

natural language, where the convex regions associated with a natural property in conceptual space is nothing but the domain of attraction of an attractor of the brain dynamics. Through this geometrical identification there would be a seamless connection between grammar and brain. This - in brief outline - is the ultimate goal of the theory; we shall see how well it stands up to experimental scrutiny.

One possible model for the dynamics of the brain is the class of attractor neural networks of Amit (1989); see also the survey paper by Changeux and Dehaene (1989). The dynamical behavior of such systems is quite complicated, and it is possible to model non-trivial cognitive functions using such networks. Above we saw how the notion of a color space could be reconstructed as a conceptual space, where colors correspond to convex region in a suitable conceptual space. Following Rosch (1978) colors have prototypical properties. The geometry of color space is thus determined by a convex triangulation based on a finite set of prototypes. This is the theory of the left line of the diagram; moving to the right line we see that color prototypes also can be interpreted as a set of fixed points for a suitable attractor neural network. In this case the prototypes are the attractors of the system and color as a property corresponds to a domain of attraction in the energy surface of the dynamics. There is a close connection between convex geometry and the dynamics of attractors; granted sufficient regularity conditions the two accounts of color prototypes are the same.

## 2 Computers and cognition: some remarks

We shall interrupt our exposition to make a few remarks on the interaction between computers and cognition. The first remark concerns computers and the brain. The complexity of the model sketched above of how grammar and the brain are connected through the geometry of the logical model space and the geometry of the phase space of the brain dynamics clearly demonstrate that the traditional AI equation  $brain = computer$ , is no longer a useful research strategy.

But this does not mean that the brain does not compute. The model sketched above is of a common type within the sciences. Every part of it can be cast in mathematical terms, such as the algebra of syntax, the formal model theory of natural language semantics, the mathematics of convexity, and the theory of dynamical systems with the associated geometry of phase space. It is part of the power of mathematical modelling that as soon as insights about nature are captured in mathematical terms there is the possibility of extracting algorithms from equations, hence of computations. Thus the brain do compute.

Let me add a further remark on language and computations. The aim of a natural language system is to connect linguistic structure and meaning. And if the activity aims toward technological applications the link between grammatical structure and meaning must be algorithmic, i.e. we must in the end be able to construct devices and write software that in real time perform some specific linguistic tasks, typical examples are systems and software as tools for transla-

tion between natural languages. We are therefore witnessing much activity in computational linguistics, ranging all the way from computerbased systems for lexical analysis to software for computational semantics; for a survey see the Handbook of Logic and Language (1997).

Computational linguistics is in a certain sense a technological partner to our main topic of brain and cognition. In the latter we aim at modelling the actual human processes. In the former we aim to model - and hence to technologically exploit - human cognitive capabilities. The two are not necessarily the same. We may be able to build systems that perform *intelligent tasks*, but they may do so in ways that are almost unrelated to human problem solving. If technology is the aim, there is no problem. But sometimes the dividing line between technology and humans becomes blurred.

In a recent article in the newsletter *elsnet* (1997. 6.3) the author has grand visions:

I propose machine awareness as a grand technological challenge which requires contributions from the fields of Human Computer Interaction, Computer Vision, Speech Recognition, Speech Synthesis, Natural Language Processing, Learning, and Artificial Intelligence. Such a technology would integrate - machine perception (vision, sound and other sensors) to enable machines to perceive, identify and follow individuals; - speech recognition, natural language understanding, speech synthesis and reasoning, to enable machines to converse with individuals, and to obtain instructions from them; - learning and reasoning, to enable machines to adapt to individuals and to master control tasks. It is safe to say that we are currently far from achieving these goals. It is also necessary to keep in mind the distinction between the scientific understanding of human cognition and the art and technology of constructing *intelligent machines*. Indeed, there may be unbridgable gaps between machine and human *awareness*.

### 3 Reductionism or emergent structures

Returning to the model and the *seamless connection* of grammar, mind and brain that it postulates, we must recognize that the picture is much too simple as a comprehensive account. What we today know about cognitive processes in the brain clearly shows that a straightforward one-to-one correspondence between cognition and neuronal activities, as assumed by almost all current work in neural network theory, simply is not correct. PET studies have revealed some of the complexity in the correspondence between cognitive functions and neuronal architecture. As one striking example we may mention the hierarchy of tasks starting with the act of passively viewing displayed words, continuing with the gradually more complex tasks of listening, speaking and generating words: Each stage of the hierarchy activates different areas of the brain. There is therefore no simple correspondence between cognitive functions and brain dynamics, and possibly different areas with different architecture may generate similar geometries and, hence, similar meaning content.

This is one, and a rather convincing, argument for an independent phenomeno-

logical theory of mind, which means a study of the geometry without presupposing a detailed knowledge of the underlying dynamical behavior; a strong argument in favor of this position can be found in the work of Freeman and colleagues; see Skarda and Freeman (1987 and 1990). This is not an uncommon situation in the sciences. We know e.g. that thermodynamics, which is a phenomenological theory, is reducible to the underlying statistical mechanics in the equilibrium case, but not so in the non-equilibrium case. In the latter case we have to write separate equations for the phenomenological stage. Other examples of importance for our theory is Turing's study of the chemical basis for morphogenesis and the Hodgkin-Huxley theory of the propagation of signals in neurons. In both cases the facts are modelled at the phenomenological level by diffusion-reaction equations. There are - of course - no new forces or substances involved, but the equations are written directly at the phenomenological level and not derived from an underlying dynamics.

In our context a similar modelling task was undertaken by Thom (1970 and 1973). The basic geometric object for Thom is an *energy surface*, which is supposed to be derived from some underlying, but not specified brain dynamics. In the 1970 paper Thom classifies spatio-temporal verb phrases in terms of singularities of the energy surface; in the 1973 paper he develops a more general approach. It is no surprise that his discussion is compatible with the model discussed above. It is entirely in line with our discussion that a noun phrase should be described as a potential well in the dynamics of mental activities and a verb phrase by an oscillator in the unfolding space of a spatial catastrophe.

The example is specific to our discussion, but the point it makes is important for the general case of interdisciplinarity. We asserted in the introduction that brain and cognition is a topic which cuts across a large number of disciplines; in addition to the neurosciences we mentioned linguistics, cognitive psychology, logic and the computational sciences. Thus the study of cognition proceeds at many levels, but the ultimate goal is an integrated account.

Ortodox methodology of science seems to postulates a *seamless connection* between all phenomena. This is a very strong form for reductionism, but it has been a belief which has been very productive in the development of modern science; we understand and hence control by reducing the complex to the simple. But we are today witnessing a possibly radical change of metaphor; to be up-to-date you have to speak of emerging structures and properties, which are not necessarily reducible to properties at a more elementary level; for a general discussion of non-linear dynamics and emerging structures see Scott (1995). There is much to discuss in this connection. For our immediate purpose we need to remind the reader that interdisciplinarity sometimes presupposes a reductionism of phenomena that is not always warranted.

## 4 Programmes for teaching and research training

Advice is better built on experience. My involvement with logic, language and computation began in 1976. In the Fall Term of that year a psychologist, a philosopher, a linguist, and a mathematical logician at the University of Oslo

decided to organize a seminar together on language as seen from their various perspectives. This was an interdisciplinary venture, the topic turned out to be fashionable, and we had an overflow audience - at least until the mathematical logician started to *explain* Montague grammar and Higher Order Intensional Logic. After that our audience was reduced to a more comfortable size. But enough interest had been generated for the seminar to continue. It has served as a useful meeting place; the *hard core* has always been a group of linguists and logicians, but from time to time we have also had the participation of philosophers, psychologists, and computer scientists.

We have some experiences to share. It is nice to be fashionable and to meet across disciplines discussing *important topics*. But durable cooperation must be based on do-able problems. And solving problems usually needs a variety of tools and experiences. This is elementary but nevertheless essential. It also has structural implications. You have to learn your trade properly; tools need to be sharp and experiences need to be broad. This means that you must not compromise on basic training. One of our experiences is that you should not mix disciplines too early. Joint basic mathematics courses for mathematicians and linguists are not a good idea; you are much too easily driven to the least common denominator. One has to learn the trade and then be willing to travel. In Norway we have had a long history of migration from mathematics to fields as diverse as economics and the exact geophysical sciences - and it has worked well. We seem to have repeated this experience with the current programme in logic, computation and linguistics.

This is but one experience, and I shall be cautious in making too bold generalisations; we need to share experiences on how to create stimulating and productive environments for interdisciplinary research. In our case the story has so far ended well. An interdisciplinary seminar has grown into an academic programme. At the University of Oslo there is now a special degree programme inside the Department of Linguistics on Language, Logic and Information. The focus is on theoretical, data intensive, and computational linguistics, but the heritage from mathematics is not entirely forgotten - three of the permanent staff have their basic degree in mathematics.

Does this and similar experiences indicate that we are witnessing the emergence of a new discipline - the science of cognition? This is not necessarily so. The Oslo experience can be interpreted as an enrichment of linguistics, adding new tools of modelling and computation to an existing field of study. Moving from cognition to computation, I have since the late 1960s witnessed and partly participated in the *construction* of computer science as an independent academic discipline, but I have also seen - and still see - the tension between the *engineering* and the *formal* parts of the enterprise. A very specific statement was implicitly made about the nature of the computational sciences when Computer Science at Stanford University moved from the School of Arts and Sciences to the School of Engineering. In comparison the Cognitive Sciences are in a much more nascent stage. Indeed, we still may have our doubts if they ever will coalesce into one science; recall the discussion above on reductionism versus emerging structures. But whether one or many sciences, the understanding of brain and cognition is the grand challenge. And this is why grammar needs

geometry more than lambda-terms.

## REFERENCES

Amit, D.J. (1989) *Modelling Brain Function*. Cambridge University press, Cambridge, UK.

Barwise, J. and Etchemendy, J. (1991) Visual information and valid reasoning, in W.Zimmermann and S. Cunningham (eds.), *Visualisation in Mathematics*, Math.Ass.America, Washington D.C.

Changeux, J.P. and Dehaene, S. (1989) *Neuronal Models of Cognitive Functions*, *Cognition* 33.

Chomsky, N. (1965) *Aspects of the Theory of Syntax*, MIT Press, Cambridge, Mass.

Churchland, P.C. and Sejnowski, T.J. (1992) *The Computational Brain*, MIT Press, Cambridge, Mass.

Crick, F. (1994) *The Astonishing hypothesis*, Touchstone Books, London, UK

elsnews (1997), *The Newsletter of the European Network in Language and Speech*, 6.3, Utrecht University, Utrecht, The Netherlands.

Fenstad, J.E. (1997) *Formal Semantics, Geometry and Mind*, in X.Arrazola et al (eds.), *Discourse, Interaction and Communication*, Kluwer Academic Publishers, Dordrecht, The Netherlands

Gardenfors, P. (1991) *Framework for Properties*, in L.Haaparanta et al (eds.), *Language, Knowledge and Intensionality*, Acta Philosophica Fennica, Helsinki.

Gardenfors, P. (1996) *Conceptual spaces as a framework for cognitive semantics*, in Clark, A. (ed.) *Philosophy and Cognitive Science*, Kluwer Academic Publishers, Dordrecht, The Netherlands

*Handbook of Logic and Language* (1997), Elsevier, Amsterdam, The Netherlands

Johnson-Laird, P.J. (1983) *Mental Models*, Cambridge University Press, Cambridge, UK

Kaplan, R. and Bresnan, J. (1982) *Lexical Functional Grammar*, in J.Bresnan (ed.) *The Mental Representation of Grammatical Relations*, MIT Press, Cambridge, Mass

Kosslyn, W. and Koenig, O. (1992) *Wet Mind*, The Free Press, New York, N.Y.

Okabe, A., Boots, B. and Sugihara, K. (1992) *Spatial Tessellations, Concepts and Applications of Voronoi diagrams*, J.Wiley, New York, N.Y.

Rosch, E. (1978) *Prototype Classification and Logical Classification*, in E. Scholnik (ed.), *New Trends in Cognitive Representation; Challenges to Piaget's Theory*, Lawrence Erlbaum Ass., Hillsdale, N.J.

Scott, A. (1995) *Stairway to the Mind, The Controversial New Science of Consciousness*, Copernicus, Springer-Verlag, New York, N.Y.

Skarda, C.A. and Freeman, W.J. (1987) *How Brains Make Chaos in order to Make Sense of the World*, *Behavioral and Brain Sciences*, 10.

Skarda, C.A. and Freeman, W.J. (1990) Chaos and the New Science of the Mind, Concepts in Neuroscience, 1.

Thom, R. (1970) Topologie et Linguistique, in Essays on Topology, Springer-Verlag Heidelberg.

Thom, R. (1973) Langage et catastrophes: Eléments pour une Sémantique Topologique, in J.Palis (ed.), Dynamical Systems, Academic Press, New York, N.Y.