

# Weather Prediction Using Case-Based Reasoning and Fuzzy Set Theory

Bjarne K. Hansen and Denis Riordan \*

New Brunswick Weather Centre  
Meteorological Service of Canada, Environment Canada  
77 Westmorland Street, Suite 400  
Fredericton, New Brunswick, Canada E3B 6Z3  
bjarne.hansen@ec.gc.ca

Faculty of Computer Science  
Dalhousie University  
6050 University Avenue  
Halifax, Nova Scotia, Canada B3H 1W5  
riordan@cs.dal.ca

## Preamble

The work described here is the only line of research that combines case-based reasoning and fuzzy set theory in a weather prediction application. This paper summarizes a master of computer science thesis of the same name (Hansen 2000), available at [www.cs.dal.ca/~bjarne/thesis](http://www.cs.dal.ca/~bjarne/thesis).

## Abstract

A fuzzy logic based methodology for knowledge acquisition is developed and used for retrieval of temporal cases in a case-based reasoning (CBR) system. The methodology is used to acquire knowledge about what salient features of continuous-vector, unique temporal cases indicate significant similarity between cases. Such knowledge is encoded in a similarity-measuring function and thereby used to retrieve  $k$  nearest neighbors ( $k$ -nn) from a large database. Predictions for the present case are made from a weighted median of the outcomes of analogous past cases (i.e., the  $k$ -nn, or the analog ensemble). Past cases are weighted according to their degree of similarity to the present case.

Accordingly, a fuzzy  $k$ -nn based prediction system, called WIND-1, is proposed, implemented, and tested. Its unique component is an expertly-tuned fuzzy  $k$ -nn algorithm with a temporal dimension. It is tested with the problem of producing 6-hour predictions of cloud ceiling and visibility at an airport, given a database of over 300,000 consecutive hourly airport weather observations (36 years of record). Its prediction accuracy is measured with standard meteorological statistics and compared to a benchmark prediction technique, persistence climatology (PC). In realistic simulations, WIND-1 is significantly more accurate. It takes about one minute for WIND-1 to produce a forecast.

## Introduction

CBR and fuzzy set theory each have their own unique well-demonstrated strengths. So when both methods are combined in one system, the system stands the chance of inheriting the strengths of both methods. A recent increase in the number of such hybrid systems attests to the

effectiveness of combining CBR and fuzzy set theory (Pal et al. 2000).

CBR is recommended to developers who are challenged to reduce the knowledge acquisition task, avoid repeating mistakes made in the past, reason in domains that have not been fully understood or modeled, learn over time, reason with incomplete or imprecise data and concepts, provide a means of explanation, and reflect human reasoning (Main et al. 2000), and these are some of the challenges faced by developers of weather forecasting systems (Christopherson 1998).

Fuzzy logic imparts to CBR the perceptiveness and case-discriminating ability of a domain expert. The fuzzy  $k$ -nn technique retrieves similar cases by emulating a domain expert who understands and interprets similar cases. The main contribution of fuzzy logic to CBR is that it enables us to use common words to directly acquire domain knowledge about feature salience. This knowledge enables us to retrieve a few most similar cases from a large temporal database, which in turn helps us to avoid the problems of case adaptation and case authoring.

A fuzzy  $k$ -nn weather prediction system can improve the technique of persistence climatology (PC) by achieving direct, efficient, expert-like comparison of past and present weather cases. PC is an analog forecasting technique that is widely recognized as a formidable benchmark for short-range weather prediction. Previous PC systems have had two built-in constraints: they represented cases in terms of the memberships of their attributes in predefined categories and they referred to a preselected combination of attributes (i.e., cases defined and selected before receiving the precise and numerous details of present cases). The proposed fuzzy  $k$ -nn system compares past and present cases directly and precisely in terms of their numerous salient attributes. The fuzzy  $k$ -nn method is not tied to specific categories, nor is it constrained to using only a specific limited set of predictors. Such a system for making airport weather predictions will let us tap many, large, unused archives of airport weather observations, ready repositories of temporal cases. This will help to make airport weather predictions more accurate, which will make air travel safer and make airlines more profitable.

---

\* Copyright © 2001 Bjarne K. Hansen, Denis Riordan

## Weather Prediction

Weather prediction presents special challenges for CBR. Weather is continuous, data-intensive, multidimensional, dynamic and chaotic. These properties make weather prediction a formidable proving ground for any CBR prediction system that depends on searching for similar sequences.

Fundamentally, there are only two methods to predict weather: the empirical approach and the dynamical approach (Lorenz 1969). The empirical approach is based upon the occurrence of analogs (similar weather situations) and is often referred to by meteorologists as "analog forecasting." The empirical approach is useful for predicting local-scale weather if recorded cases are plentiful (e.g., cloud ceiling height and visibility in a few square kilometres around an airport). The dynamical approach is based upon equations and forward simulations of the atmosphere, and is often referred to as "computer modeling." Because of computer model grid coarseness, the dynamical approach, used by itself, is only useful for modeling large-scale weather phenomena (e.g., general wind direction over a few thousand square kilometers). In practice, most weather prediction systems use a combination of empirical and dynamical techniques.

## Airport Weather Prediction Problem

An airport weather prediction is a concise statement of the expected meteorological conditions at an airport during a specified period (U.S. National Weather Service Aviation Weather Center, 1999). An airport weather prediction is, in meteorology, commonly referred to as TAF, short for Terminal Aerodrome Forecast. When pilots describe weather or weather-related delays at a destination airport, they are reading TAFs.

TAFs are made by expert forecasters. These experts have general knowledge about how large scale weather systems behave and specific knowledge about how local scale weather phenomena behave idiosyncratically at specific airports. Experts bridge the gap between simple PC and computer-model-assisted statistical forecasting on the local scale (Battan 1984).

The three types of forecasts most commonly made by forecasters are TAFs, public forecasts and marine forecasts. Of these, TAFs are the most precise and thus the most challenging type of forecast to make, both in terms of measurable weather conditions and in terms of timing. Forecasts of the height of low cloud ceiling height are expected to be accurate to within 100 feet. Forecasts of the horizontal visibility on the ground, when there is dense obstruction to visibility, such as fog or snow, are expected to be accurate to within 400 metres. Forecasts of the time of change from one flying category to another are expected to be accurate to within one hour. In comparison, public and marine forecasts can be much less precise. For example, in public forecasts, it may be sufficient to predict "variable cloudiness this morning," and in marine forecasts,

it may be sufficient to predict "fog patches forming this afternoon."

The quality of TAFs is determined by the accuracy of the forecast weather elements and the timeliness of issue. The Meteorological Service of Canada measures TAF quality in four ways: with three ceiling and visibility accuracy statistics and with a speed-of-amendment statistic (Stanski et al. 1999). The commonest cause for amendments is unforecast ceiling or visibility (Stanski 1999). Timeliness refers to the time between TAF issue and the decision deadline of the intended user, or in other words, the amount of time the TAF can be used to affect critical decisions about flight scheduling.

When ceiling and visibility at a busy airport are low, in order to maximize safety, the rate of planes landing is reduced. When ceiling and visibility at a destination airport are forecast to be low at a flight's scheduled arrival time, its departure may be delayed in order to minimize traffic congestion and related costs. An examination of the causes and effects of flight delays at the three main airports serving New York City concluded that a correctly forecast timing of a ceiling and visibility event (i.e., a significant change) could be expected to result in a savings of approximately \$480,000 per event at La Guardia Airport (Allan et al. 2001). Based on a related study, the U.S. National Weather Service estimated that a 30 minute lead-time for identifying cloud ceiling or visibility events could reduce the number of weather-related delays by 20 to 35 percent and that this could save between \$500 million to \$875 million annually (Valdez 2000).

A goal of our research is to enable an improvement in the quality of TAFs in terms of accuracy and timeliness.

## System

The weather prediction system (WIND-1) consists of two main parts, a large database of weather observations and a fuzzy k-nn algorithm, described as follows.

### Large Database of Airport Weather Observations

The database is an archive of 315,576 consecutive hourly airport weather observations made at Halifax International Airport (CYHZ, located at 44°53'N 63°30'W) during the 36-year period from 1961 to 1996.

We acquire knowledge about how to recognize similar weather cases by interviewing a domain expert, a weather forecasting expert who is experienced at forecasting for CYHZ and is thus presumably more able than anyone else to, firstly, identify the attributes to be used to indicate similarity between cases and, secondly, describe degrees of similarity between such attributes. Twelve similarity-indicating attributes are identified: date, hour, cloud amount(s), cloud ceiling height, visibility, wind direction, wind speed, precipitation type, precipitation intensity, dew point temperature, dry bulb temperature, and pressure trend. All of these attributes are continuous except for precipitation, which is nominal (e.g., rain, snow, etc.).

## Fuzzy k-nn Algorithm


The fuzzy k-nn algorithm measures the similarity between temporal cases, past and present intervals of weather observations. The algorithm is tuned by interviewing an experienced forecaster who describes various attribute-difference thresholds that are to be used to signify various *degrees of similarity* (i.e., *very near*, *near*, *slightly near*).

We design a similarity-measuring function, *sim*, that is used to find k-nn for a present weather case and rank them according to their degree of similarity to the present weather. Given two cases, each identified by unique time indexes  $t_1$  and  $t_2$ , *sim* returns a real number proportional to the degree of similarity of the two cases such that:  $0.0 < sim(t_1, t_2) \leq 1.0$

Because all weather cases are unique and because the value of *sim* is calculated to double precision, *sim* can identify exactly  $k$  nearest neighbors. There are no null search results and no ties.

The three steps to construct and use the algorithm are:

1. Configure similarity-measuring function.
  2. Traverse case base to find k-nn.
  3. Make prediction based on weighted median of k-nn.
- Step 1 is performed only once and Steps 2 and 3 are performed every time a weather prediction is made. Step 1 is performed by interviewing a domain expert, and is thus the critical knowledge acquisition step in system design. For each continuous attribute,  $x_i$ , the expert specifies thresholds for considering two such attributes to be *very near*, *near*, and *slightly near* each other (Figure 1).



**Figure 1.** Fuzzy set for comparing continuous attribute. For each continuous attribute,  $x_i$ , the expert specifies values of  $c_{i1}$ ,  $c_{i2}$ , and  $c_{i3}$  that are thresholds for considering two such attributes to be *very near*, *near*, and *slightly near* each other, respectively. To measure the similarity of two cases, all their intercomparable attributes are compared using a combination of attribute-specific  $\mu$  functions across the whole time intervals of each case (Hansen 2000).

## Experiments

Each experiment consists of a forecasting scenario. Five sets of experiments are conducted. In each set of experiments we systematically change the fixed parameters of WIND-1 and measure the resultant effects on forecast

accuracy. The fixed parameters (independent variables) are: the attribute set, the number of analogs used to make forecasts, the size of the case base, and the fuzzy membership functions. The output (dependent variables) are, for each individual forecast, forecast values of cloud ceiling and visibility, and, for each set of experiments, a summary of the accuracy of all the forecasts made.

In each set of experiments, 1000 hours are chosen at random from the 1996 weather archive and are each used as an hour to produce a forecast for. So, in each set of experiments, 1000 simulated forecasts are produced. For purposes of comparison, the same 1000 randomly-chosen hours are used in each set of experiments. This is a control so that the effect of varying other input can be tested.

In each individual experiment, a case is taken from the 1996 data and is used as a present case. It is input to WIND-1. During the forecast process, the outcome of the present case is hidden from WIND-1. WIND-1 produces a forecast for the present case based on the outcomes of the k-nn in the case base, the  $k$  most analogous past cases for the present case. After the forecast process, the accuracy of the forecast is verified by comparing the forecast with the then unhidden outcome of the present case using standard meteorological quality control statistics (Stanski et al. 1999).

## Results and Interpretation

The first set of experiments varies the attribute set and shows that prediction accuracy increases as the number of attributes used for comparison increases.

The second set of experiments varies  $k$ , the number of nearest neighbors that are used as the basis of predictions ( $k = 1, 2, 4, 8, \dots, 256$ ) and finds that maximum accuracy is achieved with  $k = 16$ . This suggests that WIND-1 is effective at identifying and ranking nearest neighbors, or, in meteorological terms, it finds the best *analog ensemble*.

The third sets of experiments varies the size of the case base and shows that prediction accuracy increases as the size of the case base increases.

The fourth and fifth sets of experiments pit WIND-1 using non-fuzzy sets against PC, and WIND-1 using fuzzy sets against PC, respectively. The non-fuzzy based prediction method is only slightly more accurate than PC, and fuzzy k-nn based prediction method is significantly more accurate than PC. The only variation between the two methods is the nature of the membership functions used to compare attributes of cases. The fuzzy k-nn method uses fuzzy membership functions that span certain ranges around the case being forecast for, whereas the non-fuzzy method uses 0-1-0 functions centered across the same ranges (and thus implements the benchmark PC forecasting method). This suggests that, compared to the accuracy of PC, the significantly higher accuracy of fuzzy k-nn based forecasts is attributable to the use of fuzzy sets to measure similarity as opposed to using crisp sets. To the best of our knowledge, all previous meteorological analog forecasting systems have used only crisp sets to measure similarity between weather cases.

## Conclusions

Based on our literature review, experiments, and the results presented in our thesis, we conclude that querying a large database of weather observations for past weather cases similar to a present case using a fuzzy k-nearest neighbors algorithm that is designed and tuned with the help of a weather forecasting expert can increase the accuracy of predictions of cloud ceiling and visibility at an airport.

Of significance to CBR: We have shown how fuzzy logic can impart to CBR the perceptiveness and case-discriminating ability of a domain expert. The fuzzy k-nn technique described in this thesis retrieves similar cases by emulating a domain expert who understands and interprets similar cases. The main contribution of fuzzy logic to CBR is that it enables us to use common words to directly acquire domain knowledge about feature salience. This knowledge enables us to retrieve a few most similar cases from a large temporal database, which in turn helps us to avoid the problems of case adaptation and case authoring.

The fuzzy k-nn algorithm, even though it is of approximate Order( $n$ ) complexity, makes superior predictions with practical speed—with less than one minute of computation. This speed is achieved by strategically ordering the steps in a case-to-case similarity-measuring test and by stopping any test as soon as a step reveals that a case is dissimilar enough to be ruled out of the k-nn set without the need for further tests. For example, suppose we have a database of  $n$  past temporal cases. And suppose each case is described by  $m$  attributes and is  $p$  time units long, thus each case is described by  $m \cdot p$  attributes. To measure the similarity of every past case, we would need to perform  $n \cdot m \cdot p$  individual tests. However, we are only interested in finding the  $k$  most similar cases, and most cases can be ruled out of contention with a single test. So, the number of tests we need to perform is much closer to the order of  $n$  than it is to the order of  $n \cdot m \cdot p$ .

Of significance to meteorology and the aviation industry: Such a fuzzy k-nn weather prediction system can improve the technique of persistence climatology (PC) by achieving direct, efficient, expert-like comparison of past and present weather cases. PC is a sort of analog forecasting technique that is widely recognized as a formidable benchmark for short-range weather prediction. Previous PC systems have had two built-in constraints: they represented cases in terms of the memberships of their attributes in predefined categories and they referred to a preselected combination of attributes (i.e., defined and selected before receiving the precise and numerous details of present cases). The proposed fuzzy k-nn system compares past and present cases directly and precisely in terms of their numerous salient attributes. The fuzzy k-nn method is not tied to specific categories nor is it constrained to using only a specific limited set of predictors. Such a system for making airport weather predictions will let us tap many, large, unused archives of airport weather observations, ready repositories of temporal cases. This will help to make

airport weather predictions more accurate, which will make air travel safer and make airlines more profitable.

We plan to pursue this research and improve WIND-1 by: testing its prediction accuracy at other airports; enabling it to learn autonomously; and incorporating additional predictive information, such as user-provided hints, projections of weather radar images of precipitation, projections of satellite images of cloud, and guidance from large-scale computer models of the atmosphere.

## Acknowledgments

Thanks to the many meteorologists in the Meteorological Service of Canada (MSC) who guided and supported us during the pursuit of this research, and thanks to MSC for providing the necessary data and computing facilities and for paying for a large part of the costs.

## References

- Allan, S. S.; Gaddy, S. G.; and Evans, J. E. 2001. Delay Causality and Reduction at the New York City Airports Using Terminal Weather Information Systems, Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, Mass., <http://www.ll.mit.edu/AviationWeather/atc-291.pdf>, downloaded May 1, 2001.
- Battan, L. J. 1984. *Fundamentals of Meteorology*, Eaglewood Cliffs, New Jersey, Prentice-Hall. 157-161.
- Christopherson, D. 1998. Artificial intelligence in the weather forecast office—One forecaster's view, Preprints of the First Conference on Artificial Intelligence, American Meteorological Society, 136-143.
- Hansen, B. K. 2000. Weather Prediction Using Case-Based Reasoning and Fuzzy Set Theory, Master of Computer Science thesis, Technical University of Nova Scotia, Halifax, Nova Scotia, Canada.
- Lorenz, E. N. 1969. Three approaches to atmospheric predictability, *Bulletin of the American Meteorological Society*, 50, 345-349.
- Main, J.; Dillon, T. S.; and Shiu, S. C. K. 2000. A tutorial on case-based reasoning; in Pal, S. K.; Dillon, T. S.; and Yeung, D. S. eds. 2000. *Soft Computing in Case Based Reasoning*. London, UK, Springer.
- Pal, S. K.; Dillon, T. S.; and Yeung, D. S. eds. 2000. *Soft Computing in Case-Based Reasoning*. London, UK, Springer.
- Stanski, H. 1999. Personal correspondence.
- Stanski, H.; Leganchuk, A.; Hanssen, A.; Wintjes, D.; Abramowski, O.; and Shaykewich, J. 1999. NAV CANADA's TAF amendment response time verification, Eighth Conference on Aviation, Range, and Aerospace Meteorology, 10-15 January 1999, Dallas, Texas, American Meteorological Society, 63-67.
- U.S. National Weather Service Aviation Weather Center, What is a TAF?, [http://www.awc-kc.noaa.gov/awc/help/what\\_is\\_a\\_taf.html](http://www.awc-kc.noaa.gov/awc/help/what_is_a_taf.html), downloaded January 29, 1999.
- Valdez, J. 2000. NWS Reinventing Goals for 2000, National Weather Service, National Oceanic and Atmospheric Administration, United States Department of Commerce, <http://www.nws.noaa.gov/npr5.html>, downloaded May 1, 2001.