NASA TM-85868

NASA-TM-85868 19840021409

FOR REFERENCE

NOT TO BE TAKEN FROM THIS ROOM

Understanding a Technical Language: A Schema-Based Approach

Pierre Falzon

July 1984

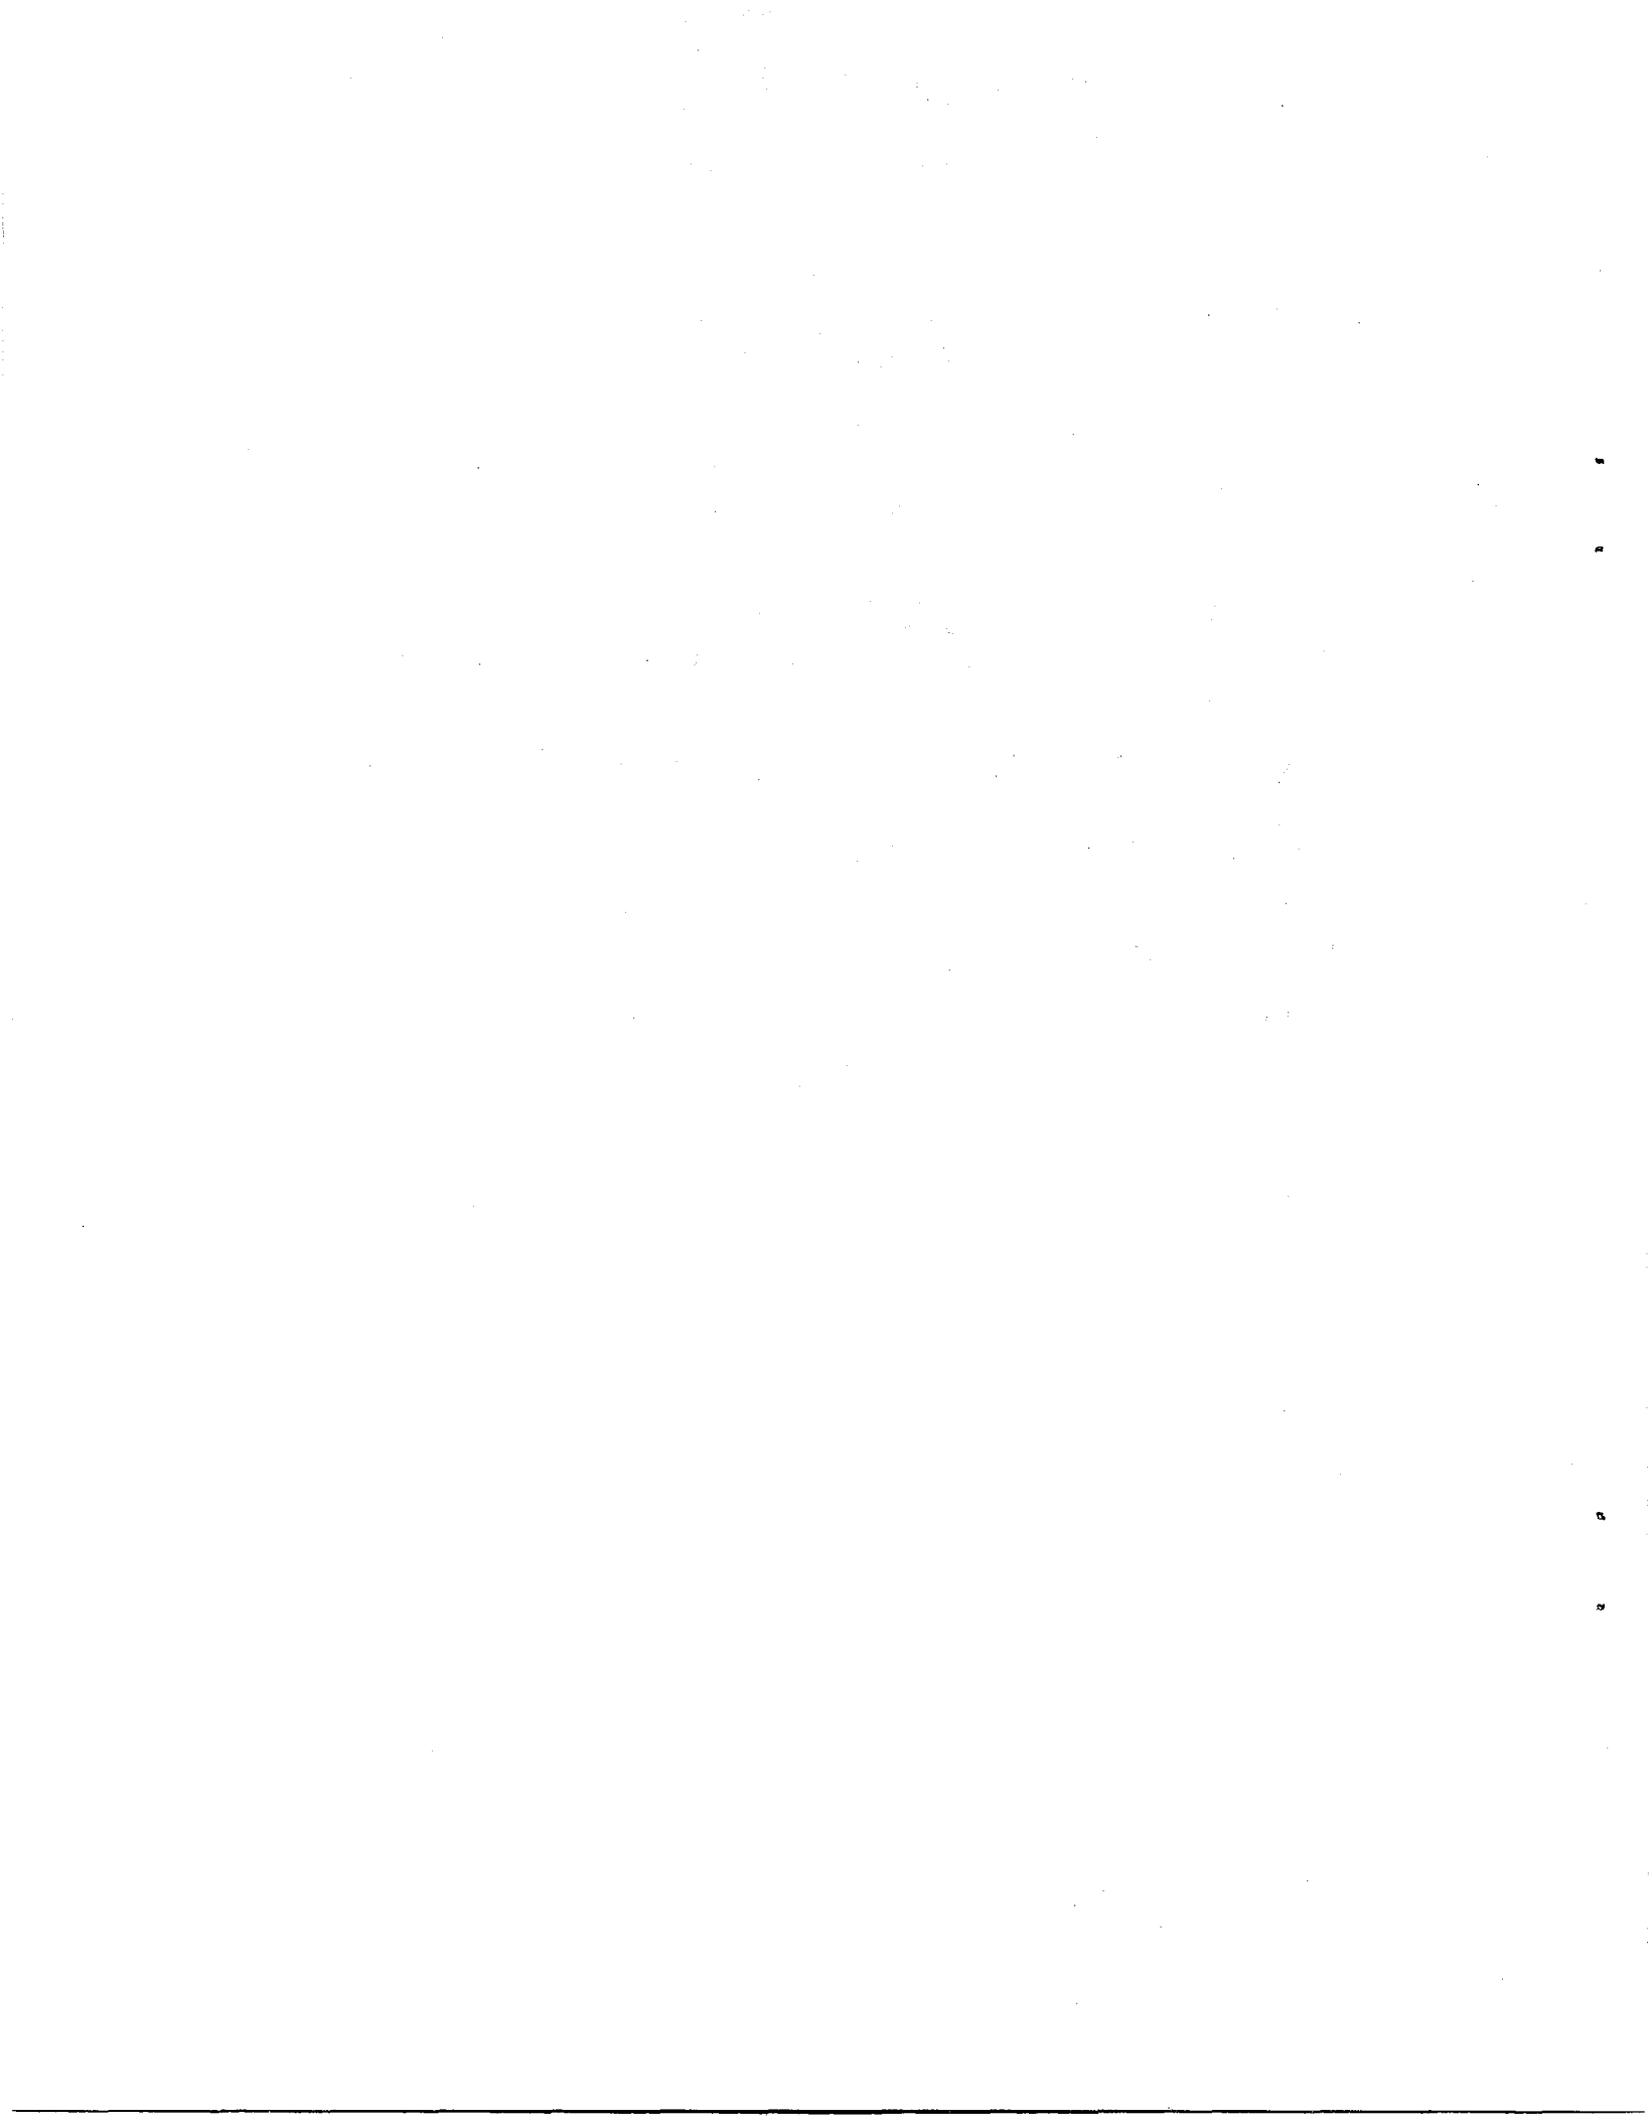# NASA

National Aeronautics and
Space Administration

Understanding a Technical Language: A Schema-Based Approach

Pierre Falzon, Ames Research Center, Moffett Field, California

N84-29478

TABLE OF CONTENTS

UNDERSTANDING A TECHNICAL LANGUAGE: A SCHEMA-BASED APPROACH

Pierre Falzon[*]

Ames Research Center

SUMMARY

Workers in many job categories tend to develop technical languages, which are restricted subsets of natural language. A better knowledge of these restrictions could provide guidelines for the design of the restricted languages of interactive systems. Accordingly, a technical language (used by air-traffic controllers in their communications with pilots) is studied. A method of analysis is presented that allows the schemata underlying each category of messages to be identified. This schematic knowledge is implemented in programs, which assume that the goal-oriented aspect of technical languages (and particularly the restricted domain of discourse) limits the processes and the data necessary in order to understand the messages (monosemy, limited vocabulary, evocation of the schemata by some command words, absence of syntax). The programs can interpret, and translate into sequences of action, the messages emitted by the controllers.

INTRODUCTION

The instructions on the shampoo bottle read: "For best results, wet hair with warm water. Gently work in the first application. Rinse thoroughly and repeat." Hill (1972) was struck by the ambiguity, lack of precision, and fuzziness of this text. Repeat from where? he wondered. He then readily proposed a "much clearer" version:

```
for best results
BEGIN
wet hair with warm water
FOR j:=1,2 DO
    BEGIN
    gently work in application (j); rinse thoroughly
    END
END
```

Hill adds that although he does not expect to see that on a shampoo bottle in his lifetime, he thinks that "it is something to be desired, far more than desiring to write plain English for computers." Hill even goes further, stating that "In my own Utopia, we shall be able to write instructions to people in programming languages, just as we do for computers."

One decade later, where do we stand? Much of the literature focuses now on the design of user-oriented systems, the language of which should be closer to natural language, although still restricted. There are two reasons for that interest in the user, neither of which has anything to do with a sudden improvement in precision and clarity of natural language. First, there is the change in user population: users are now often inexperienced with computers, and quite unwilling to learn anything

*Research Scientist, Institut National de Recherche d'Informatique et d'Automatique (INRIA), Rocquencourt, France.

about them. Computer experts now form only a small fraction of the prospective users of a system. Second, there are the problems that may arise if the interface between user and system is not easy to use. This is especially true in situations in which the main task to be accomplished is not the interaction with computers, but some other task in which some interaction with computers has become necessary. Think for example of the task of an aircraft pilot in a modern cockpit, or of the operator of a nuclear power plant.

The fundamental obstacle to a widespread use of computers was once their size and cost; now it is their lack of "friendliness" to the user. Not so long ago, Hill could assert that people had to adapt to computers, a comment that would now meet much opposition.

Nevertheless, the criticisms of natural language need to be taken into account. Some of them have long ago been stressed in a pioneer paper by Chapanis (1965). Language is currently being investigated by more and more human factors specialists, very much because of the reason already mentioned: the growing use of computers in our daily lives.

A first line of research has considered the possibility of using natural language as an interactive language. Although the last decade has seen impressive results in that field, the use of unrestricted natural language for computer interaction faces serious problems.

For example, powerful natural-language-understanding systems are major under-takings, and need a considerable amount of computer memory. This is a serious obstacle for small computer systems, but technological progress could change it.

Another problem is that very often these natural language systems look like a bulldozer trying to destroy a house of cards. There is often a huge difference between the complexity of the tool and the triviality of its application. For instance, and with all the respect due to Winograd's very sophisticated system (Winograd, 1974), it is somewhat disappointing to see that it can only function when applied to the very small world of block manipulation! And the worst is that the bulldozer sometimes cannot manage to destroy the house of cards!

A second line of research took a different approach: at least in the foresee-able future, man-machine interfaces will not use natural language, but some restricted dialect. Is it then possible to restrict the possibilities of natural language with-out greatly constraining the user? A number of authors have explored this question, studying the effects of different imposed restrictions of syntax and vocabulary (see, for example, Ehrenreich, 1981; Kelly and Chapanis, 1977). A corollary of this approach is the definition of appropriate vocabularies (see the studies on naming by Scapin, 1981, 1982; Rosenberg, 1982), and appropriate syntactic structures (Hammond et al., 1980).

However, another approach is possible. Instead of studying specific restric-tions of natural language, why not study the natural restrictions of specific lan-guages? In any work situation in which the operators have to communicate verbally, the language they use is not unrestricted natural language; the operators tend to build a specific language, molded by the characteristics of the task and its objec-tive. These task-oriented languages transform natural language into a dialect that is totally obscure to a nonspecialist, but entirely clear to the expert. Consider the following communication from an air-traffic controller to a pilot: "Intercept the 1-3-5 of Point Reyes and resume the SID and with the restrictions." This message is total nonsense to the nonspecialist, first because of the abbreviations (What is an SID?), second because of the technical meaning of some words (What are the restrictions?), and third because of the ideas involved (What does "1-3-5 of Point

Reyes" mean? I would presume it to be an object or a line, since I am supposed to intercept it.). Linguistic competence is not enough to understand a technical language.

One may argue that this is mainly a question of vocabulary, however, and that given an appropriate technical dictionary, the communications are just a sample of natural language. But there is more here than a vocabulary specialization. Several factors contribute to the modification of natural language in work situations; for example, workload, the necessity to avoid ambiguity, and the influence of a common field of work.

The workload tends to make the operators restrict the length of the messages and to concatenate several messages into one (Sperandio, 1969). The necessity to avoid ambiguity restricts the meanings of the words and the form of the messages. This is especially true in complex situations in which the risks involved are important; it often leads to the recommendation of a specific phraseology.

The influence of a common field of work makes the reference worlds and the goals of the participants the same. The restriction of the domain of discourse has two important consequences. First, given a sufficient knowledge of the domain, the possible topics are highly predictable. In air-traffic control (ATC) for example, it is not likely that one would hear "Would you mind passing me the salt?" One would expect to hear about levels, headings, and other flight-relevant matters. The operators are only interested in some of the properties of reality. Second, these topics are seen under a specific, distorted point of view (see Dupre, 1981, for an illustration of this point). Restrictions on the domain of discourse limit not only the number and the type of possible topics, but also the viewpoint from which those topics are considered.

Goal-oriented languages can then be thought of as being restricted, relative to natural language, in a number of domains: for example, vocabulary, syntax, field of discourse, and dynamics of the dialogue. A better knowledge of these "spontaneous" restrictions and of the way they are built could provide guidelines for the design of computer interfaces.

Very little work has been done in the above perspective in the human factors area, with the notable exceptions of the works of Thomas (1976, 1978), and of the series of studies conducted by Chapanis and his colleagues (see Chapanis, 1978, for a summary). In the psycholinguistic field, the research is rarely relevant; in fact, most of it has focused on noncontextual situations, trying to find general characteristics of language. However, there seems to be a recent change in this trend, with a growing interest in the influence of specific situations on the type of expressions (Gibbs, 1979, 1981; Clark and Lucy, 1975; Hupet and Costermans, 1982).

The research presented here follows these premises. It focuses on a very specific language, that used by air-traffic controllers in their communications to the pilots. A method of analysis is presented, and the results of that analysis are evaluated through the development of a language-understanding system.

Some Results and a Hypothesis

A first study (Falzon, 1982) was concerned with the vocabularies (French and English) and the forms of expression used by the controllers (if two messages are not composed of the same words in the same order, they are two different forms of expression).

Different measures on the use of words and of messages have allowed the design of restricted vocabularies, subsets of the total vocabularies the controllers use. The use of these restricted vocabularies allows the recognition of a large number of messages (a message is said to be recognized if all of its words belong to the restricted vocabulary).

This result stresses a first "natural" restriction made by the operators, concerning the vocabularies. Similar results have been observed by Michaelis et al. (1977) in laboratory experiments. The interesting point is that the vocabularies, though restricted, nonetheless allow much flexibility in the form of expression of the messages, since 60% of the different forms of expression in French (i.e., 528 different utterances), and 73% in English (i.e., 380 different utterances) are recognized.

However, the recognition performances of the restricted vocabularies vary according to the category of messages, and are very poor for some categories. The reasons for this seem to be linked first to the length of the messages ("rare" words are more likely in a long message than in a short one), second to the frequency of use of the category (in order for "conventional" expressions to appear, the category must be used frequently). Consider for example the "traffic information" category (the controller warns a pilot of the presence of another aircraft in his vicinity). The messages of this category are not very frequent, and tend to be lengthy. The restricted vocabularies are not large enough to include all the words used in these messages; according to the definition of recognition that has been given, these messages then cannot be recognized.

This last result has led to a change in our approach. Although the "traffic information" messages are not recognized, they are easy to spot: most of them include the two words "traffic" and "information." Thus, they can easily be categorized. Moreover, the type of information they mention is highly predictable: the pilot expects to hear the altitude, heading, and relative position of the other aircraft. A similar analysis can be applied to all categories of messages, each of which categories can be characterized by some "command" words and by a sequence of possible constituents. The messages of a category may be considered as a list of different instances of a common schema, as different actualizations of a single schema. As will be seen, the analysis of the different forms of expression provides the compulsory and optional elements of the messages, and the default values that are assumed in some cases.

The understanding of a message can then be (in a very approximate way) divided into two phases. In the first phase, the understanding process is data-driven: the words heard (or read) are processed and activate a previously stored schema. In the second phase, the process becomes conceptually driven: schemata are predefined representations which ask for specific information. Thus, the input is checked to see whether it can fill in the information slots of the schema. If this process is successful, the schema is validated. A more detailed presentation of schema theory is given in the next section.

Our hypothesis is that the universe of discourse is so restricted that syntax can often be neglected, provided that a sufficient knowledge of the task domain is available. This does not mean that we assume that the pilots never use their linguistic competence to understand the messages they receive; our point is only that it is possible to understand the messages without using much syntactical knowledge. The objective of this work is to test whether this assertion holds true, and to what degree. Since we are not dealing here with language in general, but with a specific dialect in use in a very specific context, we may feel free to use whatever analysis seems convenient, keeping in mind that we are not using (and certainly not building) any theory of grammar. In fact, it could even be said that a "standard" linguistic approach (if such a thing exists) would not be relevant. The interesting point is to see how a technical language differs from natural language, and not to try to analyze it through the methods that are proper to the study of natural language.

The approach will be as follows. First, for each category of messages, the underlying schema (or schemata) must be constructed. A method of schema abstraction will be presented, and a word dictionary will be built. Second, to test our hypothesis, this schema knowledge will be implemented in so-called "understanding programs," which must be able to understand messages in the technical language under study. Although the programs will be provided with typed input, there will be no punctuation whatsoever, thus keeping the input as close as possible to spoken language.

Two principal assumptions are made. First, the system will use as little syntax as possible, and there will be no grammatical parsing. The only clue that will be used is the order of the words within a message, and the order of the messages within a communication. We assume that messages begin with a command word, and that the missing elements of a message are to be found in the immediately preceding message. Second, not all of the emitted words will be found in the dictionary. Moreover, each word will be given a single definition; we assume that the restricted domain of discourse forbids polysemy.

Language Understanding and Schema Theory

Schemata are a fundamental issue for the study of memory organization. They do not stand alone in the field of cognitive science; schemata have close links with the frames of Minsky (1974) and the scripts of Schank and Abelson (1977). However, for the sake of clarity, I will only use the word schema. A thorough discussion of the schema theory can be found in Alba and Hasher (1983). The following presentation will focus on the application of schema theory to language comprehension, borrowing from Rumelhart (1978).

A schema is a data structure for representing the generic concepts stored in memory. Schemata can represent objects, situations, events, actions, and sequences of actions. A schema contains variables which can take different values; the values a variable can take are limited by variable constraints. These constraints are specified for each variable of the schema; however, the set (or the range) of possible values of a variable may also depend on the value of another variable. The constraints have two important functions: (1) they allow the correlation between the input data and the variables of the schema (this is referred to as the slot-filling process), and (2) they may be used as default values when the input does not specify any. This process of filling in the slots of the schema is generally called (in the artificial intelligence literature) the instantiation of the schema; Piaget uses the word "accommodation" to refer to a similar process.

5

Once a schema has been activated, it can be accepted or rejected; this depends on the quality of its fit to the data. This evaluation is necessary, because we do not wait until the end of a sentence to initiate a schema. We pick up the first elements and make a hypothesis, that is, we activate a schema. The following data allow us to test the hypothesis we made (in this respect, language understanding can be considered as a problem-solving activity).

Each schema is a network of subschemata; these subschemata are the conceptual components of the general concept being represented.

For several reasons, the elements of a schema must not be identified with the words of a sentence. One such reason is that the units of the conceptual level are in a way "smaller" than words. A given word may include several units of a schema (see, for example, Abrahamson, 1975, analyzing verbs of movement). An illustration of this point is given below.

The processing of paraphrases is a quite important issue in language-understanding research; in fact, many authors see it as one of the fundamental criteria in the evaluation of a language-understanding system or theory (e.g., Norman and Rumelhart, 1975; Anderson and Bower, 1973; Schank, 1975). The semantic representations must be invariant under paraphrases of the same information.

Paraphrases (and near-paraphrases) are important in that they can provide clues about the way information is stored. I will here borrow some examples from Rumelhart and Norman (1975). Consider two sentences, A and B: (A) Henry went to a store; and (B) Henry drove to a store. Sentences A and B are not paraphrases, but we feel they are closely related. In fact, the meaning of A seems to be included in the meaning of B. This means that this meaning of "drove" contains this meaning of "went." For example, we could say that the meaning of "drove" is [CHANGE OF LOCATION BY MEANS OF AN AUTOMOBILE], and the meaning of "went" [CHANGE OF LOCATION].

Now compare sentences B and C: Henry ran to a store. Sentences B and C are not paraphrases, but again we feel that they are related, although not in the same way as A and B. Neither of them is included in the meaning of the other. This points out that "ran" and "drove" probably share some common semantic elements. For example, we could say that the meaning of "ran" is [CHANGE OF LOCATION IN A QUICK PEDESTRIAN WAY]. The common elements are then [CHANGE OF LOCATION].

In A, nothing is said about the way the action is performed; nevertheless, this does not mean that we have no idea about it. Many American readers would probably think that Henry took his car; that is, they would assume a default value to that unspecified slot of the schema. It is worth noticing here that this assumption could be different in other situations. For example, if the action took place in Paris, people would probably think that Henry took the subway, or that he walked to the store. The default values are context-dependent.

METHOD

The Corpus

Two different sources of pilot-controller communications have been available. The first one is a set of recordings from a preceding study (Falzon, 1982). This corpus represented 20 hr of pilot-controller communications, and a total of

7700 messages, already categorized. The second source consisted of transcripts of four 1-hr flights (Los Angeles-San Francisco), recorded in the cockpit of an aircraft.

These two sets of recordings differ in several ways. The first recordings were made in an air-traffic control center, and focus on specific sectors of control, crossed by different flights. The second recordings focus on specific flights, crossing different sectors. The first recordings were made in France, the second in the United States. In the first, the controllers used either French or English in their communications (according to the language of the pilot); in the present study, only the messages emitted in English have been considered. The language used in the two recordings may differ for two reasons: first, because English is not the native tongue of the controllers recorded in France; and second, because the linguistic habits of French and American controllers may differ, even though the domain of discourse is the same.

The two sources have been used, with a bias toward the utilization of the American transcripts. Many messages of a single category are necessary in order to define the schema of the category; the characteristics of the corpus prevents some categories from being sufficiently exemplified (for instance, there are few data on taxiing, taking off, landing, and making final approaches). Nevertheless, the transcripts allow the analysis of the other phases of the flight.

Schemata and Categories

Previous work on pilot-controller communications has led to a categorization of the messages (cf. Janet, 1981; Falzon, 1982; Hunter et al., 1974). This categorization can be seen as an attempt to classify the messages according to the schemata they evoke. Each category thus represents a schema, and the different forms of expression are different actualizations of the schema. In some cases, however, a further classification is necessary within a category. For example, one category deals with instructions related to changes in the heading, route, or course of the aircraft (horizontal movements). There are four possible actions concerning the horizontal movements of an aircraft (speed excluded): maintain (e.g., heading, track), modify (e.g., heading, track), intercept, and depart. Four different schemata are necessary to account for these four different horizontal actions.

Each category (or subcategory) of messages may be seen as a set of paraphrases or near-paraphrases, exemplifying a single schema. The study of paraphrases and of meaning overlaps can provide us with an experimental tool in the analysis of the underlying representations. The successive steps of this analysis will now be described.

Schemata and Their Elements

For each schema, specific information appears. For example, a "depart" action will always mention a "from" position and a direction. An "intercept" action will mention a radial and a very-high-frequency omnidirectional radio range (VOR, a navigation aid), or, more generally, a track, for example, "intercept the ILS course." A "maintain" action may very well mention nothing; the value to be maintained (and sometimes even its nature) will have to be found in context. This means that the expectations are different for each schema; that is, that we are able to specify the slots we will need to fill.

7

The approach will be illustrated by an example of analysis for one category of messages, dealing with modifications to the altitude level of an aircraft. Consider these simple messages: (1) climb level 330; (2) descend level 330; and (3) leave 290 for 330.

The schema will be built step by step. First, we know that all the messages of this category deal with actions in the vertical dimension and, more specifically, with a change of level (as opposed to a change in the rate of climb or descent). These elements must be specified in the schema:

CHLVL : ((Act: VE)(Nat: LVL))

CHLVL (for change level) is the name of the schema. The abbreviation VE means that the action (Act) taking place concerns the vertical (VE) dimension; LVL (level) indicates the nature (Nat) of the action. All schemata are composed of a list of pairs of items; in each pair, the first element (in lower case) is the role, the second (in capitals) is the filler. As we will see, the filler can include several elements, among which a choice must be made.

If we now consider message (1), we see that a first element is missing in the schema: the level to be reached. Moreover, we see that messages (1) and (2) do not indicate the same type of relation between the present level and the level to be reached. The schema then needs to include these elements:

CHLVL : ((Act: VE)(Nat: LVL)(Rel: (+ -))(To: P))

In the above, Rel stands for relation; + indicates that the level to be reached is above the present level; - indicates that the new level is below the present; "To" indicates the level to be reached; and P stands for parameter.

New information appears in message (3) — the present level. This element needs to be taken into account by CHLVL, but what about messages (1) and (2)? In these messages, the present level is not mentioned, but message (1), for example, could be rewritten as "from your present level, climb level 330." We then also need to specify that if the level that is to be departed is not mentioned in a message, it must be the present level:

CHLVL : ((Act: VE)(Nat: LVL)
         (Rel: (+ -))(From: (PV P))(To: P))

"From" indicates the level that has to be departed; PV stands for "Present Value." Consider now message (4): climb level 330 at pilot's discretion.

The expression "pilot's discretion" has several implications. Without going into detail, it indicates that the action can be delayed (the pilot may wait before changing his altitude), and that the pilot has more latitude in the accomplishment of the instruction. Anyway, this information must be taken into account by the schema:

CHLVL : ((Act: VE)(Nat: LVL)
         (Rel: (+ -))(From: (PV P))(To: P)
         (Time: (def NOW PD)))

where PD stands for "pilot's discretion." NOW is a default value, assumed when the message does not mention the "Time" information. Messages (1), (2), and (3), for

8

example, do not specify when the action has to take place. In that case, a default value is specified in the schema: "def" stands for "default," and indicates that the next filler is to be chosen if the message does not specify any. The default abbreviation def is written in lower case letters because it is not really a filler, but only a flag pointing to the filler next to it.

"Time" is probably not the best name for this role. As we have seen, "pilot's discretion" indicates not only when, but also how, the action is to be executed. Another example can be found in comparing "now" to "immediately." Both words mean that the action must begin upon reception of the instruction, but "immediately" implies also a specific way to perform the action. It includes a notion of urgency, meaning that, for example, maximum thrust should be used.

In some cases, the default value will be nil. For example, compare the following two communications (dealing with modifications of the horizontal movements of the aircraft), each composed of two messages:

5a. Fly heading 230
5b. Receiving Avenal proceed direct
6a. Fly heading 230 until receiving Avenal
6b. Then proceed direct

From (6a), we can infer that this type of message (heading change) may mention a limit (until . . .). But we also notice that (5a) does not mention it. The same limit information is to be found in fact in the next message (5b). In the same way, message (5b) mentions a condition, whereas message (6b) does not (the condition is to be found in (6a)). Because of these phenomena, we need to know that a "heading change" schema may, or may not, include a limit, and that a "route change" schema may, or may not, include a condition. This is the reason why we need the possibility for the default values to be nil.

The same analysis, applied to all categories of messages, provides a dictionary of schemata.

Words Definition

Compare the following expressions (message (1) is repeated for convenience):

1. Climb level 330
7. Climb flight level 330
8. Climb to the flight level 330
9. Climb 330

These four messages are paraphrases. From an examination of them, we can infer that some of the words that are used are not needed to make the messages understandable. For example, the information given by "to," "the," "flight," and even "level" is not useful: when hearing "climb," the pilot immediately knows that the instruction refers to a modification of the flight level, and a parameter is expected. Thus, much of what is said can be discarded. Therefore, the number of words that the system will have to know is only a subset of the different words that are used; the dictionary does not need to include the words mentioned above (in fact, "level" is defined in the dictionary; it is needed to understand other messages).

The words are defined using the same elements that have been used in the definition of the schemata. It is interesting here to compare the definition of "climb" and "leave" to the corresponding schema definition (CHLVL):

```
CHLVL : ((Act: VE)(Nat: LVL)
          (Rel: (+ -))(From: (PV P))(To: P)
          (Time: (def NOW PD)))

climb : ((Act: VE)(Nat: LVL)
          (Rel: +)(From: PV))

leave : ((Act: VE)(Nat: LVL))
```

The definitions of "climb" and "leave" differ in two ways. First, "climb" expresses a relation (Rel: +), whereas "leave" is neutral in that respect. Second, a message using "climb" will not mention the present level of the aircraft; this is made clear by the presence in the dictionary definition of (From: PV), which indicates to the system that no present level is to be expected.

THE PROGRAMS

An Overview

The system is composed of two sets of programs: understanding programs and planning programs. Understanding programs are provided with ATC communications that are composed of from one to several messages. They "translate" the input, first finding an appropriate predefined schema (among several others), then filling in the slots of the schema. There are two of these understanding programs: Schematch and Dicolisp. Schematch is the main processing program; it processes the words of a communication, matching them to evoked schemata, and then stores the instantiated schemata in memory. Dicolisp includes a dictionary of words and schemata and specialized subprograms adapted to the different schemata. In fact, a schema cannot be considered apart from its subprogram; it is a data structure plus a set of operations, a representation.

Planning programs are provided with the filled schemata which are the output of the "understanding" programs. Their job is to transform these schemata into sequences of actions. The only processed schemata are those related with movements in the horizontal plane (the program is "Planho") or in the vertical dimension ("Planve"). Other schemata are produced by the "understanding" programs (for example, schemata related to frequency changes, "report" orders, politeness, questions), but are not taken into account in the planning programs. The outputs of the two planning programs differ. Planho produces a sequence of legs. Planve produces a single frame, divided in three parts: the core (the fundamental action), the rules (i.e., descent at pilot's discretion until level X), and the constraints (i.e., cross a specified VOR at a specified level).

In fact, it is somewhat inappropriate to call Schematch and Dicolisp "understanding" programs. Although the system can be said to understand, since it exhibits an "intelligent" behavior (through the production of plans of action), the output of Schematch does not meet an intuitive criterion of "intelligence." To call Schematch and Dicolisp "parsing programs" would be misleading too. For lack of a better name, though, we will continue to call them understanding programs.

All programs are written in LISP, and implemented on the UNIX system at NASA Ames Research Center.

The Understanding Programs

Schematch and Dicolisp- This presentation will try to avoid going into too much detail. Instead, I will emphasize one example, and leave aside some minor aspects of the programs.

A communication, that is, a set of messages with no punctuation, is given as input to Schematch, which processes it word by word. There are four different kinds of words: unknown words, numbers, names of places, and dictionary words.

Unknown words are words that cannot be found in Dicolisp and that are not numbers. The processor drops them and goes to the next word. Numbers are not referenced in the dictionary. They are directly recognized as parameters by the general processor; they are given in their numeric form, not spelled. Names of "places" can be, for example, VORs, airports, or ATC centers. These names can actually be composed of several words (e.g., Los Angeles, Santa Monica), which the system transforms into single words (e.g., Los-Angeles, Santa-Monica). They are referenced in Dicolisp as "Places."

Dictionary words are the words that can be found in Dicolisp. Each word has a single definition (no polysemy). This definition is composed of a list of role-filler pairs; a definition may be a single pair. For example, "now" is just (Time: NOW). Other words have more complex definitions (cf. "climb" and "leave" in the preceding section). Some of the dictionary words have a special property in that they evoke a specific schema. Words like climb, descend, fly, turn, contact, intercept are schema-associated. When the processor finds one of those, it knows that it must open a new schema. Let us consider a simple example in which the communication is "climb level 230."

The first word is schema-associated, and the processor loads the appropriate schema, called CHLVL, and opens an empty list, called BINDINGS, which will receive the instantiated elements of the schema. The first element of Bindings is the name of the activated schema:

```
SCHEMA : ((Act: (VE))(Nat: (LVL))(Rel: (+ -))
          (From: (def PV P))(To: (P))
          (Time: (def NOW ANY)))
```

```
BINDINGS : (CHLVL)
```

The rest of the process consists of creating an instantiated schema (Bindings), using both the activated schema and the information given by the words of the message. First, the definition of "climb" is called for:

```
climb : ((Act: VE)(Nat: LVL)(Rel: +)(From: PV))
```

A pattern matcher compares the definition of climb and the schema. Each corresponding element is written in BINDINGS, and deleted from SCHEMA. We obtain

SCHEMA : ((To: (P))(Time: (def NOW ANY)))

BINDINGS : (CHLVL (Act: VE)(Nat: LVL)(Rel: +)(From: PV))

All the pairs of the definition of climb have been processed and no discrepancies have appeared. The processor then considers the next word, "level." It is not a schema-associated word, and its definition is then loaded:

level: ((Nat: LVL))

This pair cannot find its match in SCHEMA. BINDINGS is then checked to see if this absence is caused by redundancy. This is the case, since the information given by "level" already existed in "climb." The next word is then considered. It is the number "230." By definition, a number is never schema-associated, and will not be looked for in the dictionary; it is internally coded as P (for parameter). The program processes numbers (and places) in a specific way (the matching process differs). Anyway, P is found in the schema, yielding the following result:

SCHEMA : ((Time: (def NOW ANY)))

BINDINGS : (CHLVL (Act: VE)(Nat: LVL)(Rel: +)
                  (From: PV)(To: 230))

So far, so good, but what is the next word? We must conclude that there is no next word. The point is, however, there is still some information in SCHEMA and we must do something about it (it is a rule that a schema cannot be abandoned unless it is empty). But there is no problem, because we are provided with a default value for Time:

SCHEMA : .....(nil).....

BINDINGS : (CHLVL (Act: VE)(Nat: LVL)(Rel: +)
                  (From: PV)(To: 230)(Time: NOW))

This does not, however, complete the process. Each schema has specific procedures attached to it, allowing different checks. For example, it could have been impossible to empty SCHEMA, because some information was missing. In that event, the procedures would have tried some heuristics to find the missing information. Although this was not the case in the present example, consider the message "leave 230 for 290." As we have already seen, "leave" does not specify the type of Relation (+ or -), but this relation can be inferred from the values of the parameters. Here, the schema procedure would infer a (Rel: +) from the pairs (From: 230) and (To: 290).

Another role of these specialized subprograms is to format the output in order to facilitate its use by the planning programs. In the present example, though, the procedures do not help much, and the instantiated schema is stored as it is in memory.

In the example we have studied, the schema was closed because there were no more words to process. There is another reason for closing a schema when the processor finds a schema-associated word; two cases may occur: (1) climb level 230 turn left heading 150; and (2) climb and maintain level 230.

12

In (1), the processor finds "turn," which is schema-associated, and tries (and manages) to close the active schema. A new schema is then loaded (CHHDG, for: change heading), and the process continues. In (2), the processor finds "maintain," which also is schema-associated, but the processor cannot manage to close the schema; it needs a parameter. In this case, a second schema is opened in parallel, and the following words are processed for both schemata.

It is here interesting to consider another situation: (3) climb level 230 and maintain. In (3), when the processor finds "maintain," it tries to close the schema and succeeds. It then stores the instantiated schema in memory, loads the schema associated with "maintain," and tries to process the message. But "maintain" is the last word; the problem is not so much that the processor is waiting for a parameter; obviously a default value could be used. The problem is that the processor does not know what is to be maintained (e.g., speed, level, or heading). In this case, the schema-associated procedure will explore the memory, assuming that the last property being talked of is the one the value of which is to be maintained.

Discussion- It is difficult at this point to give an evaluation of the "understanding" programs in terms of percentage of understood messages. In order to do so would require a larger sample of communications. However, as far as the corpus we have access to is concerned, the programs are quite successful. They are able to recognize much of air-traffic control instructions, and this despite the facts that the system has no syntactical knowledge, a limited dictionary, and a single definition for each word of the dictionary.

The simplicity of the understanding programs is of great interest and evokes a question: How is it that the programs, being so simple, are able to understand so much of ATC communications? The reason is that the messages are rarely elaborate. In a previous study, Falzon (1982) showed that especially for the categories of messages which have a high frequency, the controllers tend to stick to some standard (and simple) forms of expression. Although variations do occur, they tend to be organized along common general patterns. For these reasons, the syntactical blindness of the system is not an obstacle, because of the syntactical simplicity and stereotypy of this technical language. Most messages begin with some sort of "command" word, and this could even be considered as a characteristic of restricted natural languages (or at least of this particular one). In this respect, it may be more fruitful to compare the ATC language to a computer command language rather than to natural language, that is, to describe its syntax in terms of operators and operands, rather than in terms of generative rules.

The monosemy of the words can be a characteristic of technical languages, because of the restricted domain of discourse, and because of the necessity to avoid ambiguity. In any case, the fact that each word has only a single meaning does not seem to be a problem in decoding the ATC communications. In the same way, the fact that we are able to pay no attention to words that are not defined in Dicolisp is an interesting feature.

The very simplicity of Schematch and Dicolisp is then in itself a result. It proves something about the "natural" (i.e., user-originated) restrictions of natural language. Some of these restrictions have already been pointed out; other characteristics will be found when implementing the planning programs.

Still, the present programs are quite certainly not enough to understand all possible ATC messages. Although the human operators are willing, because they are operators, to restrict themselves to some standard phraseology most of the time,

they still have the possibility, because they are human, to switch back to the use of natural language, when they want or need to do so (for example, in case of a low workload, or when the situation is an unusual one so that there is no adequate usual phraseology). In such cases, a more elaborate system would be necessary in order to understand the communications. Does this mean that the programs are useless? Certainly not.

The language of the controllers is not homogeneous. Most of the time, the language they use is a technical dialect, in which the vocabulary and syntax are highly restricted. However, they may use other expressions and a different vocabulary in less usual situations. The fact that the programs are not able to understand all of what is said is a consequence of the use of two different modalities in the communications to the pilots, for which different analyses must be performed.

The Planning Programs

Again, a complete description of the two planning programs (Planho and Planve) will not be given; we will only outline their main characteristics and differences. The two planning programs consider that a single aircraft is dealt with. This restriction was introduced to limit the number of flight plans the system has to know of, but it can easily be changed, provided that information is given about the different aircraft. The meanings of the messages are based on the interpretation given in the Airman's Information Manual (FAA, 1982).

Planho– Planho deals with the horizontal plane. It needs two types of input: a set of instantiated schemata and the present flight plan. Instantiated schemata are messages processed by the "understanding" programs. Planho only processes those schemata dealing with horizontal movements. In order to do that, it filters the memory, retrieving only the relevant schemata.

The present flight plan (before updating) is a list of legs, the first of which is assumed to be the active one. Each flight-plan leg is composed of three elements: a trigger, a direction, and a limit. The trigger is the position at which the action begins; the direction indicates the heading, or radial, to be followed; and the limit indicates the end of the leg (if A and B are two successive legs, the limit of leg A is then the trigger of leg B).

A word or two about the definition of directions, limits, and triggers. Limits and triggers are positions, and they can be defined in four ways. First, they can be defined by a VOR, a radial, and a distance on that radial. For instance, (D-30 R-160 Avenal) means 30 miles from Avenal on its 160 radial (an interesting case occurs when the distance and radial are not mentioned, as in "receiving Avenal proceed direct": in this latter case, the point will be defined as (D-any R-any Avenal), meaning the first position meeting this condition). Second, they can be defined by the name of the VOR itself; this is in fact a special case of the preceding definition, "Avenal" for instance meaning in fact "at a minimum distance from Avenal on any radial," that is, (D-min R-any Avenal).

Third, limits and triggers can be defined by the intersection of two directions; for example, ((H-130 Pesca)(R-160 Avenal)) is the intersection between heading 130 from Pesca and the 160 radial of Avenal. Fourth, they can be defined by the present position of the aircraft (about which the program has very little information). The position is then coded "H&N," for "here and now," which makes it more intellectual.

14

We have seen how directions are coded: they always mention a reference point (VOR or not). This reference point may be H&N, in which case (H-120 H&N) is then "heading 120 from wherever you are now."

Planho creates a sequence of legs according to the input messages, and then compares this sequence of legs to the flight plan, creating new legs and inserting them in the flight plan, and deleting some legs from the flight plan when necessary. The output of Planho is then a new flight plan, which is considered valid only if it is a continuous sequence of legs (i.e., if the limit of each leg is the trigger of the following one): the validity is checked by the program.

It should not be inferred from this short summary that a single leg is created for each message. Some messages imply the creation, or the modification, of several legs ("intercept" messages, for example).

Planve- Planve, like Planho, filters the output of Schematch in order to process only the messages dealing with vertical information. The result of this processing is not a sequence of steps, but a single step, which includes three parts: the core, the rules, and the constraints.

The core indicates the fundamental general action of the step. It mentions the type of action, its limits, and its conditions. For instance, the core,

(Act: -)(From: 330)(To: 200)(Cond: PD)

means that the general action is a descent from the level 330 to the level 200, and that the action may begin at pilot's discretion (PD). The conditions may be of several different kinds; for example, they may specify a particular position (if the "from" level is to be maintained until this position before descent) or a particular speed (if the aircraft is supposed to reduce its speed before descent). The conditions can be compared to the triggers of the legs (but they are not exactly the same thing).

The rules indicate the degree of freedom of the pilots during the different substeps of the core. If a descent instruction mentions "at pilot's discretion," it means, first, as we have seen, that the action may be delayed, but also that the way to conduct the descent is not submitted to the standard rules of descent. For example, the pilot is allowed to level off during descent if he wants to or to vary the rate of descent. But "pilot's discretion" does not mean no rule at all; for instance, the pilot is not allowed to climb (i.e., the rate can be zero, but cannot be positive). The "PD" rules may be more lenient, but they still are rules. Each rule mentions the type of rule and its limit, for example, ((Rule: PD)(To: 200)).

The constraints specify the altitude restrictions the aircraft has to comply with. The constraints mention a position and its associated restrictions. For example, (Fillmore => 240) means that the aircraft must cross Fillmore at or above the level 240.

Let us examine an example. Suppose the instruction was "cross Fillmore at or above 240 maintain 200." The communication is first analyzed by the "understanding" programs, then processed by Planve. Supposing the aircraft was previously steady at 330, the output of Planve is:

```
(CORE (Act: -)(From: 330)(To: 200)(Cond: PD))
(RULES ((Rule: PD)(To: 240))
       ((Rule: ST)(To: 200)))
(CONSTRAINTS (Fillmore => 240))
```

Implied in the "cross" schema are several elements. First, there is an implicit
action: the aircraft is supposed to change its flight level (the fact that it is a
descent is inferred from the values of the parameters). The target level (i.e., the
"To" level) is temporarily set at 240. Second, if not otherwise specified, "cross"
implies that the action may begin at the pilot's discretion. This is why the Condi-
tion is PD. Third, the rules to apply to the descent are also at pilot's discretion,
until the critical level is reached. And, fourth, obviously, the constraint is to
reach Fillmore at or above 240.

Implied in the "maintain" schema is an instruction to descend to the level 200.
The "To" element of the core is then changed from 240 to 200. We need then to know
the rule to follow between 240 and 200. Nothing is specified in the communication;
in that case, the default value is "ST" (for "standard"), meaning another set of
rules of vertical movements (no leveling, constant rate, etc.). A new rule is then
added to RULES.

Discussion- The two planning programs can be analyzed using the schema theory
framework; both programs are composed of a data structure and a set of operations.
The following paragraphs exemplify this assertion.

Planho has a standard structure (Trigger/Direction/Limit) for composing each leg
of the flight plan and different operations to process the sequence of legs. For
example, it knows it has to link the successive legs (through Limit and Trigger).
When this link is not specified in the input, the program uses a characteristic of
language: namely, that unless otherwise specified, the emission order of the mes-
sages of a communication corresponds to the order in which the actions have taken (or
will take) place. Then, if message A is given before message B, Planho assumes that
the legs implied by A are antecedent (and probably immediately antecedent) to the
legs implied by B, so that the missing trigger or limit can be inferred, and the
legs linked. Of course, this is not always the case. Some messages (like the
"depart" messages) specify the position where the action will take place; these mes-
sages are processed accordingly.

Planve also has a standard structure (Core/Rules/Constraints) to describe the
vertical step, on which different operations are applied. For instance, the program
knows that rules must be specified for all levels belonging to the core. If no rule
is specified, an inference is constructed to bridge the gap, assuming a default value
("standard" rules).

In the same way, if a substep is missing, it will be built. For example, one
cannot maintain a given level unless that level has before been reached; as already
explained, Planve will infer the missing action.

Another operation specifies that whatever is the order of the messages, the con-
straints must be organized in the order in which they have to be met, that is, fol-
lowing the vertical movement of the aircraft.

In a last example, consider that the program simplifies the rules by concatenat-
ing some of them. If rule A is to be obeyed between the levels 300 and 230, and then
between the levels 230 and 200, the rules will be simplified by applying rule A from

16

300 to 200, creating a single substep. This does not mean that the pilot will apply only one of the possible modes of descent during this substep, but that any mode he chooses will have to comply with the rule of the substep.

The main criticism of the planning programs is directed at their psychological validity. There is no evidence that the structure of the pilot's representation of the future flight plan (vertically and horizontally) is similar to the structure of the outputs of Planve and Planho. these outputs can only be considered as one of the possible results of the planning activity. A good proof of that is that several different formats have been tried during the writing of the programs. For example, in the first version of Planve, the output was not a single step, but a series of successive steps, comparable to the output of Planho. The structures that were finally adopted have been chosen only because they seemed to have a good a priori probability of reflecting the structure of the pilots' plans (knowing of course that they are only structures, not complete descriptions). For example, one can think that a given leg has little influence on the following one, whereas a level constraint at a given position may have consequences on the whole vertical step, or at least on the preceding substep. This explains the structural difference between the outputs of the two programs (sequence of legs vs single step).

A second criticism (related to the first one) could also bear upon the fact that the two programs process the vertical and the horizontal instructions separately. One could argue that it is doubtful that the pilots have two totally separate plans in memory, that no correlations are made between the two. We have some excuses for such a criticism, in that the programs have no knowledge of the spatial relations between the different points of the flight plan (except for their order), and of the capabilities of the aircraft.

This has two consequences. First, some actions belonging to a single domain are difficult to process. The best example is the effect of an "intercept" schema. The planned sequence will differ according to the position of the aircraft (Will the aircraft intercept the radial if it keeps on its present heading, or is a turn necessary?). Since the program does not know the precise position of the aircraft relative to the VOR, it cannot decide between the two. Second, this makes it impossible to link the two plans. The programs cannot infer the effect of an instruction concerning one dimension upon the other one. This spatial ignorance is of course a poor "excuse," and, if correlations are made by the pilots, they should appear in intelligent planning programs.

Another point, which we will illustrate, is that in some cases the output of Planho is incomplete, in the sense that the limit of the last leg of the planned sequence is not the runway. This occurs when the controller has instructed the pilot to take some action, but has not given a complete vectoring to the runway. The pilot is then expecting some further information. It is of course quite obvious that the pilot will not passively wait several hours until something happens. First, after "some" delay, and in the absence of any message, instructions will be requested. Second, it is likely that the pilot has some default plan of action, and it may even be that the controller expects the pilot to apply this plan, defined either by the application of the formally specified default maneuvers (the "lost communications" procedures) or by common sense. The problem is that we have no idea about this common-sense knowledge. Therefore, it seems necessary to study it.

The planning programs have proved that a simple set of programs can be effective in "translating" a technical language into sequences of actions. However, further

studies of the pilot's planning activity would be necessary to improve the accuracy of the outputs of the programs.

What have the planning programs taught us? The more important result is that the schemata mean more than we expected. In a way, this is partly contradictory to what we said earlier. In the discussion of the "understanding" programs, we assumed that there was some simplicity in the dictionaries, no polysemy. We assumed further that the schemata could be described by a simple set of elements. We must realize now that this apparent simplicity corresponded to a real complexity in meaning. The data structure of each schema may be simple, but the processes attached to it are complex. They include many procedures that react in different ways in different contexts, that need to make inferences about missing elements (to use default values), and that must build some implicit actions when necessary.

Why is it so very easy to speak of the "cross" schema, the "depart" schema, the "intercept" schema, and so on? The first reason is that there is little synonymy, so that a schema can be called by the word that evokes it (this, however, is not always true; some schemata are activated by a variety of words). The second reason is that each schema of Dicolisp (evoked by these schema-associated words) finds in the planning programs its matching procedure. Quite obviously, a "cross" schema cannot be processed in the same way as a "climb/descent" schema. Specific procedures are required, both when processing the actual message ("understanding" programs) and when planning the instantiated schemata (planning programs). In fact, in many ways, the understanding programs and the planning programs share the same structure. In both sets of programs, a general processor and specialized subprograms can be found; in both sets, the general processor is independent of the type of schema (for example, the pattern matcher of Schematch applies to any schema), whereas the specialized subprograms are schema-associated (and in fact are named after the schema they are associated with).

CONCLUSION

The Need for More Data

One of the main problems with the system described herein is that the sample of communications on which it has been built is very small. The analysis of more transcripts is necessary, for at least three reasons.

First, as has already been pointed out, the system does not understand all of what is said, even in the limited sample we studied. In some cases, this is a result of the use of complex or infrequently occurring messages by the speaker. It can also result (and this is more important), however, because in order to define the schema for a given category of messages, we need to have a sufficient number of messages of this category. It is only in this way that the words used more frequently can be spotted, and that the elements of the schema can be defined. The problem is that the sample is not large enough to allow some categories to be sufficiently exemplified. We are aware that for the least-frequent categories, restricted vocabularies are probably difficult to build (because the conventionality of the surface form of a message is a function of its frequency of use), but some categories that are absent in our sample are probably not unusual (at least from a controller's point of view). Some work needs therefore to be done in that area. A selective sampling of ATC communications may be possible.

18

Second, new questions can appear when studying a larger sample. For example, because of a lack of data, there is no schema for the "clearance delivery" messages. In these messages, successive way points and airways are given by the controller. Their number is variable, and depends on the specific destination of the flight. A schema dealing with clearance deliveries would have to allow an indefinite number of elements to appear, that is, would have to be recursive, at least partly. None of the schemata presently implemented includes such a possibility. Other improvements of the system might be necessary, concerning the particular syntax. Some parsing tools may become helpful, for instance in the identification of the conditions of action. Although these conditions are successfully processed (in the limited sample of messages we studied), some syntactical knowledge could prove to be necessary in some cases.

And third, a larger sample is needed in order to validate the approach that has been chosen and to analyze its limits and weaknesses. One aspect of the validation is the study of the performance of the system in terms of the percentage of understood messages. Of greater importance, however, is the analysis of what is not understood. It is only this analysis that can really provide an evaluation of the approach we followed. Important questions in that respect are as follows: What makes a message impossible to understand by the system? Is it a matter of rarity, of syntax, of vocabulary? When and why do such messages appear? Are they necessary for the controller and the pilot? What is their function? Are some messages misunderstood? How many?

The Need for More Knowledge

Some of the problems that are met do not have their origin in the system's limited knowledge, but instead in our limited knowledge.

For example, we might experience some difficulties in "translating" a given input into a sequence of actions, because we do not really understand what was meant by the speaker (the controller). Of course, if we do not understand, the programs will not understand (unless we write some instructions based on "common sense," or on our own representation of the meaning of the message). In other words, we need a better knowledge of the meanings of some words (or messages) for both controllers and pilots and in different contexts, as well as a better understanding of the effect of these words/messages on the planning activity of the pilots.

An important issue is the mental "format" of the planned sequence, that is, the mental image, built by the pilot, of the successive steps or legs of the aircraft. It has been explained that some assumptions have been made concerning the structure of the representations. The characteristics of the actual representations are fundamental not only in order to give an accurate image of the pilot's activity, but also in the design of adapted machine representations of the flight plan. For example, in the future on-board flight-management systems, the effect of such studies is not confined to improvements of our programs, although these would indeed profit from it.

One important question in that respect is the relations that exist between the vertical and the horizontal planning activities, that is, in how an instruction in one domain modifies the planned actions of the other domain.

The Flight Plan as a Hierarchy of Representations

Finally, there is another way to improve the system. The flight plan of an aircraft must not be considered as a linear sequence of events; it can be divided into different phases, and decomposed hierarchically. A flight plan is a plan of action to reach an overall goal. This plan of action is composed of several scripts, each script using several schemata. An example will illustrate this point. For an aircraft flying from San Francisco to Los Angeles, the goal can be described as:

    GOAL : (From: San Francisco)(To: Los Angeles)

This goal can be attained using a plan specifying the different necessary scripts:

    PLAN : (Taxiing (From-gate: 15)(To-runway: 01R))
           (Taking-off on: 01R)
           (Departure procedure: Porte 5)
           (Routing: Avenal transition)
           (Approach procedure: Fillmore 8)
           (Landing on: 24L)
           (Taxiing (From-runway: 24L)(To-gate: 23)

All the slots of the above scripts are filled, but some of them may be blank at the time of departure; default values may exist here also. Each script can be expanded in its elementary schemata. Here is, for example, the expanded script of the departure procedure, Porte 5:

    SCRIPT : ((Trigger: end of runway 01R San-Francisco)
              (Dir: (H-010 San-Francisco))
              (Limit: ((H-010 San-Francisco)(R-350 San-Francisco))))

             ((Trigger: ((H-010 San-Francisco)(R-350 San-Francisco)))
              (Dir: (R-350 San-Francisco))
              (Limit: (4-DME-fix)))

             ((Trigger: (4-DME-fix))
              (Dir: (H-200 4-DME-fix))
              (Limit: ((H-200 4-DME-fix)(R-135 Point-Reyes))))

             ((Trigger: ((H-200 4-DME-fix)(R-135 Point-Reyes)))
              (Dir: (R-135 Point-Reyes))
              (Limit: (Pesca)))

             ((Trigger: (Pesca))
              (Dir: (H-090 Pesca))
              (Limit: ((H-090 Pesca)(R-116 Woodside))))

             ((Trigger: ((H-090 Pesca)(R-116 Woodside)))
              (Dir: (R-116 Woodside))
              (Limit: (Wages)))

             ((Trigger: (Wages))
              (Dir: (R-116 Woodside))
              (Limit: (Avenal)))

Some evidence of the relevance of such an approach can be found particularly in the "clearance delivery" communications, which we have begun to study. Here is an example:

(1) United one six one
(2) Cleared to Los Angeles
(3) Fly a Porte five departure
(4) Avenal transition
(5) As filed
    [...]

In this example, message (2) sets the general goal, and messages (3), (4), and (5) specify some scripts.

Such an approach (similar to the one proposed by Hammer, 1983; cf. also Rouse et al., 1983) can provide useful tools for a better understanding of ATC messages.

First, each script dictates the type of message that can, or cannot, appear. When taxiing, a pilot does not expect to be given a level change, for example, but is prepared to be told to "hold short of runway." These expectations help us understand how the pilots manage to make sense out of the garbled gibberish they sometimes hear. In the same way, these expectations can guide the comprehension of the messages in a language understanding system. (However, it must be understood that this can be the source of some misunderstandings, in real work as well as in a system, when the messages happen to differ from the expectations.)

Second, the effect of the controllers' instructions vary according to the level (goal, plan, script, schema) to which they apply. The modification of a flight plan may affect a single script of the plan and have no consequence on the next or preceding script. For example, modifications of the horizontal trajectory (shortcuts) during the departure procedure will have little effect on the other scripts. On the other hand, other modifications may affect the goal itself. For example, if the destination airport is closed because of snow, a new goal has to be set, which in turn affects the plan, the scripts, and the schemata.

REFERENCES

Abrahamson, A. A.: Experimental Analysis of the Semantics of Movement. In:
    Explorations in Cognition, D. A. Norman and D. Rumelhart, Ed., W. H. Freeman,
    San Francisco, Calif., 1975.

Alba, J. W.; and Hasher, L.: Is Memory Schematic? Psychological Bulletin, vol. 93,
    no. 2, 1983, pp. 203-231.

Anderson, J. R.; and Bower, G. H.: Human Associative Memory. L. Erlbaum Assoc.,
    New Jersey, 1973.

Chapanis, A.: Words, Words, Words. Human Factors, vol. 7, 1965, pp. 1-17.

Chapanis, A.: Interactive Communication: A Few Research Answers for a Technological
    Explosion. 86th Annual Convention of the American Psychological Association,
    Toronto, Ont., Canada, Aug. 1978.

Clark, H. H.; and Lucy, P.: Understanding What is Meant from What is Said: A Study
    in Conversationally Conveyed Requests. Journal of Verbal Learning and Verbal
    Behavior, vol. 14, 1975, pp. 56-72.

Dupre, J.: Natural Kinds and Biological Taxa. The Philosophical Review, vol. 40,
    no. 1, Jan. 1981.

Ehrenreich, S. L.: Query Languages: Design Recommendations Derived from the Human
    Factors Literature. Human Factors, vol. 23, no. 6, 1981, pp. 709-725.

Airman's Information Manual. Basic Flight Information and ATC Procedures. Federal
    Aviation Administration, U.S. Department of Transportation, Washington, D.C.,
    Dec. 23, 1982.

Falzon, P.: Les Communications Verbales en Situation de Travail: Analyse des
    Restrictions du Langage Naturel. Report CO 8211 R70, Institut National de
    Recherche d'Informatique et d'Automatique, Le Chesnay, France, 1982.

Gibbs, R. W.: Contextual Effects in Understanding Indirect Requests. Discourse
    Processes, vol. 2, 1979, pp. 1-10.

Gibbs, R. W.: Your Wish Is My Command: Convention and Context in Interpreting
    Indirect Requests. Journal of Verbal Learning and Verbal Behavior, vol. 20,
    1981, pp. 431-444.

Hammer, J. M.: A Computerized, Procedural Aid for Flight Crews. Proceedings of the
    2nd Symposium on Aviation Psychology, Columbus, Ohio, Apr. 26-27, 1983.

Hammond, N.; Barnard, P.; Clark, I.; Morton, J.; and Long, J.: Structure and Content
    in Interactive Dialogue. Report HF034, IBM Hursley Human Factors Laboratory,
    Winchester, U.K., 1980.

Hill, I. D.: Wouldn't It Be Nice if We Could Write Computer Programs in Ordinary
    English — Or Would It? The Computer Bulletin, vol. 16, 1972, pp. 306-312.

Hunter, J. S.; Blumenfeld, D. E.; and Hsu, D.-A.: Modeling Air Traffic Performance Measures. Vol. 1. Message Element Analyses and Dictionaries. Report FAA-RD-73-147, I. U.S. Department of Transportation, Washington, D.C., 1974.

Hupet, M.; and Costermans, J.: Et que Ferons-Nous du Contexte Pragmatique de l'Enonciation? Bulletin de Psychologie, vol. 35, no. 11, 1982, pp. 759-766.

Janet, E.: Analyse du Langage Employe dans les Communications d'un Systeme de Controle. Report CO 8201 R 69, Institut National de Recherche d'Informatique et d'Automatique, Le Chesnay, France, 1981.

Kelly, M. J.; and Chapanis, A.: Limited Vocabulary Natural Language Dialogue. International Journal of Man-Machine Studies, vol. 9, 1977, pp. 479-501.

Michaelis, P. R.; Chapanis, A.; Weeks, G. D.; and Kelly, M. J.: Word Usage in Interactive Dialog with Restricted and Unrestricted Vocabularies. IEEE Transactions on Professional Communication, vol. PC-20, no. 4, 1977, pp. 214-221.

Minsky, M.: A Framework for Representing Knowledge. AI Laboratory Memorandum 306, Massachusetts Institute of Technology, Cambridge, Mass., 1974.

Norman, D. A.; and Rumelhart, D. E.: Explorations in Cognition. W. H. Freeman, San Francisco, 1975.

Rosenberg, J. K.: A Psycholinguistic Model for Command Names. Report CIS-20, Xerox Palo Alto Research Center, Palo Alto, Calif., 1982.

Rouse, W. B.; Hammer, J. M.; Brown, E. N.; and Morris, N. M.: Pilot Interaction with Automated Airborne Decision-Making Systems. Third semiannual progress report, Georgia Institute of Technology, Atlanta, Ga., 1983.

Rumelhart, D. E.: Schemata: The Building Blocks of Cognition. Report 79, Center for Human Information Processing, U. of California, San Diego, Calif., 1978.

Rumelhart, D. E.; and Norman, D. A.: The Active Structural Network. In Explorations in Cognition. W. H. Freeman, San Francisco, Calif., 1975.

Scapin, D.: Computer Commands in Restricted Natural Language: Some Aspects of Memory and Experience. Human Factors, vol. 23, 1981, pp. 365-375.

Scapin, D.: Computer Commands Labeled by Users Versus Imposed Commands and the Effect of Structuring Rules on Recall. Proceedings of the Conference on Human Factors in Computer Systems, Gaithersburg, Md., 1982.

Schank, R. C.: Conceptual Information Processing. Amsterdam, North-Holland Publishers, 1975.

Schank, R. C.; and Abelson, R. P.: Scripts, Plans, Goals and Understanding. Lawrence Erlbaum Associates, Hillsdale, N.J., 1977.

Sperandio, J.-C.: Analyse des Communications Air-Sol en Controle d'Approche. Report CO 6909 R21, Institut National de Recherche d'Informatique et d'Automatique, Le Chesnay, France, 1969.
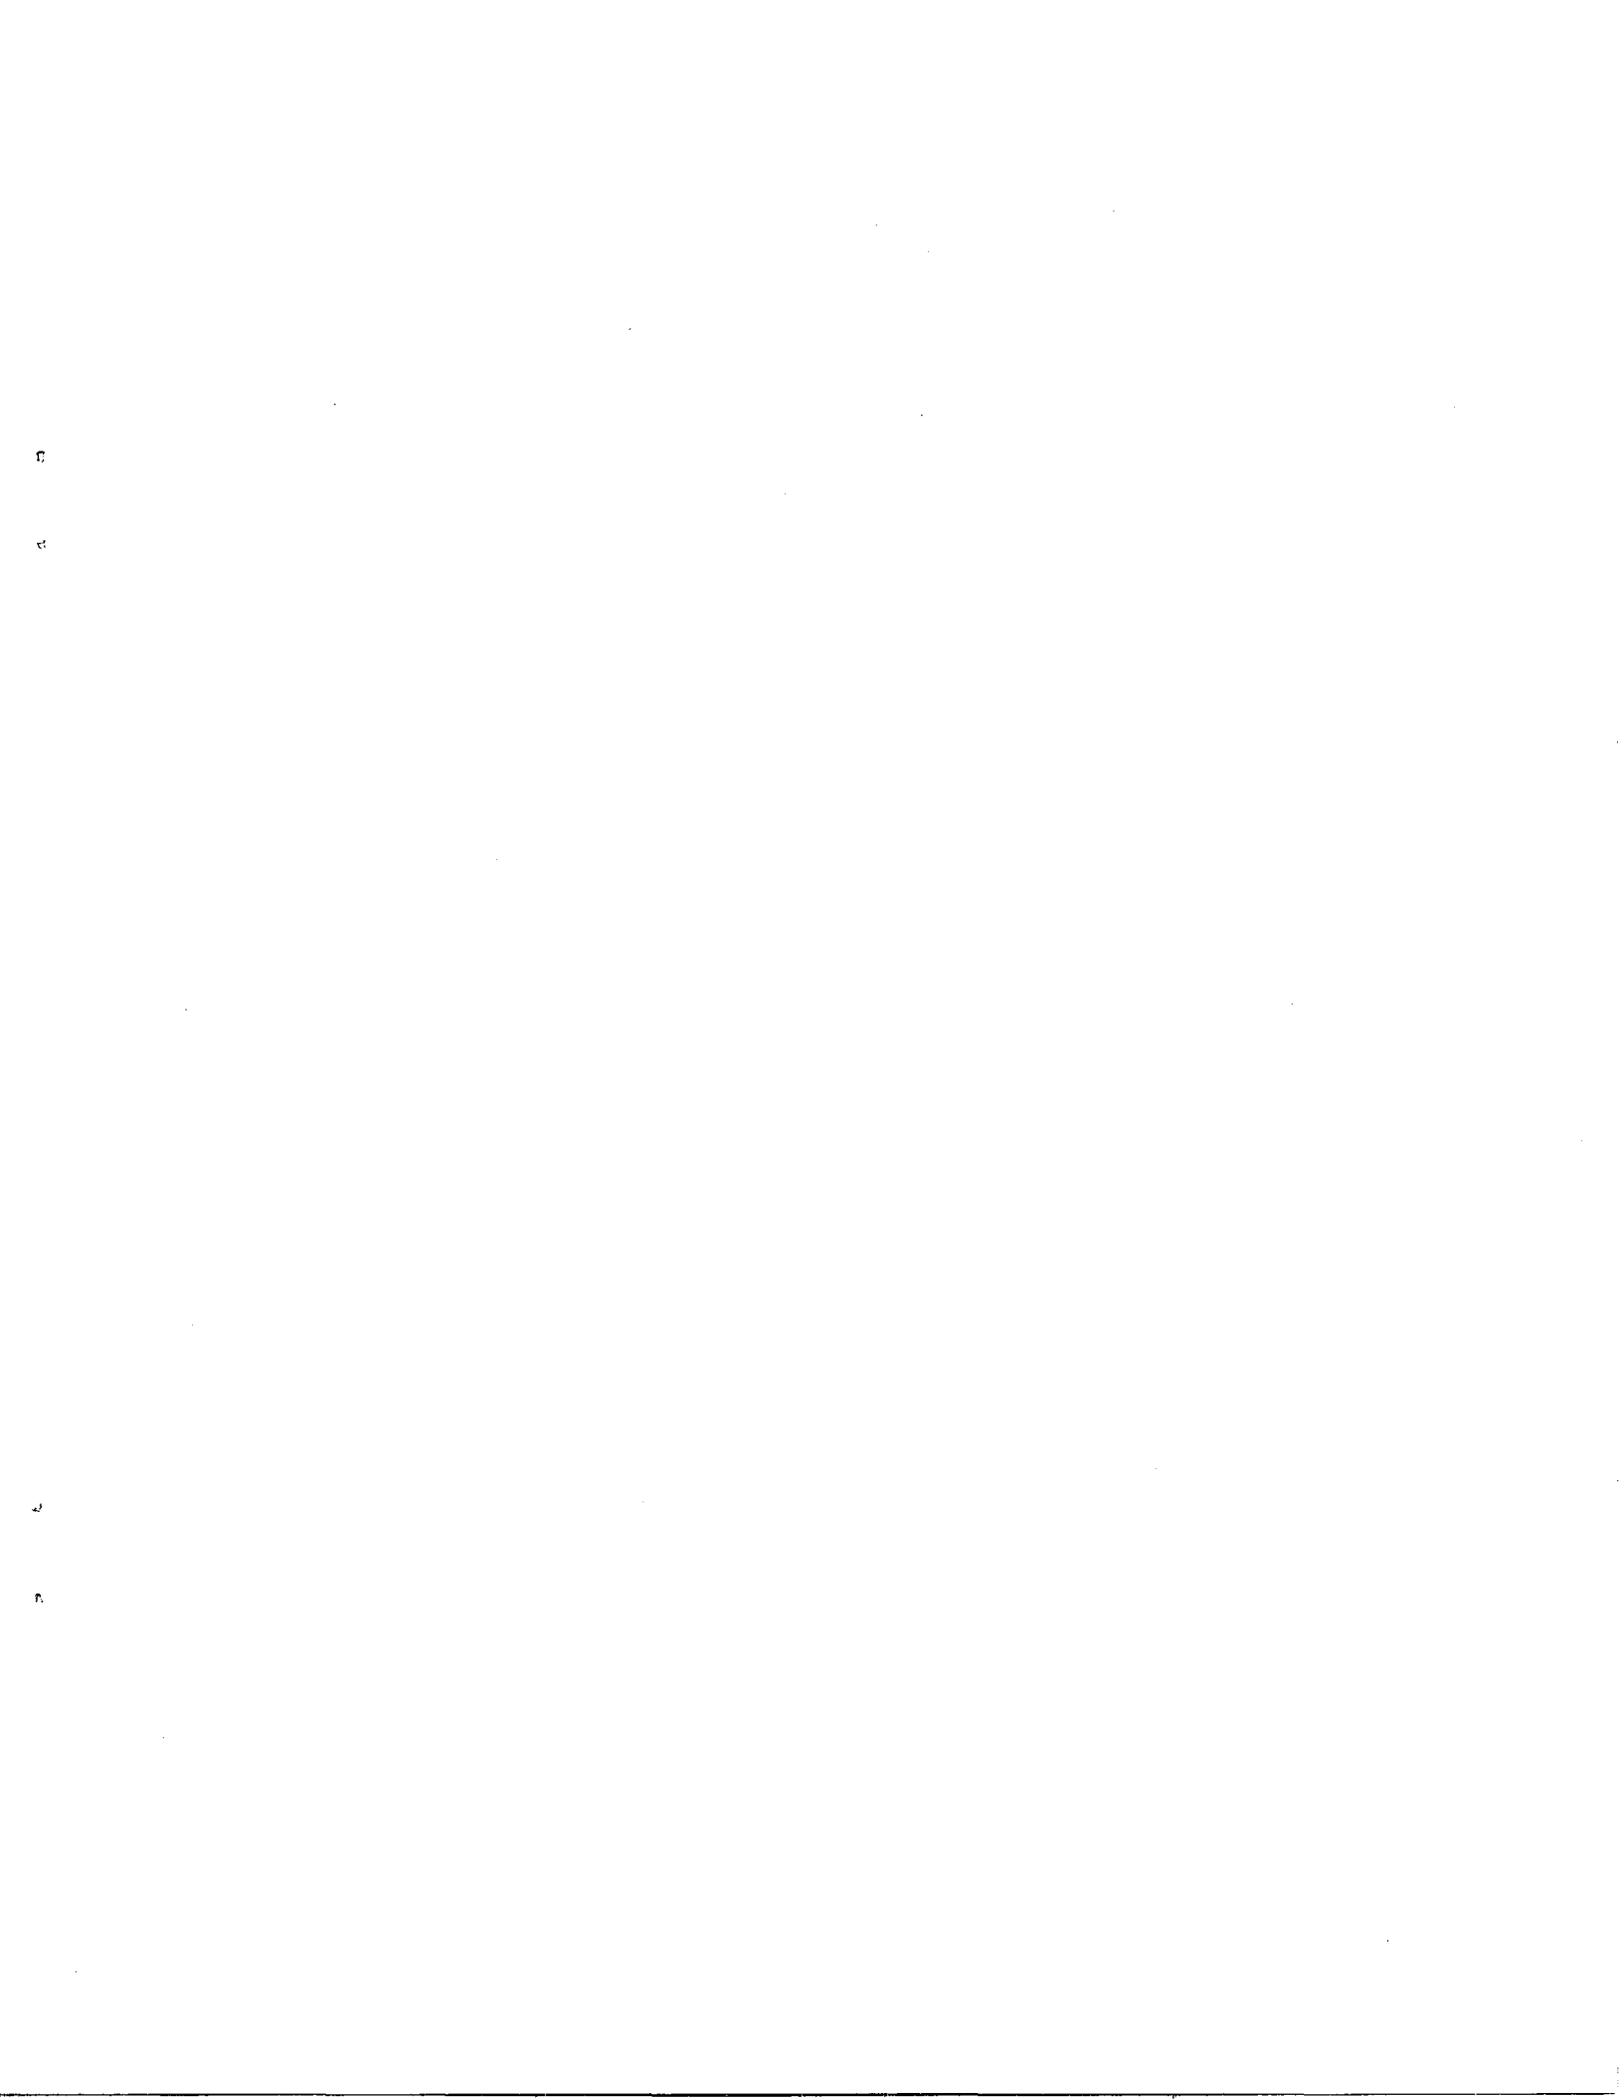
Thomas, J. C.: A Method for Studying Natural Language Dialogue. IBM Research
   Report RC 5882, Yorktown Heights, N.Y., 1976.

Thomas, J. C.: A Design-Interpretation Analysis of Natural English with Applications
   to Man-Computer Interaction. Int. J. Man-Machine Studies, vol. 10, 1978,
   pp. 651-668.

Winograd, T.: Understanding Natural Language. Cognitive Psychology, vol. 3, 1972,
   pp. 1-191.

1. Report No. NASA TM-85868	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle UNDERSTANDING A TECHNICAL LANGUAGE: A SCHEMA-BASED APPROACH		5. Report Date July 1984
		6. Performing Organization Code
7. Author(s) Pierre Falzon (Institut National de Recherche d'Informatique et d'Automatique (INRIA), Rocquencourt, France)		8. Performing Organization Report No. A-9575
9. Performing Organization Name and Address NASA Ames Research Center Moffett Field, CA 94035		10. Work Unit No. T-5694
		11. Contract or Grant No.
		13. Type of Report and Period Covered Technical Memorandum
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546		
		14. Sponsoring Agency Code 505-35-21

15. Supplementary Notes

Point of Contact: David Nagel, Ames Research Center, M.S. 239-3, Moffett Field, Calif. 94035   (415) 965-6170 or FTS 448-6170

16. Abstract

  Workers in many job categories tend to develop technical languages, which are restricted subjects of natural language. A better knowledge of these restrictions could provide guidelines for the design of the restricted languages of interactive systems. Accordingly, a technical language (used by air-traffic controllers in their communications with pilots) is studied. A method of analysis is presented that allows the schemata underlying each category of messages to be identified. This schematic knowledge is implemented in programs, which assume that the goal-oriented aspect of technical languages (and particularly the restricted domain of discourse) limits the processes and the data necessary in order to understand the messages (monosemy, limited vocabulary, evocation of the schemata by some command words, absence of syntax). The programs can interpret, and translate into sequences of action, the messages emitted by the controllers.

17. Key Words (Suggested by Author(s)) Technical language Schema Language understanding Planning		18. Distribution Statement Unlimited Subject Category - 54	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 28	22. Price* A02