Chapter 1

MINIMISING OR MAXIMISING STORAGE? AN INTRODUCTION

Sieb Nooteboom,
Fred Weerman,
and Frank Wijnen

1. Why this volume?

"The lexicon is really an appendix of the grammar, a list of basic irregularities". These words, written a long time ago by Leonard Bloomfield (Bloomfield, 1933, p. 274), have set the stage for focusing linguistic and psycholinguistic research on the compositional nature of linguistic objects, a view that has culminated, from the mid-fifties until today, in generative linguistics as proposed and elaborated by Noam Chomsky (Chomsky, 1957; 1965). This attention for the computational nature of language over many decades of linguistic investigation has proven to be tremendously fruitful. Without it, our insight in the nature of human language would be much smaller than it is. Inevitably, the continually high level of attention for computational rules has led to some neglect of the possibility of massive storage not only of irregular but also of regular linguistic objects. The idea for this volume was inspired by the observation that a growing number of linguists and psycholinguists are dissatisfied with the traditional Bloomfieldian idea that the lexicon is a list of irregularities, and all other linguistic objects are computed by rules. So the question appeared to be: if Bloomfield's view was wrong, or at least not the full truth, what then is stored in the lexicon, and what is computed by rule?

2. Why storage and computation?

When we asked a number of linguists and psycholinguists to write a chapter for this volume, reflecting on their own attempts to deal with the issue of storage and computation in the language faculty, we soon discovered that our original question can be approached in very different

S. Nooteboom et al. (eds.), Storage and Computation in the Language Faculty, 1–19.

ways. It can be applied to rather disparate aspects of the language faculty, and attempts to answer it may lead to widely different views of how the language faculty is organised and how it goes about doing its job. By way of introduction, we will briefly discuss those issues that in our perception are most relevant to what the reader will find in the different chapters of this book. Let us begin with the notions storage and computation themselves. Why would there be a need for this opposition in accounting for the human language faculty? As an analogy, one may think of the human capacity to calculate ("calculate", of course, is just another word for "compute"). The primitives of calculation are numbers. Calculation or computation is applying rules to these numbers. We need computation because there may be good reasons to want to know how much, for example, is 638 x 823, and the outcome cannot be found in a lexicon. This is similar to the need for computation in linguistics. Without computation we cannot account for our prolific capacity to generate new words, new phrases, new sentences, new interpretations, and new discourses.

Although we can imagine there is only one primitive number, the number one, and all others are derived by addition, subtraction, multiplication, division, raising to a higher power, root extraction etc., we know this is not how it works. In primary school we have learned the tables of multiplication by heart, and few of us have to do any computation to know how much is 7 times 9. So there are some real questions: Which simple and less simple "calculations" are stored, together with their outcomes, which outcomes have to be computed, and what are the strategies in doing these computations? Of course the answers to these questions may differ for different individuals, but it is of interest that it has been convincingly shown that so-called "calculating prodigies", the kind of people who tell you within a few seconds how much is 638 x 823 without using a pocket calculator, do not differ from us normal mortals in their capacity to do mental computation. What they differ in is in their memory for numbers and for outcomes of calculations. They have at their disposal a vast lexicon of outcomes of calculations plus a limited number of algebraic tricks for reducing complex numerical questions to simpler or more familiar ones (Dehaene, 1997). Their minds maximise storage and minimise computation. Of course, all 'normal' people are prodigies in using language. That is part of what makes us human. Is it possible that our minds also maximise storage and minimise computation in constructing words, phrases, sentences, and discourses?

Whereas linguists generally account for creativity in the language faculty by combinatorial rules and variables, connectionists believe that the creative aspect of language can be wholly accounted for by pattern association. Unfortunately, our attempt to include at least one chapter in this volume representing the connectionist view, has failed. Yet it may be worthwhile to discuss briefly the opposition between the rule-based approach and the neural-network or connectionist approach. It is not difficult to see why connectionists prefer pattern association to rules. They are inspired by theories about the functioning of neurons and networks of neurons, and they implement their theories of language in computational models of such neural networks (Elman, 1993; Elman et al., 1996; Fodor and Pylyshyn, 1988). Now neural networks are very good at pattern association, but it is not easy to see how abstract rules using variables could be implemented in such networks. The problem is that rules can be and are applied to objects that are widely different in all respects except in the single, possibly minute, feature that makes the rule applicable. It is a bit of a problem for connectionists that humans behave as if they use rules with variables, and currently a number of attempts are being made to bend neural networks every which way to simulate rule-like behaviour without using rules. It would seem that some 30 years of work in this domain has not yet produced a solution to this problem. Nonetheless, progress has been made, in the sense that it has become clear that the classic models — multilayer networks trained by backpropagation or similar supervisory techniques — which were once introduced as the answer to most if not all problems in cognitive science, are principally incapable of learning generalized operations on variables (Marcus, 1998). This is, however, certainly not the end of an episode. As Ray Jackendoff (this volume) and others note, other classes of models still hold the promise of successfully implementing rule- and variable-governed cognitive capacities, including language. Most likely, further developments along these lines will put the storage–computation issue in a different light.

The connectionist challenge did have a considerable impact because most of us are convinced that rules and variables are psychologically real. How, otherwise, could we learn to do simple and precise calculation, how could we learn to play chess or checkers, how could we learn to do formal logic, consisting of rule-governed operations on abstract symbols that have *no* preset associations with reality? Furthermore, there are good reasons why we have this capacity for applying rules. Steven Pinker, in his recent book "Words and rules" (Pinker, 1999), convincingly argues that in order to make optimal sense of the world around us, we need both "family resemblances" between objects, and rule-governed "good categories" that allow us to make valid predictions about aspects of objects that are not readily observable. The reason is that the world is structured that way. An animal that looks very much like a dog, may well be a dog, but an animal that does not at all look like an eagle, may still be a bird. Sometimes it is advantageous to go for resemblances, sometimes it is advantageous to generalise from single features.

Note that rules with variables may well be older than language. It seems most likely, however, that the use of combinatorial relationships within sets of symbols, most easily described in terms of rules and variables, as in language or calculation or chess, is typically human (Deacon, 1997). Nonhumans primarily work with "family resemblances". Again, the world of numbers can serve as an analogy. We humans can do precise calculation. Seven times 9 is not 62, but 63. Other animals like rats can be shown to estimate numbers, but they do not do precise calculation. Their number sense works in an analogue way, the number of things or events being represented as some value on a continuous scale. On such a scale 1 and 2 may be very different, but 5 and 6 very much the same (Dehaene, 1997). Pattern association does the trick. Digital computation, where 62 definitely is not the same as 63, where a rose is a rose is a rose, and where the plural of "wug" is "wugs" and never anything else, is a very late invention in evolution. Apparently, brains had to become very complex before they were able to work with even very simple combinatorial systems. No wonder that connectionists have such a hard time simulating this. Nature also took its time to evolve combinatorial symbol systems, and, as far as we know, did this only once. It should not surprise us, then, if humans do not rely on computation alone whenever they can avoid it. The brain has a massive capacity for storage, and it is reasonable to assume that processing of all kinds is supported by stored information whenever possible.

In his 1999 book, Steven Pinker cites extensive psycholinguistic evidence for the 'psychological reality' (to use the old-fashioned term) of both computation and storage-based processes in language. From his overview, the conclusion seems inevitable that the use of regular verb morphology (in English, but in principle in all human languages) rests on a rule-and-variable based computational mechanism, whereas irregular verb forms are, basically, stored. Many of the chapters in this volume provide corroboration for the 'dual mechanism' hypothesis, and, more importantly, extend it beyond the domain of verbal morphology. It is argued, and to our minds convincingly, that computation *and* storage (as defined by Pinker) are intrinsically involved in such diverse areas as syntax, phonology, and second language acquisition, and the strong implication is that it is no longer useful to think of the language faculty as being exclusively associated with either computation *or* storage.

3. Setting the stage

Jackendoff's opening chapter sets the stage for exploring the question of whether the human language faculty maximises computation or maximises storage by sketching a linguistic theory that would certainly allow the latter view. He observes that we cannot equate lexical units with words. We produce many words that are not lexical units and there are many lexical units both below and above word level. The latter include, of course, idioms, but also certain types of constructions in which all words are exchangeable for others, such as English resultative constructions exemplified by "Mary watered the tulips flat". This suggests that grammatical constructions, not containing a single individual word, are stored in the lexicon. In line with adherents of "Construction Grammar" Jackendoff then tentatively reduces his account of the creative aspect of language to a single rule: UNIFY PIECES, all pieces being stored in a common format that permits unification, and all necessary variables being specified in the lexical items and not in the rule.

This suggestion by Jackendoff not only sets the stage for discussing what is in the lexicon, but also for discussing the nature of computation. At one end of the spectrum we find complex rules containing many variables, at the other end of the spectrum we are asked to believe that there are no rules containing variables at all. The reader may note that, in Jackendoff's account, many phenomena that are traditionally used as arguments for computation by rule are explained by stored structures plus a very simple operation of free combination.

4. Accessing regular and irregular word forms

As soon as we have acknowledged the possibility that the lexicon is not only a list of exceptions, but can also contain regular constructions, we are faced with the possibility that retrieval from storage and computing by rule do not exclude each other. Is it possible that, when we need a particular linguistic object during processing, we simultaneously attempt to retrieve this object from storage and compute this object by rule? This sounds very inefficient, but the reader should note that the capacity of the human brain is such that efficiency in the sense of involving as few neuronal structures as possible is not called for. Current theories of how the brain might go about doing its many jobs emphasize the enormous redundancy in neuronal processes (Edelman, 1992; Gazzaniga, 1992; Calvin, 1996; Damasio, 1999). If we interpret efficiency in terms of speed of processing, the brain may be wise to follow more routes simultaneously and see which is fastest. Moreover, it is not inconceivable that accessing stored items and computing these very same

http://www.springer.com/978-1-4020-0526-8

Storage and Computation in the Language Faculty
Nooteboom, S.G.; Weerman, F.; Wijnen, F.N.K. (Eds.)
2002, 356 p. 3 illus., Hardcover
ISBN: 978-1-4020-0526-8