Computer Generation of Marine Weather Forecast Text

E. GOLDBERG

Atmospheric Environment Service, Environment Canada, Toronto, Canada

R. KITTREDGE AND A. POLGUÈRE

Departement de Linguistique, Universite de Montreal, Montreal, Canada

(Manuscript received 19 February 1987, in final form 9 December 1987)

ABSTRACT

MARWORDS is a natural language text generation system which has been developed to synthesize marine forecasts for the Davis Strait area in Northern Canada. It uses standard manually produced predictions of wind speed, air temperature, and other weather conditions to generate English language text forecasts. The resulting text compares favorably with those written by marine weather specialists. MARWORDS incorporates a detailed grammar to capture the constraints on words, sentences, and texts that are natural to the domain. This grammar was developed from a study of approximately 50 000 words of forecast text. Among the problems faced were the linguistic treatment of data-salience relations and the modulation of temporal adverbs to reflect increasing levels of uncertainty for more remote events. In addition the system was designed so that it could be extended to synthesize bilingual (French and English) or multilingual forecasts.

1. Natural language report synthesis

Natural language report synthesis (NLRS) is the term used to describe the process of creating well-formed text that summarizes formatted data in a given domain, using a style which mirrors the conventions of professional report writers for that domain. This technique was first demonstrated in the work of Kukich (1983) on knowledge-based generation of stock market reports. Kukich's ANA system produces professional-sounding stock market summaries using a daily trace of Dow-Jones' half-hourly quotations for the market average and major indices.

Most work on text generation uses a planning technique to organize the paragraph and sentence structure. This enables the system to generate a sequence of sentences which satisfy the requirements of a given goal. Changing the goal of the text utterance can affect both the organization of the text and the choice of wording, even if substantially the same raw information is embodied in the final text. Systems which have been developed with this degree of generality are still quite complex and make heavy demands on computer resources. NLRS takes a much simpler approach. It relies on a detailed grammar of text and sentence structure, coupled with a domain-specific lexicon, to capture the natural tendencies of professional report writers. Both the grammar and the lexicon is developed by analyzing a large body of reports written by people.

The work described here follows the NLRS approach, but with a more modular organization. The techniques used to select particular words (this is often termed "lexical insertion") are somewhat more general linguistically than those used in earlier works. For example, both ANA and the French version, FRANA (Contant 1986), use a phrasal lexicon (Becker 1975) approach. In these systems, the system first selects a phrase that can express the message properly. Phrases may contain variables which are instantiated at a later stage. The expressive potential of the system is therefore limited to whatever phrases the system contains. In MARWORDS, the general structure of the sentence is mapped out in a more grammatical sense, and words are selected based on linguistically motivated categories and semantic classes. That is, although it does not embody a completely formal syntax and semantics, MARWORDS recognizes that the meteorological information must be presented according to well-defined conventions with respect to paragraph and sentence structure, and that descriptive terms similar to adjectives and adverbs are used to modify the raw meteorological facts in the text. In addition, the lexicon used to describe wind variations is different from that used to describe variations in visibility. This recognition facilitates text generation, and should give MARWORDS more flexibility in expression than would be the case if a phrasal lexicon were used.

Work on MARWORDS also tests a new application

Corresponding author address: Dr. E. J. Goldberg, Environment Canada, Forecast Research Division, 4905 Dufferin Street, Downsview, Ontario M3H 5T4 Canada.

domain, which has raised a general new problem for temporal reference (cf. section 7). References to time in computer generated text have traditionally been difficult because of weaknesses in the representation of temporal knowledge.

2. Synthesis of Arctic marine weather forecasts

The MARWORDS system was developed during a five-month effort to explore the feasibility of synthesizing marine weather forecasts from formatted weather forecast data. The task was to produce Arctic marine forecasts for five forecast areas east of Baffin Island (known as FPCN25 forecasts). Marine forecasts are one of several types of weather forecast based on the same basic weather data, each type emphasizing conditions of interest to a particular community of users. In producing marine forecasts, linguistic emphasis is placed on wind direction and speed, dangerous wind and freezing spray conditions, etc. MARWORDS is designed to be sufficiently modular and flexible so as to allow easy extension and adaptation to other types of weather forecast message (e.g., agricultural forecasts, public weather forecasts). Although the current project seems to have proved the feasibility of automatically synthesizing weather forecasts, extensive testing and refinement are required before MARWORDS can be introduced into daily use.

The MARWORDS system is the natural language component of a computer-aided forecast composition (CAFCOM) project, which envisages automating the process of creating forecasts from meteorological information. In the current manual procedure all available meteorological information (observations, radar and satellite imagery, and numerical weather prediction products) is made available to the weather forecaster. The weather forecaster must diagnose meteorological processes which will affect his particular area of interest throughout the forecast period, and then translate this knowledge into appropriate textual forecasts for various users.

The current development effort will attempt to make much more meteorological information available to the forecaster by means of an intelligent graphical workstation, without causing "data overload." Much of the forecast process can then be moved into the production of detailed weather depiction charts incorporating subjectively modified numerical weather prediction products. Software such as MARWORDS could be used to draft forecast messages. Hopefully this will significantly reduce the work load on the forecasters, making it possible to focus more attention on meteorological problems.

In the normal course of events, predicted values make up a continuum in both time and space. For simplicity, values are often given at regular steps in time (e.g., hourly) and space (either at grid points or at weather observing sites). Alternatively, forecast parameters may be given in terms of significant changes only. MARWORDS is flexible enough to accept both types of data description.

3. Design of the MARWORDS system

The complete project is illustrated schematically in Fig. 1. A TRANSLATOR program mediates between the graphics editor/manager and MARWORDS. The TRANSLATOR is a traditional piece of software involving meteorological computations. For example, geostrophic winds are calculated directly from the isobars in the graphics files. The results may be adjusted for factors such as isobaric curvature, stability, and even isallobaric effects. The TRANSLATOR needs information about which forecast regions are used to generate the forecast message (such as the FPCN25) and a geographical database defining the forecast areas. Given that information, it can sample the weather depiction data and generate the input data necessary for MARWORDS. All the linguistic knowledge (e.g., paragraph and sentence structure, and the lexicon) required to generate the forecast text resides in MARWORDS.

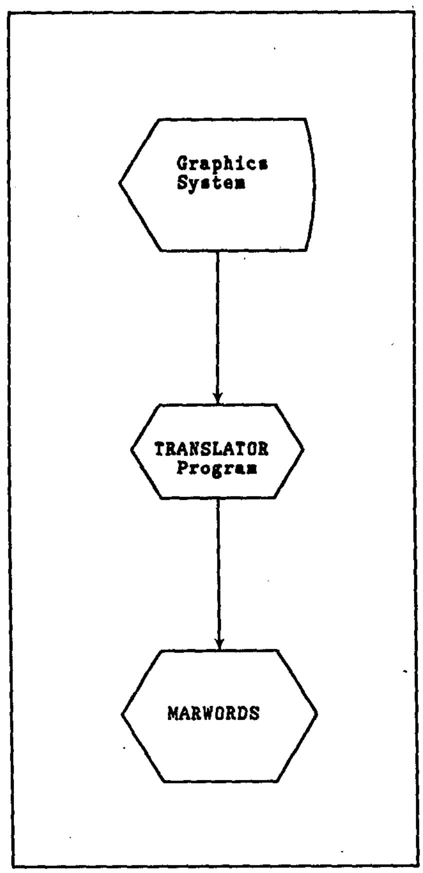

FIG. 1. CAFCOM system overview.

The definition of the data format used to communicate between the TRANSLATOR and MARWORDS was a major task in the design. This is a consequence of the variety of ways that meteorological information is presented in forecast text. A very simplistic approach would take all of the information in the graphics database and pass it on to the linguistic module. Because natural language is so expressive, this amounts to recapitulating the entire database. It becomes impractical to do this once the database grows to any size.

Some examples may help to illustrate why this is a complex task. The following texts were taken from manually produced forecasts.

1. *Northwesterly gales 35 to 45 knots diminishing to northwesterly winds 20 to 30 late this evening and to light tonight. Light winds Wednesday.*
2. *Winds southerly 25 to 30 knots except southerly 10 to 20 over the southern section of Davis Strait.*

In order to generate the first statement, the input data must be capable of expressing wind speed and direction ranges over periods of time. Modifiers such as "diminishing" and even "gales" belong in MARWORDS (i.e., not in the input data) so that terminology and warning criteria, which are specific to an individual weather office, can be isolated with the other linguistic elements.

Statement 2 illustrates that at some stage in the processing the system has to decide that the winds in the southern section of Davis Strait are significantly different from those in the rest of the area. It is not clear whether this is a decision that should be made by the TRANSLATOR or within MARWORDS itself. For the time being, this decision was left with the TRANSLATOR in order to strictly limit the scope of MARWORDS to forecast text generation. This is discussed further in section 9. In any case, the input data format has to be capable of expressing exceptions to more widespread conditions, and also attach a geographical descriptor to the data.

Further complications become evident in statements about weather phenomena relating to visibility.

3. *Patches of fog and mist becoming more extensive overnight. Visibility fair in mist and poor in fog.*
4. *Flurries over eastern areas today. Snow beginning this evening changing to or becoming mixed with rain over southern sections tonight. Flurries Monday. Visibility one half to 3 in precipitation except occasionally near one quarter in snow.*

In these cases, the variation in weather conditions (fog, mist, snow, etc.) must be given, and also linked to variations in visibility. Since MARWORDS does not incorporate any meteorological knowledge, this linkage must be expressed in the input data.

These examples illustrate some of the difficulties. The input data format must be flexible and expressive enough to provide the meteorological data, and its linkages and variations to MARWORDS.

Other criteria impacting the design were the selection of PROLOG as the programming language, and a requirement that the system be implemented on an IBM PC (or equivalent). These requirements were arbitrarily imposed so that some assessment could be made about the potential for implementation. Forcing the use of a PC-type environment is realistic in terms of hardware availability and cost when implementation at several forecast offices is considered. In addition, the program designers were forced to think in terms of program efficiency so that the computer time required to generate a forecast was not excessive. The whole system has to be able to operate within normal operational time constraints.

The selection of PROLOG as the programming language was based on previous experience with computer worded forecasts. When written using conventional programming languages (such as Fortran), these programs are very large and quite difficult to modify or maintain (Dueck 1985). Besides the fact that PROLOG is designed for symbolic computation rather than numerical calculations, it offers significant benefits in reducing the sheer volume of code for this type of problem. Ratios of 10 to 1 have been quoted for lines of Fortran code to lines of PROLOG code. Furthermore, PROLOG incorporates Definite Clause Grammar (DCG) which is a clear and powerful formalism for describing language (Pereira and Warren 1980). In fact, a portion of MARWORDS was written using DCGs.

In its current form (see Fig. 2) MARWORDS is a set of MPROLOG programs that call each other to control the information flow from data through text. (MPROLOG is a "modular" implementation of PROLOG available from Logicware, Inc.). These programs may be grouped into three components. Component I reads the formatted data (e.g., Table 1) from the input file and converts it into a set of formulas (Table 2) which can be manipulated by PROLOG.

Component II carries out a number of nonlinguistic operations. These include data consistency checks using built in databases of geographical and meteorological information, and the assignment of default values. (The only mandatory data required to generate a forecast message is wind data. All other parameters can default to climatological values.) Weather warning criteria are checked. These include freezing spray calculations based on wind speed, air temperature, and seasonally and regionally adjusted water temperature. An "archive" of data from the preceding forecast is checked to determine whether warnings are new, already in effect, or ended.

Forecasts for contiguous areas are merged whenever their forecast data is "similar" (according to criteria provided by the local weather office). This is a very important stage in generating computer worded forecasts. When similarity threshold conditions are satisfied, a single report formula is created for the merged areas under a header listing the areas. Finally, any data

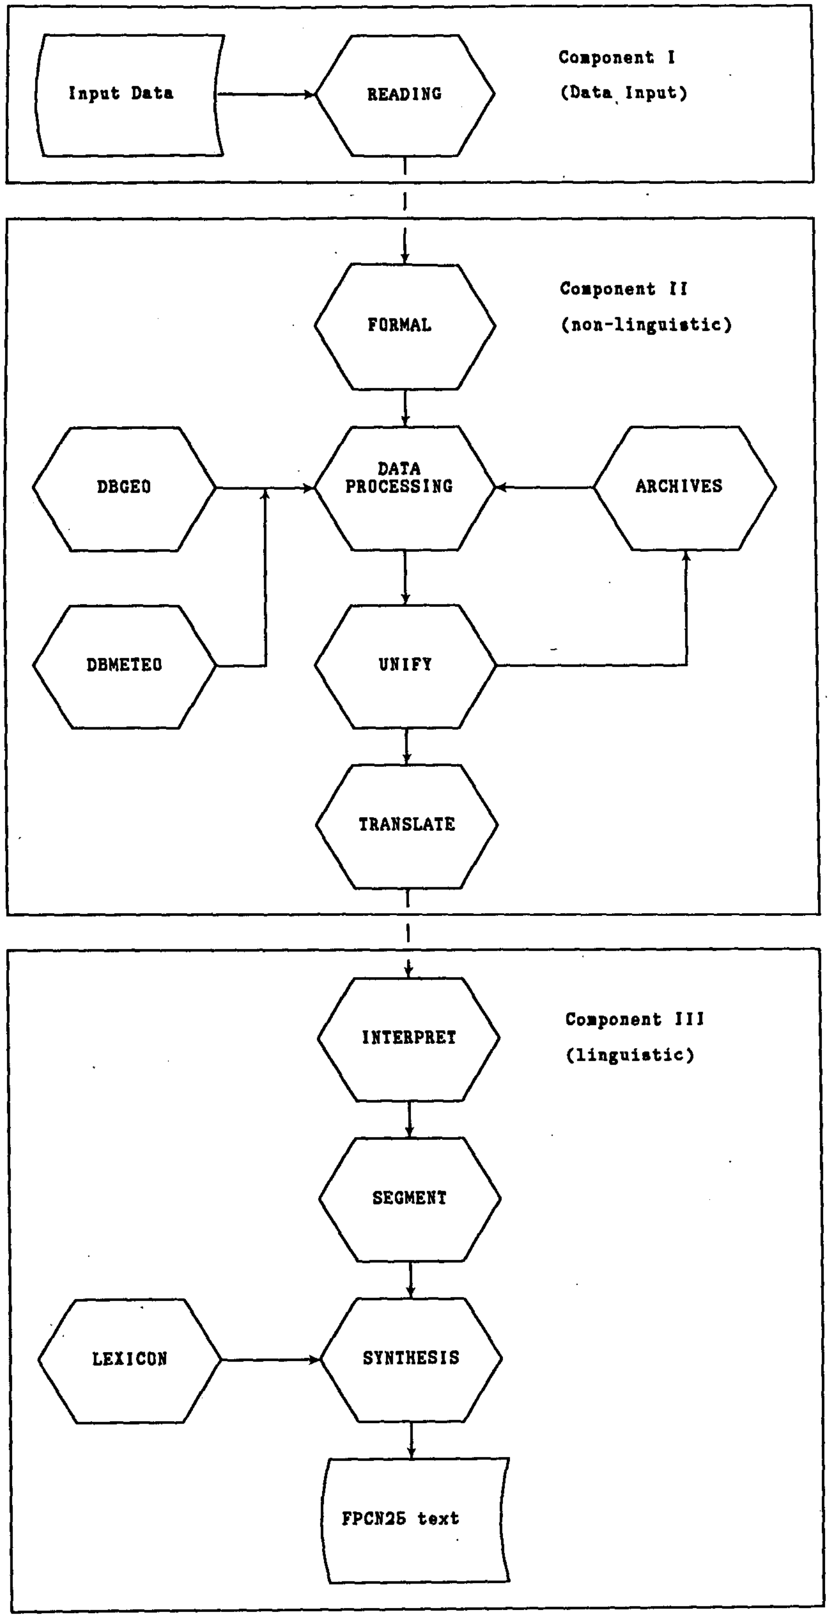

FIG. 2. MARWORDS as a set of linked MPROLOG programs.

TABLE. 1. Sample MARWORDS formatted input.

```
1000 tue 86/ 8/25 end.
brev wind 110 10
  temp -3
  sky ovc
  wea fog per & nt 6 snow per MOD
end.
davi wind 110 10
  temp -3
  sky ovc
  wea fog per & nt 6 snow per MOD
end.
frob wind 110 10 & nt 3 nl north low 180 40
  temp -3
  sky ovc
  wea fog per & nt 6 snow per MOD
end.
$
```

not explicitly used in the text is dropped. For example, air temperature is dropped after being used in the freezing spray calculation. The resulting formulas are then translated into "logical" representations (Table 3), which will eventually become separate paragraphs in the generated text.

Most of the linguistic functionality resides in Component III. For example, module *interpret* tests to see if changes in wind speed and direction are to be handled by two separate verbs. Separation of the data structure into clauses and sentences is defined, and words are selected from the lexicon based on semantic constraints. For example, terms like "backing" or "diminishing" are only applied when specific conditions are met.

TABLE 2. Initial data structure.

```
[ [hour(10),date(86,8,25)],
  [area(brev),
   wind([direction,110,speed,10]),[end_of_modif],
   temperature([-3]),]end_of_modif],
   sky([ovc]),[end_of_modif],
   weather([fog,periodic,unknown,unknown]),
     [and(nt(6),nl(_1776),
       end_of_modif]]
  [area(davi),
   wind([direction,110,speed,10]),[end_of_modif],
   temperature([-3]),[end_of_modif],
   sky([ovc]),[end_of_modif],
   weather([fog,periodic,unknown,unknown]),
     [and(nt(6),nl(_1776),
       weather([snow,periodic,moderate,unknown])),
       end_of_modif]]
  [area(frob),
   wind([direction,100,speed,10]),
     [and(nt(3),nl([north,low]),
       wind([direction,180,speed,40])),
       end_of_modif],
   temperature([-3]),[end_of_modif],
   sky([ovc]),[end_of_modif],
   weather([fog,periodic,unknown,unknown]),
     [and(nt(6),nl(_4780),
       weather([snow,periodic,moderate,unknown])),
       end_of_modif]]
]
```

TABLE 3. Data structure after nonlinguistic processing.

```
[
[brev,davi]&no_warning
   &wind(e,10)
 &sky(ovc)
 &(weather(fog,periodic,moderate)
  &weather(nt(same_24,12),
       nl(area),snow,periodic,moderate))
   &((fair,snow,periodic)&((poor,fog,periodic)
   &stop)),
frob&(begins((gale,9))
     &begins((f_s,[not_h,t(6),s(9)])))
   &(wind(e,10)
     &wind(nt(same_24,9),nl([north,low]),s,40))
   &sky(ovc)
   &(weather(fog,periodic,moderate)
     &weather(nt(same_24,12),
       nl(area)(snow,periodic,moderate)
   &((fair,snow,periodic)&((poor,fog,periodic)
   &stop))
]
```

4. A sample report

The following simplified example (Table 1) shows the input formatted data, using mnemonic descriptors, for a subset of Arctic marine areas. The first line gives details about the forecast issue time and date. Remaining lines are arranged in groups with each group describing the conditions in one forecast area. It begins with the region label (e.g., "brev" for Brevoort) and continues with the weather element information.

Input data is limited to the six most important parameters: 1) wind direction, 2) wind speed, 3) cloud cover classification, 4) precipitation types (if any), 5) precipitation frequency and intensity rating, and 6) air temperature. Further forecast parameters that are functions of the input parameters (e.g., warnings and visiblity ratings) are calculated by the first nonlinguistic module.

Changes in weather conditions are indicated with "& nt" and "& nl" where "&" is a simple connector in the input syntax, "nt" indicates a new time, and "nl" indicates a new location. Thus "nt 6" signifies that the following weather condition begins six hours after the previous one. Thus, the wind for "frob" (Frobisher Bay) is 10 knots from 110 degrees at 1000 UTC, changing to 40 knots from 180 degrees north of a low pressure center at 1300 UTC.

Table 2 shows the value of F in "*formal(Day,F)*" after the data above has been read by MARWORDS. Essentially, the data has been reorganized into a list structure which can be manipulated by PROLOG. The structure can be more easily seen if written as:

```
F → [ [date/time group],
      [  area 1 data  ],
      [  area 2 data  ],
      [  area 3 data  ]
    ]
```

In PROLOG, square brackets are used to enclose lists. Thus, F is a list structure whose elements also happen to be lists.

After reading and analysis, the data is manipulated in clausal form through data checking, area unification and data suppression stages mentioned above. It is then translated into a "logical form" just before input to the linguistic modules. Following the example above, the variable LOGIC_F in predicate *translate* has a structure:

```
LOGIC_F → [paragraph 1, paragraph 2, . . .]
```

Logical units of text have been identified (i.e., paragraphs) but clause and sentence boundaries have not yet been determined. Similarly, word selection still has to be done. Note that regions have already been grouped ([brev,davi]), and warning status determined.

Linguistic modules begin by calculating the values of significant semantic features. Specific words are selected on the basis of the meteorological data. (For example, winds that change direction in a clockwise direction will be described lexically as "veering" to the new direction.) Initially, words are selected using the most precise term available in the lexicon. Later, once the structure has been segmented into sentences it may turn out that some variant of the term would be preferable. For example, in a case where rain and snow are both causing 'fair' visibility, it would be preferable to generate:

"*Visibility fair in precipitation.*"

instead of:

"*Visibility fair in snow and fair in rain.*"

The linguistic modules also check to see if wind direction and speed are to be handled separately (i.e., using two verbs, such as "*veering and strengthening*"). The outlook sentence is synthesized for forecast issues requiring one. Module SEGMENT separates the data structure into sentences, and finally SYNTHESIS completes the forecast synthesis, including selecting particular words from the lexicon, and introducing punctuation.

Table 4 gives the final textual form of the marine forecast corresponding to the data of Table 1.

5. Knowledge sources for report synthesis

The MARWORDS architecture isolates different types of linguistic and nonlinguistic knowledge within appropriate modules. Grammatical, lexical, rhetorical and stylistic descriptions are based on an examination of all the marine forecast bulletins (manually) produced for the FPCN25 region during the 1983 and 1985 seasons (some 50 000 words in all). Examination of these forecasts had led to a fairly detailed grammar of this sublanguage (cf., Harris 1968; Kittredge and Lehrberger 1982). Linguistic knowledge is broken down into several types:

1) *Lexical semantics.* This specifies the conditions for choosing different terms (e.g. "gale" versus "storm" or "diminishing" versus "decreasing" winds). In addition, restrictions on word usage have to be specified. Thus, the term "precipitation" can be used to supersede both "snow" and "rain."

2) Frequency preferences among synonymous terms in the sublanguage of marine bulletins. For example, it may be permissible to use either "decreasing" or "diminishing" to describe the wind. If so, the word selection procedure can make use of the relative frequencies of these terms in hand-written forecasts to introduce some variety in the computer generated forecasts.

3) Syntactic patterns, including the possible and preferred sentence patterns for expressing messages of given types. That is, statements about wind follow a general pattern, and statements about visibility another pattern. A second type of syntactic knowledge concerns the rules for deleting repeated sentence constituents when two or more propositions are fused into a single report sentence.

4) Simple principles of text organization. These are specific to the variety of text to be synthesized, and

TABLE 4. MARWORDS output for data of Table 1.

marine forecasts for arctic waters issued by environment canada at 3.00 am mdt tuesday 25 august 1986.
valid until midnight tonight with an outlook for wednesday.

brevoort
davis strait
winds easterly 10. cloudy with flurries ending by midnight. fog patches. visibility fair in flurries and poor in fog.
outlook for wednesday . . . light easterlies.

frobisher-bay
gale warning issued . . .
freezing spray warning issued . . .
winds easterly 10 veering and strengthening to southerly gales 40 north of the low by mid-morning. cloudy with flurries ending by midnight. fog patches. visibility fair in flurries and poor in fog. outlook for wednesday . . . light easterlies.

hence a function of the data salience hierarchy (see below).

Nonlinguistic knowledge is of three types:

1) Geographical knowledge for each forecast area including its time zone, its limits of latitude and longitude, and the names of adjoining areas (to allow recursively merging adjacent areas in case of similar meteorological regimes);
2) Meteorological data including mean temperature values for air and water during each month of the Arctic shipping season (June through October) and record values for temperature and wind speed;
3) An "archive" of data from preceding forecasts, used to verify if wind warnings or freezing spray warnings are in effect.

Geographic knowledge is used primarily during the attempt to merge reports for adjoining areas. However time zone data is used to calculate local time associated with meteorological phenomena, and hence allow the selection of appropriate time descriptors (e.g. "by late afternoon"). Input data to MARWORDS is in Universal Coordinated Time.

6. Linguistic treatment of salience

Data salience refers to the relative significance of different types of information in the text being generated. This is important because of the impact it has on the overall structure of the text being generated (Conklin and McDonald 1982). In the case of marine forecasts, the data salience relationships are relatively simple and well defined. Consequently, they can be "hard coded" into MARWORDS. However, in public forecasts where more flexibility is allowed, these relationships will have to depend on the input data.

To illustrate this, consider how data salience affects the structure of marine weather forecasts. First, warnings of potentially hazardous conditions (strong wind and freezing spray in the FPCN25 region) constitute separate statements preceding the normal text. Only warnings are so positionally marked and informationally redundant. Within the normal text, sentence groups dealing with each forecast parameter are ordered by two principles: intrinsic interest of the data and implicit causal links between the events or states described. Thus wind direction and speed, as the critical factors in marine conditions, occupy initial position. However, visibility ratings, which should follow in order of importance, occur last by virtue of their dependence on fog/mist descriptions, which in turn are somewhat dependent on precipitation, which in turn follow cloud cover ratings. Sentence groups are therefore ordered as follows:

WINDS > CLOUD-COVER > PRECIPITATION
> FOG&MIST > VISIBILITY

Within each sentence group, sentences and clauses are first ordered according to the dichotomy "general versus local exception", and then chronologically within general and exceptional parts. Data salience also plays a role in word selection. For example, particularly strong winds are classified as "gales" (at 35 knots), "storm force winds" (at 49 knots), and so on. Also, more specialized sense verbs such as "veering" and "backing" tend to be used more for large changes of wind direction.

7. Temporal reference under increasing uncertainty

An interesting problem arises in ascribing particular time adverbials to points and intervals of (local) time. Manually written forecasts tend to "hedge" temporal descriptors slightly as reference time becomes more remote from the forecast issue time. For example, "Tuesday afternoon" or "by (Tuesday) evening" may be preferred for remote reference over the more precise "late Tuesday afternoon." This reflects the increasing difficulty in predicting onset times for remote meteorological events.

MARWORDS synthesizes time adverbials for events as a function of two factors: 1) the remoteness of the event (within or beyond 24 hours from forecast issue time), and 2) whether the event is essentially a point or an interval of time. In the case of time intervals, a more precise interval description can be chosen if all points in the interval fall within a day subpart (e.g., "Wednesday morning"), than if the interval straddles two day subparts, requiring a hyperonymic adverbial (e.g., "Wednesday").

8. Bilingual reports

The MARWORDS system was designed to accommodate the synthesis of marine weather bulletins in French as well as in English. Only the final three components in the processing sequence (*segment, synthesis,* and *lexicon*) are language-dependent (and only the last of these in a nontrivial way). Syntactic patterns and lexical entries for French must of course be furnished on the basis of independent linguistic study of the corresponding French sublanguage.

Canadian weather forecasts of all varieties are currently translated into French by the METEO system (Chevalier et al. 1978), developed at the University of Montreal. Basically, this system analyzes the English text to eventually arrive at a French equivalent. Roughly 5 to 10 percent of forecast sentences fail analysis and hence translation (Slocum 1985). This is due not only to input errors due to typing and line noise, but also to slight irregularities in the usage of English grammar and lexicon on the part of forecasters. The automatic synthesis of marine forecasts, on the other hand, should eliminate the need to deal with these irregularities by using a semantically complete and consistent subset of language to cover all foreseeable data configurations.

Work on MARWORDS may be seen as preparing the ground for an attractive alternative to machine translation of weather forecasts. The simultaneous synthesis of English and French forecasts directly from data would optimize the transfer of information to speakers of both languages. Parallel synthesis of bilingual forecasts bypasses translation altogether and minimizes human intervention, thus maximizing speed of transfer and (in principle) reliability. MARWORDS' logical structures for English forecasts are probably close to what is needed for French. Most of the system's work with a particular set of input data would therefore serve towards the synthesis of a report in either language.

9. Present status and future work

a. Bilingual MARWORDS

The CAFCOM project is proceeding at an accelerated pace. MARWORDS has already been extended to generate both French and English language forecasts. As expected, there was considerable commonality in the generation process, and the additional processing necessary to generate the second language added roughly 30 percent to the processing time. So, where unilingual generation of a forecast could take 75 seconds, complete bilingual generation of the same product takes roughly 100 seconds. In the case of marine forecasts, there are no important differences between French and English in the order of sentences in the text, and only a few cases where word order in French differs significantly from the English. Most of the differences are more localized within sentences, and have to do with word agreement (in gender and number), and the greater use of determiners in French.

b. New text generation requirements

MARWORDS is very deterministic and therefore quite efficient. That is, the program does not do much backtracking in search for solutions. This was achieved through the use of fairly complex data structures (viz. Table 2 and 3). Although this is effective in the case of marine forecasts, the same technique will not work for other forecast varieties. The reason is that marine forecasts have a very regular structure and a standardized meteorological content. This makes it feasible to develop the forecast syntax and vocabulary from archived text. However, there is a growing recognition that the structure and content of weather forecasts should be responsive to the needs of a variety of users. New kinds of forecasts and new formats are being considered. Therefore, a text generation system cannot be based on the analysis of existing archived text. It is likely that the grammar governing the text generation will have to be developed on an ad hoc basis. This means that the linguistic rules will have to be made more accessible than is presently the case in MARWORDS. Much of the efficiency built into MARWORDS was accomplished by embedding the syntactic knowledge fairly deeply in the code. However, we believe that MARWORDS is fast enough that some of this efficiency can be sacrificed without impacting the operational utility of the program.

These concerns have led to the development of a more formal linguistic treatment of weather forecasts. This will facilitate the development of a "general" text generation system (that is, one that can accommodate a wider variety of forecast products). Since weather forecasts are basically an artificial language rather than a purely natural one, this task is somewhat easier than might be expected. It has been further aided by an attempt within the Atmospheric Environment Service to standardize forecast content and terminology. Nevertheless, there is still wide latitude for regional variation, and any text generation system will have to be able to accommodate these variations.

c. Definite clause grammar

Definite clause grammar (DCG) is being used to develop this new linguistic treatment. DCGs are a powerful linguistic tool, and appear to be well suited to this particular problem. The DCG formalism is very similar to the usual method of representing context free grammars, but at the same time DCGs run almost immediately as PROLOG code. They can also be compiled for greater efficiency. As more generality is introduced, some degradation in execution time is expected, but the advantages of a more general approach should more than compensate. In addition, progress in computer systems technology is expected to overshadow this factor.

The DCG treatment of the grammar provides the means for encoding the syntax necessary to generate a sentence. For example, the structure of a simple sentence such as:

Temperatures near 8 degrees

rising to 15 later this afternoon.

can be expressed using DCGs as:

sentence → *head, value, phrase.*

phrase → *verb, value, time_phrase.*

The 'head' is simply the sentence "subject" (i.e., "Temperature"), 'value' refers to the phrase stating the temperature value ("near 8 degrees"), and 'verb' is "rising to." Other software is necessary to relate the meteorological values to the terms which are finally realized in the text. It appears that the DCGs given in the simple example above, could also be used to generate sentences about wind, sky condition, and so on. The major difference is in the selection of terms. This selection process hinges on the field being described (e.g., wind, temperature), and on the forecast variety. Thus, the verb in a corresponding wind sentence might be "strengthening" instead of "rising." This provides the thrust of the development now underway to generalize MARWORDS.

d. Semantics and the dictionary

In the revised MARWORDS, underlying differences between forecasts are concentrated in the dictionary. Each forecast variety has its own dictionary. For example, the term "veering" is included in the marine forecast dictionary, but not in the public forecast dictionary. By restricting the dictionary used in the generation process to the significant weather elements and specific terminology for a given forecast, it should be feasible to apply substantially the same grammar to many forecast products. Furthermore, by concentrating much of the text generation functionality in what is in effect a data file (i.e., the dictionary), it becomes practical to have individual users tailor their own regional or application specific terminology.

The development of a more comprehensive dictionary has had other valuable side effects. In text generation, what you do NOT say is at least as important as what you DO say. Computer generated text, which does not deal with this problem is uniformly awful. This is a pragmatic consideration that has an impact on the merging of similar forecast areas, and on the merging or lumping of weather conditions in time. Many of the restrictions that govern merging in space and in time can also be encoded in the dictionary. To put it more simply, the system cannot generate unnecessary text if it does not possess the words or terms to express that text. The phrase "winds north 10 at 3 pm becoming north 15 at 4 pm" is acceptable (i.e., grammatical) English, but nevertheless, it does not occur in written forecasts because of pragmatic considerations. The most common restriction is that the most precise term available to describe the time of day is something like "this afternoon," and it does not make sense to say "winds north 10 this afternoon becoming north 15 this afternoon." In fact, this restriction on the available terminology has its source in the dictionary used to generate the text.

This idea has been applied to the problem of merging forecast information in time. Bermowitz (1983, 1987) describes some other approaches to this problem. In the revised MARWORDS, semantic tests for significance from one time period to the next are used to determine whether the change should be realized in the text. Thus, if the dictionary only contains one term for the overnight period (e.g., "tonight"), MARWORDS will look through all the data for that time period and only realize the most significant weather condition. The determination of what is significant is itself encoded in the dictionary.

The dictionary being developed for MARWORDS incorporates both the syntactic features of the words (e.g., is the word a noun or verb with respect to the DCG), and also the semantics of the word. The term "semantics" refers to the "meaning" of the word. Thus a 'strong' wind refers to a wind speed between 20 and 33 knots in a marine forecast. These semantic features have been added to the dictionary using a form of "procedural semantics" (Woods 1968). This is a relatively straightforward approach to semantics, and appears to fit in well with the current approach.

In procedural semantics, the meaning of the sentence is identified with the action taken as a result of hearing the sentence. In our case, the meaning is identified with a time series of meteorological states, and the 'actions' taken amount to attaching values to the times and areas over which these events occur, and to the intensities associated with these states. To give an example, the dictionary entry for 'strong' is:

wind_speed(*adjective*, *strong*, *3*, [[*range*, *20–33*],
[*forecast*, *marine*]], *syntax*(. . .))

This indicates that the term 'strong' is to be used in describing wind speed, and that the speed must fall into the range 20 to 33 knots. In addition, the forecast being generated must be a 'marine' forecast. The '3' indicates a priority or significance rating. The term 'light' has a priority of 1 and 'gale' has a priority of 4. When faced with both strong winds and gale force winds during the same time period (e.g., 'tonight'), MARWORDS will select the higher priority item. The entry for a public weather forecast would very likely have a different wind speed range, if the word were to be used at all. At present, only five semantic tests are necessary to handle the terminology, which has been encoded in the dictionary.

e. Extension of the TRANSLATOR

The current CAFCOM system design involves three components (Fig. 1): the graphics system, the TRANSLATOR program, and MARWORDS. The TRANSLATOR program was expedient in the early stages of development, but has drawbacks that would make it difficult to scale up to a real implementation. Originally it was viewed as a convenient means of 'sampling' the graphics database in order to develop and format the input to MARWORDS. By necessity, it has to incorporate much of the same knowledge that is encoded in MARWORDS. Geographical areas and place names are the most obvious examples. Thus, any time a new region is added to MARWORDS, corresponding changes must be made to the TRANSLATOR. More severe complications ensue when MARWORDS needs slightly different data. For example, some areas use visibility ranges, and others do not. This means that the TRANSLATOR now needs to know specifics about the actual forecast structure. It would be preferable to embed the functionality of the TRANSLATOR in MARWORDS, so that the graphical database could be queried directly by the text generation software. In fact, this is not a new problem in text generation. Sondheimer and Nebel (1986) describe an approach known as 'inquiry semantics' to deal with a similar situation.

Work has begun on moving the TRANSLATOR functionality into MARWORDS itself. This will involve activating computational (Fortran or Pascal) software from within the PROLOG text generator. To facilitate this whole process, MARWORDS is being ported from the VAX/IBM-PC environment to the HP 9000 workstation used by the graphics subsystem.

One of the side benefits of the work on MARWORDS is that it has made us more aware of the problems of representing specific types of meteorological knowledge. There is no question that the forecast text is enriched by the inclusion of local effects. However, many of these factors (e.g., orographic effects, coastal fog, channeling of winds) are never explicitly dealt with by the computer system. The forecaster adds this level of detail in the final stages of forecast text composition. To assist in doing this, the graphics system has been designed so that the data can be manually annotated. The forecaster can note on the weather depiction chart itself that the area of fog being drawn is indeed "coastal," or that it only occurs in "onshore upslope flow." MARWORDS has demonstrated that this type of annotation can be used to good effect in the text generation process. The problem is that in a complicated weather situation, the forecaster could spend as much time annotating the chart as editing the forecast. Alternatively, a knowledge-based system could be used to infer that in a given situation, a particular local effect comes into play. This type of approach will be required eventually if we really hope to relieve the forecaster of "data overload."

By giving control of the TRANSLATOR to MARWORDS, we are preparing the way for the development of an expert system to deal with local effects that are not specifically included in the graphics. Obvious examples are the identification of coastal fog and the channeling of winds in fiords. Such a system would replace the TRANSLATOR. With both the graphics system and the text generation portion of the CAFCOM project well underway, this aspect of the project is now being actively discussed. MARWORDS will become a 'coupled' system, involving both symbolic and numerical computation. The degree of interaction between the linguistic component and the proposed expert system is still not clear. However, it seems reasonable to apply more expertise to identifying the first major frost or of the season rather than a frost which occurs in the middle of the winter. Similarly, an early frost should be mentioned in the forecast text, whereas frost which occurs in midwinter generally is not significant enough to merit special mention. This interplay between the text generator and the expert system implies that the system could be "deeply coupled" (Kitzmiller and Kowalik 1987). Initial work is expected to begin on this phase late in 1987, so that a preliminary version of the complete system can be field tested early in 1989.

f. Summary

MARWORDS is written in MPROLOG (a dialect of PROLOG), and runs on a VAX under VMS as well as on PC/XT/AT MS-DOS compatible microcomputers. It incorporates approximately 850 rules or PROLOG statements (430 linguistic and 420 nonlinguistic). Synthesis of a complete five-area forecast (about 150 words) using interpreted PROLOG code takes about half a minute for the VAX implementation and a minute for the AT implementation. Compiling the code yields about a 30% improvement.

MARWORDS is part of a larger scale research and development effort concerned with computer-aided forecast composition. The intention of the MARWORDS project was to investigate the feasibility of applying natural language techniques to the forecast text generation problem. The example in Table 4 (above) is a demonstration of the flexibility and capability of the software. Text cases, including both the weather depiction charts and the forecasts generated from them, have been reviewed by operational meteorologists at a number of sites in Canada and generally been judged to be quite acceptable. To a large extent, this can be ascribed to the richness of the input data syntax. Even so, this data is relatively straightforward to produce by computer compared to the final forecast text.

MARWORDS has succeeded in splitting the linguistic knowledge embodied in a forecast from the meteorological knowledge. By so doing, the forecaster should be able to concentrate on manipulating surface weather representations graphically, leaving the workstation to generate first the input data for MARWORDS and finally the forecast itself.

The selection of PROLOG appears to have been justified. The code is compact and accessible. It has been relatively easy to extend the original program to include differences in regional terminology. Variations in the actual structure of the text (for example, adding a sentence about wave heights in relation to wind speed) would be more difficult. However, current work on MARWORDS should help to eliminate this difficulty.

It is expected that new linguistic requirements will make new demands on the graphics system, and vice versa. At present, the TRANSLATOR samples sequential weather depiction charts, and MARWORDS assumes that weather conditions are essentially stationary from one chart to the next. Sophisticated time interpolation schemes are now being investigated for the graphics system. If these prove successful, semantics for the timing of significant changes in the meteorological fields can be made more precise. In addition, the semantics of the generated text can be refined to include descriptions of the character of the change.

REFERENCES

Becker, J., 1975: The phrasal lexicon. *Proc. TINLAP 1.* Cambridge.

Bermowitz, R. J., 1983: Computer worded forecasts. U.S. Dept of Commerce, NWS Techniques Development Laboratory, *Tech. Proc. Bull.* No. 336.

——, 1987: Combining forecast periods in computer worded forecasts, U.S. Dept of Commerce, NWS Techniques Development Laboratory, TDL Office Note 87-2.

Chevalier, M., 1978: TAUM-METEO. Groupe TAUM, Universite de Montreal.

Conklin, E. J., and D. D. McDonald, 1982: Salience: the key to the selection problem in natural language generation, *Proc. 20th Annual Meeting of the American Society for Computational Linguistics,* Toronto, 129–135.

Contant, C., 1986: Generation automatique de texte: application au sous-language boursier. M.A. thesis, Dept. de Linguistique, Universite de Montreal.

Dueck, D., 1985: CAFCOM initiation project: Preliminary analysis, Internal Rep., Atmos. Environ. Serv., Downsview, Ontario.

Harris, Z., 1968: *Mathematical Structures of Language.* Wiley-Interscience.

Kittredge, K., and J. Lehrberger, Eds., 1982: Sublanguage: studies of language in restricted semantic domains. deGruyter.

Kitzmiller, C. T., and J. S. Kowalik, 1987: Coupling symbolic and numeric computing in KB systems, *AI Magazine,* Vol. eight, (2), 85–90.

Kukich, K., 1983: Design of a knowledge-based report generator, *Proc. 21st Annual Meeting of the Association for Computational Linguistics.*

McDonald, D., and J. Pustejovsky, 1985: Description-directed natural language generation, *Proc. Ninth Int. Joint Conf. on Artificial Intelligence,* Los Angeles, 799–805.

Miller, R. L., and H. R. Glahn, 1985: Techniques used in the computer worded forecast program for zones: Interpolation and combination, U.S. Dept. of Commerce, NWS Techniques Development Laboratory, TDL Office Note 85.

Pereira, F. C., and D. H. Warren, 1980: Definite clause grammars for language analysis—a survey of the formalism and a comparison with augmented transition networks. *Artificial Intelligence,* Vol. 13, 231–278.

Slocum, J., 1985: A survey of machine translation: its history, current status, and future prospects, *Comput. Linguistics,* **11**(1), 1–17.

Sondheimer, N. K. and B. Nebel, 1986: A logical-form and knowledge base design for natural language generation, *Proc. fifth Nat. Conf. on Artificial Intelligence,* Philadelphia, 612–618.

Woods, W. A., 1968: Procedural semantics for a question-answering machine, *AFIPS Conf. Proceedings,* [33 Fall Joint Comput. Conf.], 457–471.