Words and Intelligence I

Text, Speech and Language Technology

VOLUME 35

Words and Intelligence I
Selected Papers by Yorick Wilks

Edited by

Khurshid Ahmad
Trinity College, Dublin, Ireland

Christopher Brewster
University of Sheffield, UK

Mark Stevenson
University of Sheffield, UK

Springer

A C.I.P. Catalogue record for this book is available from the Library of Congress.

ISBN-10 1-4020-5284-7 (HB)
ISBN-13 978-1-4020-5284-2 (HB)
ISBN-10 1-4020-5285-5 (e-book)
ISBN-13 978-1-4020-5285-9 (e-book)

Published by Springer,
P.O. Box 17, 3300 AA Dordrecht, The Netherlands.

www.springer.com

Printed on acid-free paper

Contents

Preface

Professor Yorick Wilks has contributed to a wide range of academic fields including philosophy, linguistics and artificial intelligence. The main focus of his work has been the fields of computational linguistics and natural language processing where his work has advanced an unusually wide range of areas such as machine translation, word sense disambiguation, belief modelling, computational lexicons and dialogue modelling. One of the distinguishing features of his work has been his ability to link the engineering of practical text processing systems with more theoretical issues about language, particularly semantics. A number of themes have run through his work and one of the aims of this volume is to show how a body of work on such a diverse range of topics also forms a coherent program of inquiry. A comprehensive record of the range and diversity of Yorick's output is beyond the scope of this volume. Rather, as part of the Festschrift organized to honour his retirement from teaching, we chose this volume to contain a selection of representative pieces including some less accessible papers.

The first paper we have chosen to include ("Text Searching with Templates") is surely one few will be familiar with. This was published as a technical report in 1964 at the Cambridge Language Research Unit, where Yorick first worked on Computational Linguistics. In this paper Yorick outlines an approach in which texts are represented using template structures and world knowledge, in the form of an interlingua, used to define the elements which could be combined into meaningful units. Later these ideas were developed and incorporated into his work on Preference Semantics.

The next paper ("Decidability and Natural Language"), published in the philosophy journal Mind, is a theoretical discussion of whether it is possible to represent the semantics of natural language in any computable way. Here Yorick argues against the accepted belief at the time that the syntax and semantics of natural language utterances should be treated independently, proposing that semantics is not an extensions of syntax but rather the other way round. He also addresses the question of whether a deterministic procedure could ever be developed to decide whether a sentence is meaningful and suggests that a suitable criterion might be whether a single interpretation of the sentence can be identified. In this paper, Yorick discusses a theme which he returns to several times: that the possible meanings of a particular word can only be defined relative to a particular sense inventory and cannot be thought of as abstract, Platonic entities.

The next paper ("The Stanford Machine Translation and Understanding Project") represents Yorick's important contribution to Machine Translation and provides detail of the English-French translation system he worked on at Stanford University. Yorick discusses how the latest advances in linguistics, particularly semantic analysis, could be used to justify another attempt at the MT problem (this paper was written only a few years after the 1966 ALPAC report damning machine translation). He also shows how these ideas could be implemented in a practical way by describing a system which made use of an interlingua approach and analysed the input text by transforming it into template structures similar to those introduced in the first paper.

One of the main outcomes of the Stanford project was Yorick's influential Preference Semantic system, various aspects of which are detailed in three of the papers ("An Intelligent Analyser and Understander of English", "A Preferential, Pattern Seeking, Semantics for Natural Language Inference" and "Making Preferences More Active"). The first paper provides an introduction and shows that, contrary to standard approaches of the day, syntactic and semantic analysis could be carried out in parallel. Preference Semantics is based on the use of selectional restrictions but, rather than treating them as constraints which must be satisfied, they were interpreted as paradigm cases, indicating normal or prototypical word usages which may be expected but could be adapted if necessary. The system represented the preferences using a set of semantic primitives which were also used to represent the possible meanings of each word (called "formulas"). These were combined, and their preferences examined to choose the correct meaning, resulting in a template representing the meaning of the text. The next paper explains how Preference Semantics can be extended to carry out reasoning about texts to perform anaphora resolution. In keeping with one of the main motivations behind Preference Semantics, that a language understanding system should always attempt to provide a usable interpretation, the approach attempted to resolve a wide range of anaphora. The system would make a best guess about the meaning of an utterance, as a human does, and act accordingly. Further experience, gained through additional knowledge about the situation, may suggest a change in interpretation but to carry out many language understanding tasks, including machine translation, requires some commitment to a preferred interpretation. The final paper on Preference Semantics provides more details about how the flexibility of the system can be used to interpret a wide range of usages. Yorick points out that word usages which are often thought of as metaphorical are common in everyday language and that the burden of interpretation should be placed on the language understanding system. Yorick argues that formal theories proposed by linguists were not flexible enough to describe the sort of language used in everyday situations. Yorick motivates this with the now famous example: "My car drinks gasoline". Yorick advocates the use of world knowledge to interpret metaphorical language, in this case we need to know that cars require the insertion of a liquid (petrol or gasoline) to run.

Preference Semantics relied on a set of semantic primitives to denote the typical, or preferred, usages although their use had been questioned. In the next paper ("Good and Bad Arguments about Semantic Primitives"), Yorick replies to these criticisms.

The main questions posed were what semantic primitives actually meant, and where these semantics derived from. Yorick proposes a position where primitives can be though of as part of the language whose semantics they represent. They form a set of building blocks within the language from which more complex statements can be formed by combination. Once again, Yorick argues that the meaning of language is found within the language itself.

Yorick's position in that paper is a theoretical one which is made practical in the next paper we selected ("Providing Machine Tractable Dictionary Tools"). This paper introduces Yorick's extensive work with Machine Readable Dictionaries (MRD) by describing several methods for exploiting the information they contain which had been developed while he led the Computer Research Lab of New Mexico State University. The first technique, the use of co-occurrence statistics within dictionary definitions, is a very different approach from Yorick's previous work on Preference Semantics and allows meaning to emerge from the dictionary definitions in an automated way. Another technique described in this paper concerns the conversion of a MRD into a full Machine Tractable Dictionary, that is a resource in which the terms used to define word senses are unambiguous and so can be readily understood by a computer. This represents a computational implementation of Yorick's view of semantic primitives. One of the main goals of this project is to identify a core set of basic terms which can be used to provide definitions and these were also identified through automatic dictionary analysis. A final application for the dictionary was to automatically generate lexical entries for a Preference Semantics system which provided a method for avoiding the bottleneck caused by the previous reliance on hand coded formulas.

The next paper ("Belief Ascription, Metaphor and Intensional Identification") represents Yorick's work on belief modeling and dialogue understanding which were implemented in the ViewGen system. His work on this area builds upon the techniques developed for understanding metaphors within the Preference Semantics framework.

In the paper entitled "Stone Soup and the French Room" Yorick returns to the topic of Machine Translation to discuss IBM's statistical approach. He is characteristically skeptical of the claims being made and controversially suggests that purely data-driven approaches could not rival mature Al-based techniques since the later represent language using symbolic structures. Yorick makes sure to point out that he does not oppose empirical approaches to language processing by reminding us that "we are all empricists" and also suggests that the roots of the statistical approach to translation could be traced back to some of the earliest work on computational linguistics. The collective memory in language processing is often short and it is important for researchers to be able to be reminded of earlier work may have been forgotten all too quickly. To a great extent Yoricks claims have been proved by recent work on statistical machine translation. During the decade or so since this paper was published work on statistical machine translation has gradually moved towards the use of increasingly rich linguistic structures combined with data derived from text.

In the final paper ("Senses and Texts") Yorick discusses recent work on semantic analysis, specifically two contradictory claims: that the word sense disambiguation problem cannot be solved since it is not well formed and another that suggested the problem had, to a large extent, been solved. Yorick points out that the notion of what is meant by "word sense" is central to these arguments but that it has not yet been adequately defined and, besides, is only meaningful relative to some specific lexicon. One of the claims Yorick discusses rests on the assumption that Computational Linguists had made naïve assumptions about the nature of meaning and he, once again, reminds us to looked to the past; "In general, it is probably wise to believe, even if it not always true, that authors in the past were no more naïve than those now working, and were probably writing programs, however primate and ineffective, that carry out the same tasks as now." Yorick points to one of the motivations behind Preference Semantics, namely that any adequate language understanding system must accommodate usages which are different from the meanings listed in the lexicon but somehow related, as in metaphorical utterances.

Khurshid Ahmad
Christopher Brewster
Mark Stevenson

Origin of the Essays

All permissions granted for the previously published essays by their respective copyright holders are most gratefully acknowledged.

1. Wilks, Y. (1964) “Text Searching with Templates” Cambridge Language Research Unit Memo, ML. 156
2. Wilks, Y. (1971) “Decidability and Natural Language” *Mind.* vol. LXXX pp. 497–520.
3. Wilks, Y. (1973) “The Stanford Machine Translation and Understanding Project” In R. Rustin (ed.) *Natural Language Processing*, pages 243–290. Algorithmics Press, New York.
4. Wilks, Y. (1975) “An Intelligent Analyser and Understander of English” *Communications of the ACM* 18(5):264–274.
5. Wilks, Y. (1975) “A Preferential, Pattern Seeking, Semantics for Natural Language Inference” *Artificial Intelligence* 6:53–74.
6. Wilks, Y. (1977) “Good and Bad Arguments about Semantic Primitives” *Communication and Cognition* 10(3/4):181–221.
7. Wilks, Y. (1979) “Making Preferences More Active” *Artificial Intelligence* 11:197–223.
8. Wilks, Y., Fass, D., Guo, C-M., McDonald, JE., Plate, T., Slator, BM. (1990) “Providing Machine Tractable Dictionary Tools” *Machine Translation*, 5(2):99–151.
9. Ballim, A, Wilks, Y., Barnden, J. (1991) “Belief Ascription, Metaphor and Intensional Identity” *Cognitive Science*, 15(1):133–171.
10. Wilks, Y. (1994) “Stone soup and the French Room” In A. Zampoli, N. Calzolari, and M. Palmer (eds.) *Current Issues in Natural Language Processing: In Honour of Don Walker* pp. 585–595.
11. Wilks, Y. (1997) “Senses and Texts” *Computers and the Humanities* 31(2):77–90.

1

Text Searching with Templates

Yorick Wilks
Cambridge Language Research Unit

1.1 Introduction

A 'template' (to be abbreviated as 'T') is a sequence of atomic marks which are either

(1) a left-hand *bracket* or a right-hand *bracket*
(2) a connective written ":" or "/"
(3) an *element* which is a member of one of the lists (p. 6) of elements selected from the elements of the C.L.R.U. interlingua "Nub"

and which are combined in the form

$$(X : p) : (m\ v\ /\ n\ (Y : q)1)$$

where p, v, q are elements chosen from given lists. X, Y, l, m, n are variables whose substitution values are *formulae* of the interlingua. A formula is defined recursively in the following manner;

(1) every element is a formula
(2) every formula not itself an element is of the form X b Y where X, Y are formulae and b is a connective.

l, m, n are only mentioned here for completeness, in that we may require their presence in the basic formula above when considering more complexities than we do now. They will not be mentioned again here, and we will therefore consider the basic formula in the simpler form.

$$(X : p) : (v\ /\ (Y : q))$$

This is itself a formula on the above recursive definition. We shall refer to X and Y as *text* variables, and p, v and q as *template* variables.

K. Ahmad, C. Brewster and M. Stevenson (eds.), Words and Intelligence 1, 1–7.

We give as section (1.5) of this note mutually exclusive lists of plausible interlingual elements as substitution values for p and/or q and for v respectively. Not all the values of the above formula, which are obtained by substituting each of the members of its appropriate list for each of p, q and v in turn, are T's.

We give next in (1.5) a table showing which of the concatenated pairs (p q), formed from this class of values for p and q, can form part of a 'permitted' T.

These lists are preceded by an outline of a preprogram (section (1.4)). This would operate upon a text coded in the interlingua, which we shall call the datum-text. A digression is required at this point on the construction of this datum-text.

1.2 The Participation of the Interlingual-dictionary Entries in the Datum-text

It is assumed that the text is provided in structured interlingual form (whether of nubs[1] or of full entries, for all the words of the original message or only for certain 'key' words is discussed below).

i) The interlingual dictionary is made by giving more than one formula, in general, for each English word; which of these are we to insert in the datum-text; since the interlingual syntax allows for only one?

   In general I think we must hope for a large enough computer to be able to insert each of the formulae in turn and subsequently to choose the 'best' output. We shall require some criterion for the best output even without this – at least, we shall until the T's are far less crude than they are at present.

ii) The question then arises as to which *part* of these formulae "participate" in the datum-text?

   If we insert only the 'nubs' we will presumably not then be able to read off semantically satisfactory values of the 'text-variables' (see below) from the processed data text. However, providing we *can* re-associate each nub with the appropriate natural language word we could leave the 'non-nub' part of each formula in the store and change the present syntactic form of a T by dropping the "bracket and connective restrictions". On the other hand it has yet to be shown that inserting the whole formula does produce *unmanageably many* contours.

(i) and (ii) are independent in so far as ambiguity is part-of-speech ambiguity (in which case interlingual entries for an English word have, in general, different nubs, and *this* ambiguity can be resolved fairly simply by syntactic considerations), which is to say that if we insert only nubs in the data text problem (i) *will not arise* at this stage.

[1] "nub" is defined in the interlingua "Nub" (a variant of "Nude") now being constructed by J. Burns. Essentially a nub is the last element of the interlingual formula for a word which also fixes some of its syntactic properties.

1.3 Contours

The program outlined in (1.4) first divides the datum-text into *strings*. A string contains the formulae corresponding to a phrase or clause in the natural language text, and the latter is divided into phrases and clauses by a modified form of the C.L.R.U. Syntax-Bracketting program. Within each string the constituent formulae are re-ordered, if necessary, to make the syntax of the string conform to the rules of the interlingua. It should be possible to stipulate a maximum length for such strings in terms of interlingual elements. The aim of the program is to match each of the strings on the datum-text with the inventory of T's that is stored in the machine as concatenated triples of elements. Output for each string is given as a list of "matched and augmented (i.e. X & Y made explicit) T's" or "*contours*". A contour is a T in which X and Y are replaced either by formulae or by nothing at all (in which case the connectives preceding the 'p' and 'q' elements are deleted also). We say there is a contour in a given string if the following marks occur in it in left-right order:-

$$p): v/q) \tag{1}$$

where p, v, q represent any of the elements in the appropriate lists (p. 6). Each pair of marks in the contour, except those forming part of the values of X or Y, may occur in the string in such a way as to be separated by other marks except that;

the 'p' element must be immediately followed by the '):'
the 'v' element must be immediately followed by the '/'
the 'q' element must be immediately followed by the ')'.

The list 1 above does not represent a necessary condition since we shall recognize cases in which the marks corresponding to values for the X, Y, v, /, marks do not occur. We actually *define* the occurrence of a *contour* in terms of a subset of the marks 1, as follows. We shall say that there is a contour in a string if it contains a substitution value for a "p):" mark followed, but not immediately followed, by a substitution value for a "q)" mark i.e. by a *pair* not a *triple*. This last stipulation requires some justification.

When searching for contours we might *treat as T's* the list of permitted concatenated triples (p v q) which we will call *full templates* by contrast with the (smaller) set of permitted concatenated pairs, or blank templates (p q). But there are troubles about this which can be readily illustrated by using the symbol '$\leftrightarrow$' (whose negation is '$\nleftrightarrow$') to mean 'can be written for' which can be roughly interpreted as 'means the same as' i.e. it denotes a symmetrical relation.

Now it is clear from the structure of the given basic T, (and will be quite transparent from the lists of suggested values for p, v, & q below), that the T's are intended to have a 'sentential feel', (though we cannot of course assume in advance that the *strings* will correspond to sentences in the text), p corresponding roughly to a subject, v to a verb, and q to an object or complement.

Generally, *in Anglo-Nude*, given values from the appropriate classes for p, q & v it is the case that

a. (p v q) ↔ (p q v) *but*
b. (p v q) ↮ (v p q) [2]
c. (p v q) ↮ (v q p)
d. (p v q) ↮ (q p v)
e. (p v q) ↮ (q v p)

Any doubt about (a) interpreted as a true statement about English word-order can be laid by *making it a necessary condition* that for a given (p′ v′ q′) to be admitted as a full template that (a) be true for *those* values of p, v, or q.

Now when a matching program is given a definite form, since the string must be searched from either the left hand end or right hand end in written order, and since the class of possible values is to be the same for p and q, equivalences rejected by (c), (d), and (e) can never arise. Since the program as given cannot distinguish such forms from forms occurring in (a) or (b). So it will be seen from (a) and (b) that the set of permitted triples is a subset of the set of possible triples, it is in fact that subset in which the 'v' separates or follows the members of the pair. We may then (having located the pair), look for a 'v' element to make it a *full template*, searching first between the elements of the pair, (since such an ordering is the more common construction in English), than after the pair in the given string. This means no more than that each triple may be considered as being written twice in the inventory, but in two forms. If no 'v' element appears in the search we record the result as a pair or '*blank contour*'.

When such a pair has been located in a string we say that the string has a substructure; we can then *read off* the values of the text variables X, Y as those elements bracketed with the p and q elements so located.

Two other things should be noted at this point. It will be seen that there may be many potential contours 'in' any given string which satisfy these requirements, and we shall require that the program locates all the contours in it.

1.4 The Form of the Program

A program to locate a single contour in a given string would operate as follows. Each element in the string would be examined in turn starting from the 'left-hand end'. In each case the examination would consist in enquiring whether the element occurred on the "p or q" list, if not we pass on to the next element (moving right on the paper); if it does so occur, we examine the next elements to the right in turn to see if one of them also occurs on this list, if one does i.e. *this* pair is designated '1' in the table on p.6 we say we have a *blank template*, if not we pass on. *For*

[2] We might refer to (b) as the "re-ordering function Q" for 'question-templates', though we shall not make use of this feature here.

each such blank template a further search for a "v" element is made as follows: the elements between the "p" and the "q" element in question are examined in *left to right* (reverse) order to see if they occur on the "v" list, if one does, *and* the triple is on the T list (p.7), we have a full template, if not we proceed to examine (again in "left to right" order) these elements in the string to the left of the "q" element and match with the 'T' list as before.

A program to extract *all* the contours would be one that continued in this way i.e. proceeding as before but *as if* the *matching* triple just found did not so match, after registering each success. It is not difficult to see that for an exhaustive search of a string of n elements the total number of comparison operations with the T list is:

$$\sum_{r=1}^{n} (n-r)\ (n-r-1)$$

which is:

$$n^2\ -\ 1\ +\ \sum_{r=1}^{n} r^2\ +\ (1-2n)\ \sum_{r=1}^{n} r$$

or:

$$\frac{n^3+6n^2+2n-3}{3}$$

i.e. for a string with 12 elements the number is 871.

A routine would be needed to ensure that the "same" contour, even if located many times, was only recorded once.

Each T may be negated by placing a 'NOT' element before its 'v' element. (In this context we ignore the logical distinction between internal and external negation, which requires separate discussion).

In order to record this in the output we should have a routine which enquired in the case of each 'accepted' "v" element whether or not the preceding element was a "NOT".

If *all* the T's are extracted it might turn out that one of them in a given string corresponded to *textual* features that could be said to be "contained in" the text corresponding to another T. But it would, I think be misleading to think of this as one template "being a value of another" or anything like that.

It must be emphasized that we cannot complain of such a program that it locates the *wrong* contours. What we *can* ask of any output are answers to the following questions:

1. Are the located contours useful for some defined purpose?
2. Is the list of templates used intuitively adequate?
3. Does it correspond with, or at least include, the list of templates found by other methods; from say, a list of contours located by phonetic stress-point analysis?

A separate note is required to indicate how this program might be achieved with punched cards, and what the appropriate method of coding both of the interlingua-word entries and of the template forms would be.

1.5 The Template List

The T's on this list are for some part wholly arbitrary, but this need not prevent the program being carried through and the value of the output assessed. The list of templates will require ultimate justification in much the same way as the choice of the basic elements and the individual word formulae. The two lists or the values of p or q, and v form a restricted sub-interlingua so that some of the elements in the templates above have idiosyneratic, though consistent, senses. These might well repay further study.

Values for p or q	*Values for v*
DO	PLEASE
FOLK	SENSE
KIND	BE
MAN	FEEL
PART	HAVE
WHOLE	USE
THING	WANT
STUFF	CAUSE
GRAIN	CHANGE
WORLD	PAIR
SIGN	

We may thus display the 121 possibilities of ordered pairs in the following table.

In it the elements written on the left hand side of the page represent the occurrence of the first element of the pair considered, those written vertically the second. An "O" indicates that the pair represented by that position is not a permitted pair; a "1" indicates that it is.

	DO	FOLK	KIND	MAN	PART	WHOLE	THING	STUFF	GRAIN	WORLD	SIGN
DO	1	1	0	1	0	0	1	1	1	0	1
FOLK	1	1	0	1	0	1	1	1	0	1	1
KIND	1	1	1	1	1	1	1	1	1	1	1
MAN	1	1	0	1	1	1	1	0	1	1	1
PART	0	1	1	1	1	1	1	1	0	1	1
WHOLE	0	0	1	1	1	1	1	0	0	1	1
THING	1	1	0	1	0	1	1	1	1	1	1
STUFF	1	1	0	1	1	1	1	1	1	1	1
GRAIN	0	1	1	1	1	1	1	1	1	1	1
WORLD	1	1	0	1	1	1	1	1	0	1	1
SIGN	1	1	1	1	1	1	1	1	1	1	1

We need to give no justification in terms of ‘logical impossibility’ or what not for the above exclusion of 21 of the 121 possibilities. We can keep which we like and simply judge the acceptability of the corresponding results. (Note that all ‘double elements’ (qq) remain at present).

2

Decidability and Natural Language

Yorick Wilks
Stanford University

2.1 Grammaticality and Decidability

It would be absurd to try to construct a procedure that could, formally and generally, decide whether English sentences were true or false. If that could be done then, at the very least, there would be no need of scientific experiment wherever English was spoken. But the absurdity of that possibility should not cause anyone to dismiss two other, more interesting, questions: one about what are usually called grammatical English sentences, and the other about meaningful English sentences. Are either of these sets of sentences decidable? Could there be procedures that would determine whether a string of English words was, or was not, either grammatical or meaningful?

Attempts have been made to construct both sorts of procedure: in the case of meaningfulness there is Carnap's theory of Logical Syntax [1]: and in the case of "grammaticality" there is Chomsky's more recent work in the field of linguistics [2]. Carnap's work rested upon an analogy between the ungrammaticality of sentences like "Caesar is and" and the apparent meaninglessness of such sentences as "Caesar is triangular". Carnap thought that, if the rules of grammar were supplemented and extended by "rules of logical grammar", then the meaninglessness of "Caesar is triangular" could be shown by the same procedures as dealt with the more obviously odd "Caesar is and". Carnap's work in this area was largely programmatic: he did not construct such a system of rules in any detail and then apply it to actual texts.

Chomsky's work in linguistics is a natural and proper successor to Carnap's. His original suggestion was that a set of "transformational grammar rules" could produce all and only the grammatical sentences of English from a number of initial "kernel sentences", that were themselves produced by rules of a different sort called "phrase structure rules". Chomsky's paradigm of a transformational rule [2:43] was one which would convert an English sentence into its "equivalent passive": *i.e.* "John loves Mary" into "Mary is loved by John".

Chomsky's work is even closer to Carnap's than is generally realised, in that it too is an attempt to "explicate" or produce, the class of *meaningful* sentences, in spite of Chomsky's initial intention to dispense with meaning and concentrate on

K. Ahmad, C. Brewster and M. Stevenson (eds.), Words and Intelligence 1, 9–27.

grammar. My claim here is not a mystical one about Chomsky's intentions. It can be supported by drawing attention to the greatly extended sense in which he has used "grammar" in the last few years: he has introduced what he calls a "semantic component" [3] into his grammars, so that grammaticality is now for him a notion quite different from syntactically correct. But to demonstrate this point in detail would be a narrowly linguistic enterprise. In this first part of this paper I want to argue that, whether or not Chomsky is now explicating meaningfulness, he cannot be explicating grammaticality as he originally set out to do, because grammaticality lacks the necessary properties for that to be possible. I shall then go on to argue that meaningfulness differs somewhat from grammaticality in this respect.

One further point: it might be objected at the outset that considerations about decidability are of a purely syntactic nature in the sense of "syntactic" in which Gödel's theorem is said to be a purely syntactic theorem, and so they can have nothing to do with questions of meaningfulness, or "grammaticality" in a linguistic sense. The premise is perfectly true, but I am not trying to introduce meaning, truth or grammaticality where it cannot belong. I am simply raising the question as to how one might interpret the notion of theoremhood in certain canonical languages of the kind described by Post [4, 5]. Chomsky himself has observed [6] that his own system of transformational rules can be viewed as a system of production rules for a canonical language in Post's sense. A canonical language has a finite alphabet, a finite number of productions, or inference rules, and a finite number of axioms. The axioms are concatenated strings of items in the alphabet, and any string which can be produced from these axioms by means of a finite number of applications of the production rules can be called a theorem. The decision problem for the language is determining whether any given string is a theorem or not. The analogy between generative grammars and this formulation of proof theory depends upon considering as "theorems" the strings produced by the operation of grammar rules, or "rules of inference". Only the last, or what Chomsky calls the terminal, string of a production is interpreted as having the property associated with theoremhood, namely, "grammaticality". The only "axiom" in Chomsky's system is then the string 'S', for "Sentence", with which all productions begin.

Chomsky does not claim that transformational grammars are complete, in that they produce all and only the grammatically correct sentences of English. Nor does he claim that they are decidable, in that their rules decide of an arbitrary string of English words whether or not it is grammatically correct. However, he is presumably trying to construct a system having either or both these properties; for that is the programme he originally set himself, whatever his subsequent disclaimers about decision procedures [2:55].

In the case of transformational grammars it is not easy to be clear about their decidability, or otherwise, because Chomsky is unwilling to give them any general form. But it does seem generally agreed that a set of transformational rules of the sort Chomsky has described characterizes a recursively enumerable set of sentences, but not necessarily a recursive set. That is to say, you can generate as many sentences as you like with the rules, but you may not be in a position to decide whether or not a given, ungenerated, sentence can be generated [7].

In this respect transformational rules differ from phrase structure rules, which are known to be decidable in a number of cases [8]. In almost all his writings Chomsky has included a number of arguments against the use of phrase structure grammars which, he contends, can produce only the sort of "surface grammatical structure" to be found in, say, school grammar books.

Yet, in this important respect of the decidability of the formal system, phrase structure grammars seem *to start with a distinct advantage*; for, whatever the practical advances made using heuristic transformational parsers (e.g.[9][10]) programmed on computers, it can never be known for certain whether or not a given transformational grammar can analyse a given sentence.

Chomsky also argues against what are called finite state [2:21] grammars, and again, by implication, for transformational grammars, on the ground that there are grammatical English sentences that no finite state grammar of a certain class can produce [11]. That is to say finite state grammars are incomplete *in the sense that* they do not cover the known "theorems". I do not want to go in detail into this linguistic dispute, for what is important here for my purposes is to point out the assumption behind Chomsky's argument; namely, that there *is* some survey of what it is to be a "theorem to be covered" by a generative grammar, which is to say, a grammatical sentence. My contention is that, on the contrary, there is not *as a matter of fact* a survey of the set of sentences Chomsky has in mind, and it is not a class of sentences about which native speakers can take reasonable decisions, in the way they can about meaningful sentences. If one asks an informant "Is 'Colourless green ideas sleep furiously' a grammatical sentence?", one tends to provoke only puzzlement; though linguists, philosophers and logicians are less unwilling to decide the question, and the diversity of their answers is an argument against Putnam's view [9:38-39] that speakers broadly agree on such questions, and hence the grammatical sentences of a language are probably a recursive set. Curry [12] thinks such sentences are grammatical; Ziff [13] thinks they are ungrammatical but not nonsensical; Jakobson thinks that they are grammatical but false [14]; Putnam [9] thinks they are at least ungrammatical, and certainly not false. Chomsky's view has, as I said, changed on this question: in the original formulation of his views, he contrasted such sentences with ungrammatical ones, but in *Current Issues in Linguistic Theory* [3:25] he argued for the existence of a set of grammatical sentences to be "explicated", and that it was a set not dependent on notions like "meaning" for its characterisation:

> "…the notion 'grammatical' cannot be identified with 'meaningful' or 'significant' in any semantic sense. Sentences (1) and (2) are equally nonsensical, but any speaker of English will recognize that only the former is grammatical.
>
> (1) Colourless green ideas sleep furiously.
> (2) Furiously sleep ideas green colourless."

Chomsky later refers to this quoted passage as a demonstration of the independence of "grammar" and "meaning", though there is no demonstration here in any strong sense, but simply an appeal to observe the difference between two sentences.

I do see a difference between the two sentences, in that I can think of a number of things that sentence (1) might mean, whereas I find it less easy to see what sentence (2) could be about. But even that might be only a question of effort: I could probably work out an interpretation for "Furiously fought men tired weary", and might well be able to do something similar for (2). I can see no other difference between (1) and (2) unless I am provided with specific grammar rules that (1) abides by and (2) breaks. By "work out an interpretation" I refer to the ability native speakers of a language have to explain what they mean by a piece of language embedding it in a larger explanatory text or conversation. In the case of sentence (1), one might try to show that it was meaningful by embedding it in some improbable story about the nature of the brain's activity during sleep, and its effect on the sleeper's behaviour. The story might also make it clear that "green" was being used in the sense of "new or untried".

This ability to explain the meaning of something, and so to show that it is meaningful, is part of the ability to write or speak a language. One thinks of Wittgenstein's "words have those meanings which we give them; and we give them meanings by explanations" [15]. Not so, however, with Chomsky's intuitions about grammaticality: I do not share his intuitions about the difference between sentences (1) and (2), yet I remain unrepentantly a native speaker of English. But in the case of grammaticality there is no such explanatory, or elucidatory, *procedure* that Chomsky can employ in order to convince anyone who fails to share his intuition of the difference between sentences (1) and (2).

Grammatical knowledge is no part of what it is to speak and understand a language, for, on the contrary, grammatical explanations and manipulations are what people *who do not speak a language well* fall back on. That is not to deny that for any actual grammar there is a set of sentences well formed with respect to it. But saying that is quite different from what Chomsky says, for it does not imply that there is such a set prior to the construction and use of the grammar, nor does it imply that there are "grammatical mistakes" except in so far as some "rule" is specified with respect to which they are mistakes.

Since he wrote the passage quoted Chomsky [3:9] has changed his view, and now considers sentences like (2) above to be "deviant" in that they should *not* be produced by a good grammar. Hence, such sentences are now considered to be ungrammatical by Chomsky. So he might seem to be in some doubt about what is and what is not a grammatical sentence; yet his whole task of explicating the set of such sentences depends on there being some independent survey or characterisation of what it is to be a grammatical sentence. Without such a survey or characterisation there is no notion of what it is to be a "theorem" for his, or any other, system of derivations to produce. Putnam [7:191] has argued that the grammatical sentences of a language *are* surveyable in that there is general agreement about the membership of the set, and that this justifies us in considering them a recursive set capable of being produced by a decidable generative grammar. The plausibility of his case comes from examples like "Mary goed home", about which there would be general agreement that it is ungrammatical. But there is not this agreement, even among experts, about interesting cases of odd, or deviant,

sentences such as "Colourless green ideas sleep furiously". Thus Putnam's case is not made.

I may have laid myself open to the charge of simply punning on the words "decision procedure" by introducing them into a discussion of natural language. Let me try to get at the main point slightly differently: as is well known, the Propositional Calculus has a decision procedure; namely, certain computations on *truth tables*. However, there is also a partial survey of what it is to be a theorem independent of the truth tables, for they are not required in order to know that "$p \supset p$" is true in the Propositional Calculus. If that were not so one could not discuss completeness or decision procedures at all. For example, when expressing Gödel's [16] theorem in the form "no consistent language can be adequate for the expression of mathematics, and at the same time be capable of proving all true propositions in elementary number theory", it is implied that there is some survey of what it is to be a true proposition in elementary number theory independent of an axiomatisation and a decision procedure. Otherwise Gödel's theorem loses its point. Yet this kind of survey is utterly lacking in the case of grammatical sentences.

Logical truths, then, can be surveyed prior to the construction of any system of explication, and, moreover, the notion of logical truth can itself be characterised in terms of other concepts. These characterisations, such as Leibniz's "true in all possible worlds", are ultimately unsatisfying but it is an important fact about logical truth that they can be proposed and sustained by argument. Similarly the notion of meaningfulness has been characterised in terms of many other concepts, and another possible characterisation of it is explored in a tentative fashion in the last section of this paper. But it is not easy to see how the notion of grammaticality can be characterised in any similar fashion. "What speakers admit as grammatical" does not seem quite good enough for, whatever the inadequacies of "those sentences that are true in all possible worlds" as a characterisation of logical truths, it is certainly better than "those sentences that speakers (or logicians) admit as logical truths".

Chomsky [11], and more recently Ziff [13], have suggested characterisations of grammaticality independent of the acceptance of sentences by some particular set of grammatical rules. Chomsky has suggested that an ungrammatical utterance is read with a falling tone on every word, and Ziff has suggested that an ungrammatical utterance is one that a native speaker "balks at". It needs no concentrated analysis to see that those suggestions will not do, and for the same reasons in each case.

On the one hand perfectly comprehensible sentences like "Mary goed home" would almost certainly be read by a speaker with the same intonation, and as little balking, as the more conventional "Mary went home", even though both Chomsky and Ziff would consider the first sentence ungrammatical and the second grammatical. On the other hand, even intelligent and well-disposed speakers balk at sentences that are perfectly grammatical by our authors' standards, but which express some particularly striking falsehood such as "An elephant isn't really an animal you know".

My conclusion from these arguments is not that many mathematical linguists are wasting their time, or are engaged in some form of linguistic circle squaring. It is rather that if their enterprise is, as it is usually, one of testing a given string of

words to see if it has a given property or not, then it would be better to call the property "meaningfulness" than "grammaticality", since the latter property does not admit of being attached to that procedure. Whereas, as I shall try and show below, meaningfulness is at least a starter in that respect. The relabelling would be quite appropriate to Chomsky's changing notion of "grammatical", which once included "Colourless green ideas sleep furiously" but now excludes it, and has at present an extension very like many people's notion of "meaningful".

If the arguments of these first and second sections are correct, then meaningfulness is not a poor relation of grammaticality, but rather the other way round. If grammaticality is to have a sense as well as an extension then it must, if it is to be anything, be a rather more general notion of meaningfulness. And that view is, I think, consistent with the traditional notion of grammar, though this is not the place to argue for that. Alice saw the point when she detected a very general meaning, or message, in the poem Jabberwocky, which has been taken to be a paradigm of "grammatical nonsense": "Somebody killed something, that's clear, at any rate".

2.2 Characterising Meaningfulness

I have discussed how one might characterise the notion of grammaticality in terms of other concepts. In this section I want to suggest a characterisation of the notion of meaningfulness: one that may seem both odd, and at the same time, obvious. I suggest that we call an utterance meaningful, in some primary sense, if and only if we can decide which of a number of things that it might mean it actually does mean. Or to put the suggestion another way: to be meaningful is to have one and only one of a number of possible interpretations.

If these two apparently different concepts, meaningfulness and sense-resolution, can be brought together then some light might be thrown on an old puzzle about meaningfulness: when grammarians deem a sentence meaningless, or when Carnap deemed a metaphysical sentence [17] meaningless on the grounds of its incorrect logical syntax, then it might not have been that the sentences *had no meaning*, but rather that each had several meanings or interpretations; though taken as single sentences in isolation from others they could not be resolved as having some particular interpretation, and so they were deemed meaningless. However, had they been put back into the context from which they came, or had other suitable context been constructed around them, each might have admitted of one and only one interpretation, as in the case of "Colourless green ideas sleep furiously" embedded in the simple story I suggested for it.

It is a trivial observation that many words have a number of meanings or senses, and that without adequate context they cannot be resolved, in the sense of being attached to one and only one dictionary explanation. If I say "I must go down to the post with these letters" then that sentence can be resolved because it constitutes adequate context to show that, for example, by "post" I mean "place for depositing mail", and not "thing to which horses may be hitched". But if I say "I found I hadn't got a jack" it cannot be resolved, because a hearer cannot resolve "jack" without knowing whether

the sentence belongs to, say, a card-playing story or a car-breakdown story. What I am maintaining is that, in some primary sense of "meaningful", the sentence "I found I hadn't got a jack" is meaningless apart from some context, or context-substitute, in just the way that "Colourless green ideas sleep furiously" is.

Before answering charges that what I have just claimed is either absurd or straightforwardly wrong, I want to say a little more to support the claim that primary meaningfulness is of resolved segments of language. I think some general justification can be constructed along the lines of Quine's discussions of synonymy where, in the course of a detailed examination and criticism of the assumptions of descriptive linguistics [18, 19, 20], he describes a situation of possible confrontation with a speaker we do not understand at all. Quine begins by distinguishing what he calls the activity of the grammarian from that of the lexicographer: the former seeks to catalogue significant sequences in a language, the latter to catalogue the synonym pairs within a language, or between languages. Quine points out that their enterprises are intimately related in that one is concerned with what it is to have meaning, while the other is concerned with what it is to have the same meaning.

Quine then directs his attention to the lexicographer's problem, which he discusses in the conventional terms of substitutions of putative synonyms within larger contexts that remain synonymous as wholes. That way of discussing "having the same meaning" is not a referential one at all, where by "referential" is meant all sense, designation and Fregean dualist theories of meaning. For on any of those theories one should determine whether or not words are synonymous by inspecting the objects or concepts (or both) to which they refer, and seeing whether or not they are the same. Quine's view is essentially a monistic, intra-linguistic, view of meaning and it concerns only relations between strings of words. I see no real difference on Quine's view of things, between saying that two utterances have the same meaning and saying that each is a meaning, interpretation, or paraphrase of the other. The problem then immediately arises of *which* of a number of possible strings of other words is the meaning under discussion, and it is here that substitution, or what Quine calls "a retreat to longer segments", comes in:

> "…a retreat to longer segments tends to overcome the problem of ambiguity or homonymity. Homonymy gets in the way of the law that if a is synonymous with b and b with c, then a is synonymous with c. For, if b has two meanings (to revert to the ordinary parlance of meanings), a may be synonymous with b in one sense of b and b with c in the other sense of b. This difficulty is sometimes dealt with by treating an ambiguous form as two forms, but this expedient has the drawback of making the concept of form dependent on that of synonymy.
>
> We may continue to characterize the lexicographer's domain squarely as synonymy, but only by recognizing synonymy as primarily a relation of sufficiently long segments of discourse."

But what other function for a "retreat to longer segments" can there be than an overcoming of sense ambiguity? What is a "sufficiently long segment" other than one that resolves such ambiguity? Quine does not say explicitly, but I think one can

reasonably infer from the quoted passage that he means a segment sufficiently long to resolve word-sense ambiguity and in particular the ambiguity of the members of a synonym pair when either of them is substituted in the segment. Quine goes on: "So we may view the lexicographer as interested, ultimately, only in cataloguing sequences of sufficient length to admit of synonymy in some primary sense" [21:58]. So the difference between Quine's primary synonymy of resolved segments and the non-primary synonymy of their parts is that the former synonymy is a context independent one. No question arises of substituting resolved segments in anything longer, for there is no more to make clear. "Resolved" means simply that all sense ambiguity has been cleared up.

Let us return to the grammarian, who was said by Quine to have the same problem as the lexicographer. If that is so, then the grammarian, too, will "retreat to longer segments". Corresponding to Quine's remark about "primary synonymy" we might expect another to the effect that "primary significance is of resolved sequences". I do not think that Quine draws this inference in the course of his arguments, but it seems to me a correct one, and a way of stating the necessary condition involved in the characterisation of meaningfulness I suggested earlier.

What is one to make of this necessary condition: the claim that a piece of language is meaningful only if it has one, and only one, interpretation, and hence that it fails to be meaningful if it has none, or two or more? The claim may sound reasonable enough for utterances whose meaningfulness is in dispute, where, as with "Colourless green ideas sleep furiously", the procedure used to show that the utterance is meaningful usually consists in constructing a narrative round the sentence so that it does have a single interpretation. I am using the terms "interpretation" loosely here, and will do something to make it more precise later on, but I think the general idea is sufficiently clear if we assume some notion of paraphrase, interpretation or synonymy between utterances. For the moment let us assume it to be Quine's "primary synonymy" of resolved utterances.

There may well be an important distinction to be made between "being ambiguous, and so meaningless, because of two interpretations" and "being ambiguous, and so meaningless, because of more than two interpretations". Poetry can often preserve two, though not usually more, interpretations over considerable lengths of text which are properly considered meaningful. But for the moment I want to consider poetry, allegories and jokes as special cases.

So then, if the utterance "He fell while getting to the ball" is embedded in a football narrative, then all proper paraphrases of it will be equivalent to "A man fell to the ground while trying to reach the object in play in the game". And the assertion that the second mentioned sentence is a paraphrase of the first would resolve the first in just the way that inserting it into a football narrative would. In the case of either procedure, inserting or giving a paraphrase, we would then know that "ball" was not being used in the sense of a "formal dance". It is also important to notice that resolving an utterance by giving a paraphrase or interpretation is not the same thing as resolving the constituent words. To know that the two sentences just mentioned are possible paraphrases or interpretations of each other is also to know that, for example, "ball" is being used in its "round object" sense. But the

converse is not necessarily true, since the interpretations of sentences are not simply computed from the interpretations of their constituent words, as anyone knows who has tried to make himself understood in a foreign language with the aid of only a dictionary, or even with a foreign grammar book as well.

But aside from sentences whose meaningfulness is *in dispute*, how reasonable is the application of the necessary condition to an everyday sentence such as "He fell while getting to the ball"? Is it not absurd to say that the sentence is meaningless just because, taken in isolation, we do not happen to know which of two likely interpretations it bears? However, if challenged to show that, or how, the sentence is meaningful, a speaker who cannot make use of gestures, and so go outside language, will certainly embed the sentence in some story or anecdote so as to give it one of its two more obvious interpretations. And that is the same procedure as the one adopted by the defenders of the meaningfulness of "Colourless green ideas sleep furiously". In other words, use is made of a procedure that does give the questioned utterance some particular interpretation. I am not taking refuge here in some highly general view such as "meaningfulness can only be discussed with respect to an entire language", or anything like that. I am calling attention to a particular *procedure* of sense resolution, in which a particular interpretation is assigned to a questioned utterance by means of telling a surrounding story, uncovering more of a surrounding book page, or perhaps simply by producing utterances with the form of dictionary entries, such as "ball means round object".

But even if the necessary condition is plausible in itself it does not shed any light on the "primary significance" that only resolved segments can have. In Quine's discussion of a "primary synonymy" he gave a quite different explication of that notion in behavioural terms. But here I think we can push the characterisation in terms of sense-resolubility a little further, and get something like a sufficient condition for primary meaningfulness as well.

The sufficient condition for meaningfulness would be that a text was meaningful if it had one and not more than one interpretation. But in virtue of what can a text be said to have a single interpretation? Why does one want to say that "I must take these letters to the post" has a single interpretation, though "He fell while getting to the ball" has two? The difference cannot be simply that "ball" has two senses while "post" has one, for "post" usually has more senses listed in a dictionary than "ball".

If an English speaker is asked to explain, in informal terms, how he knows that "I must take these letters to the post" has only one interpretation, he will probably say that the notion of "letters" is connected to only one sense of the notion "post", and so if the word "letters" is present in an utterance then it can only be the "mail" sense of "post" that is intended. But in the case of "He fell while getting to the ball" there is no such overlap of coherence of notions to disqualify either of the two natural interpretations of the utterance.

This common-sense explanation can be put in linguistic terms quite straightforwardly: if classifiers, or markers, can be attached to the senses of words so as to distinguish the senses from each other, then it is a technical matter to specify coherence rules, operating on the markers, that select certain word senses

in preference to others. So for example, if there was a marker MAIL in use, then we would expect to find it in a table of linguistic information attached to only one sense of "letters" (not the "alphabetic items" sense) and to only one sense of "post". We might then examine the sentence "He took the letters to the post" armed with the rule "if the marker MAIL is attached to senses of *more than one word* in the sentence, then select those senses". That rule would be a very simple-minded one, though it would work in this case. However, such rules can be made as complicated as necessary (e.g.[22]) and there is no more mystery about the attachment of suitable markers to word-senses than there is to the construction of the conventional entries in an English dictionary that distinguish the senses of words from one another.

Moreover, the operation of such rules as the simple one involving the marker MAIL can be equivalent, in effect, to the provision of an interpretation, or paraphrase, for the utterance under examination. If the marker MAIL pins down, or selects, one particular sense of "letters", we can suppose that sense to be expressed as a conventional dictionary entry such as "letters as papers that are mailed". That expression of the sense would be entered, in tables specifying the rules of a possible linguistic system, along with the marker MAIL, but "letters as items in an alphabet" would not. Part of the dictionary entry tagged to MAIL, namely "papers that are mailed", could then be substituted for "letters" in the original utterance. If this procedure were repeated for each word of the utterance we would end up with a new, resolved, utterance that was a paraphrase of the original one. In that sense, the operation of this sort of rule also provides paraphrases.

Now consider a different example, which I shall call a pseudo-text: "Do you like my car. I always wear a flower. The seats are leather". An utterance like that would almost certainly be said to be meaningless, even though it is not inconceivable that an ingenious person could embed it within some intelligible story, perhaps as an entry for a literary competition. It is not easy to say why the pseudo-text seems meaningless. The simplest way of putting the matter is to say that there is nothing that it is about, in that the ideas the utterance expresses do not cohere together sufficiently for there to be an interpretation that is not identical with the utterance itself.

What is claimed here by the sufficient criterion of meaningfulness, given above, is that the pseudo-text *would be* meaningful if there were sufficient coherence between its constituent concepts, of the sort expressed earlier in linguistic terms by means of the MAIL marker. Rules specifying such occurrences of markers can be as complex as necessary, and the specification is a technical matter for linguistics. What is important for the present discussion is that the meaningfulness criterion, expressed in terms of "having one and only interpretation", should refer to an overall interpretation, located by means of coherence rules of the sort I have discussed. It cannot refer simply to the sense-resolubility of the individual words of the utterance under scrutiny.

This last point can be made by looking again at the pseudo-text. Each of its three constituent sub-sentences is such that one can see, for each word in it, in which of its senses it is being used. That remains true whether the three sentences are considered separately or as parts of the pseudo-text. So there is no problem

about word-sense ambiguity in the pseudo-text, and hence, if the meaningfulness criterion were expressed simply in terms of word-sense ambiguity resolution, then the pseudo-text would satisfy the condition, and so be meaningful in terms of it.

I have given only the crudest example of the way in which a vague notion like "conceptual coherence" can be operationally expressed by means of procedures involving linguistic markers. In fact such procedures are almost always more complex than the simple co-occurrence of a single marker, and it is easy to see that a simple "threshold" notion of coherence will usually not suffice to establish meaningfulness. Consider another pseudo-text: "All kings wear crowns. All crowns are coins. All kings wear coins". That pseudo-text is like the earlier one in that each sub-sentence is resolved as regards word-sense ambiguity, although the whole pseudo-text does not seem to admit of a single interpretation. Yet, unlike the last example, it is not that there seems to be *no* interpretation, but rather an oscillation between two alternative ones, depending on the sense of "crown" selected. But this example would satisfy the very crudest standards of conceptual coherence, in that each sub-sentence would have an overlap of markers, given any reasonable choice of markers, with either of the other two sentences comprising the whole pseudo-text. Hence, any rules applying to such markers would have to be more structured than the simple one applying to MAIL that I gave earlier.

I have been defending the suggested characterisation of meaningfulness against charges of absurdity and wrongness, but is it nonetheless vague? After all, *what* precisely is being characterised as meaningful? Earlier in this paper, I described a procedure used when the meaningfulness of an utterance is challenged: a speaker defending its meaningfulness attempts to embed the utterance in a story or narrative whose overall meaning is clear. But, in terms of the characterisation, the utterance so embedded is properly deemed meaningless if it does not have one clear interpretation in isolation; and, if the whole story containing it is clear and unambiguous, then it is that whole that is shown to be meaningful by the embedding procedure. Furthermore, no inference can be made from the meaningfulness of the whole story to that of the embedded utterance, any more than one can infer that p is a theorem because it is a proper part of the theorem $p \supset p$.

If the last point is correct, then it is not proper to speak, as I did earlier in the paper, of the procedure of embedding an utterance whose meaningfulness is questioned as one giving a survey of meaningful utterances. What that discovery brings out is that the present formulation of the characterisation is incomplete: it requires an addendum "…one and only one interpretation with respect to some dictionary, or dictionary substitute". Consider again the utterance "He fell while getting to the ball". My claim was that that is meaningless in isolation, in that one could not decide whether its proper interpretation contained the "round object" or the "formal dance" sense of "ball". But that judgement assumed a conventional dictionary containing those two senses of the word "ball", even though a considerable proportion of English speakers do not know that "ball" can be used to mean "formal dance", and so to them the utterance might be said to be unambiguous and so perfectly meaningful.