

THE LOGICAL ANALYSIS OF LEXICAL AMBIGUITY

David Stallard,
BBN Laboratories Inc.
10 Moulton St.,
Cambridge, Mass.
02238

Abstract

Theories of semantic interpretation which wish to capture as many generalizations as possible must face up to the manifoldly ambiguous and contextually dependent nature of word meaning.¹ In this paper I present a two-level scheme of semantic interpretation in which the first level deals with the semantic consequences of syntactic structure and the second with the choice of word meaning. On the first level the meanings of ambiguous words, pronominal references, nominal compounds and metonymies are not treated as fixed, but are instead represented by free variables which range over predicates and functions. The context-dependence of lexical meaning is dealt with by the second level, a constraint propagation process which attempts to assign values to these variables on the basis of the logical coherence of the overall result. In so doing it makes use of a set of polysemy operators which map between lexical senses, thus making a potentially indefinite number of related senses available.

1 INTRODUCTION: LEXICAL ASSOCIATION IN A COMPOSITIONAL SEMANTICS

A tenet now held with some force among formal semanticists is that the meaning of a complex natural-language expression should be a function of just two things: the meanings of the parts of the expression and the syntactic rule used to form the expression out of those parts. Systems such as Montague Grammar [9] give phrases like "former senator" compositional treatments by first translating them to an expression of intensional logic, and then giving this expression a model-theoretic interpretation in the usual way. The practical relevance of this goal to work in natural language processing is clear: for any application domain, maximum coverage could be obtained from the same domain-independent set of rules, needing only to add the relevant entries, with their primitive associations of meaning, to the lexicon.

An obvious technical issue for this program is raised by the phenomenon of lexical ambiguity. This problem is not one that has been particularly addressed in the Montague Grammar literature. The most obvious approach is simply to make alternative lexical senses separate entries in the lexicon, and to allow these disambiguated lexical items to give rise to separate syntactic and semantic analyses. The computationally unattractive consequences of this are quite clear: the same work must be done over again for each variant.

An alternative class of proposals defers the lexical part of the analysis until the rest is done. Hobbs [5] has presented the most detailed general treatment of this type to date. This treatment simply associates each ambiguous lexical item with the logical disjunction of its separate senses. Standard reasoning techniques (such as theorem proving) can then be applied. The problem with this approach is that it is simply not correct. This may be most straightforwardly seen in yes/no questions that contain an ambiguity. For example, suppose the the ambiguous verb "have" is to be treated as the disjunction of the predicates POSSESS, PART-OF, etc. Then the answer to the question "Does the butcher have kidneys?" must always come out "yes", because the second alternative is (assumably) true regardless. This method goes wrong because the issue in resolving ambiguity is determining which possibility was intended, not which possibility is true.

A more correct approach is due to Landsbergen and Scha [8] and implemented in the PHLQA1 system. There, the result of semantic interpretation is an expression of an ambiguous logical language called EFL (for English-oriented Formal Language). During semantic interpretation each lexeme is assigned to one (possibly ambiguous) descriptive constant of that language, which is later mapped, via local translation rules, to one or more expressions of an unambiguous logical language called WML (for World Model Language). The result is a set of complete WML translations of the entire EFL expression, from which sortally anomalous alternatives are subsequently eliminated.

The PHLQA1 system, while handling homonymy acceptably, does not address the problem of *polysemy* - the presence of an indefinite number of related senses for a single word. Consider the polysemous lexeme "mouth", which is used differently

¹The work presented here was supported under DARPA contract #N00014-85-C-0016. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States Government.

in the phrases "mouth of a person", "mouth of a bottle", "mouth of a river", and "mouth of a cave". Surely the same logical relationship is not involved in each of these cases. Generalizing the meaning of the word will not help either, for if we tried to re-define "mouth" to mean just any aperture, we would lose our ability to refer to human "mouths" independently of other parts of the body. Enumerating these separate senses with separate translation rules does not look like a very promising approach either, since it is not at all clear that the list above could not be continued indefinitely. The problem with such an approach to meaning seems to be that it is too discrete: in linguistic terms, it does not "capture a generalization".

This paper presents a method of dealing with lexical meanings which does seek to capture the generalizations implicit in polysemy. The complexes of meanings associated with polysemous lexical items are generated, structured and extended by a kind of "grammar" of word meaning: a set of operators which take descriptive constants of a meaning representation language onto other descriptive constants or expressions of that language. These operations include not only metaphorical and metonymic extension of the word sense, but "broadening", which allows a word to refer to a wider class of items than before; "exclusion", which removes from the denotation of a word the members of a particular subset thereof; "narrowing", which narrows the denotation down to a particular subset. Each word is assumed to have a core sense (or in the case that it is homonymous, several core senses) from which extended senses can be derived by recursive application of the operators.

Related to the issue of lexical ambiguity, if traditionally studied apart from it, are the problems raised by nominal compounds and metonymies. Here the problem is determining the binary relation which has been "elided" from the utterance. This could in principle be any relation; a translation rule approach cannot help here. Novel metaphorical uses of a word, such as the substitution of an individual for a whole class, will also escape such an approach. The point about all three of these phenomena is that they essentially create new lexical senses. The productivity of this process suggests that the established senses of polysemous lexemes may be generated in the same way.

A key innovation of this work is to treat *every* non-logical word as being potentially ambiguous. Thus semantic interpretation initially assigns to each such lexical item not an ambiguous constant, but a free variable capable of ranging over the appropriate type of a higher-order intensional logic [4]. These free variables are restricted to range not over an explicitly enumerated set of logical expressions, but over a potentially infinite set of them which is recursively enumerable by the polysemy operators. Obviously, the core sense itself (and other established senses) are not excluded as candidates. A separate con-

straint propagation stage then assigns appropriate descriptive constant values to these variables based on the sortal coherence of the whole expression.

This two-stage method of semantic interpretation will be seen to have an advantage over one not discussed so far: a single stage method which not does not allot a separate role to lexical semantics or pay close attention to compositionality, but rather seeks to interpret distinct patterns like "mouth of a cave" as a whole. Besides suffering from the same lack of generality criticised above this latter method encounters difficulty when an ambiguous word-form and a pronoun or trace are combined together. A second constraint propagation stage enables the dependence of word meaning on context - specifically, on the meanings of other words and the referents of anaphors and deixis in the utterance - to be captured. The computational effect is that search can be cut down in a space that is essentially a cartesian product over the ambiguous elements of an utterance.

2 THE NOTION OF A "LOGICAL VOCABULARY"

Lexical association cannot be considered apart from a notion of "logical" or "conceptual" vocabulary - the set of descriptive constants of a logical language which are available for making such associations. This notion may be identified with the "domain model" or "conceptual model" of such systems as PHLIQA1 [11], TEAM [3] and IRUS [1]. Logical vocabularies, or "domains", are what the polysemy operators work with. The present section lays down the representational structure which the next, dealing with the polysemy operators themselves, will make use of.

Let a "domain" be defined as a set of descriptive constants and axioms involving them, subject to three conditions: (1) The descriptive constants are such that a specification of each of their extensions gives a "state of the world" relevant to the domain (2) The axioms are such that they constrain which states of the world are possible or allowable (3) The axioms do not define the constants with the biconditional, but with one-way implication only, thus leaving the constants primitive. If complete definitions of constants via lambda-abstraction is allowed it is only as a technical convenience; these are to be regarded as "extra".

The latter condition (3) captures the important fact that domains are not definable in terms of other domains. Thus expressions cast in logical vocabulary D_A cannot be directly used to refer to states of affairs, etc. expressible only in terms of logical vocabulary D_B . This has an impact for natural language question answering systems in which D_A is the notions of ordinary language and D_B the logical vocabulary of some technical domain. In this case, only lexical items specially invented for the technical domain (such as "JP-5", a particular kind of military jet fuel)

have an unproblematic lexical association in terms of D_B . Obviously not all the words a user employs will have this characteristic, nor will all the constants of the technical domain be lexicalizable in this way. In other cases the notions of D_A will have to be mapped to those of D_B , in some way that is not yet specified.

A common occurrence is for lexical items available in regular English to be employed to bridge the gap, in such a way as to multiply their effective ambiguity. Consider a question seeking to find ships with a certain offensive capability: "What ships carry Harpoon missiles?". On a literal interpretation of the word "carry" the predication of the sentence is satisfied whether the ships "carry" the missiles as weaponry or as incidental cargo, yet of these only the first alternative is the desired one. If the query were instead "What ships carry oranges?" the second alternative is the preferable one. The resultant "splitting" of lexical senses can be regarded as a form of ambiguity generated by the contact between logical domains.

Other kinds of mapping between notions of different domains are more complex, not taking place along the lines of greater or lesser specificity, but involving instead another kind of mapping that is really tantamount to metaphor. A phrase like "in Marketing", for example, is not locative in the literal sense of location in space but rather makes use of a metaphor having to do with this notion. Here the initial domain is that of space and spatial inclusion, while the final one is that of, say, fields of employment or expertise.

The formal representation of metaphor used in this work is that of Indurkha [7]. Indurkha identifies a metaphor with the formal notion of a "T-MAP": a pair $\langle F, S \rangle$ where F is a function mapping descriptive constants from one domain to another and S is a set of sentences which are expected to carry over from the first domain to the second. A metaphor is "coherent" if the transferred sentences S are logically consistent with the axioms of the target domain, "strongly coherent" if they already lie in the deductive closure of those axioms.

Depending on the formal language used to represent the statements S , one may encounter computational difficulties (i.e. decidability) with this program. One way around this is not to use predicate calculus (as Indurkha does) but a language that is more restrictive than predicate calculus. For the price of surrendering complete expressive power one gains the advantage of deductive tractability.

One system which may be used for this purpose is the NIKL [10] system, in which only a few types of axioms can be encoded. A descriptive constant *subsumes* another of the same complex type if its extension is always a superset of the other. Two constants are *disjoint* if their extensions are always disjoint. (Note that respective subsumees of the two constants "inherit" this disjointness.) Relations of more than one argument have sortal (one-place

predicate) restrictions on their argument, thus stipulating that the extension of the relation will always be a subset of the cartesian product of the extensions of the sorts. Finally, a one-place predicate P *restricts* a binary relation R to be Q if the image under R of each member of P 's extension is a member of the extension of the second one-place predicate Q . In what follows I will treat this operation as restricting the form that the extension of the relation R may take on, so that the placing of constraint P on the first argument results in a propagation of the constraint Q on the second argument.

3 THE LEXICAL CONSTRAINT MODULE

3.1 Overview

In this section I present a solution to the multiple problems of ambiguity posed by a natural utterance. Added to the architecture of semantic interpreter, discourse model, lexicon and domain model is a new component - the lexical constraint module. It accepts from the semantic interpreter a logical form containing free variables of higher-order and constructs from it a constraint graph structure in which such variables are connected in accordance with the syntactic structure of the expression. This structure is then used in a constraint-propagation process that attempts to assign descriptive constant values to the expressions. The lexicon in this scheme stores for each non-logical word an extendable polysemic complex (or complexes, in the case of homonymy) of logical associations. I shall describe assumptions about the semantic rule set-up as I go along.

In making these assignments, the module applies a "maxim of coherence". That is, we assume that the user will not deliberately speak nonsense to us, use terms redundantly, or make use of elaborate means to refer to the null set. A coherent outcome is one where the descriptive constants being applied to the same terms (bound variables and individual constants) are not sortally disjoint. This may not always be achievable with the core sense of words. When it is not, a set of "polysemy operators" is invoked to re-interpret a lexical assignment in such a way as to make sense of the expression.

I will first consider an example where no such re-interpretation is required. For the utterance "John has a car", the following logical form is given as input to the constraint module:

$(\exists x (\underline{\text{car}} x) \ \& \ (\underline{\text{have John}} x))$

The underlined symbols are the free variables. Suppose the main verb "have" to be homonymous between the various predicates PART-OF, OWN, AFFLICTED-WITH. The last of these is eliminable because the argument sorts it requires and the the sorts given to it do not agree: physical objects and

diseases are disjoint sets. Such surface inspection of argument sorts is not the only source of constraint, however. For some relations a particular constraint on the first argument causes a constraint on its second argument. Thus, the alternative PART-OF is eliminable because the parts of an organism must themselves be organic material, something clearly disjoint with artifacts like cars. The constraint graph is now satisfied, and we are left with:

$(\exists x (\text{CAR } x) \ \& \ (\text{OWNS JOHN } x))$

3.2 The polysemy operators

We now proceed to overconstrained cases in which potential assignments are in conflict, and re-interpretation by the polysemy operators is required. For the first pair of such operators, *generalization* and *exclusion*, we will make use of the Montague Grammar notion of *universal sublimation* [2]. A universal sublimation of a concept A is just the set of properties which are true of all A's members, or:

$(\lambda P (\forall X A(X) \rightarrow P(X)))$

Generalization and exclusion operate upon lexical senses by modifying their universal sublimations and looking for the alternative meaning (if any) of the word that most closely corresponds to this new set.

As an example of *generalization*, consider the phrase "plastic silverware". While in literal terms this is oxymoronic, one often sees it used to refer to plastic eating utensils, and in situations where only these items are available, the word "silverware" alone may be used to denote them. Obviously for such speakers the class EATING-UTENSIL is available as an extended and generalized sense of "silverware". The initial representation would be:

$(\lambda X (\text{and } (\text{plastic } x) \ (\text{silverware } x)))$

A portion of the sublimation of the concept SILVERWARE is the set {MADE-OF-SILVER, EATING-UTENSIL}. Of these, it is the first property that is disjoint with PLASTIC and a new sublimation is constructed which excludes it. In the partial representation above, this new sublimation is just the class EATING-UTENSIL itself.

Exclusion takes a lexical sense onto one from which particular sub-senses have been explicitly excluded. Consider the sentence "The Thresher is not a ship, it's a submarine", or, to be free about its logical form:

$(\text{CONTRAST } (\text{ship Thresher}) \ (\text{submarine Thresher}))$

If we assign the core meanings to these words this is nonsensical, since SUBMARINES are, by definition, SHIPs as well. The expression coheres if whatever is assigned to ship excludes SUBMARINE. We form a partial sublimation {SHIP, ~SUBMARINE}, and find

corresponding to it the alternative sense of "ship", SURFACE-SHIP.

A surprising number of words have such alternative exclusionary senses, among them "axe", where HATCHET is excluded; "animal", where HUMAN is excluded; and "blue", where TURQUOISE (and other off-color shades) is excluded. The phenomenon seems to be that a specialized term for some distinguished subset of a concept comes to be the preferred term for members of that subset. The all-embracing word can still be used, but it comes to have a sense which is contrastive with these distinguished subsets. From the impression made by a Venn diagram of the set and its excluded subsets we might call this "cut-out" polysemy.

One wonders if certain phenomena which have been described as ill-formedness might not in fact be instances of this sort of polysemy. Goodman, for instance, uses the actual word pair "blue" - "turquoise" as an example of "miscommunication" (Goodman85). What seems more plausible however, is that the speaker describing a turquoise object as "blue" is not really misspeaking, but is rather using the word "blue" in the more inclusive sense which embraces all shades of the color.

Metonymic extension re-interprets a predicate by interposing an arbitrary, sortally compatible relation between an argument place of the predicate and the actual argument. An example can be seen in the command "Highlight C3 tracks", where "C3" is a predication made of ships and "tracks" are trajectories of ship positions, traced out on a screen. Obviously, on literal interpretation, this utterance does not cohere, since physical objects (ships) and graphical objects (tracks) are disjoint. We have:

$(\text{HIGHLIGHT } \wedge (\lambda X (\text{c3 } X) \ \& \ (\text{track } X)))$

The categories SHIP and TRACK have too many clashing properties for generalization or exclusion to prevail. Instead, the two clashing elements are reconciled by finding a function or relation reaching between SHIPs and TRACKs (or subsuming categories) and metonymically extending one of the items with it. The extended meaning of "C3" can be expressed by:

$(\lambda X (\exists Y (\text{and } (\text{SHIP } Y) \ (\text{SHIP-TRACK } Y X) \ (\text{C3 } Y))))$

In any usage of the metonymy operation there is a choice about which of two clashing elements to extend. In this case it would also have been possible to have metonymically extended "track" instead of "C3" in this example. The resultant expression would then be a set of ships instead of tracks - clearly not what is wanted here. It would moreover not be an immediately coherent one itself, since "highlighting" can only be done on graphical objects. More importantly, it would seem to be that metonomies are less likely to

shift the head noun meaning, since this changes the sortal category of what is being referred to and operated upon by the utterance. This seems to be particularly strong when the head noun's meaning has an underlying functional role, as does "track" in this case.

Note that many words which at first appear to have unitary senses are actually better described in terms of metonymic complexes. Thus, "window" can be used to refer to its constituent pane of glass, its sash, or the opening around it. Similar examples can be seen in "light", which can be used to refer to the actual electromagnetic radiation or the device for producing it, and "bank" (in the fiscal sense), which can be used to refer to the building or the financial institution itself.

Metaphorical extension operates not by shifting an argument place of a predicate, but by shifting the predicate itself. Capturing the generality in the meaning of "mouth" in the example of section 1 involves capturing a class of metaphors involving that concept. Classes of metaphors are described by the notion of a *parameterized T-MAP*, in which the mapping function *F* and set of sentences *S* are not completely specified, but may instead have missing elements which must be solved for. Let "mouth of the cave" be given by:

(mouth (iota x (cave x)))

The functional constant MOUTH is restricted to operate on individuals of the class ANIMAL, so the above is incoherent on literal interpretation. A metaphorical re-interpretation must select certain constants for the mapping function *F* and certain facts *S* which carry over to the new domain. Two such facts are:

**SUBSUMES (ENCLOSES-SPACE, ANIMAL)
SUBSUMES (OPENING, MOUTH)**

In this use of the word "mouth" it is operating on individuals of the class CAVE instead of ANIMAL. One element of the mapping function *F* is thus the pair (ANIMAL,CAVE). In order to determine the relationship that the word "mouth" really means in the example we must solve for a function variable *P* which MOUTH is mapped to. This function must be sortally coherent with CAVE; it is the righthand member of the second ordered pair of the mapping function *F*.

The sentences to be transferred are:

**SUBSUMES (ENCLOSES-SPACE, CAVE)
SUBSUMES (OPENING-OF, P)**

Of these, the first is not only not inconsistent, but true. One descriptive constant of the geological domain which is obviously not incoherent with CAVE is the function CAVE-ENTRANCE. If this function is used in place of *P* the second sentence is satisfied as well.

An important metric of metaphorical plausibility is how much structure in *S* is transferred from source to target domain versus how many descriptive constants

are mapped via the function *F*. In the present example the ratio is one. Clearly if this ratio is high the metaphor is stronger and more plausible; if it is low the metaphor is less so.

3.3 Nominal Compounds

Nominal compounds are treated by assuming that the semantic rules formulate their interpretation with a free binary predicate variable standing in for the relation which must be determined to complete the interpretation of the compound. Interpreting the nominal compound thus becomes solving for this predicate variable. This variable is initially unconstrained except by the sorts of the noun meanings it connects.

A problem with some nominal compounds is that they seem to violate the restrictions imposed by their component parts. For example, a "staple gun" is not a weapon at all, and would thus on some treatments have to be treated either idiomatically or as a completely incoherent expression. With the approach presented here, however, the polysemy operators can be invoked to find a re-interpretation of the words for which a solution does exist. The word "gun" can be re-interpreted to discard the clashing property of shooting bullets only, and to denote in this case the wider class of devices that eject objects of whatever type.

An important point about nominal compounds is that they cannot be treated extensionally. A *soup pot* is still such whether it currently contains something different from soup, or indeed whether it contains anything at all. Clearly, the relation to be solved for in a nominal compound may in general be a non-extensional one between *kinds*. Such a relation may in turn have a meaning postulate which dictates which actual entities (such as the actual soup) may be related at which indices of time. This phenomenon would seem to pose a problem for Hobbs and Martin [6], who view as a sub-problem resolving the reference of the "lube oil" in the compound "lube oil alarm". One can imagine a "lube oil alarm" which only sounds when all the lube oil is gone.

3.4 Effect on Anaphora Resolution

Even after syntactic and pragmatic considerations have been taken into account, the decision on the correct referents for anaphora cannot take place independently of considerations of word meaning choice. Consider the following two sentences:

- (1) The table is in that building
- (2) It is a bank.

The proper referent of "it" in (2) is constrained by the predication made by the ambiguous lexical item "bank", namely that it either be a RIVER-BANK or a BANK-BUILDING. Neither is sortally coherent with TABLE, the referent described by "the table" is elimin-

able. The only thing left is the individual described by "that building" and since BUILDING, being an ARTIFACT, is disjoint with RIVER-BANK, the proper sense of "bank" is BANK-BUILDING and the referent of "it" is "that building".

3.5 Algorithm and Heuristics

The algorithm used by the lexical constraint module is a search loop consisting of just three parts - *tentative assignment*, *constraint propagation* and *re-interpretation*. On the first iteration tentative assignment constrains each word-variable with its core logical sense, or the set of its core senses if it is homonymous. These serve as entry points to the polysemy complexes. Variables associated with anaphors are initially constrained by whatever pragmatic and syntactic (such as C-command) considerations are seen to apply. The variables associated with nominal compounds are initially left unconstrained.

Thereafter, constraint propagation may end up in one of three states: *satisfaction*, in which case the module returns a single logical expression; *underconstraint*, in which case there is an ambiguity with which the user must be presented; *overconstraint* in which case re-interpretation is invoked to search for an interpretation which coheres.

The most important issue in performing re-interpretation is controlling the process that the system does not "hallucinate" arbitrary meanings into an expression. The control heuristics include:

1. consider overconstrained variables for re-interpretation first
2. prefer generalizations and exclusions which modify a small number of properties
3. prefer metaphorical extensions with a high ratio of plausibility (as in Sec 3.2) and minimize the number of "augmentations" and "positings" [7]
4. avoid multiple re-interpretations of the same item
5. prefer re-interpretations to already established polysemous senses instead of creating new ones

In Hobbs' work [6] control turns on a notion of a "cost-function" associated with the lengths of proofs. The notion of "minimality" in that work has some similarity to the heuristics above, which seek to avoid arbitrary re-interpretations of lexical meanings by preferring conservative re-interpretations and discouraging multiple ones.

The creation by the polysemy operators of new sense for a word can effectively be regarded as a kind of "learning". Thus, given the sentence "That's not a lion, that's a lioness" the system could deduce (via the

exclusion operator) that an alternative sense of "lion" means a male lion only. One should not presume, however, that the discovery of new lexical senses will occur on a constant basis. The last heuristic above is therefore an important one.

4 CONCLUSIONS

This component will be implemented in a future version of BBN's JANUS natural language understanding system. Included in this system will be a unification parser with a large grammar and a new and improved semantic interpreter.

I have tried to show how a compositional semantics need not be incompatible with a context-dependent notion of word meaning by making a division of labor between the rule-to-rule translation of syntactic structure and the complex semantics of lexical items. I shall even go so far as to say that such a division of labor is necessary for the compositional program to succeed. A component which takes into account the creativity of lexical meanings and which utilizes knowledge representation and limited inference not only gives word meaning its proper place in a modular system but also has the potential of extending coverage and flexibility beyond what is currently available in natural language systems.

Acknowledgements

I would like to thank Remko Scha for his many useful comments on this work. I would also like to thank Erhard Hinrichs and Bob Ingria for their comments and encouragement, and Jessica Handler for valuable linguistic data.

References

- [1] Bates, Madeleine and Bobrow, Robert J. A Transportable Natural Language Interface for Information Retrieval. In *Proceedings of the 6th Annual International ACM SIGIR Conference*. ACM Special Interest Group on Information Retrieval and American Society for Information Science, Washington, D.C., June, 1983.
- [2] David R. Dowty, Robert E. Wall, and Stanley Peters. *Introduction to Montague Grammar*. D. Reidel Publishing Company, 1981.
- [3] Barbara Grosz, Douglas E. Appelt, Paul Martin, and Fernando Pereira. *TEAM: An Experiment in the Design of Transportable Natural-Language Interfaces*. Technical Report 356, SRI International, Menlo Park, CA, August, 1985.

- [4] Hinrichs, Erhard W.
A Revised Syntax and Semantics of a Semantic Interpretation Language.
1986.
- [5] Hobbs, Jerry R.
Overview of the TACITUS Project.
In *Proceedings of the DARPA 1986 Strategic Computing Natural Language Workshop*, pages 19-25. The Defense Advanced Research Projects Agency, May, 1986.
- [6] Jerry R. Hobbs and Paul Martin.
Local Pragmatics.
In *Proceedings, IJCAI-87*. International Joint Conferences on Artificial Intelligence, Inc., August, 1987.
To appear.
- [7] Indurkha, Bipin.
Constrained Semantic Transference: A Formal Theory of Metaphors.
Technical Report 85/008, Boston University, 1985.
- [8] Landsbergen, S.P.J. and Scha, R.J.H.
Formal Languages for Semantic Representation.
In Allen and Petofi (editors), *Aspects of Automated Text Processing: Papers in Textlinguistics*. Hamburg:Buske, 1979.
- [9] Montague, R.
The Proper Treatment of Quantification in Ordinary English.
In J. Hintikka, J.Moravcsik and P.Suppe (editors), *Approaches to Natural Language. Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, pages 221-242. Dordrecht: D.Reidel, 1973.
- [10] Moser, Margaret.
An Overview of NIKL.
Technical Report Section of BBN Report No. 5421, Bolt Beranek and Newman Inc., 1983.
- [11] Scha, Remko J.H.
Logical Foundations for Question-Answering.
Phillips Research Laboratories, Eindhoven, The Netherlands, 1983.