METEOR: An Artificial Intelligence System for Convective Storm Forecasting

RENÉE ELIO

Department of Computing Science, University of Alberta, Canada

JOHANNES DE HAAN

Computing Department, Alberta Research Council, Canada

G. S. STRONG

Atmospheric Sciences Department, Alberta Research Council, Canada

(Manuscript received 15 May 1986, in final form 12 December 1986)

ABSTRACT

An experienced forecaster can use several different types of knowledge in forecasting. First, there is his theoretical understanding of meteorology, which is well entrenched in current numerical models. A second type is his "local knowledge," gained over years of experience, of how weather is likely to form in his forecast area. This kind of local familiarity is not easily captured with traditional numeric techniques, but might provide additional insights for prediction that someone unfamiliar with the area might not have. A third type of knowledge is how to interpret forecast tools already in use. This might include knowledge of the tool's limitations and how it works in a particular locale. Capturing these types of knowledge is important in building computing systems that can serve as intelligent consultants to forecasters. This paper describes a prototype system, called METEOR, that incorporates all these types of knowledge to predict the location, severity, and motion of convective storms in Alberta; METEOR interprets contoured maps of a synoptic-based instability index and of surface equivalent potential temperature. It also gathers additional information about a variety of ongoing weather conditions from three portions of surface aviation reports: the cloud cover section, the obstructions visibility, and the observations provided in the "remarks" section. Interpreting remarks made by human observers, while useful to a forecaster experienced with local weather conditions, can be too time consuming for people to do in real-time and too complex for traditional computing methods to handle. However, METEOR interprets these remarks and keeps track of where various weather activities are occurring and how they are changing over time. At present, METEOR's final forecast is a prediction of likely areas of storm initiation, direction of motion, and intensity, plus summaries of current conditions and their implications for storm development.

1. Introduction

The Alberta Research Council's Atmospheric Sciences Department has conducted a research program on weather modification, particularly of hail suppression, for several years. The ability to predict the location, severity, and motion of hailstorms is important to the cloud seeding experiments that are conducted during the summer. Meteorologists assigned to forecasting for the research program, called the Alberta Hail Project, have considerable knowledge about how severe storms form in Alberta. In this paper, we describe an artificial intelligence system called METEOR, which uses the meteorologist's heuristics, strategies, and statistical tools to forecast severe hailstorms in Alberta.

Although it is not a "general" forecasting system, METEOR's design and approach features have wide applicability. It illustrates the combination of many types and sources of meteorological knowledge and data within an artificial intelligence framework. This includes "understanding" surface aviation reports of current conditions and their implications for storm development. From these surface aviation reports, METEOR abstracts information on cloud cover and interprets the observer's verbal remarks describing ongoing conditions. METEOR also uses a particular statistical model developed especially for storm prediction in the foothills east of the Rockies (Strong & Wilson, 1983). It applies heuristics and meteorlogical knowledge to interpret the statistical model and to determine the implications of current weather conditions. In order to do this, METEOR keeps track of what the local conditions are, how they are interacting, how they are changing over time *and what those changes mean.*

First, we will describe the METEOR system, concentrating on the information and knowledge it uses

TABLE 1. Definitions of Convective Day Categories (CDC's) for the AHP operations area, with qualifying notes (after Strong, 1979).

CDC	Definition	Qualifying notes
+5	LARGER THAN GOLFBALL (>52 mm diameter) maximum size reported	1) All HAIL categories are based on a climatological day 0600–1800 MDT (1200–1200 UTC). For forecasting and verification purposes, the day begins at briefing time, 1700 UTC.
+4	GOLFBALL (33–52 mm) maximum size reported	
+3	WALNUT (21–32 mm) maximum size reported	2) At least two hail reports are required to qualify in each HAIL category.
+2	GRAPE (13–20 mm) maximum size reported	
+1	PEA (5–12 mm) or SHOT (≤4 mm) maximum size reported	3) Showers must be accompanied by cumulus congestus clouds (TCU) to qualify as CDC = −1 or 0.
0	WIDESPREAD SHOWERS at three or more of the synoptic stations BA, YC, QF, RM, CT, and EG, or the occurrence of a THUNDERSHOWER at or near either station after 1700 UTC	4) Three reports of showers from volunteer reports (telephone surveys) or forestry towers are considered the equivalent of one synoptic station reporting showers.
−1	SCATTERED SHOWERS (at one or two of the six synoptic stations after 1700 UTC), with NO HAIL or THUNDERSHOWER	5) The surface synoptic reports used include all "specials" and also make use of "remarks"
−2	ACC, TCU, CB reported by two or more of the six synoptic stations, but no showers after 1700 UTC	6) Additional restrictions are placed on the use of remarks from those synoptic stations which lie just outside the AHP project area (YC, BA, EG, and CT); eg., "SHWRS W CT" are counted but "SHWRS E CT" are ignored
−3	NO DEEP CONVECTION (no ACC, TCU, CB)	

to mimic the forecasting procedure of our expert meteorologist. Then, we will discuss METEOR as an artificial intelligence system, emphasizing the ways in which it is qualitatively different from algorithmic or statistical approaches to prediction. These differences will become more apparent as we present some features of METEOR's design and the artificial intelligence techniques for representing meteorological knowledge and for reasoning and inference. Finally, we will conclude with some observations on designing and implementing intelligent consultants for meteorological applications.

2. Forecasting hail storms in Alberta

METEOR's task is to forecast the initiation region, intensity, and motion of severe hail storms in Alberta. To do this, METEOR uses two sources of information. The first is a statistical model, called the Synoptic Index of Convection (Sc4). The second is information about current cloud conditions and other weather activity, obtained from surface aviation reports gathered from about 100 weather stations throughout Alberta and adjacent areas. Although human forecasters have access to much more information than these two data sources, the ability to intelligently interpret the Sc4 model and the surface aviation reports was both interesting from an artificial intelligence standpoint and useful from a forecasting standpoint. This section will describe each type of information and how it is analyzed to produce a forecast.

a. Interpreting the Sc4

Over the past several years, one of us (G.S.) has developed a statistical predictor, the Sc4, to assist in predicting areas of storm initiation and severity. Details on the Sc4 development and use are explained by Strong and Wilson (1983) and Strong (1986). We will summarize the main points of the model here. The Sc4 is computed from two synoptic scale upper air predictors and two subsynoptic-scale surface instability predictors. These variables are 1) the 24-h, 12–1200 Universal Coordinated Time (UTC), forecast 700 mb temperature change; 2) the 24-h, 12–1200 UTC, forecast 1000–500 mb thickness change; 3) the 1500 UTC EGDEX forecast[1]; and 4) the 1500 UTC surface parcel lifted index. Each of these predictors depends on the output from the Canadian Numerical Weather Prediction Spectral Model. The predictors are each linearly correlated to a discrete-valued predictand called the Convective Day Category (CDC), defined in Table 1. Basically, the CDC represents a scale of maximum storm severity within the AHP operations area. Its values range from −3 to +5. Higher numbers correspond to severe hail-producing storms. The negative/positive distinction corresponds to a no-hail/hail distinction, with −3 to 0 representing an increase in convective activity up to, but not including, the occurrence of hail. The occurrence of walnut-sized or larger hail (CDC ≥ +3) has been designated as a "severe convective storm." This threshold value is considered a "severe

[1] The EGDEX is defined as the difference between the wet-bulb potential temperature of the surface parcel, corrected for diurnal changes, and the forecast 2100 UTC 500 mb wet-bulb temperature (assuming saturation). It was developed by S. M. Checkwitch, 1972: "A short-term forecast technique for summer precipitation forecasting in the Edmonton area", an unpublished internal manuscript by Alberta Weather Center, AES, Edmonton, Alberta, 32 pp.

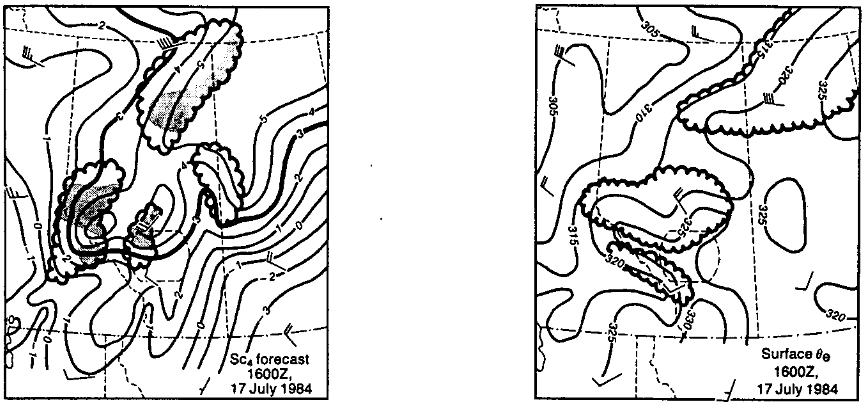

FIG. 1. (a) Sc4 forecast and (b) surface theta-E analysis at 1600 UTC 17 July 1984. In each case, scalloped regions represent predicted regions of convection. In (a) shaded regions are the intersection of scalloped regions from (a) (above the +3 "threat" threshold) and (b). These shaded regions represent forecast convective complex initiation regions.

hailstorm" in Alberta and introduces the possibility of tornados.

An Sc4 value is computed for each station in the forecast area and each value is valid for a radius of 130 km. Each value is normalized to the range of the CDC values so that, in theory, it can be interpreted in the same way. The final output of the Sc4 analysis is provided in the form of a contoured map. An example Sc4 map is shown in Fig. 1a. The Sc4, by itself, does not indicate where convective complexes will form, their size or intensity, or how they will move. It only provides an indication of maximum expected convective intensity for the day.

To interpret the Sc4 accurately, the forecaster applies his specific experience with the Sc4 model and his general experience with convective complexes. The strategy for interpreting the model using surface moisture fields and advection principles is given by Strong and Wilson (1983) and more recently detailed by Strong (1986). This strategy is summarized as follows. By considering the advection of complexes by the mean wind, it has been found that these systems usually form on the upwind side (at 700 mb) of the strongest Sc4 gradients, i.e., within regions of strong negative Sc4 advection, indicated by the gray areas in Fig. 1a. "Severe" storm candidates are those lying within the +3 Sc4 contour. These regions therefore give a first guess as to the probable size of the convective complex. As a first guess of storm motion, storms will tend to propagate toward lines of maximum Sc4, based simply on the higher convective potential there.

Operational experience has shown that the forecast regions of convection and their motion can be refined using the surface moisture field. An example is given in Fig. 1b. This approach follows from the fact that the propagation and intensity of a storm is largely controlled by its ability to intercept the greatest amount of moisture. To incorporate surface moisture in the technique, the forecaster looks for regions of strong negative advection of surface theta-E or mixing ratio,[2] using the 850 mb flow. These regions, again indicated by scalloped areas in Fig. 1b, define slightly differently different regions of maximum convective potential. The intersection of the regions of negative Sc4 advection and negative theta-E advection (scalloped areas in Figs. 1a and 1b) then narrows down both the choice and the size of regions of highest potential. Once individual storms have formed, their short-term (1–3 h) propagation tends to be towards lines of maximum theta-E. For longer periods, and to estimate when and where storms might start to decay, propagation towards the line of maximum Sc4 is preferred.

The Sc4 model has undergone considerable evaluation and is still being improved today. Like most statistical models, it *implicitly* embodies expert knowledge about how measurable meteorological conditions can be quantitatively related to infer weather dynamics. To design an artificial intelligence system that interprets the model as an expert might, there are several important pieces of knowledge that the system must apply. First, the model must be interpreted using knowledge about surface moisture, low- and midlevel winds and, in most cases, knowledge of local topography. Second, a meteorologist familiar with this model knows the factors that can "fool" its predictor variables. For example, ongoing or recent precipitation often inflate surface moisture predictors. Yet without knowledge about current weather conditions, a statistical model can only treat all surface moisture measurements the same. This knowledge of the index's limitations, assumptions, and simplifications is an important part of the forecaster's expertise. Finally, most experienced forecasters have rules-of-thumb they use to make qual-

[2] The forecaster prefers to use mixing ratio (for forecasts made before noon,) and theta-E for afternoon forecasts when surface temperatures are not changing so rapidly.

itative judgments about the likelihood and nature of severe weather. The significance of current weather conditions, evaluated with this qualitative local expertise, could serve as a framework for interpreting a statistical model like the Sc4. The role of this qualitative knowledge and data is described in subsection 2b.

b. Interpreting surface aviation reports

Some information about current weather conditions comes from surface aviation reports made at weather stations. An example surface aviation report is given in Fig. 2. Two portions of this report provide the qualitative information that is important to an experienced forecaster: cloud cover, marked a in the figure; and the remarks section, marked b. The cloud cover information indicates which types of clouds are observed and what observable portion of the sky they cover. In this example, the first cloud layer the observer sees, cumulonimbus (CB), covers seven-tenths of the sky. A second cloud layer of altocumulus (AC) covers the remaining three-tenths of the sky that the observer can see. Other data in the report describe the state of this cloud cover. For example, the CB clouds are "broken" (BKN). If an earlier report from this station had indicated that this layer was overcast, the *shift* to broken might be significant for storm forecasting.

The remarks in Fig. 2 reflect the human observations that are not easily captured in coded format. What these remarks report is: "There is lightning in the clouds, from cloud to cloud, and from cloud to ground in the southwest direction [relative to the station]. Showers are heavier in the north and west directions. The pressure [at the station] is unsteady."

Both cloud cover information and the informal remarks in the surface aviation reports have gone unanalyzed in the past because there is no simple way for traditional methods to "understand" this information and codify it in a form usable by statistical models. However, these qualitative conditions are meaningful to an experienced meteorologist familiar with this particular forecast area. For example, our expert has a number of informal heuristics based on cloud formations. An example is "If there is a widespread deck of cirrostratus cloud, then it's not likely a storm will occur." Like most heuristics, this rule comes with a set of qualifications and hedges. Other observations are more direct evidence of the kinds of dynamic processes that the Sc4 model is indirectly trying to measure. They can also alert the expert to situations that can fool the statistical model's predictor variables into yielding an inaccurate intensity estimate. The case of ongoing precipitation, noted earlier, can be detected by a system that can read the surface aviation reports, discover where it is raining, and evaluate the impact of this on the accuracy of the model.

2006 YEG SA 0600 E50 BKN 90 OVC 15+TRW–
148/16/13/3010/996/CB7AC3
(a)
LTGIC–CC–CG SW QUAD. SHWRS
HVIER N,W PRES UNSTDY 3012
(b)

FIG. 2. An example of a surface aviation report.

3. METEOR's forecasts

METEOR is a noninteractive system. It does not rely on a person to describe or interpret weather conditions or the index's contour maps in order to apply what it knows about severe storms in Alberta. It is completely automatic, accessing other computers to obtain the same contour maps and surface aviation reports that the human forecaster uses. METEOR takes about 50 min to analyze this information and produce its predictions. Our main interest to date has been to design a framework for representing knowledge for forecasting and the exact form of the output has been secondary.

Currently, METEOR's forecasts consists of a set of maps with accompanying remarks about their interpretation. One map shows storm initiation regions, which METEOR determines by analyzing Sc4 and theta-E maps. An example is shown in Fig. 3. The remarks that accompany this map are shown in the same figure. Each storm area outlined on the map has an identification number (e.g., STORM-05) so that the forecaster can find the associated remarks. The remarks include METEOR's predictions for storm intensity and direction. The intensity prediction for a particular storm is initially based on the level of the closest Sc4 maximum. METEOR examines the four predictor variables that constitute the Sc4 model, checking to see if each one is "in line with" the overall Sc4 rating. The notion is that, like the expert, METEOR does not treat the model as a "black box." METEOR uses qualitative heuristics to evaluate each predictor variable, rating each one as positive, negative, or neutral evidence for the initial intensity rating. Thus, if the two instability variables (SPLI and EGDEX) are both neutral, while the overall intensity rating is high (e.g., 4 or 5), then METEOR may downplay the final intensity.

There are a few important points to note about this. First, METEOR's analysis of these predictor variables is based on its ability to understand the spatial locations of these measurements with respect to the individual storm it is considering. Second, METEOR incorporates the forecaster's knowledge of the index's limitations and the factors that can "fool" the predictor variables. An example of this is shown in Fig. 3. In this case, METEOR has found areas of precipitation and alerts the forecaster to the fact that this activity may be inflating theta-E ratings.

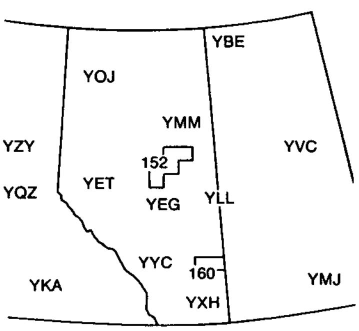

Map 1: Storm Initiation Regions at 1215

a) Date and time for this forecast is 1215

Storm Initiation Region: STORM-152
Predicted direction: SOUTH-EAST
Sc4 CDC rating for this region: 5.0
Evidence in support of this intensity level ...

SPLI	WEAK-NEGATIVE
EGDEX	POSITIVE
D24TH	POSITIVE
D24T7	POSITIVE
THETA.E	NEUTRAL

Final predicted intensity: 4.0

Storm Initiation Region: STORM-160
Predicted direction: SOUTH-EAST
Sc4 CDC rating for this region: 5.0
Evidence in support of this intensity level ...

SPLI	POSITIVE
EGDEX	POSITIVE
D24TH	POSITIVE
D24T7	POSITIVE
THETA.E	POSITIVE

Final predicted intensity: 5.0

b) SC4 MODEL COMMENTS

For this forecast, theta.e gradients were intersected with Sc4 gradients to find storm initiation regions.

It is possible that precipitation occurring at the time the Sc4 model was run caused theta.e values to be inflated. This suggests that the above theta.e interpretations should be tempered.

The 850 winds used for this forecast were ((SPEED (YYE 18)) ((DIRECTION (YYE 329)) and were based on data taken at time 1212.

The 700 winds used for this forecast were ((SPEED (YQF 17) (YYE 18)) ((DIRECTION (YQF 342) (YYE 315)) and were based on data taken at time 1212.

FIG. 3. METEOR's predictions of storm initiation areas and interpretation of the Sc4 index.

METEOR's knowledge about the location of precipitation, or any other current weather activity, is derived from its analysis of surface aviation reports. In subsection 5b, we describe the role of artificial intelligence techniques in analyzing these reports and inferring general regions of weather activities from isolated reports. METEOR has a set of particular observations and long-term events that it routinely looks for. If it finds these observations or events, it presents their potential significance for storm development. This analysis constitutes the second part of METEOR's output and also comes in the form of maps. An example of a portion of this analysis is shown in Fig. 4. The knowledge METEOR has to interpret the implications of these conditions is based on the meteorologist's local knowledge and heuristics about how storms form in Alberta. For example, one event METEOR looks for is labeled "midlevel stabilization." To recognize that this event is occurring, METEOR examines surface aviation reports (both the cloud cover and remarks section) for mention of, say, altostratus clouds. If METEOR notes that, *over time,* a large midlevel cloud cover area is increasing in size and changing from altocumulus type to altostratus, or from a broken condition to overcast, it infers that this indicates a stabilization of conditions. It outlines this area on a map and reports the nature of the change and its implication to the forecaster. On the other hand, it also recognizes the occurrence of altocumulus castellanus clouds as evidence of existing midlevel instability. Keeping track of where such activities are in space, and how particular features of these activities are changing over time, is critical for METEOR's ability to forecast.

4. A knowledge-based system approach to forecasting

METEOR differs from traditional statistical or algorithmic approaches to problem solving in its representation and use of nonnumeric task-specific knowledge to produce a forecast. This knowledge includes not only the formal training a meteorologist has, but more importantly, the strategies, heuristics, and conceptual understanding gleaned from experience. These qualitatively different types of knowledge are not easily represented and manipulated as algorithms, decision tables, or other numeric forms. For example, it is difficult to understand a surface aviation report like the one given in Fig. 2, particularly when its implications for the meteorologist have a general form like "A and B tend to suggest C most of the time." For this reason, systems like METEOR are called "knowledge based" because they derive their problem-solving power by using artificial intelligence formalisms for representing and manipulating the types of knowledge that distinguish a human expert from an inexperienced practitioner on a particular task.

In many scientific domains, human experts have developed sophisticated numerical models and software that, in a very real sense, represent their knowledge about a problem or domain. How is this different from what we have just described as knowledge-based systems? In numerical models, some of the expert's knowledge is *indirectly* represented in a set of equations. While numerical formalisms and statistical models are one way to represent relationships in a well-understood domain, such models often require assumptions, simplifications, and sometimes gross approximations that may not be appropriate. In addition, an expert often uses hunches, rules of thumb, and considerations that cannot easily be represented as a set of equations. While general statistical models of weather development can be very useful tools, much of the local forecaster's knowledge can become diluted in a statistical model. This is particularly true when a number of data sources are involved. In forecasting, those sources may include

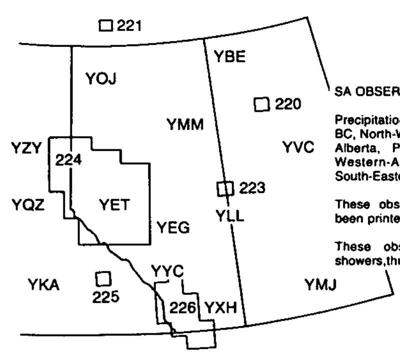

SA OBSERVATIONS

Precipitation sighted in North-BC, Parkland-BC, North-Western-Alberta, Parkland-Western-Alberta, Parkland-Eastern-Alberta, Central-Western-Alberta, Central-Eastern-Alberta, South-Eastern-Alberta, and North-Sask.

These observations, sighted at 1215, have been printed on Map 4.

These observations include sightings of showers,thundershowers and virga.

Map 4: Precipitation at 1215

FIG. 4. A portion of METEOR's analysis of current precipitation areas, located by analyzing human remarks in the surface aviation reports.

weather maps, surface aviation reports, numerical forecasting models, knowledge about what happened an hour ago, how conditions are changing over time, and finally, experience in combining them all to make a forecast. Thus, the forecaster engages in a problem-solving *process* that involves the integration of many diverse types of knowledge and data. A knowledge-based system tries to represent *directly* what a forecaster knows (both the formal training as well as "tricks-of-the-trade") and how it is used during each stage of the forecasting process.

A critical aspect of building a knowledge-based system is eliciting the knowledge an expert uses to solve a problem. Our expert for this problem (G.S.) has had 20 yr of forecasting experience, with the last 10 yr focused on convective storms in Alberta. As a first step in learning about forecasting and extracting some of the expert's heuristics, one of us (R.E.) attended a 2-week training session given by the expert for new Alberta Hail Project forecasters. However, the real insights into how the meteorologist interpreted local conditions as well as the Sc4 model came from extensive interviews and discussions of case studies that took place over about an 18-month period. As is often the case with knowledge-based systems, METEOR's deficiencies in what it "knows" and how it uses that knowledge became apparent during its field test (described below). These problems are due in part to the system's developers, who sometimes misunderstand the exact nature of a particular heuristic or strategy, and in part to the expert forgetting to mention some critical qualification or other piece of knowledge. Thus, the creation of a knowledge-based system is an iterative process of extracting the knowledge, representing it with the appropriate artificial intelligence formalisms, and having the expert critique the system's performance on particular case studies.

Forecasting represents some challenging aspects for knowledge-based systems. One is the integration of existing statistical models and utilization of radar, satellite and data collection systems. Such methods and tools, like the Sc4, are often an integral part of a forecaster's reasoning and prediction. Indeed, what distinguishes an expert forecaster from a less experienced forecaster may be the knowledge that allows him or her to interpret "general" statistical models with greater accuracy, understanding how idiosyncratic features of the forecast area interact with a model's assumptions or limitations. This is precisely the type of knowledge that we tried to capture by giving METEOR the task of interpreting the Sc4 model.

Another difficult problem for knowledge-based forecasting systems is the need to reason about activities in spatial and temporal dimensions. To forecast severe convective storms using the Sc4 model and surface aviation reports, METEOR not only requires meteorological knowledge, but the ability to recognize patterns on contoured maps and to make inferences about changes over time. The need for spatial and temporal reasoning make forecasting a particularly difficult task for knowledge-based systems. These abilities, which humans do rather handily, are considerably more painstaking for knowledge-based systems. Section 5 outlines some features of METEOR's knowledge base that address these problems.

5. METEOR's knowledge-base

To interpret the Sc4, METEOR recognizes and manipulates features on contour maps in much the same way that the human forecaster does: It locates areas of Sc4 maxima and strong contour gradient, and uses knowledge of 700 mb winds to select storm initiation regions. METEOR consults contour maps of surface moisture and considers the relative surface moisture in areas immediately surrounding the storm initiation areas. Because this analysis is based on spatial reasoning, METEOR has an internal representation of the "world" onto which it can project these contour maps and identify spatial relationships among map features. However, it also locates ongoing weather conditions in this world so that it can make inferences about their relative spatial proximity and possible interactions. Finally, METEOR keeps track of the conditions from hour to hour, so that it recognizes what the "significant" changes are. Thus, it has the ability to relate meteorological activities temporally, as well as spatially. In this section, we give a general description of the kinds of artificial intelligence formalisms we used and developed to represent the knowledge METEOR needed to perform these tasks. Details on METEOR's representation of meteorological knowledge, how it makes spatial and temporal inferences, and how it operates as a system, are presented more fully from an artificial intelligence perspective in Elio and de Haan (1985).

a. *Knowledge representation and organization*

Figure 5 presents a schematic overview of how METEOR's knowledge is organized. METEOR's knowledge is represented using several artificial intelligence formalisms and methods. The first of these is production systems. These systems have three main parts: a set of IF-THEN rules, a working database that reflects the current information about the world and the problem at hand, and an interpreter for deciding when a particular IF-THEN rule matches information in the working database. The rules are independent, modular constructs that represent knowledge of the form "IF a particular situation or set of features is occurring, THEN take the following action, or make the following conclusion." Unlike IF-THEN statements in a conventional programming language, production rules are independent chunks of knowledge. The application of any particular piece of knowledge is dynamically determined by the current features of the problem, not by a preprogrammed order. A rules's action typically changes information in the working database, but it can also access information elsewhere in the system, or start processes on other computers. Rules can represent goals that must be achieved, or heuristics. The following are simplified renditions of some of METEOR's rules:

IF I am beginning a forecast
THEN set goals to check for new Sc4 maps and a new set of surface aviation reports.

IF there is a region of cirrostratus cloud
THEN this suggests the dynamics are not favorable for severe storm development

IF the Sc4 model is predicting a very severe storm but both instability factors are weakly positive or neutral with respect to this prediction
THEN downplay the model's severity rating by 2 categories

The first rule exemplifies how rules can "drive" METEOR through the various subtasks associated with the process of making a forecast. Other rules would respond to the availability of the Sc4 contour maps to initiate map interpretation. Whether or not the Sc4 model was run, METEOR would turn its attention to understanding the latest surface aviation reports. This involves parsing the cloud cover and remarks sections, and then looking for particular regions of meteorological activities (e.g., an area of lightning reports, reports of virga or towering cumulus clouds.) The second rule illustrates a heuristic that might be included in METEOR's final report. The third rule is one of many heuristics that the expert uses to interpret the Sc4 model by considering whether each of the predictor variables is giving positive, neutral, or negative evidence for the index's final overall rating. Basically, production systems are a useful formalism for representing heuristics, moving the system through a well-defined series of tasks and subtasks, and implementing reasoning strategies.

Not all knowledge is best represented in this IF-THEN formalism. This is particularly true of taxonomic relationships, definitions, and categories. METEOR's knowledge about particular meteorological conditions, clouds, and long-term changes (like mid-level stabilization) is represented by another formalism called frames. Frames are essentially packets of associated knowledge about concepts and their features. Special mechanisms are designed to reason about these concepts and draw inferences about their associations to each other. Thus, METEOR knows that cirrostratus clouds are a kind of high-level cloud, that they are composed of ice crystals, and that they usually signal conditions that are not conducive to storm development. This kind of knowledge is represented in a frame for cirrostratus clouds, which in turn is associated with a frame for high-level clouds. This association allows a concept like cirrostratus clouds to share all the knowledge that is true about high-level clouds, which in turn can share knowledge that is true of clouds in general. On the other hand, facts and rules-of-thumb that are true of only cirrostratus clouds can override more general knowledge that is associated with its superclasses. Frames are one formalism that easily captures this type of taxonomic relationship, which can be exploited for efficiently representing knowledge and for making certain kinds of inferences.

Both the production system and the frame system rely on an internal representation of the forecast area that allows METEOR to locate activities in space. All the contoured maps as well as the observations and events inferred from an analysis of surface aviation reports are "projected" onto this internal representation, which is essentially a large grid. This grid of map elements, labeled "mapel map" in Fig. 5, corresponds to approximately 2.3×10^6 km^2. The size of each mapel, which defines the resolution of the world, is a parameter that can be easily modified. All meteorological activities that METEOR discovers each day, as well as fixed geographical features (such as cities, weather stations, large areas like "central Alberta," and mountains), have a spatial location defined by this mapel map. By knowing where particular conditions are rel-

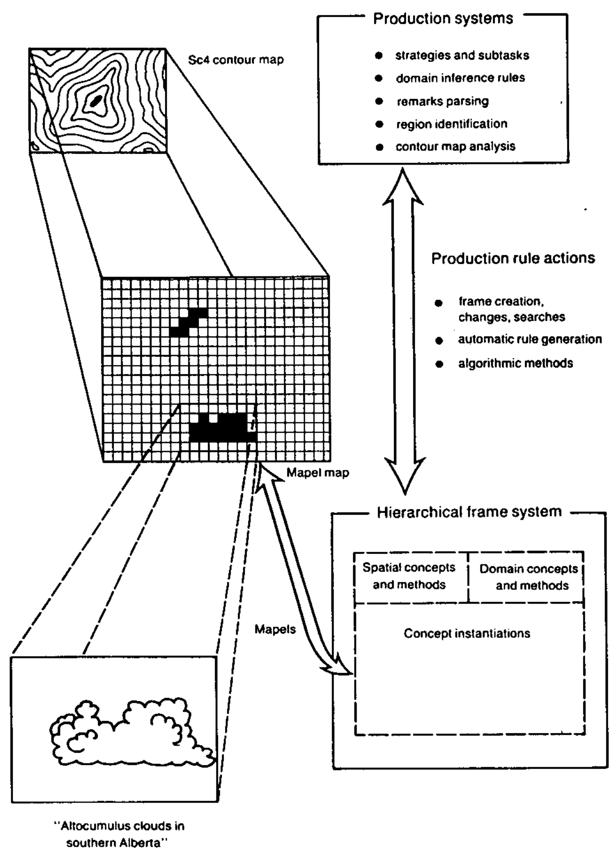

FIG. 5. A schematic view of knowledge organization and interaction in METEOR.

ative to topographical features, METEOR can reason about the effects of topography on weather development in a qualitative manner similar to a human forecaster. Special geometric methods allow METEOR to make spatial inferences and answer questions of the sort "Was there a region of overcast altostratus clouds southwest of STORM-005 two hours ago?" and "Are there lightning reports in the same area currently reporting virga?". Note that these questions are not posed by a human (although they could be). These are questions that METEOR seeks to answer on its own, based on the strategies represented in its IF-THEN rules and on its knowledge of weather conditions and their implication for storm development.

b. *Interpreting surface aviation reports and current weather conditions*

The process of interpreting surface aviation reports is done in two stages. First, each word or abbreviation in the remarks section is matched against a corresponding concept in METEOR's knowledge base. For example, the words in a portion of the remark in Fig. 2 "SHWRS HVIER N,W" are matched against words in METEOR's lexicon. The closest matches are taken to be the full-word "translation" of these abbreviations: SHOWERS HEAVIER NORTH, WEST. Second, using specific knowledge about these concepts and a set of heruistics about the syntax of remarks in general, METEOR groups the concepts in the Fig. 2 remark into noun phrases. IF-THEN rules represent the syntactic heuristics that METEOR relies on to do this task. For example, one such rule might be: IF punctuation occurs between an adjective and a noun THEN the punctuation signals the start of a new phrase. Another rule might be: IF punctuation occurs between two adjectives THEN the two adjectives are part of the same phrase. Exceptions to these general rules are detected by other rules. METEOR bundles each recognizable phrase into a standard format so that they can easily be used later on.

Using the locations of reporting stations, information about current activities gleaned from the remarks and cloud cover sections of the reports is spatially located on the mapel map. When METEOR is looking for a particular observation (e.g., as evidence for certain events), a description of the event is matched against observations associated with particular stations. Having found a set of stations with the required observations, METEOR decides how to join these individual observations into one or more regions of activity. Several heuristics come into play, such as the absolute distances between observations or whether intervening stations report a conflicting pattern of activity. The net result is that complete sets of related surface aviation observations are coalesced into regions which can easily be used for further inferencing. In short, the maps of current weather activity, like the one in Fig. 4, require a good deal of intelligence to make—METEOR cannot simply connect isolated reports of "ACC" or use them intelligently in isolation. Furthermore, METEOR's knowledge about concepts like *virga* or *towering cumulus* allows it to understand the implications of these regions, something that goes beyond rewriting *tcu* as *towering cumulus.*

c. *System organization*

METEOR runs on a Xerox 1100 LISP machine, a special-purpose workstation for developing artificial intelligence systems. This workstation is connected via an Ethernet to a VAX 11/780 in Edmonton, which in turn is connected to a VAX 11/730 in Red Deer, Alberta. (Red Deer, the center of AHP operations for the Alberta Research Council, is about 150 km from Edmonton). Both VAXs run the VMS operating system. The 11/730 receives weather data over a dedicated circuit line. METEOR automatically monitors the arrival of new data on the 11/730, transfers it to the 11/780 and then to the workstation. METEOR then transmits its predictions (maps plus accompanying commentaries) to the Red Deer office.

METEOR is written in Interlisp-D and a production system language called OPS4 (Forgy, 1979). The production system rules that represent heuristics and strategies can access the general meteorological knowledge represented in the frame system. Production rules can also initiate processes on the VAX, such as the transfer of new data or maps to the workstation. METEOR currently begins with 140 production rules and 300 frames that represent its meteorological knowledge for this task. During a forecast, new frames are dynamically created to hold knowledge about concepts that are idiosyncratic to this particular forecast, such as "the area of *virga* reported in southern Alberta at 1000 MDT." METEOR automatically creates as many as 300 new frames as it goes through the forecast process.

Some features of METEOR are configured by information that resides on the VAX 11/780. These features include names and location of weather stations, geographical details about the forecast area, the division of the area into GEOREGIONS, and the resolution of the mapel map (i.e., the size of each mapel). The station data that arrives hourly are automatically collected by processes that monitor a particular set of stations in the forecast area (currently around 100). Other processes wait for the forecaster to run the Sc4 model (during development and testing, METEOR operated in tandem with the forecasters in Red Deer). When the Sc4 analysis is completed, processes automatically create the Sc4 and surface-moisture contour maps and organize them in a special directory on the 11/780. The important role of this directory is to coordinate the transfer of data from the VAX to the workstation. By using this directory, METEOR finds newly arrived station reports and maps that it has not yet analyzed.

6. Field testing

A short field test of the METEOR prototype was conducted for 2 weeks in August 1985. Before starting the field test, it was recognized that there was much that METEOR did not know about forecasting hailstorms. However, it was important to begin the iterative process of knowledge refinement by allowing forecasters to critique METEOR's predictions and analyses. In this way, METEOR's mistakes, as well as its correct interpretations, would indicate flaws in its reasoning or knowledge base. Second, insight was required on how forecasters would integrate knowledge of current weather conditions over a large area with the interpretation of the Sc4 model. Although METEOR has some knowledge about the significance of current conditions for interpreting the model, the forecasters themselves never had the opportunity to combine these two sources of knowledge in real time during a forecast. Finally, it was desirable to test the operational aspects of METEOR as a system. This included automatic data collection from Red Deer to Edmonton, and the transfer of predictions back to Red Deer.

During the 2-week test, 3 days had to be omitted because of data communication problems. Of the remaining 11 days, METEOR correctly predicted that no severe storm would occur on 9 days. The prototype was not configured to forecast the intensity of weak-intensity storms, only that there would be no severe storms (CDC $\geqslant +3$) on those nine days. Within this period, METEOR successfully predicted the intensity (CDC = +5) and location of the most severe storm of the season (3 August 1985). However, it made its most notable error on the following day, when it over-predicted the intensity of a weak convective storm. The forecaster that day had the advantage of knowing that the short wave trough involved had already passed by, so that conditions negating convective storms were already influencing the region at forecast time. METEOR does not currently have access to this information, but similar types of information could be incorporated in an improved version.

7. Observations and conclusions

From a meteorological standpoint, METEOR's current knowledge about the implications of local weather conditions needs to be refined and expanded. For example, its knowledge about local weather conditions can be integrated better with its interpretation of the statistical model. This was a primary goal of the METEOR project, but was complicated by the fact that the human forecasters often do not have the time to integrate the available data on current weather conditions with their interpretation of a numerical or statistical model. The surface aviation reports are a good example of this: much of the useful information contained in these reports, particularly the English-like remarks provided by the observer, cannot easily be understood by traditional computing methods and are too numerous for humans to interpret in real-time. The field test was the first concrete opportunity for investigating the interaction between quantitative data from the statistical model and the information contained in the surface aviation data on local weather conditions.

METEOR's performance would improve significantly by increasing its access to, and knowledge about, other types of information. This includes more accurate indications of wind patterns at 700 and 850 mb, information on synoptic patterns such as approaching troughs and ridges and especially information on cloud systems from satellite imagery. (Knowledge about synoptic patterns enabled the forecaster to downplay the storm that METEOR over-predicted.) In addition, the experts have improved their interpretation of the Sc4 model since the original knowledge base was incorporated into METEOR. Because METEOR uses a particular numerical model, it may seem to be a rather limited forecasting system. However, it is precisely this kind of "local" knowledge that knowledge-based systems can exploit. Rather than applying the same numerical model to all geographical areas, knowledge-based systems can represent the forecaster's field experience, rules of thumb, and even topographical influences, rather than dilute (or ignore) local experience, as numerical models generally do.

The implementation of METEOR highlighted several important features of applying artificial intelligence techniques to meteorological applications. First, any "expert system" in meteorology will have a significant "system" component as well as "expert" component. Communications issues will arise when statistical models or data sources resident on other computers are integrated with a knowledge-based system. METEOR demonstrates that automatic data acquisition is doable, but a better scheme than the one METEOR currently uses is also possible. At present, METEOR explicitly consults the VAX to see if new surface aviation reports or maps have arrived and if so, transfers them. This is one of METEOR's agenda items that are processed by the production systems and initiating these processes is the result of a rule's action. Ideally, METEOR should be interruptable. That is, if new data arrive when METEOR is in the middle of another task, METEOR should recognize this and possibly suspend the current task to deal with it. This requires a qualitatively different kind of artificial intelligence system architecture—one that responds to interruptions—as well as a different problem-solving strategy—one that allows priorities to change dynamically and tasks to be suspended and resumed. It seems that many meteorological applications would require this approach for optimal integration with data-collection systems.

Second, the ability to reason about spatial and temporal relationships is critical to any meteorological application. There are special-purpose artificial intelligence techniques, representations, and algorithms that are best suited for this kind of reasoning (Allen, 1983; McDerrmott, 1982; Schubert et al., 1983). These techniques and representations must be part of a coherent system that includes other artificial intelligence formalisms such as rules and frames. In METEOR's case, the geometric methods and spatial inferencing techniques are an integral part of knowledge representation framework. A concept like "altocumulus clouds in southern Alberta" inherits meteorological features as well as spatial inferencing methods from associated concepts.

Third, artificial intelligence techniques can capture the knowledge and forecasting strategies that are not easily represented in equations or statistical models. However, artificial intelligence should not be seen as an alternative or replacement for numerical models. Rather, it provides meteorology with the means for designing computer systems that know about both the idiosyncracies of a particular locale, and how these impact the interpretation of quantitative forecasting tools.

Acknowledgements. This work was conducted at the Alberta Research Council, and supported in part by a Natural Science and Engineering Research Council of Canada Industrial Research Fellowship to Renée Elio. Portions of the material in this article were presented at the Workshop on Operational Meteorology held in Winnipeg, February 1986, and sponsored by the Atmospheric Environment Service of Environment Canada and the Canadian Meteorological and Oceanographic Society. We would like to thank Chris Sackiw for his cooperation during the field test and the knowledge refinement processes. We are also grateful for the suggestions of two anonymous reviewers which helped the organization and presentation of this paper considerably. Correspondence about this article should be sent to Renée Elio, Department of Computing Science, University of Alberta, Edmonton, Alberta. T6G 2E7.

REFERENCES

Allen, J. F., 1983: Maintaining knowledge about temporal intervals. *Communications of the ACM,* **26,** 832–843.

Elio, R., and J. de Haan, 1985: Knowledge representation for a severe storm forecasting system. *Proceedings of the Ninth International Conference on Artificial Intelligence,* Los Angeles, 401–409.

Forgy, C. L., 1979: The OPS4 Reference Manual. Technical Report, Department of Computer Science, Carnegie-Mellon University.

McDermott, D., 1982: A temporal logic for reasoning about processes and plans. *Cognitive Science,* **6,** 101–155.

Schubert, L. K., M. A. Papalaskaris and J. Taugher, 1983: Determining type, part, color, and time relationships. *IEEE Computer,* **16,** 53–60.

Strong, G. S., 1986: Synoptic to mesoscale dynamics of severe thunderstorm environments: A diagnostic study with forecasting applications. Ph.D. thesis, University of Alberta, Edmonton, 346 pp.

Strong, G. S., and W. D. Wilson, 1983: The Synoptic Index of convection: Application to the Fort Collins hailstorm of 30 July 1979. Paper presented at the Amer. Meteor. Soc. Conf., Tulsa.