TEACHMET: An Expert System for Teaching Weather Forecasting*

NATHAN M. REISS AND JAMES C. HOFMANN**

Rutgers—The State University of New Jersey, New Brunswick, New Jersey

18 May 1987 and 25 August 1987

ABSTRACT

Students of weather forecasting need to learn to identify efficiently the information relevant to the elements they predict. One way students learn these skills is by engaging in discussions of interesting weather situations as they occur. In these discussions, the mentor explains his reasoning in developing the forecast, and his rationale for using particular data. This process is very time consuming and relies on the student's willingness to ask questions. In this project, we developed a computer-assisted instructional system that interactively guides the student of weather forecasting in choosing and interpreting the wide variety of available data. The system is capable of giving detailed responses to the student's questions and is available whenever the student wishes to use it. The system incorporates the expertise of several faculty members and graduate students, as well as two National Weather Service forecasters. A description is given of the creation and evaluation of the system as well as difficulties that were encountered.

1. Introduction

Students in undergraduate meteorology curricula typically receive instruction in the interpretation of various types of weather observation displays (e.g., map analyses, satellite photographs, weather radar displays) and in the assessment of forecast model output. Weather forecasters are continually inundated with a huge variety of satellite photographs, weather maps depicting various levels of the atmosphere, and printed observations and forecasts from individual weather stations. Perhaps the most important skill for a weather forecaster is to be able to identify efficiently the information that is relevant to the elements that are being predicted. Because weather features differ so much from place to place, developing general techniques applicable to a wide range of locations is virtually impossible. Students usually learn these skills by engaging in discussions of interesting weather situations as they occur. In these discussions, the mentor (usually a professor or experienced graduate student) explains his reasoning in developing the forecast and the rationale for using particular subsets of the data. This approach has two problems:

1) This method of teaching is manpower-intensive. In order for it to be carried out properly, a skilled individual should be available to answer questions and provide guidance whenever the student is working on a forecast. Ideally, the student should be able to work on the forecast at any time, taking advantage of the arrival of new data or of breaks in his class schedule.

2) This method of teaching relies heavily on the student's willingness to ask questions and request clarification. Many students have inhibitions about asking questions for fear of appearing ignorant or disturbing the faculty members.

The objective of this work was to develop a computer-based tutorial system for teaching weather forecasting that deals with these two problems by providing

1) expert tutorial assistance whenever needed by the student;

2) answers to the student's questions without judging the student's abilities. It also allows itself to be subjected to extended questioning without becoming impatient;

3) facilities by which several students at one time can receive independent tutorial assistance.

Racer and Gaffney (1984) have noted that the use in meteorology of expert systems that have sufficient explanatory capabilities has the potential to enhance weather forecasting by

> . . . forcing codification of relevant procedures, thought processes, etc. in a precise, thorough, consistent manner. Also, when available on-line to the forecasting

---

* An earlier version of this paper appeared in the *Preprints of the Third International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, New Orleans, LA, January 1987, sponsored by the American Meteorological Society.

** Present affiliation: Dresdner, Robin & Associates, Jersey City, New Jersey.

Corresponding author address: Professor Nathan M. Reiss, Rutgers-The State University, Dept. of Meteorology & Physical Oceanography, P.O. Box 231, New Brunswick, NJ 08903.

community, such algorithms should enhance the skills of the forecaster . . .

To use the system, the student enters the tutorial by stating what type of forecast he wants to make. The system then engages him in a dialogue about the specifics of the forecast that he is preparing. In reaction to the student's responses, it directs him to seek out particular items of information and type them in at the terminal. The computer then examines these responses and requests more information on the basis of what has been reported until the computer finally suggests a categorical forecast on the basis of the data that the student has entered. At any point during the tutorial, the student may ask the computer why particular information has been requested. The computer responds from its knowledge base as to why the information is important to the forecast. The initial tutorial that we developed is confined to giving probabilistic forecasts of whether or not precipitation will occur. However, the system is open-ended in the sense that tutorials for other aspects of precipitation forecasting (e.g., precipitation type and intensity) and forecasts of other weather elements can easily be added.

The tutorial system is designed to discourage the student from forecasting by rote. In contrast to forecaster checklists, which have been widely used for many years, the system does not ask redundant questions or questions that do not pertain to the specific forecast problem. For example, if the user indicated that he is making a 48-h forecast, the system will not ask a series of questions pertaining to the present weather conditions because the rules state that model output results should take precedence over present weather in the preparation of 48-h forecasts. Because of the very wide range of responses that its extensive logical base is capable of producing, repetitiousness and readily anticipated behavior on the part of the system are minimized. Questions are posed that the student has to answer before proceeding to the next step.

2. Expert systems

The approach makes use of the rapidly developing technology of expert systems, a computer-based technology intended to solve a narrow domain of problems with expert proficiency (Shortliffe and Duda, 1983). Expert system software differs from the more usual types of computer programs in that it is directed primarily to decision making rather than to computational uses. The knowledge base (that is, the set of rules and definitions governing the system's operation) of the expert system may be entered a priori by the system developer. Alternatively, some expert systems can be programmed to deduce rules from actual cases. The former method is most useful when the rules are known in some detail, while the latter method is best suited when the rules are ill defined. Since the purpose of the system described here is to impart the experts' knowledge to the students, it was decided to enter the rules beforehand. Because the rules of weather forecasting are certainly ill defined, an expert system designed for use in an operational setting (e.g., at a forecast office) might well be of the self-learning variety.

Some of the early expert systems were developed for providing medical consultations (Shortliffe, 1976), performing spectral analysis (Buchanan and Feigenbaum, 1978), and structural studies (Bennett and Engelmore, 1979). More recently, some applications of expert systems in the field of meteorology have begun to appear. These have included systems for convective storm forecasting (Elio et al., 1987) and recognizing wind shear hazards (Campbell and Olson, 1987).

3. Hardware and software

The weather forecasting tutorial system consists of an IBM PC-XT computer, with a 20-megabyte fixed disk capable of storing a large base of information (the knowledge base). It is equipped with a Classic PC/XT Multiuser System[1] that enables it to support four users simultaneously. Access to the computer is provided by four video terminal work stations placed in the area of the meteorology department where current weather information is received. The system makes use of a commercially available expert system software development package and inference engine, ES/P Advisor.[2] This package is capable of processing user requests for information during the decision making process, rather than only after the process has been completed. Examples of such information include a trace of the logic that led the system to ask the current question or to arrive at the current conclusion, textual material explaining the reasoning behind a particular question or conclusion, or bibliographic references.

4. Creation of the knowledge base

The existing literature pertaining to expert systems stresses the importance and difficulty of developing an appropriate knowledge base (see, for example, Hayes-Roth et al., 1983; Doyle, 1985). Creation of the present knowledge base has been a recursive process involving extensive one-on-one interviews with meteorology faculty members, graduate students, and forecasters of the U.S. National Weather Service. Although the ES/P Advisor software includes an excellent expert system development package, for logistical reasons we used a consultation system, EXPERT, developed by Weiss et al. (1985) at Rutgers University and residing on the

[1] Classic Technology Corp., 2090 Concours Dr., San Jose, CA 95131.

[2] Expert Systems International, 1700 Walnut St., Philadelphia, PA 19103.

University's mainframe computer, for early work in developing the knowledge base. The knowledge base was later ported to ES/P Advisor.

Initial interviews generally followed the outline given in appendix A. Formal letters were sent to each expert participant in the project. Accompanying each letter was a short introduction to expert systems in general and to this project in particular. With the permission of the expert participants, interview sessions were tape recorded. Recordings later were reviewed and detailed notes made.

After a first-draft knowledge base was developed, the participants were given instructions on how to access and run the software. They could then access the system, run the consultation program, and relay any comments to the model developers. Because not all of the participants were at the same location, it was advantageous for this phase of the process to have the system operating on a computer with remote job entry capability. The EXPERT software also allowed participants to save their consultation sessions in a file for later examination and comparison with their comments.

Additional tools were used besides the interviews to obtain specific information for the knowledge base. The first was a questionnaire designed to determine how forecasters quantify certain parameters, such as the strength of upper-air flow patterns, vorticity and thermal gradients. The quantifications desired in each case were "strong," "moderate," and "weak." The participants were each given 15 850-mb charts and 15 500-mb height/vorticity charts and were asked to circle areas of strong, moderate, and weak values of the above quantities. Their responses were tallied, and representative contour spacings for each quantity were determined. The definitions of "strong," "moderate," and "weak" were made a part of the explanatory material that the expert system offers its users.

A second tool used was a five-page forecast worksheet that was filled out before, during, and after precipitation events. The sheets were used by one expert who was a professional forecaster at a nearby National Weather Service office. The first page of the worksheet provided space where the forecaster could note the magnitudes of the various precipitation predictors and the forecasting product used to obtain each predictor. The second page provided space for a discussion of the weather situation, a forecast, and a discussion of how the parameters influenced the forecast. The third and fourth pages were local and regional base maps on which the forecaster was asked to draw frontal location, high and low pressure centers, and other important weather features. The fifth page was a blank skew T–log$P$ diagram, to be used only when the potential existed for severe weather events. The worksheets were completed during May through August 1986 and in December 1986 and January 1987.

Additionally, for three semesters, tape recordings were made of map discussions held weekly at the Rutgers Meteorology Department. During these map discussions, weather events of the previous week are discussed in the context of the map analyses and forecast products. The discussions include appraisals of the accuracy of the model runs and conclude with the discussant presenting a forecast for the next 2 days. Generally, there is vigorous audience participation, providing an excellent atmosphere for developing an understanding of the forecasters' deductive processes. Each map discussion recording was reviewed twice. Notes were made of the reasoning used in making the forecasts and of any rules of thumb that were stated. Weather maps used in the discussions were recovered and used as an aid in interpretation of the recordings.

Figure 1 is a flowchart giving a general overview of the precipitation forecasting process. The logic of the program is modularized into sections, each one dealing with one aspect of the forecast problem. As in the case of the section dealing with use of existing weather conditions, sections may be skipped altogether, depending on information that is received during the consultation. Figure 2 is a flowchart of the section that evaluates if sufficient moisture is available. The diagram illustrates only the original operational version of the section and obviously does not at this time represent the ultimate precipitation forecasting technique. When the relative humidity predicted by the National Meteorological Center's Forecast Output Statistics (FOUS) model is less than or equal to 50%, none of the other questions are asked and the conclusion is that moisture is lacking. Figure 3 gives examples of some of the screens displayed to the user.

Initial tests of the tutorial, a small segment of which is shown in appendix B, were directed primarily to confirm the technical integrity of the overall system. Experienced forecasters, including some not involved in formulation of the knowledge base were asked to "exercise" the system (including documentation) for real and/or simulated weather situations of their own choosing, and to give an assessment of its technical integrity. Minor changes were made as needed. Acceptance by the experts of the system as it currently exists does not imply that the knowledge base is complete. We view the knowledge base as a constantly growing and evolving body of knowledge. Continual feedback resulting from use of the expert system provides the stimulus for further changes.

5. System evaluation

The purpose of this stage of evaluation was to locate any portions of the tutorial that were confusing to the students or were being interpreted incorrectly. A small group of students, chosen for their variety of experience and talent, were asked to use the system (including

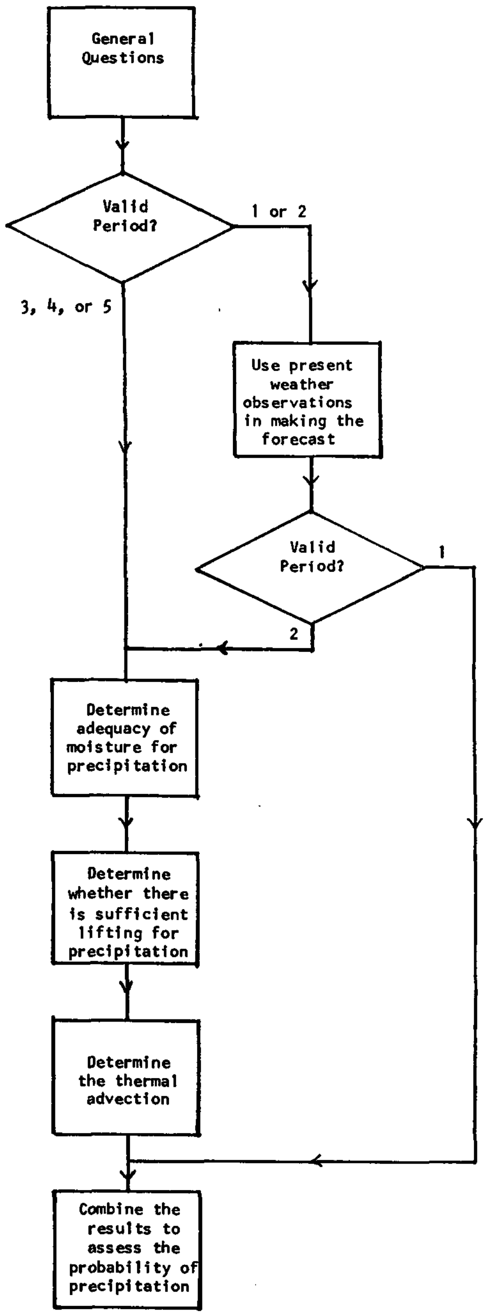

FIG. 1. Flow chart depicting a general overview of the precipitation forecasting process.

documentation) as part of their weather forecasting classes and were provided with forms to record their comments. Since the Meteorology Department of Rutgers is small, one-on-one discussions also took place with most students who had used the system, often while they were using it. Again, changes were made as needed. Since the students were only just beginning to learn forecasting, they had few if any comments relating to forecasting procedures, but were in some cases able to expose ambiguities in the system's requests for information, and suggested subjects on which explanatory information would be useful.

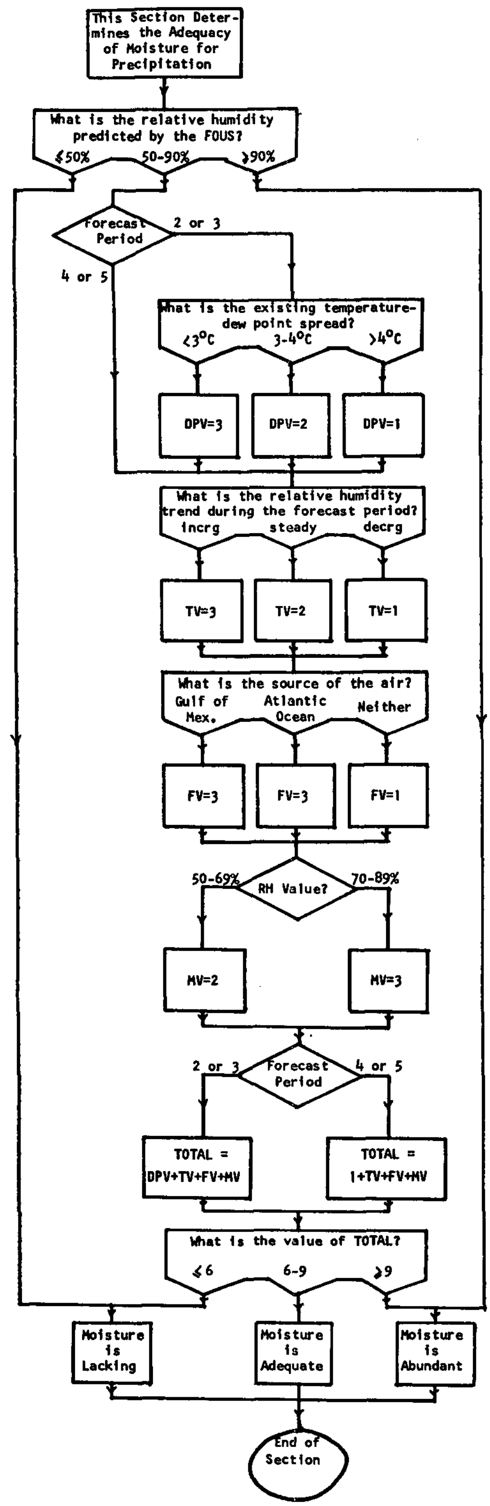

FIG. 2. Flow chart of the section that evaluates if sufficient moisture is available.

```
GOAL:  rh_trend           SECTION: humidity
--------------------------------------------------------------------------------

Please indicate the trend in relative humidity during the
period for which you are forecasting.

   (1)  -  positive
   (2)  -  steady
   (3)  -  negative

--------------------------------------------------------------------------------
Enter the number of the relevant entry:
--------------------------------------------------------------------------------
```

```
GOAL:  rh_trend           SECTION: humidity
--------------------------------------------------------------------------------

I am asking because I wish to establish the state of
the relative humidity trend during the forecast period (rh_trend)

which is a necessary pre-condition for the display of the
following paragraph:-

'NOTE:
 -----

Generally with a relative humidity between 50 and 70 percent, regardless of
other humidity parameters, a forecast of increasing cloudiness should be
made.  This is especially true after an upper-level (500 mb) ridge has
passed over the area and when there is a positive or steady relative humidity
trend for the forecast period.'

--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
                    PRESS ANY KEY TO CONTINUE
```

```
GOAL:  rh_trend           SECTION: humidity
--------------------------------------------------------------------------------

FINAL DETERMINATION OF SUFFICIENT MOISTURE:
-------------------------------------------

Upon evaluation of the parameters used for determination of the
availability of moisture for precipitation, it appears that there
will be sufficient moisture for precipitation to form.

--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
                    PRESS ANY KEY TO CONTINUE
```

FIG. 3. Examples of screens presented to the user during a consultation session with the expert system. The top screen prompts the user for information needed in the section outlined in Fig. 2 that determines the adequacy of moisture. The middle screen appears when the user types WHY in response to the question. The bottom screen informs the user of the determination that was made as soon as the section has been completed.

6. Effectiveness of the system

We found that the system, even in its early stages, quickly provoked the interest of users. Experienced forecasters using the system often remarked that they had never seen so much information about precipitation forecasting concentrated and systematized in this way; they found their experience with the system very informative and useful. The nature of the impact on the students is a greatly enhanced opportunity to learn weather forecasting through dialogues with an "expert," thereby exposing them in an orderly way to the complex rules and guidelines that experienced forecasters use to assimilate the large quantities of data that are available. Students are able to obtain assistance in preparing their forecasts at any hour of the day, without regard to the presence of a professor or teaching assistant. It should be emphasized that the proposed system does not replace any classroom contact that students have had with faculty.

Students and faculty rightfully have approached the system with a certain amount of skepticism—probably in much the same way that they would approach a human otherwise unknown to them who claimed to be an expert in precipitation forecasting. Before relying on him, they would wish to gauge his knowledge for themselves. As with a human expert, full acceptance of the system's expertise will have to be earned after it has had the opportunity to prove itself over a number of seasons and variety of weather situations.

A pedagogical weakness of the system that we hope to remedy in the future is that if the dialogue that it provides is followed to its conclusion, it eventually leads to a forecast made by the system, rather than by the student. Implicit in this procedure is the concept that the system is capable of providing a "correct" answer. Of course, the student is not required to use the system's forecast. In fact, by making available a diagram of its reasoning, and by allowing a manual override of any of its conclusions, the system in fact enhances the student's ability to examine its logic and to substitute his own reasoning. However, this aspect of the expert system's operation is not at present being advertised to the student strongly enough, so he is not likely to take full advantage of it.

7. Future enhancements

We expect that the system described herein will continue to be refined and expanded. The work described only scratches the surface of an exciting array of potential applications of expert systems in teaching weather forecasting. The most obvious extension of the system is the addition of knowledge bases for other weather elements, particularly temperature and wind. There are two other important enhancements to the system that we hope to incorporate in the future. The first of these would allow the system to "remember" the details of its consultations with students and provide summary reports. This information would be useful to faculty members in determining where students' forecasts may have gone wrong. The ES/P Advisor software already has the ability to store consultation files, but does so only at the request of the user. It would be useful to be able to automate this function. The second enhancement would provide a means for entering the actual weather that occurred and comparing it to the forecasts that were made. Not only would this allow for automatic scoring of the students' forecasts, but more importantly, it would set the stage for allowing the system to learn by experience. That is, the system could automatically compile a record of the success rate for the use of each rule or combination of rules and report this information either to the students or to the system developers. A system of this sort would likely be of the type that is capable of deducing rules from actual cases, as discussed earlier.

Acknowledgments. The bulk of the forecasting knowledge contained in the system described here was kindly supplied by Drs. Robert Harnack and Mark D. Shulman and Mr. Keith Arnesen of the Rutgers Meteorology Department; Mr. Harry Woodward of the U.S. National Weather Service Office in Newark, New Jersey; and Mr. Tom Morgan of the U.S. National Weather Service Forecast Office in New York, New York. This work was partially supported by Contract 87-990640-2048 from the New Jersey Department of Higher Education.

APPENDIX A

Outline Used for Initial Interview with Experts

I. Introduction
   A. General project description
      1. What is an expert system
      2. What is knowledge engineering
      3. Descriptions of some expert systems currently in use
      4. Brief description of the precipitation tutorial project
      5. Explanation of the interviewing scheme
   B. Introduction to the expert system building tool, EXPERT
      1. Types of problems handled by EXPERT
      2. Knowledge representation
         a) Types of questions that EXPERT can ask
         b) Types of rules that EXPERT can use
II. Knowledge acquisition
   A. Problem identification and definition
      1. Participants and their roles
      2. Sources of knowledge other than experts
      3. Scope of the problem
      4. Nature of the input data
      5. Type of result that the system is supposed to provide
   B. Problem conceptualization
      1. General structure of the reasoning

2. Relationships between the various types of observations
3. Inferences that can be drawn from the observations
4. Relationships between observations and conclusions
5. Important intermediate conclusions

APPENDIX B

Example of Tutorial System Dialogue

The following is an example of a brief portion of the dialogue produced by the tutorial system.

Please indicate the present season:
1) Winter
2) Summer
3) Spring or Fall
*3

Please indicate the period that contains your valid time:
1) Period 1 (3 to 6 hours from the initial NGM or LFM)
2) Period 2 (7 to 12 hours from the initial NGM or LFM)
3) Period 3 (13 to 24 hours from the initial NGM or LFM)
4) Period 4 (25 to 36 hours from the initial NGM or LFM)
5) Period 5 (37 to 48 hours from the initial NGM or LFM)
*3

You chose period 3 as your valid time. For that period we use the 12 and 24 hour NGM prog charts, the 12, 18, and 24 hour NGM FOUS output, the MOS, the 12 and 24 hour 850 mb prog, and the LFM as our sources of forecasting guidance. Press RETURN to continue.
*

Use the NGM map indicated at the beginning of the consultation to make a forecast. Go to the 500 mb vorticity panel. Follow the 500 mb flow upstream from New Brunswick and look for a vorticity trough or ridge. What did you find?
1) A trough, indicating PVA over New Brunswick
2) A ridge, indicating NVA over New Brunswick
3) Neither a trough nor a ridge, indicating neutral vorticity advection over New Brunswick.

[Note: At this point, if the student types "EXPLAIN", the system will provide the following amplification of the question:

A vorticity ridge is the point where there is a change from PVA to NVA following the 500 mb height contours.

A vorticity trough is the point where there is a change from NVA to PVA following the 500 mb height contours.

If the vorticity GRADIENT crosses the 500 mb contours at an angle from 81 to 90 degrees, consider the vorticity advection to be neutral.

Which of these do you encounter first as you travel upstream from New Brunswick? Use the NGM prog map that is closest in time to your valid time. If your valid time occurs at or near the middle of a period use the prog maps at the beginning and end of the period and estimate what the conditions would be like at your valid time. Take into consideration how fast systems are moving, or how fast the NGM and LFM are forecasting for systems to move.]
*1

Use the NGM FOUS output or the NGM prog map (700 mb panel) to determine the relative humidity (average over all layers) for New Brunswick. Use the NGM prog that is closest to your valid time. NOTE: Choose the Relative Humidity value that is the highest given for your period.
*45

Please indicate which is true:
1) New Brunswick will experience warm advection at the valid time.
2) New Brunswick will experience cold advection at the valid time.
3) New Brunswick will experience neutral advection at the valid time.
*1

REFERENCES

Bennett, J. S., and R. S. Engelmore, 1979: SACON: A knowledge-based consultant for structural analysis. *Proc. Sixth Intl. Joint Conference on Artificial Intelligence*, IJCAI, Tokyo, 47–49.

Buchanan, B. G., and E. A. Feigenbaum, 1978: DENDRAL and METADENDRAL: Their applications' dimensions. *Artif. Intell.*, **11**, 5–24.

Campbell, S. D., and S. H. Olson, 1987: Recognizing low-altitude wind shear hazards from Doppler weather radar: An artificial intelligence approach. *J. Atmos. Oceanic Technol.*, **4**, 5–18.

Doyle, J., 1985: Expert systems and the "myth" of symbolic reasoning. *IEEE Trans. on Software Engineering*, **SE-11**, 1386–1390.

Elio, R., J. de Haan and G. S. Strong, 1987: METEOR: An artificial intelligence system for convective storm forecasting. *J. Atmos. Oceanic Technol.*, **4**, 19–28.

Hayes-Roth, F., D. A. Waterman and D. B. Lenat, 1983: *Building Expert Systems.* Addison-Wesley.

Racer, R., and J. Gaffney, Jr., 1984: Interpretive processing/expert systems: An initiative in weather data analysis and forecasting. *Nat. Wea. Dig.*, **9**, 31–45.

Shortliffe, E. H., 1976: *Computer-based Medical Consultations: MYCIN.* American Elsevier.

——, and R. O. Duda, 1983: Expert systems research. *Science*, **220**, 261–268.

Weiss, S., K. B. Kern, C. A. Kulikowski and M. Uschold, 1985: *A Guide to the EXPERT Consultation System.* Tech. Rep. CBM-TR-94, Dept. of Computer Science, Rutgers University, 123 pp.