A Knowledge-Based System for the Diagnosis and Prediction of Short-Term Climatic Changes in the North Atlantic

SERGEI RODIONOV AND JAMES H. MARTIN

Computer Science Department, University of Colorado, Boulder, Colorado

(Manuscript received 4 April 1995, in final form 15 January 1996)

ABSTRACT

Understanding and predicting climate change is the key problem in climatology. The most well-accepted current approach to this problem involves the development of general circulation models (GCMs). This approach is based on modeling fundamental physical principles in large computer programs. At the same time, however, an increasingly large proportion of the available information regarding the climate system exists in the form of heuristics, or empirical rules of thumb. The objective of the CESNA (Climatic Expert System For the North Atlantic) project is to develop a practical system that can manipulate this qualitative information in such a way as to facilitate insights into observed and anticipated climate changes. The methods used to reach this objective are based on concepts and techniques derived artificial intelligence research on representing and reasoning with uncertain knowledge. A recently completed evaluation of the prototype CESNA measured how well it could predict the sea temperature of the Kola section of the barents sea for the period 1965 to 1991 with a one-year lead time. The system's predictions paralleled the observed temperatures with remarkable accuracy. Similar results were obtained for two other regions, the northwest Atlantic and the southeastern United States. Qualitatively, these experiments show that even though some rules may be poor predictors in a given year, the combined evidence from the remaining rules results in an accurate prediction.

1. Introduction

Understanding and predicting climate change is the key problem in climatology. The most well-accepted current approach to this problem involves the development of general circulation models (GCMs). This approach is based on modeling fundamental physical principles in large computer programs. At the same time, an increasingly large proportion of the available information regarding the climate system exists in the form of heuristics, or empirical rules of thumb. Numerous articles presenting statistical analyses, detailed case studies, and more anecdotal personal experiences comprise a source of this sort of information. Unfortunately, for both practical and theoretical reasons, it is quite difficult to incorporate such information into numerical models.

The objective of the CESNA (Climatic Expert System For the North Atlantic) project currently underway at the Computer Science Department of the University of Colorado at Boulder is to develop a practical system that can manipulate qualitative information in such a way as to facilitate insights into observed and anticipated climate changes. The methods used to reach this objective are based on concepts and techniques from artificial intelligence (AI). Such methods usually materialized in the form of knowledge-based systems, or expert systems. Expert systems are software systems designed to "reason" as one or more human experts would within their own area of expertise in order to solve a problem or give advice. This is accomplished through the use of declarative representations of expert knowledge (instead of explicit procedures). Such systems are typically used to deal with problems that are often poorly understood, for which there is no crisp algorithmic solution, and that can benefit from some sort of symbolic reasoning (Buchanan and Shortliffe 1984; Rolston 1988; Tsai and Weigert 1993).

There have been a number of previous efforts to apply AI techniques to problems in meteorology. These efforts have been directed toward the development of expert systems for weather predictions (operational forecasts). Applications areas have typically involved situations such as thunderstorm forecasting, where there is little time to provide a more detailed analysis (Conway 1989; Moninger 1990). In cases like this, an expert system plays the role of a single human expert who might not be available at the moment. Moreover, such systems make their judgements on the basis of an extremely limited and carefully circumscribed amount of information.

The general framework of CESNA is quite similar to these systems. Knowledge about the climate is rep-

Corresponding author address: Dr. James H. Martin, Department of Computer Science & Institute of Cognitive Science, University of Colorado, Campus Box 430, Boulder, CO 80309-0430.

resented as a large set of interrelated macroclimatic objects, such as centers of action, upper ridges and troughs, jet streams, temperature anomalies in key regions, extension of polar ice cover, precipitation patterns, etc. Empirical rules describing the relationships among the macroclimatic objects comprise the knowledge base of the system. However, the CESNA system differs from these meteorological systems and more traditional expert systems in two ways: 1) the individual rules used in the system have a high degree of uncertainty when used for prediction and 2) the knowledge base represents a superset of the information that would normally be used in creating a prediction or explanation by any single climate expert. It is through the careful use of these two facts that CESNA operates.

In preparing a forecast CESNA exploits several features of the climate system and external factors that make it possible to project the current state of climatic parameters into the future.

- *Cycles,* for example, the 11-year cycle of solar activity, or quasibiennial oscillation (QBO) of stratospheric winds over the equator. There are also numerous other cycles found in variations of climatic characteristics, but their confidence factors are much lower (Burroughs 1992).
- *Persistence,* which can be caused by enormous thermal capacity of the ocean or by positive feedback loops in the large-scale air–sea interaction and may exhibit itself in so-called climatic regimes (Namias 1982).
- *Teleconnections* with a significant time lag. One of the most appealing predictors here is the El Niño signal. The observed lag correlations between sea surface temperature (SST) anomalies in the equatorial Pacific during the middle months of the year and atmospheric teleconnection patterns in the Northern Hemisphere during the following winter are strong enough to be of some practical value in predicting wintertime temperature anomalies one or two seasons in advance (Wallace and Jiang 1992).

It is known that climatic forecasts based on empirical relationships are not very reliable. Correlation coefficients between climatic variables may change over time, amplitude of a cycle may reduce, and an obvious trend may halt or change its direction. Nevertheless, each of these rules gives small but valuable information about future state of the climate system. What is important is to collect as many such rules as possible and put them into one system to provide a necessary ''critical mass'' of evidence. The goal is that even if some of the rules are quite weak, the overall conclusion can be quite strong if the evidence is combined in the right way. A related, and no less important, goal is that the system must then be able to display and explain the chain of reasoning used to arrive at its conclusion.

2. An overview of knowledge-based systems

The problem of capturing an expert's knowledge of a domain has traditionally been approached by encoding the knowledge in the form of a collection of antecedent–consequent rules. In their purest form, such rules correspond to a set of logical implications. Within this approach there are two natural ways of encoding the necessary domain knowledge: causal rules and diagnostic rules. Causal rules model the relationship of events or states to their results in a direct fashion: inferring the results of known causes. Diagnostic rules work in the opposite fashion: allowing the reasoner to infer possible causes from observed states.

Once a body of domain knowledge has been collected, a logically correct and computable form of inference is used to conclude useful new facts as needed. For efficiency reasons, inference in such systems is typically implemented as some form of simple forward or backward chaining through the rules.

Unfortunately, in practice the vast majority of expert rules are heuristic rules of thumb. They are true often enough to be useful but they are neither logically correct nor complete. What is needed is a way to capture the inherent degree of uncertainty in the bulk of expert domain knowledge. Over the past 20 years of AI research, a number of formalisms have been developed to deal with this problem. These formalisms all attempt to capture uncertainty through the use of what amounts to probabilistic reasoning. In addition, however, they all attempt to do so in a way that remains computationally tractable. The most common and well-studied methods include Bayesian Belief Nets (Pearl 1988) and Certainty Theory (Gordon and Shortliffe 1984). CESNA uses a variant of Certainty Theory.

In this approach, both ordinary atomic propositions and implication rules are augmented with numeric certainty factors (CFs) representing an expert's degree of belief. In addition to these numeric values, a CF-based system must specify a set of combination rules that parallel logical inference rules. The most important of these involve chaining and evidence combination.

To make this discussion more concrete, consider the following CESNA rule linking central pressure and geographical position of the Icelandic low (Glowienka 1985):

```
if central_pressure = low

then position = shifted_north, CF = 60.
```

This rule is based on statistical analysis between these two variables and reflects a general tendency for the Icelandic low to be shifted northward when its central pressure is anomalously low. Since this rule does not work in every case, some degree of uncertainty must be associated with this rule. In this case, the certainty factor 60, attached to the conclusion of the above rule, indicates that the conclusion is drawn with 60%

confidence. In our system, confidence factors vary from 0 to 100. A certainty factor with a value of zero indicates that we do not believe in the rule at all, that is, even if we know the premise of the rule, our forecast based on this rule will not differ from a climatological forecast based on statistical distribution of the conclusion of the rule alone. CF = 100 indicates knowing the premise of the rule gives us 100% confidence that the conclusion will occur.

Continuing with our example, the rule assigns a confidence factor of 60 to the conclusion value `shifted_north` if the variable `central_pressure` is assigned the value low. But the certainty factor in this conclusion assumes that a confidence factor of 100 has been assigned to the value of the premise. Suppose instead that for some reason, we do not know for sure that the atmospheric pressure was anomalously low and we assign a confidence factor of 50 to this value. In such a case, the confidence factor of the conclusion is derived by multiplying the CF of the premise by the CF of the rule itself: $(50 \times 60)/100 = 30$. This technique is the fundamental means for propagating CFs through a chain of rules.

In the above rule, this CF = 60 is based on an observed correlation between the central pressure of the Icelandic low and its geographical position. However, it is important to note that such certainty factors may reflect a subjective, rather than statistical, value; it may reflect a personal belief in correctness of the statistical analysis, whether or not it has been confirmed by other independent studies, model experiments, and so forth.

In addition to chaining, CF-based systems must be capable of computing the certainty of a proposition that is supported by two or more distinct chains of rules. Certainty theory provides a simple algorithm for these computations, which is clear from the following example. Suppose that a conclusion (c) is known to be supported by two rules with the certainties CF1 and CF2:

```
Rule 1
      if (e1), then (c), CF1 = 70;
Rule 2
      if (e2), then (c), CF2 = 30.
```

The combination rule in the certainty theory produces a combined answer as follows:

```
CFcomp = CF1 + CF2
                          − (CF1 × CF2)/100 = 79.
```

Now, suppose a third new piece of evidence is discovered that activates a third rule in support of the same conclusion:

```
Rule 3
      if (e3) then (c), CF3 = 30.
```

An updated estimate of the certainty of the conclusion (CFcompnew) will be

```
CFcompnew = CF3 + CFcomp −
                   (CF3 × CFcomp) = 85.
```

The remaining issue to be dealt with is the issue of where the CFs come from. There are two aspects in assigning confidence factors to rules and propositions. First, the available data on climatic variables are often imprecise, incomplete, vague, or uncertain. For example, there are several indices that can be used to characterize the strength of the midlatitude westerly winds over the North Atlantic. Therefore, the result can be somewhat different depending on which index was used. In addition, some data might not be available for the moment, and only indirect evidence can be used, or it may be a predicted, rather than an observed, value. In all these cases, the CFs reflect the degree of certainty one has in the value of the climatic variable considered.

Another situation occurs when there is enough quantitative information to determine the value of a climatic variable. However, the variable is not well pronounced and just slightly deviates from normal. In this case we also cannot assign 100 for the corresponding confidence factor. We consider the variable as being fuzzy (Zadeh 1965), with the confidence factor reflecting its degree of membership. The membership function is determined by the range of variation of a climatic variable, being 100 for all the values exceeding one standard deviation. Intermediate values of the confidence factor are calculated by linear interpolation between zero and the standard deviation.

Generally, CFs for climatic variables represent a mixture of these two interpretations. In practice, however, those climatic variables that have already been observed by the time of issue of the forecast are interpreted as fuzzy sets. In contrast, for climatic rules, CFs characterize our belief in them, and for forecasting rules, such as those based on cycles and lag relationships, CFs usually do not exceed 15 on a scale from 0 to 100. The result of this fact is that few of the rules can by themselves lend a high degree of confidence to any variable. It is only by chaining rules and combining multiple sources of evidence that a conclusion reaches a significant level of confidence.

3. Climatic Expert System for the North Atlantic (CESNA)

A prototype of CESNA has been developed using the VP-Expert system. VP-Expert is an expert system development tool that provides an inference engine, evidence combination methods, a graphical user interface, and rudimentary database tools.

The rules in CESNA are divided into seven separate sets, or knowledge bases. The first six are 1) solar activity, 2) global characteristics, 3) El Niño, 4) time series, 5) lag relationships, and 6) the North Pacific. These comprise the system's knowledge about climatic factors that can effect the North Atlantic climate. The final, seventh, set contains the rules describing relationships between climatic variables within the North Atlantic region itself. The North Atlantic region includes eastern North America, the North Atlantic and adjacent Arctic seas, and Europe. The system currently contains 420 such rules. The following sections provide brief overviews of the contents of each rule subset.

a. Set 1: Solar activity

Solar activity is one of the most fundamental quantities in relation to the terrestrial climate. It can be characterized by such parameters as sunspot numbers, phase of the 11-year solar cycle (ascending or descending branch), magnetic polarity of sunspots (even or odd cycle in Zurich numeration), umbral–penumbral ratio, and general level of solar activity (Coffey 1989). The latter, for example, can be measured by the length of the solar cycle (Friis-Christensen and Lassen 1991). Forecasts of solar activity are based on the existence of the 11-year cycle in its fluctuations and are available on the Internet from the National Geophysical Data Center.

The rules in this set describe the effect of solar activity on midlatitude westerlies (Arora and Padgaonkar 1981), winters in central Europe (Lamb 1972), precipitation (Girskaya 1987), sea ice (Nikiforov and Shpaiher 1980), temperature and cyclonic activity in the Barents Sea (Bochkov and Seliverstov 1978), and other climatic variables.

As an example of a rule from this subset, consider the effect of solar activity on El Niño described as following:

```
  if solar_activity = decreases and
     solar_activity = low and
     sunspot_gradient = small
then El_Nino = yes, CF = 6.
```

This rule is based on findings of Mendoza et al. (1991), who analyzed the occurrence of El Niños in coastal Peru as related to solar activity. They found that El Niño events tend to occur for small negative gradients (differences in annual sunspot numbers) and low sunspot numbers, conditions that in the sun correspond mainly to the descending phase of the solar cycle and around the minimum.

b. Set 2: Global characteristics

This set consists of those rules that have to do with climatic variables that are global rather than regional in nature. Variables of this class include stratospheric winds, circumpolar vortex, and surface and tropospheric temperature over the Northern Hemisphere among others. Of particular importance here are winds in the stratosphere over the equator that exhibit quasi-biannial oscillations (QBO). The regular reversal of these winds appears to be the best-established periodicity in climatic variables for frequencies lower than one cycle per year. Westerly and easterly winds alternate in an oscillation of around 28 months that dominates all seasonal and lesser variations.

In spite of the ubiquity of the QBO in surface weather records, there is no adequate physical explanation as to how the QBO in the stratosphere is linked with all the fluctuations of similar duration. There are, however, some empirical forecasting rules supported by high statistical correlations that allow us to use information on QBO to estimate atmospheric circulation and temperature patterns in the lower troposphere. For example, during the eastern phase of equatorial stratospheric winds (at 50 mb) frequency of tropical storms is low (Gray and Sheaffer 1991).

Modulation of effects of other forces is also important. Labitzke and van Loon (1990) proposed the following forecasting rule: when the sun is at its most active and the stratospheric wind at 45 mb is in the west phase, the pressure will be higher than normal over North America and lower than normal over the Pacific and Atlantic Oceans. Such anomalous pressure patterns play a major role in extreme seasonal weather. When pressure is high over North America in winter, cold northerly winds will sweep down the eastern seaboard. So in west-phase years we should expect to see cold winters on the East Coast when the sun is most active. In the CESNA this is encoded as

```
  if solar_activity = high and
     stratospheric_winds = west_phase
then SE_US = cold, CF = 12.
```

c. Set 3: El Niño

One of the most prominent sources of interannual variation in weather and climate around the world is the El Niño–Southern Oscillation (ENSO) phenomenon. All the rules in this set can be divided into two groups. The first group of rules is used to establish whether or not El Niño (or La Niña, which is considered here as just negative El Niño) will occur later in the year of forecast issue. Some judgement about possibility of El Niño has already been made based on information on solar activity (set 1). Other rules link occurrence of El Niño with preceding temperature and precipitation patterns in the western equatorial Pacific and regions surrounding the Indian Ocean (Kiladis and Diaz 1989). CESNA also takes into account various El Niño forecasts that can be found in the Climate Diagnostic Bulletin regularly published by the Climate Prediction Center. The performance of five ENSO prediction systems is examined by Barnston et al. (1994).

As soon as possibility of El Niño is established, the system considers what effect this can have on climate in the Northern Hemisphere extratropics. The effect is more clear in the Pacific North American sector and has been discussed in numerous observational and experimental studies. It is particularly noticeable for the Aleutian low, Pacific/North American (PNA) circulation index, sea surface temperature and upwelling along the American seaboard, and surface air temperature in western Canada and the southeastern United States.

El Niño's effect on the North Atlantic climate is less obvious. However, some helpful relationships have been established. For example, Fraedrich and Muller (1992) analyzed the ENSO signal for the winter seasons (December to February) at the end of the year of the event when the midlatitude circulation reveals the strongest response. They found that the mean surface temperature anomaly fields show negative (positive) temperature deviations over northern and northwestern Europe as a response to warm (cold) ENSO events.

d. Set 4: Time series analysis

This set consists of rules that make use of internal regularities, such as cycles, in the historical time series. There seem to be no periods in the range from 2 to 20 years that have not been found in climatic variables. Unfortunately, for the overwhelming majority of them their energy of fluctuations is very low. Some cycles, however, are strong enough to make contribution into assessment of future climatic state. Among those are alternation of cold and warm winters in central Europe (Lamb 1972) and the United States (Dettinger et al. 1995), quasibiennial periodicity in the North Atlantic Oscillation (van Loon and Rogers 1978), 9–12-yr cycle in fluctuation of sea ice in the Labrador Sea (Deser and Blackmon 1993), and others. In addition to rules extracted directly from the published literature, this set includes rules based on Box and Jenkins models that proved to be useful in analysis and prediction of times series of different climatic variables (Privalsky 1985).

Quite often climatic cycles can be asymmetric, with phases of different longevity. For example, one of the basic features of iceberg number variability in the Northwest Atlantic is a tendency for high annual iceberg numbers to occur in groups of 3 or 4 consecutive years, with local minimum counts tending to occur at intervals of 4 to 9 years (Marko et al. 1994). Some climatic variables demonstrate fluctuations in the form of abrupt transition from one relatively stable regime of fluctuations to another. Thus, Dickson and Namias (1976) noted an existence of climatic regimes on the order of about 10 years in surface air temperature fluctuations in the southeastern United States. These climate regimes have much in common with those in the Barents Sea (Rodionov and Krovnin 1992) and the Bering Sea (Rodionov and Krovnin 1991). It implies that if a transition has occurred and a climatic regime has established, one may assume that it is more likely for temperature anomaly to be of the same sign next year.

The system also makes use of decadal and longer-term climatic variability in terms of increased probability of a particular event. Thus, the rule

```
if year > 1992 and year < 2000,
then circulation_pattern
                         = meridional, CF = 5,
```

describes the results obtained by Shirley (1988), who found that because of the position of the sun, the period after 1992 and until the end of the century will be characterized by predominance of the meridional atmospheric circulation pattern. Decadal climate regimes (background memory) are now part of the Climate Prediction Center's effort in seasonal forecasting.

e. Set 5: Lag relationships

This set is comprised of those rules that describe relationships between climatic variables when one of them (predictand) has an appreciable time lag in its changes compared to another one (predictor). For example, Deser and Blackmon (1993) found a link between winter sea-ice concentration anomalies in the Davis Strait–Labrador Sea region and a time series of the second EOF of winter (November–March) SST. The latter is characterized by a dipole pattern with anomalies of one sign east of Newfoundland and anomalies of the opposite polarity off the southeast of the United States. If the two time series are superimposed, the maxima in sea-ice concentration precede the maxima in SST by 1 to 2 years. The correlation between the two time series is -0.26 at 0-lag; -0.62 when sea ice leads SST by 1 year; -0.76 when sea-ice leads by 2 years; and -0.62 when sea ice leads by 3 years. The strong lag correlations indicate that winters of heavy sea ice in the Labrador Sea precede winters of colder than normal SSTs east of Newfoundland.

In many cases these lags in relationships are caused by gradual transportation of anomalies of temperature, salinity, or other properties along the ocean currents. It is plausible in the above example that sea-ice anomalies in the Labrador Sea advected southward, resulting in colder than normal SSTs east of Newfoundland in the following (or second) year. Another prominent (and, at the same time, controversial) example is the "Great Salinity Anomaly" that traveled along the subpolar gyre from 1968 until 1982 (Dickson et al. 1988).

f. Set 6: The North Pacific

The importance of North America for the climatology of the North Atlantic is well known (Dickson and Namias 1976). North America, in turn, is under strong influence of the processes in the North Pacific. Therefore, it is not surprising that climatic processes in the North Pacific have a remote effect on the processes in the North Atlantic. Lamb (1972), for example, showed that if waters in the Namias region (30°–45°N, 155°–175°W) are warmer than normal in winter (DJF), there is a tendency for above-normal pressure over broad belt across Arctic Canada and Greenland toward the British Isles. This rule was confirmed later by Esbensen (1984). His one-point correlation map shows that 500–700-mb height anomalies in the east central North Pacific are positively correlated with 500–700-mb height anomalies west of the United Kingdom. Apparently, this high pressure cell west of the United Kingdom is indicative of blocking situations in this region.

Another good example is an opposition in sea level pressure (SLP) fluctuations between the Aleutian and Icelandic lows. Wallace and Gutzler (1981) found this opposition so strong that they suggested to consider the difference between SLP in these two centers of action as another index of the North Atlantic Oscillation (NAO).

g. *Set 7: The North Atlantic*

This is the largest (236 rules) and, presently, most developed set of rules. Five groups of rules can be distinguished in this set.

Atmospheric circulation—this group includes the rules describing the relationships between such macroclimatic variables as the Icelandic low and Azores high, central pressure in these centers of action and their geographical position, westerlies and stormtracks, cyclonic activity and position of the atmospheric polar front, upper trough over eastern seaboard of North America, and frequency of blocking situations in the Northeast Atlantic, and so forth.

Thermal advection—rules in this section describe changes in temperature over Europe and eastern North America caused by advection of cold and warm air.

Oceanic processes—this group comprises the rules for such oceanic processes as position and intensity of major ocean currents, oceanic polar front, sea surface temperature anomalies, ice cover, and others.

Ocean–atmosphere interaction—the rules describe large-scale interaction between oceanic and atmospheric processes in the North Atlantic.

Teleconnections—the most well-known teleconnection is the seesaw in winter temperatures between Greenland and Northern Europe. Some other teleconnection patterns (such as in-phase temperature fluctuations between European Arctic seas and the Sargasso Sea, between southeastern United States and northern Africa, and others) are also included.

4. How the system works

A typical consultation session begins with a human–computer dialog where the system asks questions and the user answers those questions based on information provided by the system and his–her personal experience. The problem-solving method used by the CESNA inference engine is called "backward chaining." In each rule set the inference engine starts by identifying the goal variable(s) and then moves through a sequence of rules until it finds a value that can help it assign a value to the goal variable. In each rule set there may be one or more such goal variables. Whenever the inference engine cannot find a variable in the conclusion of a rule, it first tries to get information from the database. If the information is not found or it is marked as preliminary, CESNA asks the user. In turn the user can ask why the question was asked. As part of this explanation, CESNA provides the user with a detailed description of the rule under consideration, along with other relevant text and graphical information that can help the user to answer the question.

To be more specific, the system system begins by asking the user for the year for which a prediction is going to be made. Currently, a lead time for the forecast is one year. CESNA then systematically processes the rules in the rule sets in the order given above. When the consultation session comes to the last, North Atlantic, set of rules, CESNA asks the user to specify the climatic variables and region that they are interested in and comes up with the final assessment of this variable. Since the climatic system is very complex and there are so many rules involved, the final result usually contains opposite categories of the variable with different confidence factors attached. For example, if the user is interested in thermal conditions in the southeastern United States, the possible answer may look like this:

```
''in [the year of forecast], thermal conditions in the southeastern United States
are expected to be:
COLD CF = 25,
WARM CF = 7.''
```

An important advantage of the system is the ability to trace back the line of reasoning and find out how one arrives at each of the possible conclusions. During the consultation, CESNA records the path of the inference engine and stores it in two files. The first file is a graphic file and represents a logic tree that was used during the current consultation. The second file is a text file and represents a report prepared by the system. The report contains references to all the rules that were used during the consultation, preliminary results for each set of rules, and the final result. Having this report, the user can run CESNA again in the regime of information retrieval. This regime has four options: a rule description, climate of the year, time series, and a rule inventory. The first option gives the user detailed textual and graphical descriptions of each rule in the system. The second option provides information about characteristic features of climatic situation in each year. The third option allows the user to retrieve historical time series of various climatic variables, along with analyses of their interannual variation. The final option gives the user the ability to easily manipulate the inventory of rules in a number of ways. Users can experiment with the system by modifying, adding or deleting rules from the knowledge and observing the effect (or lack of effect) on the final conclusion.

To test our approach, we used CESNA to produce experimental forecasts of winter and annual climatic conditions in the Barents Sea. An acknowledged indicator of these conditions is sea temperature in the upper 200 m of the Kola section. The forecasts were given with a 1-yr time lead for each year from 1965 to 1993. Since the mean winter and annual sea temperatures are highly correlated (correlation coefficient $r = 0.73$), we will consider only annual values. As a forecast temperature index we used the difference between confidence factors for warm and cold gradations.

As seen in Fig. 1, there is strong parallelism between the index and observed annual temperatures in the Kola section. The correlation coefficient between these two variables is $r = 0.63$. If we take into account only the sign of anomaly, the skill score of the forecast, calculated as the relationship of the number of correct forecasts to the total number of forecasts given, will be 89%, which significantly outperforms forecasts based on persistence (skill score 73%). In fact, there were only two years when the sign of forecast did not coincide with the sign of observed temperature anomaly: 1971 and 1989. Both of these years were characterized by a significant intraannual variations, so that winter and annual temperature anomalies had different signs.

We also tested CESNA for two other regions, the Northwest Atlantic and southeastern United States. The forecasts were given one year in advance for the period 1965–1995. The results were compared with observed air temperature anomalies at Gander (Newfoundland) and Charleston (South Carolina), respectively. Although the skill score of the forecast for the Northwest Atlantic was high (83%), the low-frequency temperature variations in this region were so strong that the forecast based on persistence had the same skill. As for the southeastern United States, CESNA had a skill score of 77% and outperformed persistence (skill score 65%).

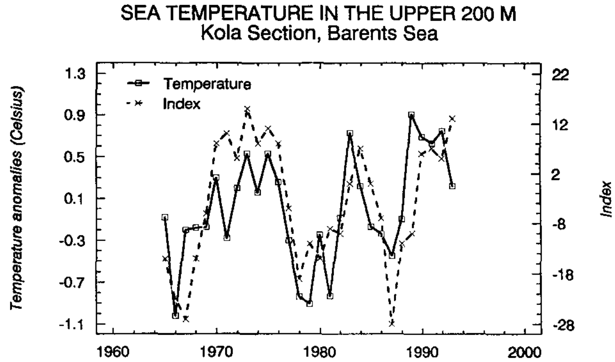

FIG. 1. Observed and predicted changes in the sea temperatures in the upper 200 m of the Kola section of the Barents Sea. These forecasts were made with a 1-yr lead time.

5. Conclusions and future plans

The major objective of our early efforts with the CESNA prototype has been to demonstrate a practical system that can manipulate qualitative information in such a way as to facilitate insights into observed and anticipated climate changes. This has been accomplished through the use of declarative representations of expert knowledge. An ever increasing amount of this type of information exists in the form of heuristics, or empirical rules of thumb. The principal technical issue in using this kind of knowledge is dealing with the inherent high degree of uncertainty of any given rule.

The fundamental idea behind CESNA is to overcome this problem by combining various forecasting methods. This is, of course, not a new approach, but in the past it has been limited to a few methods developed for the same region. The existence of teleconnections between climatic variables allows us to make use of empirical rules about the climate system from remote areas, virtually from all over the globe. Each of these rules, no matter how uncertain, carries valuable information that can be used for predictions.

Fortunately, techniques derived from research in artificial intelligence and expert systems provide a excellent basis for combining exactly this kind of information. Our experiments with CESNA have shown that even though some rules may be poor predictors in a given year, the combined evidence from the remaining rules can result in accurate predictions. The system's predictions from the Barents Sea temperature evaluation paralleled the observed temperatures with remarkable accuracy. Similar results were obtained for two other regions, the Northwest Atlantic and southeastern United States.

Our experience with the prototype system has lead us to focus on a number of areas for improvement. One major goal is to expand the system's core knowledge base to include more of the globe. CESNA system still has a fairly limited number of rules. For example, the knowledge-base for the North Pacific contains only those rules that are particularly important for the North Atlantic. In the future, we plan to develop this knowledge base so that it can be used not only as a supplement for the North Atlantic but to predict climatic processes within the Pacific–North American sector itself.

A second area concerns the system's ability to create true forecasts. Note that the temperature forecasts presented here are actually "hindcasts" since the temperature for the forecast years was already observed. True forecasts require a significant amount of current information on observations about various climatic variables around the globe be readily available with a minimum delay. For systems like CESNA, it is important that this operational information becomes more and more available through the global Internet. We are currently modifying CESNA so that it will easily incorporate such information. Among the information of that kind, we intend to access the monthly and seasonal forecasts of temperature and precipitation over the United States that have recently started to be issued by the Climate Prediction Center.

Acknowledgments. We wish to thank Roland Madden and George Kiladis for their valuable comments and suggestions regarding CESNA. We would also like to thank two anonymous reviewers whose critical comments helped improve an earlier version of this article. This work is supported by the National Science Foundation under Grant ATM-9320455.

REFERENCES

Arora, B. R., and A. D. Padgaonkar, 1981: Time-variation of solar influence on tropospheric circulation. *J. Atmos. Terr. Phys.*, **43,** 91–95.

Barnston, A. G., and Coauthors, 1994: Long-lead seasonal forecasts—where do we stand? *Bull. Amer. Meteor. Soc.*, **75,** 2097–2114.

Bochkov, Yu. A., and A. S. Seliverstov, 1978: Solar activity and herring population dynamics (in Russian). *Rybnoe Khozyaistvo [Fisheries],* **1,** 17–20.

Buchanan, B. G., and E. M. Shortliffe, 1984: *Rule-Based Expert Systems.* Addison-Wesley, 748 pp.

Burroughs, W., 1992: *Weather Cycles, Real or Imaginary?* Cambridge University Press, 210 pp.

Coffey, H. E., 1989: Geomagnetic and solar data. *J. Geophys. Res.*, **94,** 7020–7022.

Conway, B. J., 1989: Expert systems and weather forecasting. *Meteor. Mag.*, **118,** 23–30.

Deser, C., and M. Blackmon, 1993: Surface climate variations over the North Atlantic Ocean during winter: 1900–89. *J. Climate,* **6,** 1743–1750.

Dettinger, M. D., M. Ghil, and C. L. Keppenne, 1995: Interannual and interdecadal variability in the United States surface-air temperatures, 1910–87. *Climate Change,* **31,** 35–66.

Dickson, R. R., and J. Namias, 1976: North American influences on the circulation and climate of the North Atlantic sector. *Mon. Wea. Rev.*, **104,** 1255–1265.

——, J. Meincke, S. A. Malmberg, and A. J. Lee, 1988: The "Great Salinity Anomaly" in the North Atlantic, 1968–1982. *Progress Oceanography,* Vol. 20, Pergamon, 103–151.

Esbensen, S. K., 1984: A comparison of intermonthly and interannual teleconnections in the 700-mb geopotential height field during the Northern Hemisphere winter. *Mon. Wea. Rev.*, **112,** 2016–2032.

Fraedrich, K., and K. Muller, 1992: Climate anomalies in Europe associated with ENSO extremes. *Int. J. Climatol.*, **12,** 25–31.

Friis-Christensen, E., and K. Lassen, 1991: Length of the solar cycle: An indicator of solar activity closely associated with climate. *Science,* **254,** 698–700.

Girskaya, E. I., 1987: Droughts and severe winters in the Northern Hemisphere and solar activity (in Russian). *Tr. Gl. Geofiz. Obs. im. A. I. Voeikova,* **513,** 144–149.

Glowienka, R., 1985: Studies on the variability of the Icelandic low and Azores high between 1881 and 1983. *Contrib. Atmos. Phys.*, **58,** 160–170.

Gordon, J., and E. H. Shortliffe, 1984: The Dempster–Shapher theory of evidence. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project,* B. G. Buchanan and E. H. Shortliffe, Eds., Addison-Wesley, 272–292.

Gray, W. M., and J. D. Sheaffer, 1991: El Niño and QBO influences on tropical cyclone activity. *Teleconnections Linking Worldwide Climate Anomalies,* M. H. Glantz, R. W. Katz, and N. Nicholls, Eds., Cambridge University Press, 256–284.

Kiladis, G. N., and H. F. Diaz, 1989: Global climatic anomalies associated with extremes in the Southern Oscillation. *J. Climate,* **2,** 1069–1090.

Labitzke, K., and H. van Loon, 1993: Some recent studies of probable connections between solar and atmospheric variability. *Ann. Geophys.*, **11,** 1084–1094.

Lamb, H. H., 1972: *Climate: Present, Past and Future.* Vol. 1. Methuen, 872 pp.

Marko, J. R., D. B. Fissel, P. Wadhams, P. M. Kelly, and R. D. Brown, 1994: Iceberg severity off eastern North America: Its relationship to sea ice variability and climate change. *J. Climate,* **7,** 1335–1351.

Mendoza, B., R. Perez-Enriquez, and M. Alvarez-Madrigal, 1991: Analysis of solar activity conditions during periods of El Niño events. *Ann. Geophys.*, **9,** 50–54.

Moninger, W. R., 1990: Summary report on the third workshop on artificial-intelligence research in environmental science (AIRIES-89): 2–4 May 1989, Tysons Corner, Virginia. *Bull. Amer. Meteor. Soc.*, **71,** 672–679.

Namias, J., 1982: Case studies of long-period air–sea interaction relating to long-range forecasting. *Physical Basis for Climate Prediction,* Scientific papers presented at WMO/ICSU Study of the Physical Basis for Climate Prediction: Seasonal, Annual and Decadal Time Scale, Leningrad, 13–17 September, 1982, WMO, WCP-47, 293–325.

Nikiforov, E. G., and A. O. Shpaiher, 1980: *Regularities in Formation of Large-Scale Hydrological Regime in the Arctic Ocean* (in Russian). Leningrad, Gidrometeoizdat, 248 pp.

Pearl, J., 1988: *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufman, 522 pp.

Privalsky, V. E., 1985: *Climatic Variability, Stochastic Models, Predictability, and Spectra.* Nauka, 248 pp.

Rodionov, S. N., and A. S. Krovnin, 1991: Interannual variability of thermal conditions in the Bering Sea. *Proc. Int. Symp. on Bering Sea Fisheries,* Khabarovsk, Russia, Alaska Fisheries Science Center, 49–55.

——, and ——, 1992: The 1980's context of climate changes in the North Atlantic region. *ICES Marine Science Symp.*, **195,** 86–95.

Rolston, D. W., 1988: *Principles of Artificial Intelligence and Expert Systems.* McGraw-Hill Book Company, 257 pp.

Shirley, J. H., 1988: When the sun goes backward: Solar motion, volcanic activity and climate. *Cycles,* **39,** 113–119.

Tsai, J. J.-P., and T. J. Weigert, 1993: *Knowledge-Based Software Development for Real-Time Distributed Systems.* World Scientific, 234 pp.

van Loon, H., and J. C. Rogers, 1978: The seesaw in winter temperatures between Greenland and Northern Europe. Part I: General description. *Mon. Wea. Rev.*, **106,** 296–310.

Wallace, J. M., and D. S. Gutzler, 1981: Teleconnections in the geopotential height field during the Northern Hemisphere winter. *Mon. Wea. Rev.*, **109,** 784–812.

——, and Q. Jiang, 1992: On the observed structure of the interannual variability of the atmosphere/ocean climate system. *Atmospheric and Oceanic Variability,* H. Cattle, Ed., Royal Meteorological Society, 17–43.

Zadeh, L. A., 1965: Fuzzy sets. *Inf. Control,* **8,** 338–353.