SIGMAR: A Fuzzy Expert System for Critiquing Marine Forecasts

Bjarne Hansen
Maritimes Weather Centre
Environment Canada
1496 Bedford Highway
Bedford, Nova Scotia
Canada B4A 1E5

bjarne.hansen@ec.gc.ca

Hansen, B. K., 1997: SIGMAR: A fuzzy expert system for critiquing marine forecasts,
AI Applications, Vol. 11, No. 1, 59-68. [This paper was reformatted for *AI Applications*.]

Abstract

Meteorological information and knowledge are often uncertain, ambiguous, or vaguely defined. Fuzzy logic lets expert systems perform optimally with uncertain or ambiguous data and knowledge. With a fuzzy logic framework, one can efficiently implement linguistically expressed rules derived from experts. Operational meteorology is therefore treated as a fuzzy environment. An argument is made for the applicability of methods based on fuzzy logic for the optimal solution of problems related to the evaluation of meteorological data and forecasts. An expert system, SIGMAR, has been designed which uses fuzzy methods to interpret meteorological data. The system automatically evaluates the significance of actual wind reports. Two activities that challenge weather forecasters are coping with information overload and maintaining accuracy of forecasts. Both tasks can be performed more easily and consistently with SIGMAR. The system efficiently identifies significant information contained within huge amounts of data. Forecasters using the system can more consistently and easily monitor the accuracy of weather forecasts. Systems such as that described here are bound to become more common as time goes on.

Introduction

Forecasting meteorologists have to interpret huge amounts of data. Forecast production is complicated by the fact that meteorological information is often uncertain, ambiguous, or vaguely defined. Meteorologists will benefit from the development of new tools that use techniques from the fields of fuzzy logic and expert systems. This paper demonstrates how one of meteorologists' tasks – weather watch – can be performed more easily and consistently with the use of a fuzzy expert system. The system is called SIGMAR, short for "Significant Information Generated from Marine Area Reports." The system continuously monitors wind data, and alerts a forecaster if wind data indicate that a marine forecast is tending to become inaccurate.

Marine Forecasts

Winds at sea can pose a serious danger to those who are unprepared, or those whose vessels are not suited for rough conditions. The best way to deal with the danger is to avoid it. Marine forecasts give people some idea of what to expect at sea and are depended on by mariners who plan to go to sea safely. Mariners make decisions based on marine forecasts. Accurate forecasts of strong winds enable mariners to plan to avoid dangerous conditions. Likewise, forecasts of light winds let them to plan to profit from ventures at sea. Users of forecasts are a diverse community including fishermen, container shipping lines, weekend sailors and surfers. Each group of users has to respect its own safety thresholds where wind is concerned. The cost of exceeding safety limits is high. Lives are lost and property damaged in waters around the Maritimes provinces of Canada every year due to weather related incidents.

The Maritimes Weather Centre (MWC) works to maintain a high standard of accuracy for marine forecasts. Accurate forecasts equate to protection of life and property. When people or property are lost at sea, forecasts and conditions are closely reviewed. Three strategies are used to maximize the accuracy of marine forecasts: skilled forecasting; surveillance of current weather; and periodic verification of forecasts (i.e. quality control). All three of these activities can benefit from expert systems based on fuzzy logic. This paper describes a prototype expert system developed to assist forecasters with the second task – weather watch.

Hansen, B. K., 1997: SIGMAR: A fuzzy expert system for critiquing marine forecasts, *AI Applications*, Vol. 11, No. 1, 59-68. [This paper was reformatted for *AI Applications*.]

Fuzzy Reasoning with Meteorological Information

Meteorologists use heuristic reasoning all the time. The process equates to using series of "if-then" rules. We can also reason effectively with fuzzy syllogisms. For instance:

(a) *if* snowfall is <u>heavy</u> *then* transportation will be seriously affected

Most people will agree this is true. The question is: How true? How will we define *heavy*? A simple approach is to pick one value of snowfall and test the truth of the statement. We can make the assertion more specific:

(b) *if* <u>15 centimetres</u> of snow falls *then* transportation will be seriously affected

If we survey 100 people, perhaps 90 people will agree with this statement. But when we select arbitrary values we limit the flexibility of our reasoning; we artificially restrict the meaning of the word *heavy* as it pertains to snow. Statement (a) represents a unit of knowledge that people can use to reason about the implications of various amounts of snowfall. The truer it is that snow is heavy, the truer it is that transportation will be seriously affected. With statement (b), we are limited to drawing conclusions based on only one particular value of snowfall – "15 centimetres or more of snow implies transportation will be seriously affected."

People's perception of *heavy* depends on who they are and where they live. Among 100 people in Halifax, there may be 10 elderly or disabled people who find even a trace of snow limits their mobility. In Washington D.C, a light snowfall will cause traffic problems. In Anchorage, people are prepared to cope with snow. If we seek to define the word heavy as it pertains to snow, we could survey 100 people in three different places and get results as shown in Figures 1, 2, and 3. These figures represent *fuzzy sets*. The functions show the membership of given amounts of snow in the set of snow amounts we call *heavy*. Because the term *heavy* is fuzzy, the functions do not jump abruptly from 0 to 1.

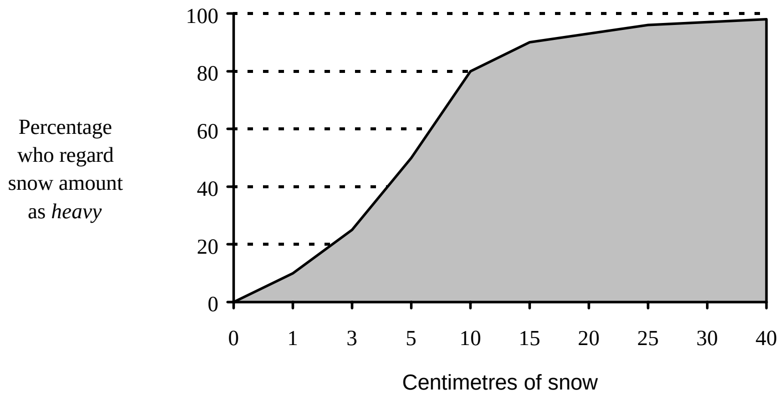

Figure 1. Percentage of people in Halifax who think a given amount snow is *heavy*.

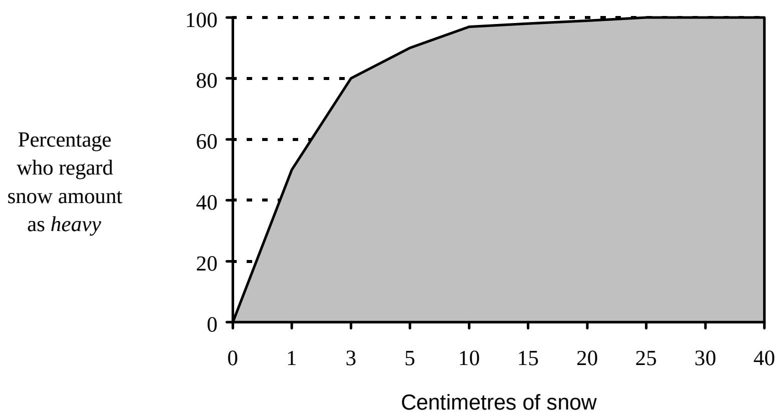

Figure 2. Percentage of people in Washington., D.C. who think a given amount snow is *heavy*.

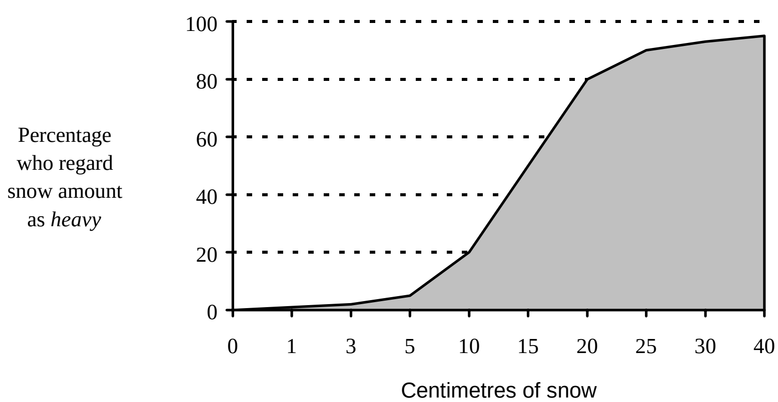

Figure 3. Percentage of people in Anchorage who think a given amount snow is *heavy*.

Hansen, B. K., 1997: SIGMAR: A fuzzy expert system for critiquing marine forecasts, *AI Applications*, Vol. 11, No. 1, 59-68. [This paper was reformatted for *AI Applications*.]

Fuzzy Analysis of Wind Data

Before considering how we can categorize winds with fuzzy sets, we will consider how winds are normally categorized. The conventional way of verifying, or measuring the accuracy of a marine forecast is to examine records of forecast and observed winds and check whether the forecast and observed winds are within the same crisp categories. For instance, winds equal to or greater than 34 knots belong to the category of *gales*. Gales imply hazardous conditions for many mariners. When forecasters predict that gales will blow, they issue warnings to that effect. Periodically, the accuracy of the forecasts is examined and tabulated. The case of an accurate forecast is called a “hit.” An inaccurate forecast is called a “miss.” The membership function of a wind speed in the category of gales is shown in Figure 4.

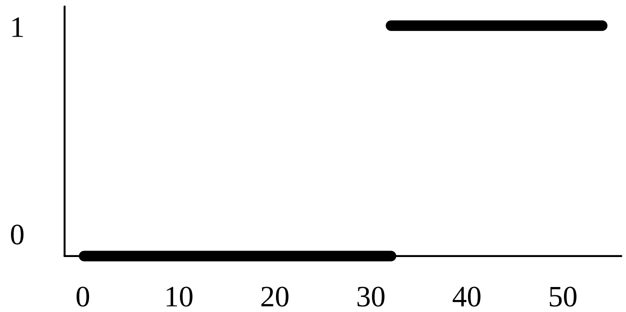

Figure 4. Membership function of a wind speed in the crisp category of gales.

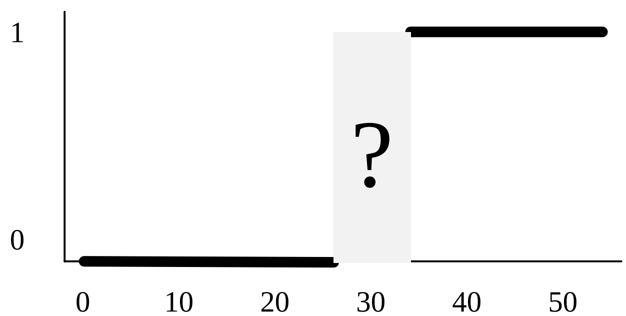

Figure 5. Perceived membership of a wind speed in the category of gales.

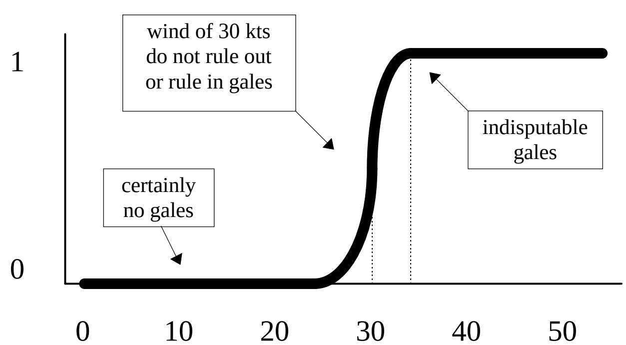

Figure 6. Fuzzy membership function of a wind speed in the category of gales. Given a single wind observation, one forms a degree of belief in the presence of nearby gales.

Hansen, B. K., 1997: SIGMAR: A fuzzy expert system for critiquing marine forecasts,
AI Applications, Vol. 11, No. 1, 59-68. [This paper was reformatted for *AI Applications*.]

The rules of verification are simple: winds of 33 knots or less are not gales; winds of 34 knots or more are. The implication is that 5 knot winds and 33 knot winds are one type of wind, and 34 knot winds and 50 knot winds are another type of wind. This is absurd; there is hardly any difference between a 33 and a 34 knot wind, yet they fall into opposite categories.

In actual practice, meteorologists avoid the trap of categorizing winds in a simple binary way. At the MWC, winds in the range of 28 to 33 knots are treated as “near gales” and are regarded separately from those winds under 28 knots and those of 34 knots or more. This is a reasonable practice. Forecasts apply to thousands of square miles of ocean. Over the course of six hours, meteorologists typically receive only several actual wind observations with which to verify a forecast. If one of the reports is of 30 knots, it is not unreasonable to suppose that nearby winds may have reached 34 knots. One cannot be certain that gales blew or that they did not. This trivalent treatment of winds is akin to Dempster-Schaffer treatment of uncertainty (MacNeil and Freiberger 1993). For purposes of practical verification, one can replace Figure 4 with Figure 5. A trivalent categorization of wind speeds is more realistic than the “gale or no gale” categorization of winds implied by a crisp definition of gales. One can distinguish between 5 knot and 30 knot winds. The winds are different, and this difference in wind is significant.

Fuzzy methods are enabling advances in a rapidly increasing number of data processing and expert system applications. Meteorological data is amenable to treatment with fuzzy methods. For instance, consider the trivalent treatment of winds shown in Figure 5. Although it is better than a binary system, it still wastes information. 28 knot winds are lumped together with 33 knot winds and regarded as “uncertain” or “near gales.” Intuition says that a 33 knot wind is more suggestive of gales nearby than a 28 knot wind is. All antecedent weather data imply consequent facts with varying degrees of certainty. When people attach degrees of belief to implied facts, they are in effect mentally performing a fuzzy operation. The trivalent function shown in Figure 5 can be replaced by a fuzzy set as shown in Figure 6.

The function shown in Figure 6 models the intuitive decision making behavior of a meteorologist. In plain English: A wind measurement of 10 knots very strongly refutes the presence of gales. A wind measurement of 34 knots or more definitively confirms the presence of gales. A wind measurement between 28 and 33 knots is suggestive of gales, and the closer the speed is to 34 knots, the stronger one’s belief is in the presence of gales.

The fuzzy gale function is a simple, illustrative example. It shows how meteorological data yields to fuzzy analysis to produce results that correspond closely to a person’s intuition. Fuzzy analysis brings with it several advantages:

Tunability - Different data sources have different assumed reliability. Sources include ships, buoys, land stations, and visual estimates. When reliability is low, functions can become correspondingly broader.

Adaptability - Different users’ sensitivities and circumstances may necessitate different interpretation of linguistic values. “Strong” means different things to pleasure boaters than it does to container ships. One system can serve both users by making reference to different fuzzy sets. In different seasons, inshore fishermen (with small boats) extend their range far offshore. At these times, the expression “strong winds in offshore waters” takes on a different meaning. One system can adapt to different user requirements by using different fuzzy sets.

Computational efficiency - Many developers report the high speed with which data can be processed using fuzzy logic.

Heuristic system design - Zadeh (1994) says that fuzzy logic simplifies the application of rules by enabling a form of data compression. He refers to the data compression as *granulation*.

Suppose we need to initiate one of three types of responses depending on whether wind is light, moderate or strong. Wind speeds may assume a hundred different values, but we do not want to implement a hundred different rules. The usual simplifying strategy is to arbitrarily lump ranges of values into rigid categories and to assign a meaning to each category. A problem persists: how should we handle "grey cases" that straddle categories?

The advantage of Zadeh's granulation is that it mimics the way in which people interpret linguistic values. Linguistic data and rules are processed to produce linguistic results that are consistent with human reasoning. For instance, if a wind speed is on the threshold between moderate and strong, should a system respond as if the wind is moderate or as if it is strong? Fuzzy systems respond smoothly to transitions by applying multiple rules in an appropriate balance. This is how fuzzy systems achieve smooth and robust behavior.

There is a precedent for treating wind data with an intuition-based, granular system of measurement. The Beaufort wind scale assigns numbers to intervals of wind speed ranging from 1 to 12. Wind speeds can be translated between knots and Beaufort force values. For example, Beaufort force 8 refers to winds ranging from 34 to 40 knots. The Beaufort system is the standard system used to express wind speed in Europe. The salient point here is that its users are more concerned with the *effect of the wind speed* than they are with the exact value of the wind speed. Measurements of wind in the Beaufort system are commonly made by subjective estimates based on waves, spray, movement of a vessel and other loose objects. Users of the system find it an intuitive and user-friendly way to convey wind speed information. The Beaufort system's prevalence proves that people operate effectively with linguistic approximations of wind speed.

Critiquing System for Marine Forecasts

Further examination of the fuzzy gale membership function could lead to the development of a new method of marine forecast verification. Such an examination would refer to a lot of statistics. There is a huge literature dealing with conventional methods of verifying forecasts. Before we attempt to convince meteorologists to try fuzzy methods, it will help to show that practical systems can be based on fuzzy methods. Therefore, this section will describe a new type of wind surveillance system based on fuzzy methods. Surveillance and verification are two methods used by meteorologists to achieve accurate forecasts. One can think of surveillance as real-time verification. The main purpose of the system developed, SIGMAR, is to monitor winds continually, and to alert the forecaster if winds are significantly higher than forecast. SIGMAR also performs fuzzy analysis of time series of wind data and produces reports about a variety of other wind characteristics.

Background

There is a common perception that weather forecasting is a static problem that only needs periodic attention. Many wonder why the weather must be continually monitored, and forecasts often revised. The author has received questions such as: "If you people are so good at forecasting, why do you have to change the forecast so often?" It is a good question. Weather is a dynamic, non-linear, and chaotic process. In other words, you cannot predict it accurately well ahead of time. The cause of errors in weather predictions is well illustrated with the tale popularly known as "The Butterfly Effect." Suppose a butterfly in Brazil flaps its wings. Ripples in the air spread outward through space and through time. Conceivably, the resulting cascade of events could lead to the development of a storm two weeks later in North America. The triggering causes of major events can never be perfectly identified.

Realizing that weather forecasts tend to become inaccurate, forecasters remain vigilant and often revise forecasts to maximize accuracy. Weather forecasting is comparable to driving a car. People monitor their position as they drive and make continual error corrections in their course. This strategy for safe driving is so basic that everyone takes it for granted. Imagine if drivers simply pointed their cars, closed their eyes, and pressed the gas pedal. This is not an ideal way to take advantage of a continual stream of vital information. The continuous forecasting-monitoring cycle works the same way. At the moment a forecast is issued, the forecast is the best possible prediction of future weather developments. A forecast is based on current available information and on models that assume the future weather will evolve through a precisely specified series of states. Just as milk sours when left unnoticed on the table, forecasts go bad when left unattended. Information arrives after issue time and unpredicted sequences of events unfold that compel one to change forecasts. There is a dividend for rapid response. When forecasts are kept accurate, public safety and the public's appreciation for the weather service are increased. When forecasts are inaccurate, avoidable weather-related accidents occur, and the perceived value of the weather service is reduced.

The task of monitoring wind data to ensure maintenance of forecast accuracy can be difficult. There is a huge amount of wind data demanding the forecaster's attention. The data arrive in different forms and at irregular times. Weather forecasts can be quite detailed. Forecasters have other responsibilities that pull their attention away from the task of monitoring for periods of time. Expert systems can help forecasters to cope effectively with growing volumes of information. The recent success of fuzzy systems in various domains indicates the use of fuzzy methods for the processing of meteorological databases. The analogy between driving and weather forecasting is germane. Cars are being built with new types of fuzzy expert systems that help people to drive more safely and smoothly.

The MWC issues marine forecasts and warnings for waters around the Maritimes. Because forecast accuracy drops as forecasts age, new forecasts are issued about every eight hours. The MWC's area of responsibility includes Atlantic waters within two hundred miles of the Maritimes as well as the Gulf of St. Lawrence. The MWC's area of responsibility is shown in Figure 7. The MWC's marine district is divided into nineteen named areas. Forecasters use many tools and lots of data to help them assess and predict the weather. When they compose forecasts, they use an editor called COMAR, short for "Coded Marine." (COMAR was developed at the MWC and is undocumented elsewhere.) A forecast is entered into COMAR in a concise, coded form. COMAR translates the code into worded marine forecasts in English and French. Forecasters prepare long forecasts efficiently by typing only a few essential characters.

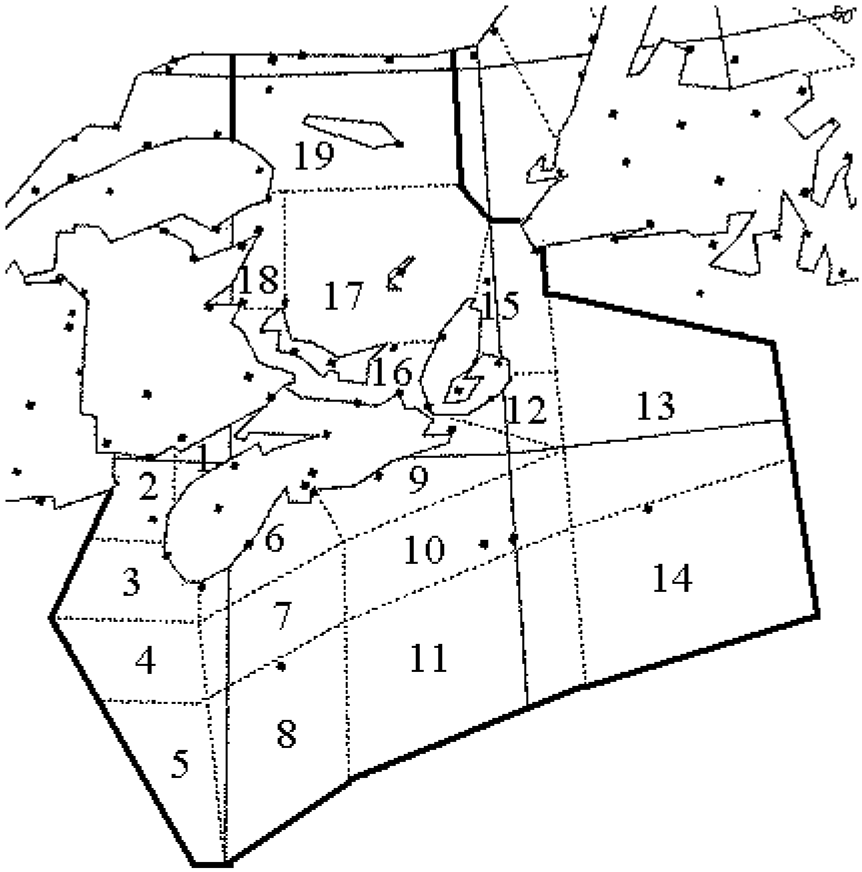

Figure 7. Marine forecast area of responsibility of the Maritimes Weather Centre. The district is has nineteen named areas. About seventy sites provide regular wind information.

A marine forecast in COMAR code resembles a terminal forecast for an airport. A particular wind strength or wind range is forecast for the initial period, and different wind ranges are forecast to begin at particular hours. For instance, a segment of COMAR code may be: "W20-30. 15h W15-20." This means winds are forecast to be initially westerly 20 to 30 knots. At 15:00, winds are forecast to diminish to westerly 15 to 20 knots. The actual forecast is not disseminated in this form. COMAR takes the code and translates it into English phrases; precise hours are translated into vague time ranges. The segment of code is translated to: "Winds westerly 20 to 30 knots diminishing to westerly 15 to 20 this afternoon." Marine forecasts have used this form of grammar for decades.

One can assume that at the start of the afternoon, winds will be in the range of 20 to 30 knots, and by the end of the afternoon, they will range from 15 to 20 knots. Exactly when and how winds are predicted to diminish during the afternoon is not specified. Based on the worded forecast alone, one may expect winds anywhere in both forecast ranges during the middle of the transitional period. Winds at 3 p.m. can range from 15 to 30 knots, and not contradict the forecast. This sort of interpretation is used to preprocess COMAR coded forecasts for use as input into SIGMAR.

System Design

SIGMAR is a diagnostic system, not a predictive system. SIGMAR provides information about the accuracy of the current forecast for given current marine forecasts and actual measured winds. Although the current marine forecast may raise false alarms, SIGMAR cannot. Users of SIGMAR are more likely to detect false alarms. When actual winds contradict the current forecast, SIGMAR consistently identifies and reports on such events.

SIGMAR continually critiques a marine forecast. It monitors winds, compares them to the range of forecast winds, and informs the forecaster when a marine forecast's accuracy deteriorates. SIGMAR requires as input a current valid marine forecast and continually updated wind information. The forecast wind strengths that SIGMAR uses include three values that are functions of time: minimum wind, maximum wind and average wind. The ambiguity of these values near transitional times is resolved by smoothing forecast winds across six hour time intervals centred about the hour specified in the COMAR code. Figure 8 shows how marine forecasts are taken from COMAR, and preprocessed for use as input for SIGMAR.

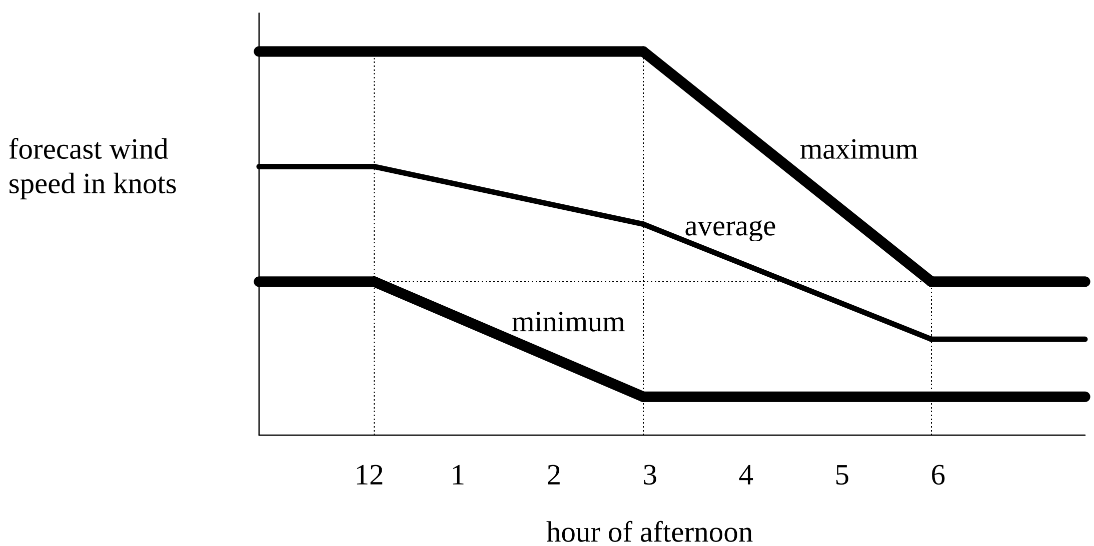

Figure 8. Forecast wind range as a function of time. Forecast corresponds to COMAR code of "W20-30. 15h W15-20" and worded forecast of "westerly 20 to 30 diminishing to westerly 15 to 20 this afternoon." Implied minimum and maximum values of forecast wind are depicted with bold lines.

The way the marine forecast is represented in SIGMAR corresponds to the intended meaning of the worded marine forecast. The forecast range is 20 to 30 knots at the start of the afternoon, and 15 to 20 at the end of afternoon. At mid-afternoon, when the forecast is most ambiguous, the forecast used by SIGMAR is a total union of early and late afternoon forecasts. The marine forecasts for all marine areas are taken from COMAR, processed in a way similar to that shown in Figure 8, and given to SIGMAR as input.

Hourly winds are available from about seventy stations around the MWC's marine district. Wind data are fairly unambiguous, and passed directly to SIGMAR as input. SIGMAR performs fuzzy analysis using the variables of forecast wind and series of recently observed wind data. Every set of wind observations is compared to the forecast for the marine area in which the observations are taken. In addition to being compared to the marine forecast, the observations are compared with each other to see whether significant trends are present. The ten factors evaluated for each station's set of wind observations are listed in Table 1.

1. Is the current wind speed strong compared to the forecast for now?
2. Is the current wind speed light compared to the forecast for now?
3. Is the current wind direction different from the forecast for now?
4. Is the projected wind speed strong compared to the forecast for three hours from now?
5. Is the projected wind speed light compared to the forecast for three hours from now?
6. Is the projected wind direction different from the forecast for three hours from now?
7. Have winds recently increased?
8. Have winds recently decreased?
9. Have winds recently changed direction?
10. Are recent reports for the station unexpectedly missing? (Malfunction?)

Table 1. Conditions evaluated by SIGMAR.

The answers to the questions in Table 1 are not expressed in "yes or no" terms, but with fuzzy numbers. This way one can prioritize the results. By attaching a priority to wind data, significant values can be flagged, and when indicated, an alarm can be triggered. Conversely, if SIGMAR gives a low value of significance to wind reports, the implication is that the forecast is remaining accurate (no news is good news).

Design of Fuzzy Sets

The behavior of SIGMAR is totally determined by the shape of the fuzzy sets used to evaluate wind data. Fuzzy sets are used to perform operations on variables of forecast and observed wind values. Forecast wind values used by SIGMAR include four variables: minimum, average, maximum wind speed, and direction of wind. The observed wind data consist of a recent series of reports, each report containing values of: average wind speed, direction of wind, and (sometimes) gust strength. If a gust value is reported, it is taken to represent the wind strength; otherwise, the average wind speed is used to represent the wind strength.

Fuzzy sets are designed to accommodate forecast and observed wind data. The fuzzy sets make reference to the *difference* between forecast and observed wind, not to the absolute values of wind. It is difference that is significant. By focusing on difference alone, we need only design one fuzzy set to accommodate any possible combination of absolute values of forecast and observed wind values. Sets are designed so that the *significance value* of a difference ranges from 0.0 to 1.0. The author has chosen two particular values to correspond to linguistic terms of significance: values above 0.0 and below 0.2 imply "slightly significant"; values above 0.8 imply "quite significant." These arbitrary thresholds enable output to screened. The objective is to isolate reports that contain critical information about the marine forecast. For this purpose, observations scoring below 0.2 are of little interest; they imply the observation fits the forecast. Reports scoring above 0.8 are significant enough to warrant the forecaster's attention.

Figure 9 shows how a fuzzy set used to evaluate to what degree an observation is *strong* when compared to a forecast range. Suppose a forecast says "20 to 30 knots." With reference to Figure 9, min=20 and max=30. The average wind speed forecast is 25 knots. The degree to which an observed wind speed is thought of as *strong*, compared to the forecast, can be evaluated along a continuum from 0 to 1. The number can be translated into linguistic terms. A set of wind observations is evaluated this way in Figure 10.

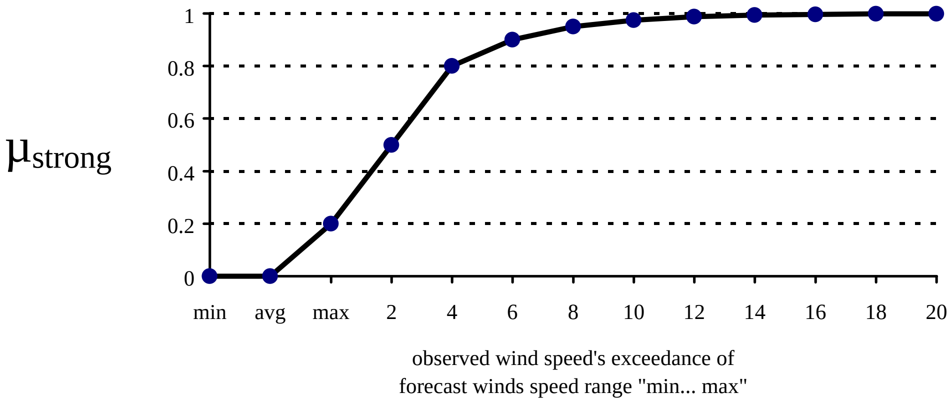

Figure 9. Function for membership in the category of strong winds for observed wind speed's exceedance of forecast range "min...max."

Speed	μstrong	Strong?
15	0.0	no
25	0.0	no
30	0.2	slightly
32	0.6	↓
34	0.8	quite
36	0.9	...
40	0.975	...

Figure 10. Degrees to which wind speeds are thought of as *strong* when forecast is "20 to 30 knots."

Suppose a forecast says “20 to 30 knots.” With reference to Figure 11, min=20 and max=30. The degree to which an observed wind speed is thought of as *light*, compared to the forecast, can be evaluated numerically and linguistically. A series of wind observations are evaluated this way in Figure 12.

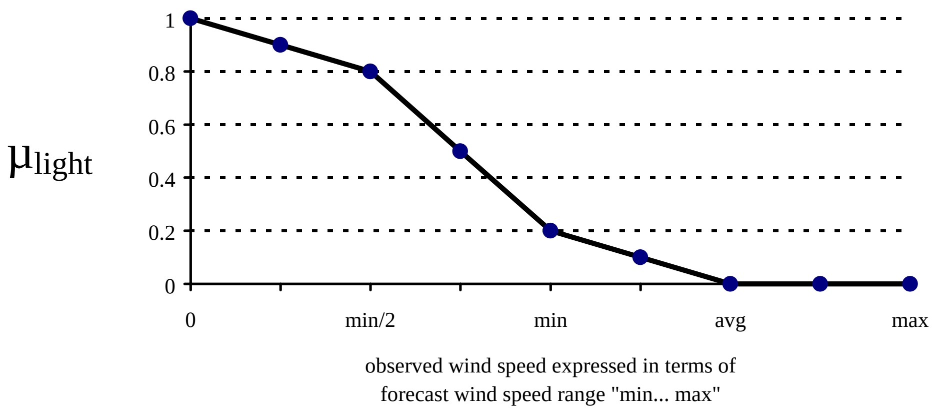

Figure 11. Function for membership in the category of *light* winds for an observed wind speed with forecast wind range "min...max."

Speed	μ light	Light?
5	0.9	
10	0.8	quite
15	0.5	↑
20	0.2	slightly
25	0.0	no
30	0.0	

Figure 12. Degrees to which wind is thought of as *light* when forecast is “20 to 30 knots.”

Suppose a forecast says “Northerly.” The degree to which an observed wind’s direction implies that the forecast wind direction is *off* (i.e. wrong) can be evaluated numerically and linguistically. A series of wind observations are evaluated this way in Figures 13 and 14.

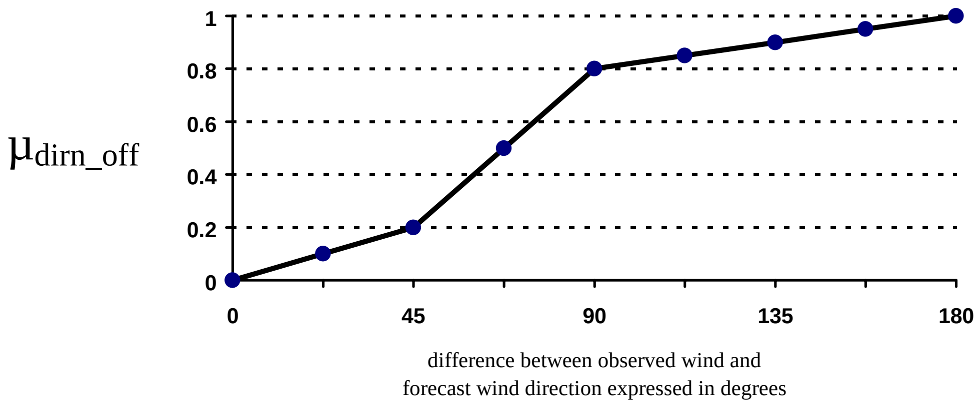

Figure 13. Function for membership in the category of winds whose direction suggests the forecast direction is *off*.

Direction	μ dirn_off	Direction off?
Northerly	0.0	no
northeast	0.2	slightly
easterly	0.8	quite
southeast	0.9	↓
southerly	1.0	

Figure 14. Degrees to which forecast wind directions are thought of as *off* when given directions are observed and forecast is “northerly.”

The sensitivity to discrepancies between forecast and actual winds is determined by the shape of the fuzzy sets shown in Figures 9, 11, and 13. These fuzzy sets are used to evaluate wind's present strength, lightness, and direction – factors 1, 2 and 3 in Table 1. SIGMAR also evaluates *trends* in these parameters. Simple linear extrapolations of observed winds are used to evaluate probable problems with the forecast three hours after the current time – factors 4, 5, and 6 in Table 1. The appropriate marine forecast for three hours after the present time is used. The same fuzzy sets are used for comparison of present and presumed future conditions and forecast.

A series of recent wind data in itself may contain significant information. If winds at a station are increasing rapidly, or have shifted direction sharply, this is worth knowing about. If wind reports have recently gone missing, it is significant (data outages require attention). Several other fuzzy sets are designed enable evaluation of these qualities of wind data – factors 7, 8, 9, and 10 in Table 1.

SIGMAR's results are presented to the user via three types of display:

1. Wind Alert. On a forecaster's Hewlett Packard (HP) workstation, there is a background menu option called "Strong Wind Alert". The forecaster selects this option and a window opens to display SIGMAR's reports about winds that are "quite strong" compared to the forecast. The report updates itself automatically.

2. Stations can be listed in order of decreasing significance. To access this display, a forecaster opens an HP command line window, and types "**sig** *number*" (*number* is the number of stations to display results for). Results for each station are reported in rows, and each row contains:

   significance index (2 $\Rightarrow$ slight, 8 $\Rightarrow$ quite)
   brief description of problem (e.g. wind strong)
   station identifier
   number of marine area (e.g. 5 $\Rightarrow$ George's Bank)
   station's observed wind direction and speed
   current valid forecast wind direction and speed
   hour extrapolation of station's wind direction and speed
   forecast wind direction and speed for three hours from now

3. A one-line summary is given for each of the marine areas. The type of information is basically the same as in the second type of display. Detailed information is only displayed for a single station – the station with the highest significance score among all the stations in the forecast area.

System Results

When SIGMAR was first introduced to the weather office, it was programmed to provide detailed analyses of wind data in tabular form. All ten factors listed in Table 1 were evaluated for each actual wind observation. The factors were sorted, and output in order of decreasing significance. The developer wanted the output to be detailed enough to show why specific significance was attached to the features of specific weather observations. The output demonstrated that the system performed reliably. For every actual weather observation, the ten potential causes for concern (Table 1) were evaluated correctly.

When presented with the system's comprehensive output, users commented: "Too many numbers." Users are already challenged with having to interpret screens full of data – SIGMAR appeared initially like more of the same. A streamlined version of SIGMAR was developed: the Wind Alert program. Wind Alert only gives concise reports of winds that are significantly stronger than forecast. If there are no such winds, the program reports that “winds are light enough.”

The Wind Alert version of SIGMAR is regarded by users as useful. Forecasters are most concerned when winds that are stronger than forecast occur. By design, Wind Alert’s output presents information that either causes concern or reassures. The output is concise: out of data that would fill several screens, one or two lines of significant information are extracted and displayed. The display is perceived by users as non-intrusive.

Wind Alert was first used during late July and early August 1996. During this trial, winds in the MWC’s district were generally light. This is unfortunate as far as testing the program is concerned. Wind Alert’s single purpose is to alert users when unpredicted strong winds occur. When forecasters received reports of strong winds, they found the reports useful. But there were only a few such instances during the trial period. Another factor hindered the testing of the system: when forecasters are confident that winds will not be very strong (as in late summer), concern about strong winds tends to lower. As a result, forecasters were less motivated to use SIGMAR. However, forecasters have continued to use SIGMAR since the initial trial. From October through December 1996, as winter and strong winds returned to the Maritimes, forecasters described the Wind Alert system as increasingly useful.

Forecasters say they find SIGMAR most useful at the start of their shifts. At these times, forecasters are unfamiliar with the details of the current forecast and actual winds. They must determine if, where, and how the current forecast is inaccurate; SIGMAR provides relevant information in a glance. Forecasters see which areas’ marine forecasts require their attention first.

SIGMAR’s detailed fuzzy analysis presents novel opportunities for forecast verification. For instance, it is quickly apparent that observed winds tend to be light compared to forecast winds. This has been a general perception, but SIGMAR’s output makes it quite obvious. More detailed examination of SIGMAR’s fuzzy analysis could yield a number of results that are not available via conventional forms of forecast verification. For instance, during the trial period, the ratio of “light winds” to “strong winds” was 3910:757. This ratio can be interpreted more deeply. One can examine the “degree” to which winds are *light* or *strong*. Special attention may be given to winds that are “strong” to a degree of 0.95 or more.

Acknowledgements

I thank Dr. Denis Riordan, who helped me to appreciate the principles of critiquing expert systems; Allan MacAfee, who gave me valuable help with computer programming; and all my colleagues at Environment Canada, who have supported my studies in various ways.

References

MacNeil, D., and Freiberger, P. 1993. Fuzzy Logic – The Discovery of a Revolutionary Computer Technology and How it Changing our World. Simon & Schuster, New York.

Zadeh, L.A. 1994. Fuzzy Logic, Neural Networks, and Soft Computing, Communications of the ACM, Vol. 37, No. 3, March 1994, pp.77-84.

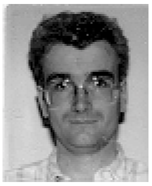

Bjarne Hansen received a B.Sc. in biology (1986) and Diploma in Meteorology (1987) from Dalhousie University in Halifax, N.S., Canada. He is currently a graduate student in the School of Computer Science at the Technical University of Nova Scotia. He has worked as a meteorologist at the Maritimes Weather Centre of Environment Canada in Bedford, Nova Scotia since 1989. He is currently on a research assignment with Climate and Atmospheric Research Directorate of Environment Canada. His research objective is to develop new types of expert systems for meteorological problems. He is particularly interested in fuzzy logic and artificial neural network techniques.

Hansen, B. K., 1997: SIGMAR: A fuzzy expert system for critiquing marine forecasts, *AI Applications*, Vol. 11, No. 1, 59-68. [This paper was reformatted for *AI Applications*.]