

# Automatic Generation of Multimodal Weather Reports from Datasets

Stephan M. Kerpedjiev\*

Institute of Mathematics  
Bl.8, Acad. G. Bonchev Street  
1113 Sofia, Bulgaria

## Abstract

Weather reports are created in various modes — natural language text, specialized language text, tables and maps. The system presented allows the user to define his needs of weather information and requirements on the form of presentation. The system analyzes a dataset obtained through specific procedures of forecasting or observation, plans the product according to the user requirements and generates its components. Special emphasis is placed on the coherence of the report by investigating the rhetorical structures observed in this kind of text and the coordination between a map and a text specifying it. The method of generation is a knowledge-based one with three types of knowledge employed in the system — terminological, rhetorical and grammatical. A prototype has been implemented and tested with original datasets.

## 1 Introduction

The generation of information products stepped into a new phase characterized by the intensive application of artificial intelligence, computational linguistics and other modern information technologies. Currently, various data are collected into databases and specific procedures are applied for processing those data into forecasts, analyses, surveys and other types of information products. Usually, those products are in numerical form which is unsuitable for the general audience and even for many specialists. Therefore this data has to be converted into a human-oriented mode such as natural language (NL) text, tables, maps, diagrams. The automatic conversion requires formalizing the process — a problem which nowadays cannot be attacked successfully except by gathering, coupling and employing various types of knowledge — common sense, about the subject domain, grammatical, etc.

In this paper, we report on a study of the automatic generation of multimodal weather reports from observed or predicted data. This particular problem is significant both from a practical point of view (various weather reports are to be made every day in many weather centers all over the world) and for its scientific aspects (it manifests the basic features of the generation of verbal

reports from data). Our work relates closely to three areas: NL generation, multimodal documents and weather information processing.

The communicative act performed by the system is the description of an observed or predicted situation. Other works that consider analogous communicative acts are (Davey, 1979) on the description of tic-tac-toe games, (Kukich, 1983) on the generation of market reports, (André *et al.*, 1988) about simultaneous commenting on a soccer game recorded as a sequence of digitized video frames. In our case the situation to be presented is coded into a dataset obtained through routine procedures of weather forecasting or observing.

Our approach to NL generation follows the basic steps as described by McDonald (1987), viz. selection of the content portions that are to be communicated to the user, planning the text by adoption of the most suitable rhetorical schemas, realizing the discourse plan as a surface structure and its rendering as a text. The goal and the context of the utterance are specified by the user together with parameters concerning the precision of the information and the message length. We place special emphasis on the content production component which scans the dataset and extracts assertions from it, as well as on the rhetorical structures observed in weather reports.

Recently an increasing interest has been observed in the processing of multimodal documents, the research being focused on the coordination between the different modalities (NL, graphics, video images, pointing). Some projects with intensive research in this area are XTRA (Algayer *et al.*, 1989), COMET (Feiner and McKcown, 1990), ALFresco (Stock, 1991). To a large extent this aspect of our project was inspired by the WIP project (Wahlster *et al.*, 1991) in which the coherence of multimodal discourse is investigated and common sense knowledge is employed in the coordination between the textual and the graphical components of instructions for the use of domestic appliances.

We consider the case of supplementing a weather map with a verbal note specifying those content portions that cannot be presented on the map or whose graphical presentations distort the original information. The system discovers such deficiencies of the graphical presentation and generates a verbal comment on the map.

There are various projects concerning the production of weather reports, each of them setting specific goals

---

\*This work has been partially supported by the Ministry of Education and Science and the Bulgarian Academy of Sciences.

The RAREAS, RAREAS-2 and FoG series of systems (Bourbeau *et al.*, 1990) developed by one of the most successful groups in weather report generation shares many concepts with the current project, the main differences lying in the specification of the product and in the modality of the generated documents. Thus the Canadian group deals exclusively with NL forecasts while our project considers the generation of multimodal reports and employs diverse means of specifying what one needs and in what mode he wants to receive the information.

Our previous work (Kerpedjiev, 1990; Kerpedjiev and Noncheva, 1990) concerns the conversion of weather forecasts from textual form to weather maps or texts in another language. This is a translation problem rather than a generation one. Another feature that makes the previous work different from the current one is the lack of coordination between the graphical and the textual parts. Finally, in this project we employ the knowledge-based approach which allows a higher degree of flexibility and easy adoption to various types of products.

## 2 Architecture of the system

The architecture of the system is shown in Figure 1. The initial data for the system are the dataset and a specification of the final product prepared by the user in the form of a template. The system works as follows:

1. The monitor interprets the template and successively calls the scanner, the planner, the text or map generator, and the formatter with parameters extracted from the template.
2. The scanner analyzes the dataset and extracts assertions about the weather situation.
3. The planner applies rhetorical and grammatical knowledge to convert the extracted set of assertions into a surface structure or a map plan.
4. The text generator makes use of the lexicon to linearize the surface structure into a text; the map generator creates a cartographical presentation from the map plan using the visual library.
5. The monitor evaluates the generated text against parameters specified in the template (e.g. length of the text), and if the result does not meet the requirements of the user, re-activates the previous processors with adjusted parameters for providing alternative solutions.
6. The formatter assembles the product.

## 3 Weather report analysis

In this section we describe the initial dataset, the assertions extracted from the dataset, and the form of the final product as well as its specification.

### 3.1 Structure of the initial dataset

The initial dataset obtained through observations or numerical forecasting techniques is compiled in tabular

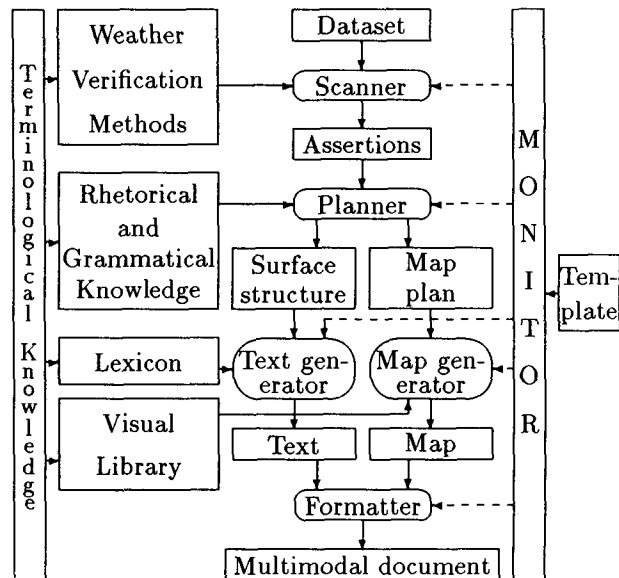


Figure 1: The system architecture

form with lines corresponding to the locations, columns — to the weather elements considered in the report, and subcolumns — to the time instants to which the data refer. The locations are either the stations where data is collected or the nodes of a regular grid in which the numerical forecast is computed.

In our experiments we used weather data collected through observations made at the main synoptic hours (00, 06, 12, 18 GMT) in 50 weather stations dispersed over the territory of Bulgaria. Ten weather elements have been considered: cloud amount, precipitation type and amount, wind speed and direction, min and max temperatures, and the phenomena fog, frost, thunderstorm.

### 3.2 Intermediary representation

An intermediary representation is necessary because the initial dataset describes the weather in terms of a scientifically based model which may not meet the user conceptions. It is intended to accommodate in a language independent form those facts that will be conveyed to the user.

What are the major differences between the initial data and the intermediary representation? Firstly, they pertain to different territory and time models. While the locations in a dataset are weather stations or grid nodes, in the intermediary representation they are administrative and geographic areas known to the audience. The dataset contains data referring to time instants, whereas the facts of the intermediary representation refer to parts of the day (such as morning) and whole days. Hence, the facts in the intermediary representation summarize the initial data with respect to time and space.

The second difference concerns the weather models. In addition to the basic weather elements employed in the

initial dataset, the intermediary representation makes use of some derived attributes. So, the basic numerical quantities *wind speed* and *precipitation amount* are converted into qualitative characteristics — *wind strength* and *precipitation intensity*, respectively. Particular examples of other derived attributes are given in section 4.

We call the facts from the intermediary representation *assertions* and denote them as quintuples:

$(w\_attribute, w\_value, region, period, precision)$ .

The weather attribute and the weather value represent the goal of the assertion; the region and the time period form its context; the last component denotes the precision of the summarization both over time and space and in the case of facts with derived weather attributes.

### 3.3 Structure of the final product

The final product is a natural or specialized language text, a table<sup>1</sup> and/or graphics. The basic constructs of those modes are oriented towards the expression of assertions — the atomic content portions extracted from the dataset. A NL sentence or clause, an icon placed in a certain position on the map, and a lexical or numerical weather value put in a cell of the weather table are all constructs of this type.

Figure 2 illustrates some modes. For example, the assertion (*cloud-amount, overcast, Nor-Bul, morn, high*) expressible through the NL sentence “In the morning it will be cloudy over North Bulgaria” is represented as a weather map (Figure 2d) and in the upper left cell of the weather table in Figure 2c. Weather reports can be structured in different ways. The text in Figure 2b is an *enumeration* type of text with four independent segments labeled by the regions they pertain to, and the text in Figure 2a is a sample of a *narrative* text.

### 3.4 Specification of the final product

The user’s requirements on the final product are specified by means of a template. It defines the mode, goal and context of the product, as well as various parameters concerning the precision of the information, the length of the message, and the style of text or map. The template consists of two types of statements: statements defining the modal structure of the document and content production statements.

There are four statements defining the modal structure of the final product: *narration*, *enumeration*, *table* and *picture*. The general format of a modal structure statement is given below:

```
<modal_struct_statement>{<external_context>}
  <sequence_of_content_production_statements>.
```

The following examples of statements are intended to generate a product with the modal and content structure of the forecast in Figure 2.

<sup>1</sup>We should distinguish the standard tabular report representing the initial dataset from the user tables.

```
narration{}
  text{clouds,precip, wind, phen, temp;
    Bul,whole_day; precision=0.6};
enumeration{Nor_Bul,East_Bul,Sou_Bul,West_Bul}
  text{clouds,precip,wind,phen,temp;
    whole_day; length=[20,100]};
table{Nor_Bul,East_Bul,Sou_Bul,West_Bul}
  value{clouds;morn}; value{clouds;noon};
  value{clouds;aftern}; value{clouds;even};
picture{Nor_Bul,morn}
  map{clouds;;maxasrt=2}
```

The content production statements are *value*, *text* and *map*. The first type of statement produces the lexical presentation of a single weather value (e.g. *overcast*, c 15°C); the text production statement makes complete sentences linked in a coherent text; and the map production statement generates a cartographical presentation of the assertions by placing icons of the particular weather values in certain positions of the map. The format of content production statement is as follows:

```
<content_production_statement>
  {<goal>;<context>;<parameters>}
```

The goal is the set of weather attributes in which the user is interested. The context contains the region and the time period for which weather information should be extracted. The part of the context given as an external context in the modal structure statement makes a heading of the corresponding section, and therefore this context may not be explicitly mentioned in the text. The parameters specifying the produced content portions are divided into three groups: precision rate, length and style of the message.

The *precision* parameter defines the minimum precision rate that must be guaranteed by the generated message. By specifying a high precision value, we rule out vague sentences like “it will be cloudy in some portion of North Bulgaria” and force the system to retrieve more precise assertions from the dataset.

The parameters restricting the length of the message are of three types:

- *maxasrt* – determines the maximum number of assertions generated for each attribute from the goal;
- *length* – restricts the length of the final text by specifying the minimum and maximum number of characters in it (applies to text production only);
- *detail* – specifies the level of detail of the produced message on a three-element qualitative scale {*concise, normal, full*}; *concise* detail implies that only a summary information for each goal should be extracted; *full* detail makes the system extract complete information; and *normal* detail produces a text with a level of detail in between the two extremes.

The *style* parameter defines the message language. In the case of text production, the language could be a sul-

South and East Bulgaria will be mostly sunny. Clouds with showers are expected in North Bulgaria and in the afternoon. In East Bulgaria the wind will increase. High temperatures 25–30°C. Low temperatures 18 – 20°C.


(a) A narrative text

North Bulgaria: Mostly cloudy weather with showers in the afternoon. High temperatures 25°C. Low temperatures 18°C.  
East Bulgaria: Clear in the morning and cloudy in the afternoon. Increasing of the wind. High temperatures 25 – 27°C, low temperatures - 20°C.  
South Bulgaria: Mostly sunny weather. In the mountains the afternoon will be cloudy with showers. High temperatures 27 – 32°C. Lows 18 – 22°C, in the mountains 8 – 12°C.  
West Bulgaria: Cloudy sky will prevail in North-West Bulgaria. In South-East Bulgaria mostly sunny weather but the afternoon will be cloudy with showers. High temperatures 25 – 30°C. Low temperatures 18 – 20°C.

(b) An enumeration type of text

Region	Cloud amount			
	morning	noon	afternoon	evening
North Bulgaria	ov	pc	pc	cr
East Bulgaria	ov	ov	pc	cr
South Bulgaria	cr	cr	pc	cr
West Bulgaria	pc	pc	pc	pc

(c) A weather table (cr – clear, pc – partly cloudy, ov – overcast)



(d) A weather map

Figure 2: A multimodal weather document

set of a NL, a telegraphic type of language, or a special-purpose language conformed with specific users needs. In the case of map production, the style determines what types of icons should be used and how the time will be presented (through several maps, by explicitly indicating the time periods on the map, etc.). For reference purposes, each style is given a unique identifier, e.g. *english*, *teleg-bul-report*, *avionic*.

## 4 Terminological knowledge

The terminological knowledge-base (TKB) represents the weather, territory and time models.

The weather model consists of the set of weather attributes, their domains, relations between some domains, and rules for calculation of derived attributes. So the qualitative weather element *wind strength* with a five-element ordered domain is calculated from the numerical basic attribute *wind speed* by means of the rule:

$$w\_strength = \begin{cases} calm & \text{if } w\_speed \in [0, 2] \\ light & \text{if } w\_speed \in [3, 6] \\ moderate & \text{if } w\_speed \in [7, 14] \\ strong & \text{if } w\_speed \in [15, 20] \\ gale & \text{if } w\_speed > 20 \end{cases}$$

The derived weather attribute *cloud change* with a four-element nominal scale is calculated by means of a rule based on the properties *monotonicity* and *amplitude* of the basic attribute *cloud amount* with a three-element

ordered scale  $\{clear, partly\_cloudy, overcast\}$ . Similar rules allow the system to calculate summary weather attributes. For example, the *clouds* attribute unifies the domains of the attributes *cloud amount* and *cloud change* into the domain  $\{clear, partly\_cloudy, overcast, clouds\_increase, clouds\_decrease, variable\}$ .

Two weather values  $v_1$  and  $v_2$  are considered related if they represent co-occurring weather characteristics (e.g. *overcast* and *rain*) and opposite if the characteristics are associated as contrary (e.g. *clear* and *overcast*). The two relations are defined in the TKB by means of the predicates  $related(v_1, v_2)$  and  $opposite(v_1, v_2)$ .

The territory model represents the set of regions, their carriers and certain logical links between them. The function  $carrier(r)$  returns the set of stations that belong to  $r$ , thereby allowing us to treat the regions as sets. The predicate  $path(r_1, r_2, \dots, r_n)$  indicates that there is a path starting from region  $r_1$ , passing through  $r_2, \dots, r_{n-1}$ , and reaching  $r_n$ .

The time model defines the time periods as intervals of time instants through the functions  $begin(t)$  and  $end(t)$ . Two relations between time periods supported by the TKB are partial order ( $t_1 \leq t_2$  iff  $end(t_1) \leq begin(t_2)$ ) and inclusion ( $t_1 \subset t_2$  iff  $[begin(t_1), end(t_1)] \subset [begin(t_2), end(t_2)]$ ).

The relations between weather values, regions and time periods are employed in the selection of rhetorical schemas (cf. section 6).

## 5 Scanning the dataset

The scanner determines the content portions of the message by computing relevant assertions from the dataset. The monitor calls it with two types of queries specifying the goal (a single weather attribute), the context and a parameter concerning either the precision rate or the maximum number of assertions to be produced:

```
scanp(clouds, Bul, whole_day, 0.8)
scana(clouds, Bul, whole_day, 3)
```

The first query makes the scanner extract assertions about the *clouds* attribute applied to Bulgaria and the whole day, and with a precision rate greater than or equal to 0.8. The second query restricts the maximum number of assertions that should be extracted to three.

The scanning is carried out in three steps: generation of a full set of assertions, pruning the full set of assertions and selection of the final set of assertions.

In the first step, the scanner applies weather verification techniques (Kerpedjiev and Ivanov, 1991) to generate an assertion for each context that belongs to the query context. Such an assertion contains the weather value that approximates the data subset corresponding to that context with the highest precision rate.

In order to avoid a combinatorial explosion during the selection, the set of assertions is pruned by removing all assertions that can be inferred from other assertions. (An assertion  $a_1$  can be inferred from  $a_2$  if both assertions convey the same weather value, but  $a_1$  relates to a sub-context of  $a_2$  and its precision rate does not exceed that of  $a_2$ .) The average reduction rate of the pruning is 70%.

The selection of a combination of assertions is first made independently for each weather value of the goal attribute. A combination of two assertions  $(w, v, r_1, t_1, p_1)$  and  $(w, v, r_2, t_2, p_2)$  is evaluated by means of the formula  $\min(p_1, p_2, 1 - p', 1 - p'')$  where  $p'$  and  $p''$  are the precision rates of the assertions  $(w, v, r - r_1 - r_2, t, p')$  and  $(w, v, r_1 + r_2, t - t_1 - t_2, p'')$ , respectively,  $r$  and  $t$  being the query context. Then the scanner selects the most precise combinations for the different weather values and returns them as a response to the query.

## 6 Planning the report

The planner assimilates a set of assertions into a surface text structure or a map plan. Since planning is essentially a process of arranging the information in a coherent way, we will consider at first the coherence in weather reports and then will elaborate the planning techniques.

### 6.1 Coherence

Coherence can be ensured for the portions created only by the text production and map production statements since the modal structure statements combine the constituent parts mechanically without caring for the consistency between them.

Coherence of a text portion is achieved by selecting a rhetorical schema that suits best the current set of assertions. The main vehicle for ensuring proper organization of the text content is the employment of existing relations in the TKB. Indeed, those links represent common associations and orderings of the objects, and following any of them while reading or hearing the text will enable the user to assimilate the information easily with minimum cognitive effort.

Based on the analysis of a number of textual weather forecasts and reports, we have extracted and collected seven types of rhetorical schemas:

*Presentation by weather attributes.* An assertion about a given attribute cannot interpose a sequence of assertions concerning another attribute.

*From a summary to details.* An assertion with a context which includes the context of another assertion is conveyed before the second assertion.

*Temporal progression.* The assertions are ordered by the successive time intervals they pertain to.

*Spatial progression.* The assertions are arranged in such a way that their regions form a conceptually existing path.

*Coupling related values.* Assertions with related values and intersecting contexts are rendered in a group.

*Contrast.* Two assertions with opposite values are conveyed together to contrast with each other.

*Presentation by weather values.* The assertions about a given attribute with an ordered domain are conveyed in successive groups relating to the particular weather values.

The problem of supplementing a graphical portion (created by the map production statement) with a verbal comment may arise when the situation presented or the map is dynamic, imprecise or uncertain. Due to the lack of proper graphical means of expression for such properties, a text has to be created that specifies the information available on the map. The following example illustrates the problem.

Suppose that the assertion (*phen, fog, Nor-Bul, morn high*) has to be shown on a map created for the whole day. A presentation with the pictograph for fog placed in one or more positions dispersed uniformly over the region specified may prove misleading because the important information about the time period is absent. To resume the correctness of the map the following concise text message should be created:

“The fog in North Bulgaria will clear by noon.”

It consists of the reference part “the fog in North Bulgaria” and the specification part “will clear by noon”. The reference part identifies the phenomenon through elements expressed on the map while the specification part conveys the missing or distorted elements.

## 6.2 Text planning

The conversion of a set of assertions into a surface structure poses two main problems:

- How to find the most suitable rhetorical structure of the text?
- How to realize this structure into the surface structure of cohesive sentences?

We employed rhetorical and grammatical knowledge embedded in rules to cope with those problems. For each rhetorical schema, a rule is formulated whose condition part evaluates how well the set of assertions is stratified by the corresponding schema. For example, we regard a set of assertions as well stratified by a path of regions if all assertions pertain to the same attribute and time period and there exists a one-to-one correspondence between the regions of the path and the regions of the assertions, or a set of assertions is well stratified chronologically if all assertions pertain to the same region and there is no overlap between their time periods.

Since the conditions of the “temporal progression” and “spatial progression” rhetorical schemas as described above are too rigid and so they are rarely satisfied by the assertions produced by the scanner, we loosened them by allowing partial instead of full coincidence between the regions. The grade of similarity between two regions  $r_1$  and  $r_2$  is defined by the formula:

$$d(r_1, r_2) = 1 - \frac{|carrier(r_1) \cap carrier(r_2)|}{|carrier(r_1) \cup carrier(r_2)|}$$

and they are considered coincident if  $d(r_1, r_2) > 0.7$ . Thus the set of three assertions concerning the regions “the lowlands of West Bulgaria”, “Central Bulgaria” and “the Black sea coast” can be successfully mapped out along the path “West Bulgaria”, “Central Bulgaria”, “East Bulgaria”.

There are certain priorities among the rhetorical schemas. The schemas “presentation by attributes” and “coupling related values” have priority over the others; the schema “from a summary to details” has priority over the temporal and spatial progressions, etc. A rule with a higher priority than another is applied first, and only if it fails, then the second rule is tried.

The action part of the chosen rule breaks the set of assertions into a chain of chunks. The link between two chunks represents the conversational move that takes place when the discourse passes from the source to the target chunk. Then each chunk is broken down into a subchain, and so on until a hierarchical discourse structure is obtained, at the terminal nodes of which are the assertions of the initial set (cf. Figure 3b).

The conversion into a surface structure proceeds by applying the rules embedding the grammatical knowledge. They analyze the discourse structure by means of patterns. The matching of a pattern with a discourse substructure leads to a transformation of the latter into the surface structure of a sentence, clause or phrase and

its bounding to the surface structure of the text. In addition to indicators of the elements of the discourse structure, the patterns may contain conditions on the contents of the assertions and on the types of the preceding sentences. Figure 3c shows a portion of the surface structure realizing the discourse structure in Figure 3b.

The following features characterize the creation of the surface structure of a text:

- A good deal of sentences are constructed on the basis of impersonal verb phrases typical for weather description.
- The tense of the verbs is determined by the type of the report. If it is a forecast, then future tense is adopted, otherwise — past tense.
- Where appropriate, function words are inserted that indicate the type of conversational move (e.g. “but” for contrast, “also” for addition, etc.).
- Certain elements of the context (the region and/or the time period) are omitted, if implied from the preceding text or the external context, or are replaced by adverbial or relative adverbial phrases (“there, then, where, when”), if the corresponding element is implied but the grammatical structure requires such a phrase.
- The precision rate of the assertions, if lower than *high*, is indicated by inserting proper modifiers, such as “at many places of ...”, “mostly” etc., which warn the reader to accept the information with some reservations.
- The word order of the sentences is selected in such a way that the elements constituting the topics and the focuses of the consecutive sentences alternate (Hajičová, 1987). For example, if the region is the focus of one sentence, it is good to generate the next sentence with the region being its topic. Thus the text will flow rhythmically and at a proper pace.

## 6.3 Supplementing a map with a text

A technique of converting a set of assertions into a weather map has been described in (Kerpedjiev, 1990). Here we concisely recall the technique and extend it to allow the generation of text supplements.

The conversion of a set of assertion into a map is based on the existence of a set of visual objects (pictographs) and two functions —  $f$  and  $g$ ;  $f$  assigns a pictograph to each weather value; and  $g$ , for each region, determines the positions where the icons related to a given attribute should be put in. The algorithm of conversion scans the selected set of assertions and generates the map plan by replacing each assertion  $(w, v, r, t, p)$  with a list of statements  $\{(q, x_i, y_i)\}_{i=1..n}$ , where  $q = f(v)$ ,  $\{(x_i, y_i)\}_{i=1..n} = g(w, r)$ . A statement  $(q, x, y)$  of the map plan drives the formatter to place icon  $q$  in the position with coordinates  $(x, y)$ .

Some problems arise with this technique. Firstly, two pictographs may occur to overlap and distort each other.

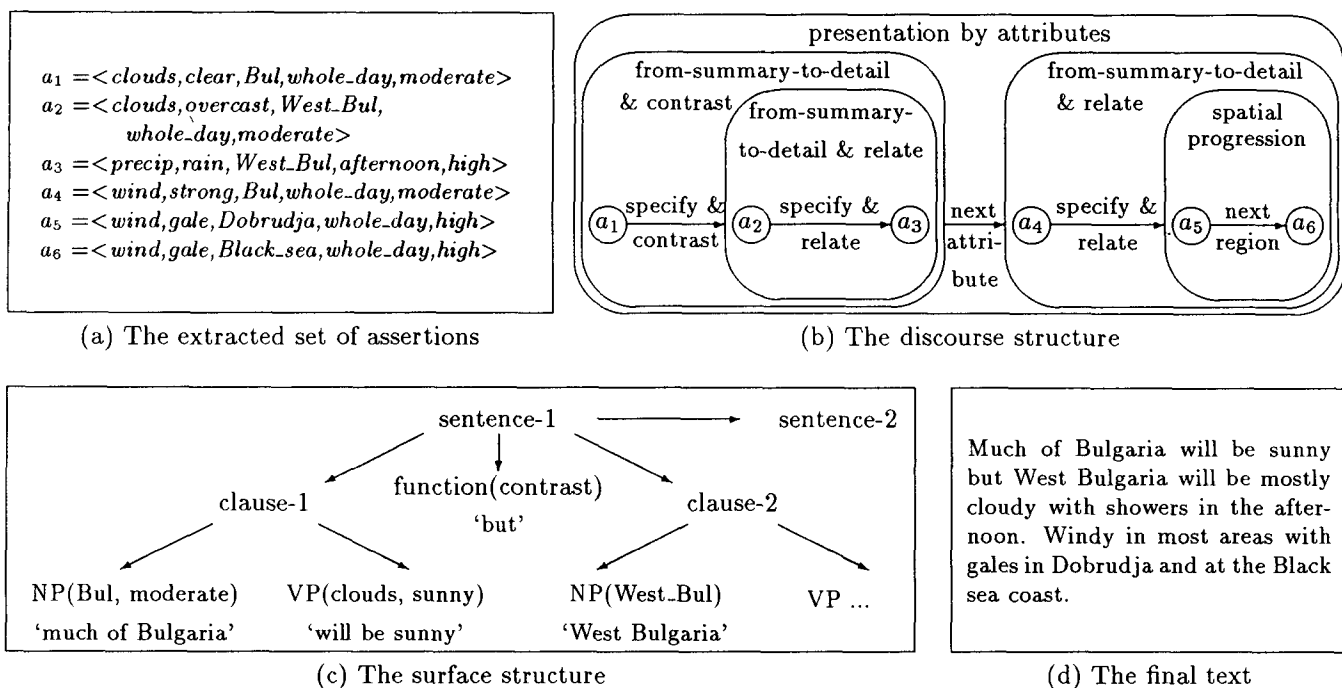


Figure 3: The successive steps in the conversion of a set of assertions into a text

Secondly, certain geometrical relations between the icons of related assertions should be ensured. Thirdly, the information concerning the time period and the precision rate is completely ignored by the conversion technique. The first two problems are resolved by carefully designing the function  $g$ . Information about a time period and/or precision rate, when necessary, is provided by a verbal comment as described below.

Suppose that the assertion  $(w, v, r, t, p)$  is visualized on a map representing the weather situation in context  $(r', t')$ . The corresponding map plan will represent the assertion correctly only if  $t' \subset t$  and  $p = high$ . If any of these relations is violated then the system has to impart this information to the user. We call it *residual information* or a *residue*. It consists of a *reference part* determined by the weather value (the user associates it with the corresponding icon) and possibly by the region (if  $r \cap r' \neq r'$ ; the user should identify it with the locations where the icons are situated) and a *specification part* determined by the elements  $t' \cap t$  (if different from  $t'$ ) and  $p$  (if not *high*). The grammar for rendering a residue as a sentence is available in the grammatical knowledge base. Furthermore, the residual information can be formulated either as a characteristics of the reference part (e.g. "the rain in East Bulgaria will be scattered") or as a process (e.g. "the rain in East Bulgaria will stop by noon").

The planning of a text supplement may face the following problem. Consider the example in section 6.1. If another assertion occurs about a fog in North Bulgaria

at noon, then the residual information must be adjusted to the following message:

"The fog in North Bulgaria will clear by the afternoon"

In order to avoid any inconsistency in the generated message, the system collects the residues relating to the same weather attribute and region, unifies their specification parts whereby some content portions may partially neutralize each other, and then generates the message.

## 7 Generation and composition

The text generator converts the surface structure into a text by making use of the phrasal lexicon. For some of the terminals of the surface structure it provides ready-made strings, while other terminals have to be refined further. For example, if the regions "South Bulgaria" and "East Bulgaria" occur in the same phrase, the generator will combine them into the contracted form "South and East Bulgaria". Whenever possible, the text generator takes care of the diversity of lexical forms by providing different phrases for terminal nodes of the same type (compare "much of Bulgaria" and "most areas" in the example in Figure 3).

The map generator interprets the map plan and converts it into an image by rendering the successive state elements (cf. section 6.3). The maps contours as well as the various icons are prepared in advance and stored in the visual library. The text supplements, if any, are prepared in the same way as ordinary text portions and attached to the bottom of the map.

The formatter composes the document from components delivered by the generators. An enumeration structure is created from the constituting text portions by inserting the lexical representation of each external context as a heading of the corresponding portion. A table is composed of phrases produced by the text generator from a series of values extracted by the scanner.

## 8 Conclusion

The system presented in this paper interprets a weather dataset and generates a multimodal report. The following features distinguish the method from similar projects carried out elsewhere:

- Various weather documents have been analyzed to determine their structures and to define a language that allows the user to specify the generated product with respect to mode, content, style, length and precision of the information.
- A knowledge-based technique for selection of the discourse structure of the generated document has been devised on the basis of typical rhetorical schemas and relations defined in the TKB.
- The system controls the precision of the assertions extracted. Information with a low precision rate is rendered as a sentence with lexical indicators of imprecision (e.g. *in some portions, possibly*), which warn the user to accept the corresponding assertions with a certain degree of reservation.
- The cartographical presentation, though superior in many respects to the textual presentation, still suffers from the lack of proper means of expression for certain elements such as time and precision of the information. Therefore, a map may be supplemented with a concise verbal comment on the underspecified elements. Thus the two modalities, NL and graphics, complementing each other offer a highly expressive and efficient weather report.

The system has been implemented in Pascal on an IBM PC. The TKB is filled in with models of the general-purpose short-range weather forecasts for Bulgaria. The grammatical knowledge base and the lexicon contain styles corresponding to subsets of Bulgarian and English. Experiments have been performed with datasets compiled at the National Weather Service in Sofia.

In order to make the system practical it has to be coupled with a layout manager. Thus the user will be able to specify the arrangement of the different units in the plane.

Another point of future work is the enrichment of the weather model with attributes summarizing the weather over a longer period (say five days) taking into account climatic data. Thus the system will be able to extract and render more interesting facts about the weather.

A promising research area which may contribute to the further development of the system is user modeling.

Experimentation in this area can be combined with the various applications of weather reports.

## References

- J. Algayer, K. Harbusch, A. Kobsa, C. Reddig, N. Reithinger, and D. Schmauks. XTRA: A Natural-Language Access System to Expert Systems. *International Journal of Man-Machine Studies*, 31:161-195, 1989.
- E. André, G. Herzog and Th. Rist. On the Simultaneous Interpretation of Real World Image Sequences and their Natural Language Description: The System SOCCER. In *Proc. 8th ECAI*, pages 449-454, Munich, 1988.
- L. Bourbeau, D. Carcagno, E. Goldberg, R. Kittredge, A. Polguere. Synthesizing Weather Forecasts in an Operational Environment. In *Proc. 13th Int. Conf. COLING*, vol.3, pages 318-320, Helsinki, August 1990.
- A. Davey. *Discourse Production*. Edinburgh University Press, Edinburgh, UK, 1979.
- S. Feiner and K. McKeown. Coordinating Text and Graphics in Explanation Generation. In *Proc. 8th National Conf. of AAAI*, pages 447-454, 1990.
- E. Hajičová. Focussing - a Meeting Point of Linguistics and Artificial Intelligence. In *Artificial Intelligence II: Methodology, Systems, Applications*, eds. Ph. Jorrand and V. Sgurev, North-Holland, Amsterdam, pages 311-321, 1987.
- S. Kerpedjiev. Transformation of Weather Forecasts from Textual to Cartographic Form. *Computer Physics Communications.*, 61:246-256, 1990.
- S. Kerpedjiev, D. Ivanov. Automatic verification of general-purpose short-range weather forecasts. Internal Report, Institute of Mathematics, 1991.
- S. Kerpedjiev and V. Noncheva. Intelligent Processing of Weather Forecasts. In *Proc. 13th Int. Conf. COLING*, vol.3, pages 379-381, August 1990, Helsinki.
- K. Kukich. Design of a Knowledge-Based Report Generator. In *Proc. 21st Annual Meeting of ACL*, 1983.
- D. McDonald. Natural Language Generation. In *Encyclopedia of Artificial Intelligence*, ed. S.C. Shapiro, pages 642-655, 1987.
- O. Stock. Natural Language and Exploration of an Information Space: The AIFresco Interactive System. In *Proc. IJCAI*, Australia, August 1991.
- W. Wahlster, E. André, W. Graf, and Th. Rist. Designing Illustrated Texts: How Language Production is Influenced by Graphics Generation. In *Proc. 5th Conf. European Chapter of ACL.*, Berlin, Germany, April 1991.