A Near-Fregean Framework for Discussing Language-Understanding

by

J. G. Laski

February 1977 TR-67

Department of Computer Sciences
The University of Texas at Austin

A Preliminary draft of:--

A Near-Fregean Framework for Discussing Language-Understanding (Programs)

J. G. Laski February 1977

index:

The Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712

Prologue

The following work is a first version of a set of ideas that seem to me cogent and worth further exploration. I therefore solicit comments from workers in a field where I am a green beginner. In particular, I would appreciate pointers to where I have been unnecessarily obscure, mistaken, or confusing, and references to other work that this may shed light on, or, conversely, be inconsistent with; contrariwise, to other work that, itself, may stimulate related ideas that could, with advantage, be included.

I am particularly worried that the argument, from the Fregean sense/reference distinction to the data-structures involved in a language-understanding system and back to an objectification, in terms of these data-structures and their access patterns, of a sense/reference distinction, may present an appearance of circularity. It may help the reader to view the original Fregean distinction as the explicandum, and the latter data-structure-based distinction as an explicatum, to use Carnapian terminology.

The genesis of this work lies in patient explanations and repeated clarifications of my misunderstandings of Yorick Wilks' language-understanding system by Margaret King, a draft paper of hers: "Several ways of being Suggestive" on presupposition, and her remarks that, consequently, a) she wished to be able to include very much more deduction, whether using probable general knowledge of the world, or using necessary particular knowledge of what is the case, already gleaned from the text or previously read texts, and b) she was not at all clear how to relate this activity to the algorithms and data-structures presently employed in the system.

The next stimulus was my observation that that system would represent its understanding, say, of the two texts "Mary was sitting on the table. She

has red hair" and "Mary has red hair, She was sitting on the table" by two totally (apparently?) distinct data-structures,in which the sequence in which the information given by the text is preserved, whereas the knowledge of what is the case in the world that an understanding system should have obtained from the two texts for subsequent possible use in deductive reasoning, it seems intuitively to me, should be the same---or at least very nearly so.

Apart from this prologue and the subsequent epilogue, which were written after the body of the paper was typed in early February, the body of the paper was written in mid-November, and is reproduced from what was written then without change, apart from tidying up small points of poor grammar or replacing obviously obscure expressions by what I hope will prove to be less so.

A Near-Fregean Framework for Discussing Language - Understanding (Programs)

Introduction

We state 7 doctrines of Frege concerning sense and Reference and show that, with some modification, they can be used to provide a framework to understand natural language - understanding (programs). In particular we show that we can make a precise distinction between (the data-structures that realize) information in the realm of sense and (those that realize) information in the realm of referents, based on accessing paths to that data.

For the benefit of non-computing people, a data-structure is that combination of data that enables programs acting on it to access such data as these programs need. If these programs answer questions about what the data-structure realizes, then these programs realize information-yielding algorithms and the data they access will realize the information they will need in order to answer the questions they ask.

Initially we will separate (data-structures that realize) information in the realm of sense from (those that realize) information in the realm of referents on the basis of a pragmatic choice of what kind of information is useful in the realm of sense and what kind in the realm of referents; the choice is pragmatic since: --

(i) What is available in the realm of referents is available to algorithms working in the realm of sense.

(ii) Information about a referent justifies corresponding information in the realm of sense.

We will argue pragmatically against using corresponding information in the linguistic realm.

However, we will go on to separate information in the realm of sense from information in the realm of referents on the basis of where (how it

may be accessed) it is held rather than what kind of information it is.

This allows information of whatever kind to move from the realm of sense to the realm of referents (or vice-versa) and we shall argue that this happens on efficiency grounds of storage of and access to the data that can realize the information and examine the consequence of this migration.

As Frege did, we shall distinguish between objects and first-level properties (or functions from objects). Thus "the spade" refers to a simple object, "$\lambda\xi$. ξ is black" refers to a first-order property. Similarly "the spade is black" refers to the true or the false. Properties (or other functions) on first-order properties and simple objects are called second-order. "$\lambda\alpha$. α (a spade) = the true or α(a spade) = the false" refers to a second-order property. We will call the level in this functional hierarchy in the realm of referents the <u>type</u> of the property or function.

If objects in the realm of referents have such a functional hierarchy, this will be reflected back to a similar functional hierarchy in the realm of sense. However, the sense of a phrase that refers to the true or the false may embody information about the position of the referents of its component sub-phrases in the referent type-hierarchy and therefore, by reflection back to the similar type hierarchy in the realm of sense, about the position of the sense of its component subphrases in the functional hierarchy in the realm of sense. We shall also call this position in the hierarchy in the realm of sense the type of the sense.

Similarly, the functional hierarchy in the realm of sense reflects back into a functional hierarchy among the phrases of the object language and we shall call this position in the hierarchy among the phrases of the object language the type of the phrase.

If the application of a referent at some position in the referent type

hierarchy to any and every set of arguments that satisfy the required type levels in the functional type hierarchy yields a referent, the system of referents is called non-sorted; otherwise the simple objects can be divided into sorts, the first level properties into functions from some sort to the truth values, such that the result of applying them to any other object not of the right sort is not defined; and similarly for functions and predicates of each type and each number and type of arguments. The sort distinctions in the realm of referents reflects back, similarly to the type distinctions to sort distinctions in the realm of sense and among the linguistic phrases of the object language.

If the set of arguments that a referent can be applied to are not of a single type, or the results of so doing are not of a single type, but, perhaps, satisfy some relations among themselves, the referents are said to be untyped; if the referents are such that no referent can be applied to itself, or to a referent that can be applied to the original referent, or to a referent that can be applied to a referent that....can be applied to the original referent, the referents are said to be stratified.

The properties of being stratified and of being sorted reflect back from the realm of referents to the realm of sense and thence to the linguistic phrases of the object language.

Frege thought that the realms of referents, senses and phrases were essentially typed, stratified and non-sorted; we hold conversely that they are, in general, untyped, unstratified and sorted.

Frenge's Doctrines

1) To every word or phrase corresponds a sense, and to most a referent.
2) It is the sense of a phrase that enables us to discover the referent from the phrase.

3) The sense of a phrase is a combination of the senses of its component subphrases.

3a) It follows from 2) and 3) that the referent of a phrase is a combination of the referents of its component subphrases.

4) Senses are elements of a "third realm", neither objective, as referents are, nor subjective, as ideas are (nor linguistic, as phrases are, which Frege did not mention).

5) There is a strict type-hierarchy of functional level of applicability among simple objects, their properties or functions upon them, properties of or functions on these first level properties and functions. This type-hierarchy is reflected back to a similar type-hierarchy among the linguistic phrases that refer to them, and also to a similar type-hierarchy in the realm of sense.

6) Phrases that do not refer to simple objects (including truth values) are essentially incomplete (unsaturated), as are their senses and referents.

7) Given the functional type of some phrase (or a sense or a referent), and any collection of phrases (or senses or referents) of the functional type required for its arguments, the first may always be combined with the rest by application to yield a meaningful phrase (or sense or referent).

I claim that these provide a useful framework to discuss natural language understanding (programs). For our purposes: --

1-4) carry over unchanged

5) Carries over except for funnies like Twice ($\equiv \lambda x.xf.f(fx)$) This is parametrically polymorphic of type $(\alpha \rightarrow \alpha) \rightarrow (\alpha \rightarrow \alpha)$, so

will not fit into a typed hierarchy; it is also applicable to itself and therefore will not fit into a stratified hierarchy.

6) requires elaboration:--

a) a property is not the same as its extention, i.e. the set of what it holds of, nor a function the same as the subset of ordered pairs that relate each element of its domain to the corresponding element of its range.

b) By Frege's use of greek letters, e.g. ⌜ξ is a man⌝, or, better by the λ calculus, e.g. ⌜λξ . ξ is a man⌝ or, better still by a sorted λ-calculus, e.g. ⌜λα: humans . α is a man⌝ incomplete phrases, senses or referents may be finitely given in a form amenable to calculation.

c). Notwithstanding a), if the domain of application of a function or predicate is finite, an algorithm to calculate it can be given by the set of corresponding pairs, <u>together with</u> a lookup algorithm.

Notice that these remarks apply equally - with anysuitable changes - to referent combination, sense combination and the combination of linguistic phrases.

7) is clearly wrong; the realm of referents is clearly many-sorted.

(i) It is incoherent to apply (the referent of) ⌜ran⌝ to (the referent of) ⌜Everest⌝. If there were a metaphorical coherence, which in this case I doubt, a preference logic and preference list of metaphorically coherent sorts can deal with such secondary coherences.

(ii) Since sense determines the referent, a sort-coherence restricting what can be combined in the realm of referents

entails a sort-coherence in the realm of sense, which, in turn entails a syntactic sort-coherence among linguistic phrases that can be put together in a larger phrase.

(iii) In whichever realm, the sort structure is very complicated with one sort included in another, union of sorts, overlap of sorts, product sorts, etc. among non-functional elements of the sort domain. Moreover, the sort of some functional elements of the realm may be parametrically polymorphic, not only with respect to the type of what it applies to and the result of applying it, as we showed above (5) for Twice, but also with respect to sort.

Consider for example the sort of the phrase: -

[$\lambda a{:}\alpha.\lambda b{:}\beta$. a and b] or its sense or its referent.

Consider the cases: -

a) α,β some subsets of simple sorts, e.g. ⌜men and women⌝ is of sort the union of the sorts α,β(say α = male , β = female)

b) α,β some subsets of predicative sorts, e.g. ⌜λx.x is a man and x is a woman⌝ is of sort the meet of the sorts α,β

c) α,β the sort truth values, e.g. ⌜true and false⌝ has a result of sort truth values.

d) α,β two simple sorts not both truth values, e.g. ⌜the man and the woman⌝ is of sort the cartesian product of α,β. . We could go on exploring possible cases for a long time.

It may be argued that these alternative cases, in the realm of sense and the realm of referents, are between distinct senses and referents that the linguistic phrases in isolation fails to disambiguate. It seems to me that this alternative way of looking at the problem comes to the same thing when looking at what happens during the understanding process. I do not think that every phrase necessarily loses all ambiguity, so that polymorphic senses and referents are needed; however, I will agree that such senses or referents are structures of alternative component senses or component references amongst which selection takes place during the understanding process.

I also agree that there is overlap between the idea of what sort a referent is and the idea of what properties it has; indeed I see a sort potentially corresponding to each property and vice-versa. How this affects the understanding process will be further treated below but I will use sort to discuss whether or not phrases, senses or referents can coherently be combined.

To sum up: elaborate sort distinctions and coherence rules can be formulated and used, without any bound I can see, and are powerful tools to select which of alternative components [1] to combine.

---

1) I am using component roughly to mean alternative senses of a word or phrase, reserving sense for its Fregean use. However, we may be speaking of component senses, component referents, or component structures of data structures that represent these.

Sense and Referent in Data and Algorithm

As we build up the data-structures that realise our understanding of a text, the consequence suggested by the preceeding are four-fold:--

I. Incomplete senses and referents may be refered to by finite, i.e. complete, phrases of the meta-language, one out of which will be an algorithm. Therefore, they can be realized by finite data-structures.

II. To each word or phrase will correspond a data structure, and we may think of this as divided into substructures, one of which realizes the sense of the word or phrase, the other its referent. How to distinguish between what realizes the sense and what realizes the referent, and what information each data-structure realizes is the principal question we will be treating below.

III. Since, by 2), the sense of a phrase enables one to reach its referent, any algorithm in the realm of sense may access data from the data structure that realizes the realm of referents.

IV. Since the properties of the sense of a phrase are justified by and correspond to the properties of its referent, any data that realizes information in the realm of sense implies that, at least in principle, there will be corresponding data in the realm of referents. In practice, to avoid storing duplicate data, we may allow algorithms in the realm of referents to access the corresponding data from data-structures in the realm of sense.

III and IV appear to suggest that we are mistaken to distinguish between the realm of sense and the realm of referent. From, the point of view of III, since information realized by data in the realm of referents is available to

algorithms in the realm of sense, why distinguish it? Why not consider it all as information in the realm of sense? Conversely, since by IV all information realized by data in the realm of sense can also be realized by equivalent data in the realm of referents, why not consider all data as realizing information in the realm of referents?

I fully accept the force of both these arguments. They lead one to the conclusion that any distinction between the realm of sense and the realm of referents, and what information is held in the data-structures that realizes the sense of a phrase and what in that that realizes its referent can only be justified on pragmatic grounds.

However, I think that these arguments are not to the point. Rather, I think, the point is whether or not we find it useful to make a distinction between parts of the data structure that corresponds to each word or phrase, what information is realized in each part, and what algorithms use the data in each part, and, finally, whether the Fregean distinction between sense and referent is appropriate to illuminate the distinction we choose to make.

Now we will want to make a distinction, if, by doing so, we can more readily consider, study, design the data structures and algorithms that work on them to realize our language understanding system, or model a theory of how language can be understood.

My belief is that the distinction is helpful and properly can be called Fregean and I now proceed to demonstrate this and discuss the implications of how I do so both for that design process and for the theory of meaning.

I take as example the Yorick Wilks language understanding program since I am acquainted with it and it works quite well. I would expect the discussion to apply to any other system mutatis mutandi.

As far as I can see, apart from input, output and the strategic

algorithm that determines which algorithm to apply to what data structure next, the algorithms can be divided up into three classes:--

A) Component Selection by Coherency Criteria

B) Individuation and Cross-Reference

C) General Laws and Deduction

though I do not think that he or his co-workers make the distinction between A and B; those that perform B seem to be in the middle of those that do A, without any indication that they are different in nature.

Component Selection by Coherency Criteria

The main concern of this section is to argue that the algorithms concerned with this operate with data that can best be considered as realizing information in the realm of sense.

In Wilks' system there are the so-called template matching and paraplate matching algorithms. For the benefit of those not familiar with the system we will attempt a very brief oversimplified and probably inaccurate description.

As the text is read, each word is associated with a data-structure, and the text is split up into short phrases separated by punctuation marks or conjunctions or prepositions together with, for the latter two, their associated data structure. Within these phrases the words are, presumably, nouns, verbs, adjectives, adverbs. The data structure associated with each word consists of a set of one or more components each of which realizes information about one possible meaning for that word. For each such component, one sub-component realizes what might be called the principal meaning of that meaning. For example ⌜run⌝, ⌜walk⌝ might have a component with principal meaning MOVE(the first might have subsidiary meaning FAST, both that

the mover was probably MAMMAL). ⌜ball⌝ and ⌜chain⌝ have components with principal meaning PHYSOB. We now consider the set of sequences of principal meanings that would correspond to the selection of some component from the data-structure corresponding to each word in the sequence, and match each of these to the allowable so-called templates. A template consists of an actor-action-acted upon triple. For example we may have:--

Allowable triples			Disallowed triples			
PHYSOB	SUPPORT	PHYSOB	IDEA	SUPPORT	PHYSOB	
HUMAN	THINK	IDEA	HUMAN	THINK	HUMAN	
HUMAN	THINK		PHYSOB	THINK	IDEA	etc.

In the last we have an empty acted upon third.

A match occurs if an allowable triple is a subsequence of a sequence of principal meanings. Then, unless there is no match, we restrict our attention to those sequences that match an allowable triple.

Presumably, the unmatched principal meanings come from data-structures that correspond to words that qualify some element of the triple. We now examine the subsidiary meanings of the selected components and the principal and subsidiary meanings of the unmatched components to see what they qualify and which alternative components to select, on the basis of how many of the preferences of the subsidiary meanings are satisfied.

The next stage is to put these phrases together. Consider a preposition or conjunction. To this is associated a data structure consisting of alternative so-called paraplates. A paraplate tells us what relationships are to be preferred between the data structures associated with the following phrase and with the master phrase, and word within that phrase to which they should attach. Again those combinations are selected which satisfy by their principal

and subsidiary meanings the most preferences.

In doing so, pronouns will often become related to what they refer to and, hopefully, a best component will have been selected to realize the meaning of each word.

If more than one selection of components or alternative attachments remains, the application of general rules about the world is used to select which to finally choose. For example in 'The soldiers shot the women and I saw some of them fall', 'them' is linked to 'women' since 'shot' yields, in part, CAUSE to PUSH, fall yields MOVE, and we can apply the general rule that if a PHYSOB is PUSHed, it tends to MOVE.

This brief description is, I hope, adequate to convey the flavor of the data and algorithms of this system.

I think that it is obvious that the component selection performed by template matching is applying a coherence rule to a relation or property and its arguments, where the sorts are given by the principal meaning of each component. Selecting the components of the qualifying words, and attaching these to what they qualify, is again maximizing the sort coherence with the subsidiary meanings playing a part. Finally, each paraplate describes the complicated sort properties of,and relations that should presumably hold, between the data-structures it joins the one corresponding to the phrase following it, the other to the phrase and word within that phrase to which the former attaches. These properties and relations are based on information realized, as before, in the components of the data-structures that get joined.

Hence, apart from the cross-reference and deduction algorithms, the algorithms are bases on rules of coherence, and the majority of the data realizes sort information.

Should these coherence algorithms and the sort data be held to be in the realm of referents, or in the realm of sense, or as syntactic algorithms on linguistic data?

For seeing them as in the realm of referents, there is one argument: sort distinctions initially arise in the realm of referents; in seeing them elsewhere, one is using distinctions that are secondary, derived from the realm of referents.

For seeing them as in the realm of sense, there are two arguments.

(i) In the case of the good old 'present king of France,' it seems reasonable to find it coherent that he may 'be bald' or 'sing for his supper' and find it incoherent that he may 'be circular' or 'parturate,' on the assumption that he is of the suggested sort, viz. human, male. But, since the phrase has sense but no referent, there can be no data structure realizing information in the realm of referents.

Against this, one may consider a data structure realizing a non-existent referent, but this trick has led to a lot of paradoxical nonsense.

(ii) Before individuation (see below), even though a phrase may have a referent, the data structure that realizes that referent has not yet been created in the system. Hence, if, as is the case in the system described above, and as appears likely for any system, sort data and coherence algorithms are applied before individuation, then these must be working either on data that is not in the realm of referents, that is to say that is in the realm of sense, or as syntactic algorithms working on linguistic data. (See below) This latter explanation should be rejected, since the algorithms are justified by how the world is, rather than by the allowable forms of linguistic phrases.

For seeing them as syntactic algorithms working on linguistic data, there is one, basically ontological argument. This claims that all data is derived from linguistic phrases and, therefore, all algorithms that work on it are syntatic. This has the advantage that we are saved the need, in the metalanguage, to be able to distinguish between referring to a phrase, the sense of a phrase and the referent of a phrase.

I find this argument unpersuasive. Of course we can see all data as linguistic and all rules as syntactic, justified by how we use words and phrases, and not by properties of how the world is and what their meanings are. But then we could only find some phrase confusing or in error on formalist grounds, and not because its meaning is not in accord with how the world is, which seems not to be how we talk and think about how language is used.

Moreover, we would be unable to talk about meanings that our object language cannot express. Tarski has shown that, for suitably powerful languages, the property of phrases of the object language that, in the metalanguage is expressed as: '$\lambda\eta$ for all valid interpretations, the interpretaion of η is true', is not expressible as a phrase of the object language. Hence we are able to talk about a meaning inexpressible in the object language by talking in the metalanguage. This, a purely formalist way of talking would forbid.

Once the distinction between linguistic phrases and what they mean is acknowledged, the distinction between sense and referent follows by Frege's original arguments.

All these arguments, and, in particular, argument (ii) in favour of the realm of sense against the realm of referents, led me to regard sort-realizing data and coherence-realizing algorithms as being in that part of the data-structure that realises information in the realm of sense.

Incidentally, this would accord with Frege's dictum that the sense of a phrase depends on the sense of its component subphrases.

After individuation and cross-reference (see below), possibly as a result of deduction from general laws, we may discover properties of some referent that locate it in a narrower sort than that given by a sense of one of the phrases that refer to it. Since the coherence rules of sense are derived from those in the realm of referents, we may as a result be able to reject components of the realm of sense that, before cross-reference, were still acceptable.

It may even happen as a result of this process that no component of the realm of sense is left as the sense of some phrase that still leads to this reference. In this case, we must either reject the validity of the deductions we have made, or reject that this and some other phrase(s) refer to the same referent, and reinstate some component sense for this and other phrases that refer to this referent.

We do not discuss the backtracking algorithms that realize such changes, but we do pose the question of whether these various later algorithms should be seen as operating in the realm of referents or in the realm of sense.

Individuation and Cross-reference

All the previous discussion is empty. We have not yet specified what data-structures are associated with the phrases of the text by our language understanding system, whether these data-structures can be seen as divided into two substructures, if so, whether or not the distinction between these two substructures can approximately be made in terms of the Fregean distinction between sense and referent, and, finally what kind of information the data in the various parts of the data-structure realizes.

In order to do so we give a preliminary sketch of what individuation and cross-reference algorithms do.

Individuation associates an individual of an appropriate type (simple, predicate, functional, etc.) with each phrase, which we see as the referent of that phrase. In terms of the data-structure, therefore, the individuation creates a part of the data-structure that realizes the referent of the phrase from which it is initially derived, with the data within this part of the data-structure realizing information in the realm of referents.

We see this referent data-structure as initially empty of information, although, perhaps, it will immediately come to contain data to realize what type it is. Also, just as some phrase may be realized as a structure of its various subphrases, its sense realized as no more than a similar structure of the data-structures that realize the sense of various subphrases, its referent may be realized as no more than a similar structure of what realizes the referents of its various subphrases. We would prefer, however, to see these latter structures as realizing information about the sense and referent of the phrase.

Since what sort (the sense of) a phrase is determines necessary properties of its referent, the sort data used by the coherence algorithms we have been discussing must either be held in the referent data-structure itself, or in a part of the data structure that leads from the phrase to the referent data-structure. We call this part of the data-structure the sense data-structure and see the data within it as realizing information in the realm of sense.

The first, admittedly weak, argument for locating the sort information in the sense and not the referent structure of some phrase is that there is no particular reason to perform this referent-creating indiviuation while we are engaged in rejecting data-structure components on the basis of local

coherence requirements of a particular phrase. But I think the force of this argument is more to suggest that the sense/referent distinction is not useful unless more than one phrase can refer to an individual, so that more than one sense can lead to the same referent.

Cross-reference coalesces the referent of some phrase with some other referent in the system. In order that his may be validly done, the sort data must imply properties for the referent of this phrase that are consistent with the properties of this other referent, so that the properties of the referent are, at the very least, the union of those implied by the sort of the various phrases that refer to it.

More properties may be deduced to hold of some referent in two ways. First as the example 'the observer was bearded,' which implies that the referent of 'the observer' was bearded, shows, these may be derived from the fact that a phrase that refers to a referent is a subphrase of a phrase used assertively. Second, this referent is probably male rather than female, since bearded humans are usually (though not invariably) men, and this example shows that we may use our general knowledge about the world to deduce further properties.

Thus the properties of a referent may imply that its sort is more restricted than can be known by considering a single phrase that refers to it in isolation. If so, we can, perhaps, exclude components from the data-structure associated with the phrase that could not be excluded by considering the phrase in isolation.

As mentioned above, all components may be excluded as a result. In this case either deductions that led to this must be rejected or one of the phrases that were believed to refer to the referent must be detached.

If this is done, we need to remove properties of the referent that depended on the fact that this phrase referred to it; its remaining properties are those that are derived from the sort of those phrases that still refer to it. Also, we must recover the sort of the detached phrase in order to consider the properties of its referent and whether it can be cross-referred to some other referent.

This implies that the sort-data associated with a phrase is not a part of the data-structure that realized its referent but is, rather, a part of the data-structure that leads from the phrase to its referent, which earlier we called the sense data-structure.

It is striking how very Fregean this distinction between the sense data-structure and the referent data-structure has turned out to be.

The following diagram, in which the first two of three phrases turn out to have the same referent may be helpful. Backtracking does not occur and it should not be thought that the algorithms of one step are all completed before those of the next are begun.

Linguistic Data Sense Data Referent Data

Read-Lookup:

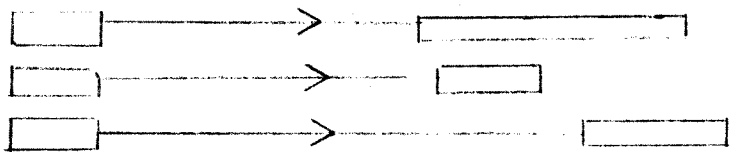

1

Coherence of Sorts:

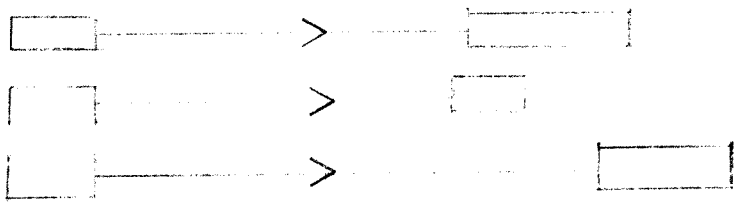

2

Individuation:

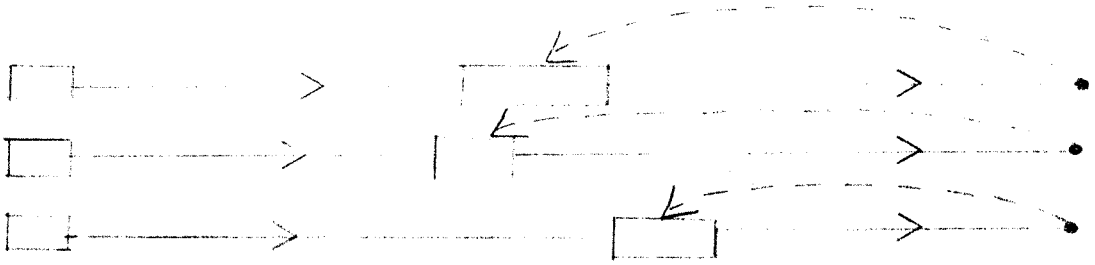

3

Cross-Reference:

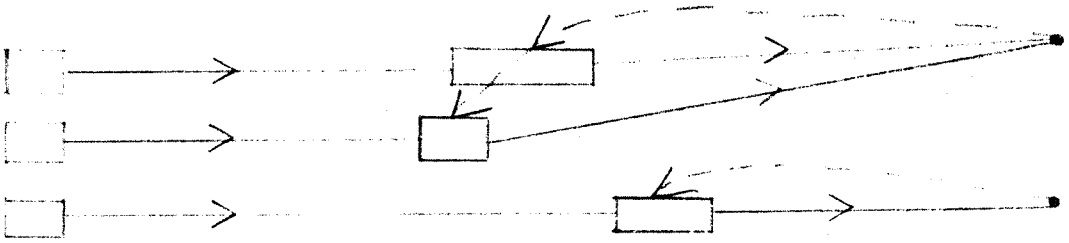

4

Sort Coherence from X-reference:

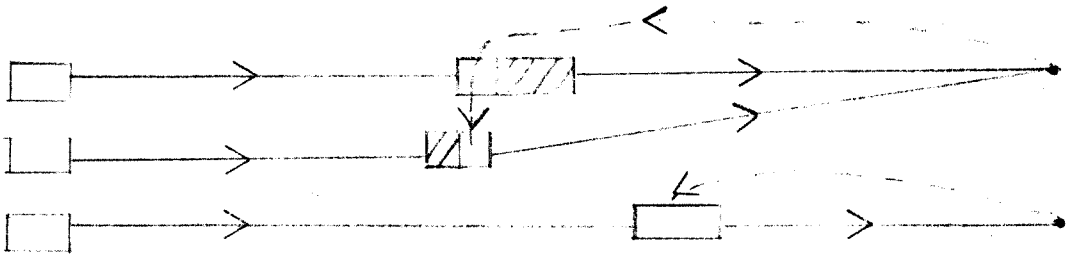

5

Data Migration:

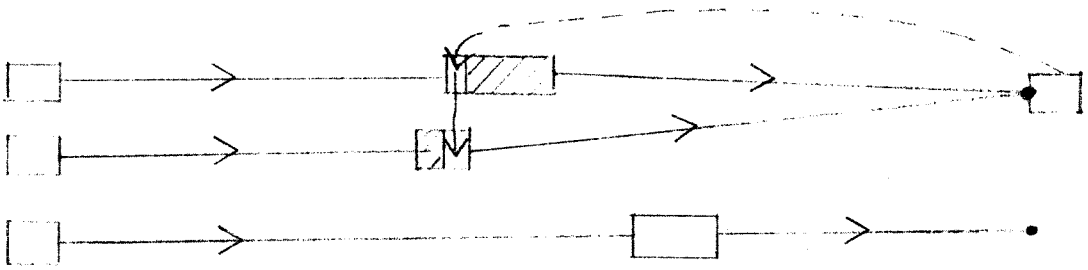

6

Deduction:

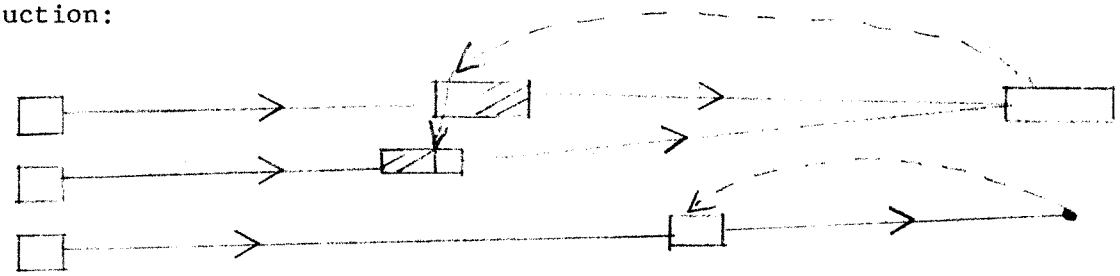

7

Sort Coherence from Deduction:

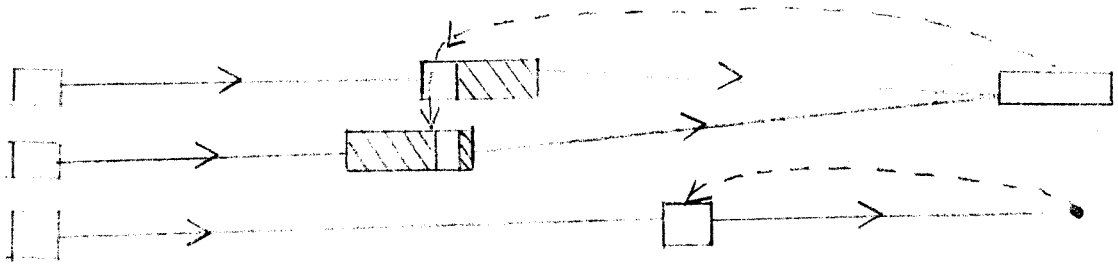

8

Final Data Migration and Sense Deletion:

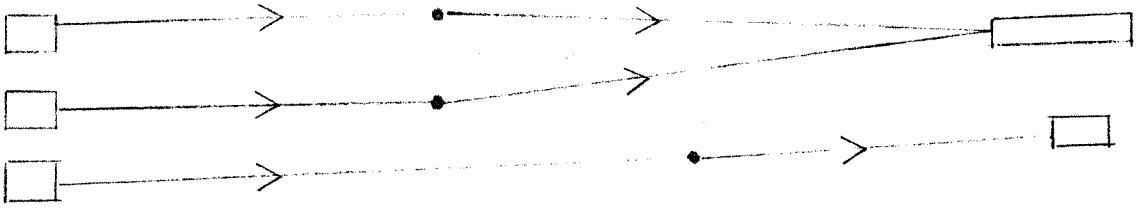

9

Separation of Text and Referential Model:

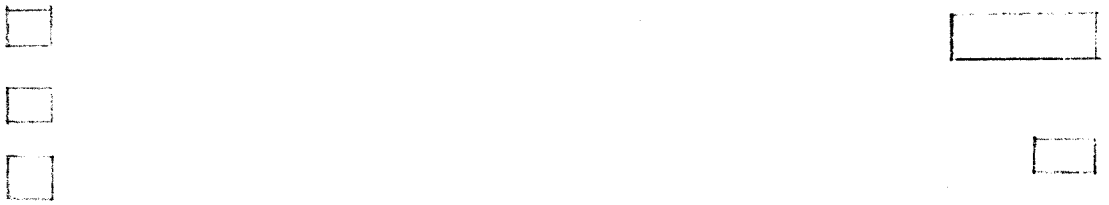

10

More about Individuation

We see the individuation algorithms as marking the key divide between working entirely with sense-data and the possibility of working with referent-data as well, which allows the use of cross-referencing and deductive algorithms, since we see these latter, unlike coherence algorithms, working on data entirely in the realm of referents.

The examples and discussion above may have given the impression that individuals corresponded solely to referring phrases. Also, unlike Pat Hayes, perhaps, I do not see referents as without structure except that given by what sort they are.

Consider, for example, "the referent of 'λ x.x is green or x is yellow,'" which must, of course, be distinguished from "either the referent of 'λ x.x is green' or the referent of 'λ x.x is yellow.'" [Note to M. King: I am curious as to how you would make such a distinction without using the distinction between the object language and the metalanguage or some analogue thereof.]

All three are individual referents of first level predicative phrases, the first complex, the latter two simple. Their sort is that of being a colour property.

Being some particular sort is equivalent to having some particular property, and it seems obvious that identical data can be interpreted to realize either that a sense is of a particular sort or that the referent of the phrase which has that sense has the corresponding property. (Query: is this more than another way of saying that the referent is of the corresponding sort? I think not.) Rather than taking a copy of the sort data to realize the properties of the referent, it seems more economic of storage simply to allow access from the referent to the data structure that realizes the sense of the individuated phrase, which I have illustrated by the dotted arrows of line 3 of the diagram.

Cross-Reference and Data Migration

In line 4 of the diagram, I have shown the coalescence of two referents as enabling both senses to be accessible to that referent. In the diagram, one sense is accessible only via the other. I think this is misleading and the referent would have access to a ring of senses. I am doubtful whether marking consequently incoherent sorts (illustrated in line 5) can be separated from taking a copy of the sort data common to both senses (illustrated in line 6). Indeed it would seem space efficient to move a copy of the consistent sort data to the referent data-structure, and add a pointer from the sense data-structures to access it as sense-data.

It would seem sensible to store with this data a note of which phrase or phrases it is derived from, in case the application of rules of deduction--which I see as taking place in the realm of referents--led to an inconsistent referent. If it was then deemed advisable to break a cross-reference, the data derived solely from the sense of the phrase concerned can thus be recognized and removed from the referent data structure. In any case, as illustrated on line 7, the data structure that realises the referent will grow. If this is incompatible with some components from one or more sense, these can be returned to the sense data structures they came from, as is illustrated in line 8.

Without detailed experiment and design, it is not possible to be more than tentative about what kind of data is stored in the realm of sense, what in the realm of referents and how the linkage between the two should be set up.

I am sure however that, in any implementation, storage and access efficiency will encourage the migration of data that is duplicated in more than one sense data-structure to the single referent data-structure. I am also sure that this migration will not take place before there is any need for it,

i.e. as a result of some successful cross-reference algorithm rather than as the result of individuation.

I am less certain that components of the sense data-structures are actually lost and replaced by pointers to corresponding components of the referent structure. It seems to me that we must distinguish between or-alternatives which realize the residual ambiguity and the and-components that realize, as it were, properties a narrower sort. Further we must consider what data structures should be used to realize paramaterised sorts properties. It is far from clear that the data-structures that are convenient as sense data-structures are precisely those that are convenient to realize the same information as referent data-structures; though the more this was so, the easier it would be to handle the data migration and avoid duplication.

Data Migration and the Referential Model

Although data only migrates to the referent data-structures when it has good reason to do so, which, above, we proposed was after successful cross-reference, we envisage that ultimately all will do so. As we say: at last, the text is completely understood, even though the implications may not be understood. By this stage, any referent referred to by more than one phrase will be realized by a data structure that realized all that is known about that referent by the processes described in the previous section. I envisage that the referents, hitherto empty data-structures, that have not been involved in cross-reference will now have transferred to them the sense data-structure, or whatever transform of this realizes the same information as a referent structure as the sense from which they were individuated. This is shown in line 9 of the diagram, which also shows the sense data-structures deleted.

Since the understanding process is completed, any cross-references

that have been set up will not be revised, at least without rereading the text, and therefore we have no need to revise some proposed cross-reference.

Indeed, all information that embodies our understanding of the text is now in the realm of referents. We can, line 10, break the links between the text and the referential model that we understand by it, and the elements of this referential model provide possible referents to cross-reference to when reading further text.

In the description above we spoke as though the various algorithms were successively applied to the text in the order shown in the diagram. This is most unlikely. How we advance our understanding of a text is an apparently random process, and this, I feel sure, would be reflected in the strategy with which the various algorithms were applied to the various parts of the evolving data-structure.

Historical and Philosophical Afterword

When I began to write this paper, I thought that what information was realized in the realm of sense and what in the realm of referents was determined by what kind of information it was; that this distinction was fixed throughout the understanding process, albeit on a pragmatic basis that would vary according to the precise system in use. I was, also, not clear how to distinguish the data-structures that realized information in the two realms.

I then isolated the notion of individuation as a necessary process between the understanding processes that could be done based on information that could be found locally in some particular phrase, and those that use what could only be found outside that phrase. These later processes use information available either as the result of cross-reference, or as a result of the immediate context within which the phrase occurs. Thus, it became clear that the data structures could be seen as divided into two parts, the one,

produced by individuation that could, at least potentially, be accessed from more than one phrase, the other, initially attached to each phrase and leading nowhere, but by individuation and cross-reference leading to a data-structure of the first kind. This implies that more than one phrase via a data-structure of the second kind particular to that phrase, could reach a data structure of the first kind.

It thus seemed that, if a data structure of the second kind were seen as realizing information about the sense of the phrase and a data structure of the first kind were seen as realising information about the referent of the phrase, the distinction between the realm of sense and the realm of referents could be based on whereabouts in the data-structures it was realized and not on what kind of information was being realized. Moreover, this structural distinction very closely paralleled Frege's original distinction between sense and referent.

Considerations of efficient storage and retrieval of the data that realized this information then led to the introduction of the notion that this data would migrate from the sense data-structures to the referent data-structures. I see this as the original and interesting idea in this paper, both as to how the data-structures of a language understanding system should be organized and for the philosophical implications of this organization.

These implications are threefold:

1. The distinction between the realm of sense and the realm of referents can be made not by what kind of information is held in each realm, but by the access paths to that information; moreover, this distinction very precisely mirrors Frege's original formulation of the distinction between sense and referent.
2. Consideration of efficient storage and access to this

information leads to the notion that this information will tend to migrate from the realm of sense to the realm of referents, so that in what realm some particular kind of information can be found is not fixed throughout the understanding process, but changes as the understanding process advances.

3. While the understanding process is based solely on local information within the phrase being considered, this information is held in the realm of sense. This accords with Frege's views that it is in the sense of a phrase that depends on the sense of its component sub-phrases.

Cross-reference causes this information to migrate to the realm of referents so that deductive processes based on the truth of assertions of which the phrase is a component sub-phrase, or from general knowledge of the properties of the world, will take place in the realm of referents.

Finally, all information reaches the realm of referents and we are left with a purely referential model that embodies our understanding of the text.

At this stage we can lose the information in the realm of sense and then sever the links from the phrases of the text to this referential model.

It is, at the moment, an open question what data structures can realize the information involved. Practical design of experimental systems and examination of existing systems is needed to throw light on the many problems involved and how, let alone how best, they can be solved. The two that seem to me to be both pressing and interesting are: first, how information about a phrase can be related to information about its sub-phrases,whether in the realm of sense or in the realm of referents; second, whether information realized in the realm of sense, should, when it

migrates to the realm of referents, be held in the same form in the realm of referents.

The second question leads to the further interesting question of what data-structures can be used to realize a purely referential model. I suspect that this may well be very different from how this same information is realized when scattered in the various sense data-structures of the phrases that have led to the building of this model. Indeed, the catalogue of what individuals are a part of this model does not seem to me to have any analogue in the realm of sense.

I suspect that, after migration of information from the realm of sense to become attached to the various individuals in the realm of referents, there is a further process that rearranges how this information is realized and builds the catalogue needed.

However, the testing of this and other hypotheses that may be formulated must wait on the result of further work.

Epilogue

A) My description of Wilks' language understanding system maligns it in a slight, but conceptually important way; I imply that individuation, and thus the transition from sense to reference is totally missing in the system. This is not so; without individuation, cross-reference and deduction could not take place, and, in that system, algorithms to do so exist and are put to work, as in the example given above about the soldiers and the women.

Nevertheless, I would claim that these are a very peripheral and conceptually murky part of that system whose importance, to my mind, is vastly underrated by those concerned with it.

B) I may inadvertently have given the impression that I hold that there exists some essential referential model that embodies the understanding of some text. This is very far from my present opinion, which is that the meaning/understanding of a text is a vague notion, which is made more precise by the pragmatically based notion of what that meaning/understanding is to be used for. For any particular use, some particular referential model will be, pragmatically, adequate and appropriate; again, how best to realize this by particular data-structures within some machine so that the algorithms that embody that use may be realized by efficient programs that access that data efficiently is a further pragmatic problem to be solved.

Nevertheless, I am of the opinion that these various pragmatically determined referential models are related quite closely one to another, and that the relation: "λ x, y : Ref Models. x can be used for all purposes for which y can be used (and possibly more)" is a partial ordering on these models.

Moreover, tentatively, I think that the space of referential models is closed under finite union so that:

there exists z: referential models such that $z > x$ and $z > y$.

where $>$ is the partial ordering of the previous paragraph; further that if $x > y$, then in some sense y is a homomorphic image of x.

The essentialist view that I rejected at the outset of B) can be reformulated within this terminology. I am skeptical of the existence of (let alone the construction, of or usefulness of construction of) a maximal referential model, of which all other referential models are homomorphic images.

This last paragraph should be understood metaphorically; there are many philosophical traps that it opens, and not all the monsters there released are paper tigers.

C) From the reasoning in B) and the fact that, in the body of the paper, language-understanding is considered without a pragmatic context of intended use of that understanding, it follows that there are three important gaps in this paper.

i) No particular data structures are proposed for realizing the senses of the phrases of some text, nor to realize the referential model that embodies the understanding of a given text. Consequently:

ii) No algorithms are proposed to create such data structures from a given text, nor is their capacity for doing so demonstrated.

iii) No algorithms are proposed to enable that intended use to be performed by the use of the referential model realized by these data-structures to be demonstrated.

Given an intended use, proposals must be formulated for (i)-(iii) and their adequate performance demonstrated. This paper attempts to lay out a general methodological and conceptual structure for doing so.

D) Speculatively and entirely intuitively, I think I understand in general terms

the conceptual structures that are realized by the trees that are associated with words and phrases in Wilks' system and also, tentatively, it seems to me that in principle they are adequate, at least as a starting point, for representing the information needed in the realm of sense in a usable way.

Contrariwise, I am completely bewildered by what conceptual structure it is that is realized by semantic nets, which seem the current front-runner for representing information in the realm of referents; indeed, (cf What's in a link), it seems to me that these directed graphs realize different conceptual structures in the intent of the various proponents of their use.

I am very skeptical indeed,to the limited extent that I can perceive them, of the adequacy of these conceptual structures for representing the information needed in the realm of referents, though clearly, they almost certainly have contributions to make in discovering adequate conceptual structures as a starting point to represent the information in a referential model.

I am less unhappy about scripts, frames, et al.to represent the information needed to embody our general knowledge of the regularities of the world we know.

I must again emphasize that these views derive from subjective intuitive introspective imagination and not from objective experiment and should be valued accordingly.