Fire Research Report

Medium Range Forecasts of Fire Weather Indices

An assessment of canonical correlation based forecast systems from an ensemble prediction system

MetService

December 2005

Rural fire authorities use indices of fire weather and fire danger to assess current conditions, and also to anticipate future fire risk. This report describes and evaluates a forecast system that extends the lead-time at which forecasts of fire weather indices are available to beyond day-7. The forecasts are based on a statistical approach that matches patterns in how computer models see the weather around New Zealand, to patterns in the observed weather at individual sites around the country.

New Zealand Fire Service Commission Research Report Number 64
ISBN Number 1-877349-36-4

Medium Range Forecasts of Fire Weather Indices

An assessment of canonical correlation based forecasts from an ensemble prediction system

Tony Simmers
Advanced Technology Division
MetService

1 December 2005

Medium Range Forecasts of Fire Weather Indices

An assessment of canonical correlation based forecasts from an ensemble prediction system

Tony Simmers
Advanced Technology Division
MetService

1 December 2005

Issued by

Dr Neil Gordon
Chief Meteorologist

Executive Summary

A method of producing medium range forecasts of fire weather indices is described. Pattern recognition techniques based on canonical correlation analysis (CCA) are used to relate forecasts from the US National Centers for Environmental Prediction's (NCEP) ensemble prediction system to observations at 137 sites from the New Zealand Fire Service's fire weather climatology. The forecasts show significant skill to day-10, and if implemented operationally, the system could provide daily charts showing the spatial variation of fire weather indices. The method is designed to be transportable to forecasts from other ensemble systems, although this ability was not tested.

Beyond day-10 any remaining skill is due to the persistence of initial values and the influence of seasonally varying factors. This information is not intrinsically due to the forecast system, however the significant absolute level of variance explained in the Weekly Severity Rating summarises these effects. Fire managers may find a chart showing the spatial variation of WSR at day-14 useful.

The anticipated ability of the ensemble forecasts to add weather dependent probabilistic information was not found. Probabilistic forecasts could still be made; but they would be based on historical errors, rather than a day-to-day assessment of whether the atmosphere was more or less predictable.

If the forecast scheme was to be implemented operationally, improvements to the NCEP ensemble, and availability of data from the more skilful 51-member European Centre for Medium Range Forecasting ensemble, may result in slightly increased skill. More ensemble members may also allow variability in the predictability of the atmosphere to be described.

1. INTRODUCTION 1

1.1 Fire Weather Indices 2

2. REANALYSIS DATA 4

2.1 Principal Component Analysis of NCAR/NCEP Reanalysis data 6

3. OBSERVATION DATA 9

3.1 Principal Component Analysis of Observation Data 12

3.2 Estimating Data for Missing Stations 13

4. CANONICAL CORRELATION ANALYSIS 15

5. VALIDATION OF CCA RELATIONSHIPS 18

5.1 Reconstructing observations with CCA 18

5.2 Treating Weather Elements Individually 19

6. CCA BASED FORECASTS OF FIRE WEATHER 21

6.1 NCEP Ensemble Prediction System Forecasts 21

6.2 Forecasts from a Single Ensemble Member 22

6.3 Comparison with Other Forecasts from Single Models 23

7. ENSEMBLE FORECASTS OF FIRE WEATHER 25

7.1 Impact of Persistence 26

7.2 Correction for Bias 30

7.3 Forecasts for 2003-04 Fire Season 34

7.4 The Spread Skill Relationship 35

8. CONCLUSIONS 38

ACKNOWLEDGEMENTS 39

REFERENCES 40

APPENDIX: DETAILS OF STATIONS USED 42

1. Introduction

Rural fire authorities use indices of fire weather and fire danger to assess current conditions, and also to anticipate future fire risk. Typical conditions leading to dangerous fire weather are strong winds, combined with high temperatures and low humidities. The risk can become extreme when these conditions are experienced following a period of prolonged drought. Fire managers are generally well aware of the underlying risk arising from past conditions through the fire weather indices generated from observations of the weather. In order to anticipate fire risk, indices derived from weather forecasts may be used. The Meteorological Service of New Zealand Limited (MetService) has been providing 3-day forecasts of fire weather indices to the New Zealand Fire Service since 2001. These forecasts are based on output from the 12 km resolution mesoscale weather prediction model run at MetService.

This report describes and evaluates a forecast system that extends the lead-time at which forecasts of fire weather indices are available to beyond day-7. The forecasts are based on a statistical approach that matches patterns in how computer models see the weather around New Zealand, to patterns in the observed weather at individual sites around the country. Unsettled situations where a series of fronts cross the country every few days do not usually lead to severe fire risk, and are inherently more difficult to predict accurately. More settled situations are both easier to predict further into the future, and tend to lead to increased fire risk. The pattern recognition technique combines a long record of computer weather analyses with observational data from an updated version of Pearce et al.'s (2003) climatology of fire weather.

As the lead-time of a weather forecast increases, the skill and accuracy of the forecast inevitably decreases. Advances in numerical models of the atmosphere have gradually extended the lead-time at which model forecasts show skill, but on any particular day the skill of a model depends to a large extend on the volatility of the weather situation, and on how well the initial conditions can be determined. Over the last 10 years ensemble prediction systems have been developed to provide estimates of the uncertainty inherent in a forecast, and consequently identify days when the forecast is likely to retain skill for a longer forecast period. An ensemble system comprises a number of numerical weather models that are each started from very slightly different initial conditions. In some cases the physics schemes of the models that make up the ensemble themselves are also slightly different. As each forecast evolves, small differences in the models amplify in a chaotic manner. By taking the average of all the ensemble members, it is hoped to separate the 'signal' of the underlying processes driving the weather, from the 'noise' of day-to-day variability. Ensemble average forecasts typically show skill at 1-2 days beyond forecasts based on a single model. The range, or spread, of results at a particular forecast time is often used to estimate the likely variability in the forecast. If the estimates of variability are reliable, ensemble systems should be able to recognise the stable, predictable, weather situations that can lead to high fire danger, makes them particularly suited to forecasting fire weather indices. However, as will be shown later, this assumption needs to be tested.

This project describes the methods used to create forecasts of fire weather indices using output from the NCEP ensemble prediction system, and reports on validation of the forecasts over a 12 month period. The major steps taken to produce this report were:

1. Principal Component Analysis (PCA) of 17 years of data from the US National Centers for Environmental Prediction/National Center for Atmospheric Research (NCEP/NCAR) reanalysis project was used to describe typical atmospheric patterns over New Zealand.
2. A similar procedure was performed on an updated set of the data used by Pearce et al. (2003) to describe patterns of observed weather at 137 sites around New Zealand.
3. Canonical Correlation Analysis (CCA) was used to form a relationship between the patterns in the reanalysis dataset and those from the fire weather climatology.
4. This relationship was applied to archived forecasts from a single numerical weather model to produce and evaluate forecasts of fire weather indices for the period June 2003 to May 2004 inclusive.
5. The technique was then extended to the multiple weather models that make up the NCEP ensemble prediction system, and the resulting ensemble forecasts evaluated.

1.1 Fire Weather Indices

Fire-weather forecasting systems arise from the need to manage fire risk in both wild land and commercial forests. The following list (Reifsnyder and Albers, 1994) illustrates some of the issues authorities need to consider:

a) Protection of life and property in wildland environments;
b) Protection of the commercial value of the forest resource;
c) Protection of watersheds from fire-accelerated erosion;
d) Protection of high-valued plantations;
e) Controlled use of fire for forest regeneration;
f) Controlled use of fire for fuel reduction;
g) Controlled use of fire in land clearing;
h) Controlled use of fire in vegetation type conversion.

New Zealand fire managers use the New Zealand Fire Danger Rating System (NZFDRS) to assess the potential for ignition and subsequent development of forest and rural fires (Fogarty

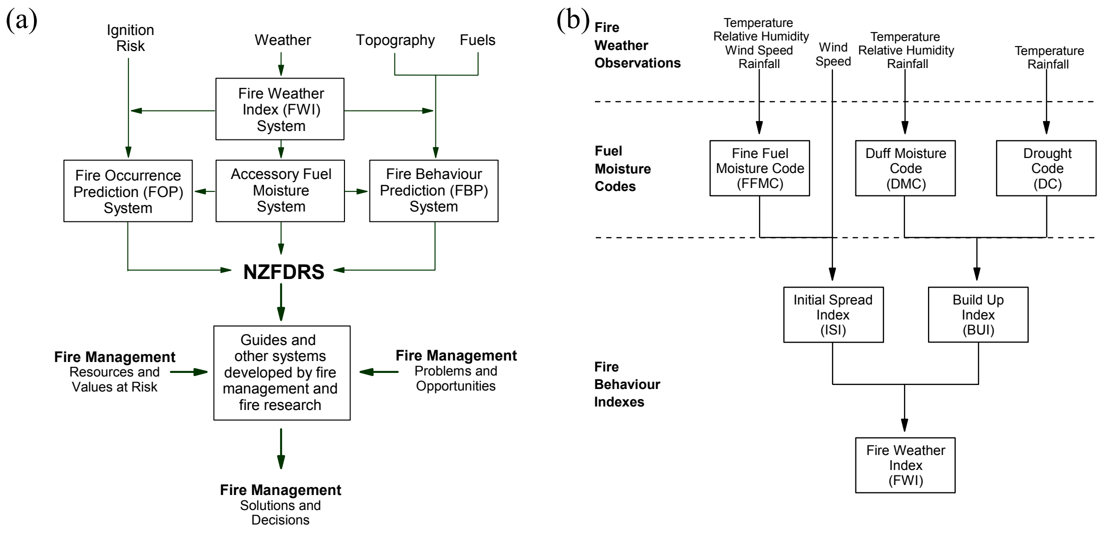

Figure 1. Simplified structure diagrams for **(a)** the New Zealand Fire Danger Rating System (NZDRFS), illustrating the linkage to fire management actions and, **(b)** the Fire Weather Index (FWI) System (after Pearce et al., 2003).

et al., 1998). The system is an adaptation of the Canadian Forest Fire Danger Rating System (Stocks et al., 1987) and the complete system takes into account fuel type, topography, and present and historical weather conditions. The fire weather component of the NZFDRS is derived from either observations or predictions of midday temperature, humidity, wind speed, and rainfall over the preceding 24 hours. The overall Fire Weather Index (FWI) is a means of rating wildfire potential for a standard fuel type (mature pine forest) on level terrain. It is calculated from a number of component indices, some of which are used in more specific algorithms to estimate fire risk for forest, grassland and scrubland (Alexander, 1994; Pearce, 2001). Figure 1 gives an overview of the NZFDRS, and the component parts of the FWI System. The implementation of the FWI used in this project is based on the algorithms of Van Wagner and Pickett (1985) with minor adjustments to suit New Zealand conditions (Alexander, 1992).

When describing the fire risk aggregated across a number of stations or from one station over a period of time it is more usual to use the Daily Severity Rating (DSR). The DSR is derived by increasing the weight placed on high values of FWI, and was designed to reflect the difficulty of controlling a fire more directly than the FWI (Van Wagner and Pickett, 1985).

$$DSR = 0.0272 FWI^{1.77}$$

Heydenrych et al. (2001) used monthly and seasonal averages of DSR, known as MSR and SSR respectively, when describing climatic influences on severe fire seasons at a number of locations in New Zealand. The weekly severity rating is introduced later in this report.

This report maintains the distinction between the use of the averaged severity ratings as measures suited to evaluating the performance of forecast fire weather over a period or over a number of stations, and the other indices, which are more likely to be used by fire managers in the field.

2. Reanalysis Data

In devising a scheme to produce weather forecasts at individual locations it is usually necessary to downscale the coarse spatial resolution gridded output from numerical weather prediction (NWP) models to the observed weather elements at each location. There are two main approaches to assembling an archive of gridded model data. In the first approach, known as model output statistics (MOS), data from each lead-time of a forecast model is saved, and separate equations to predict weather elements at each lead-time are developed. This has the advantage of tailoring the equations to the specific traits of the model, including accounting for decreased skill with increasing lead-time by the equations tending towards forecasting climatology as lead-time increases. A significant disadvantage of MOS is that the equations should be redeveloped each time the model is upgraded or changed.

For this work the second approach of a 'perfect prognosis' (PP) technique was chosen. In the PP method the archive of gridded data comes from the initial analysis of the numerical model. The centre running the model creates an analysis by assimilating observations into their model, and producing what is effectively a forecast with a lead-time of zero hours. A single set of equations is developed based on the assumption that if the model prognosis at a given lead time is 'perfect', then the equations based on the analysis will apply. The principal advantage of the PP method is that there is usually a longer historical record of analyses, enabling more stable equations to be developed. Also, as improvements to the forecast model are made its predictions should become closer to perfection and, rather than needing to recalculate the statistical relationships, the existing equations should show more skill.

Forecasts based on a perfect prognosis system do not tend towards climatology as lead-time increases. When data from an ensemble of forecast model runs are put through a perfect prognosis scheme a range of values will result, even at long lead times. This is in contrast to the MOS approach, where at long lead times each member of the ensemble will produce the same, climatological, forecast. Bertrand and Verret (2004) cite similar reasons for their choice of a perfect prognosis regression scheme to forecast temperatures from the Canadian ensemble prediction system.

To build the training set of gridded data, fields from the NCEP/NCAR reanalysis project (Kalnay et al., 1996) with a spatial resolution of 2.5°×2.5° latitude/longitude were obtained for the region surrounding New Zealand (56 grid points from 165°E-180°E, 30°S-50°S; see Figure 5). Analyses are available at six hourly intervals, but because fire weather indices are calculated from midday observations, only 0000 UTC analyses were used. The reanalysis project contains data extending back to 1958; however, until widespread observations from satellites and drifting buoys became available in the 1980's much of the southern oceans were a data void. For this reason data from 1985 to 2004 were selected to describe the synoptic situation around New Zealand.

There is a known problem in the reanalysis data relating to the inclusion of psuedo-observations (PAOBs) of surface pressure between 1979 and 1992. The Australian Bureau of Meteorology generated PAOBs in an attempt to improve analysis over data sparse regions, but unfortunately they were included in the wrong locations. Although the impact on the reanalysis data set was found to be greatest in the latitude band between 40°S-60°S, the region between Australia and New Zealand was much less affected due to the availability of conventional observations. The size of the errors introduced also decreased rapidly with height. Figure 2 shows the difference in 850 hPa temperatures when the PAOBs were included in the correct and incorrect locations. There is an obvious minimum in variation over the New Zealand region. By way of comparison, there was typically a difference of 1°C between analyses from the NCEP/NCAR reanalysis and a similar reanalysis project using data

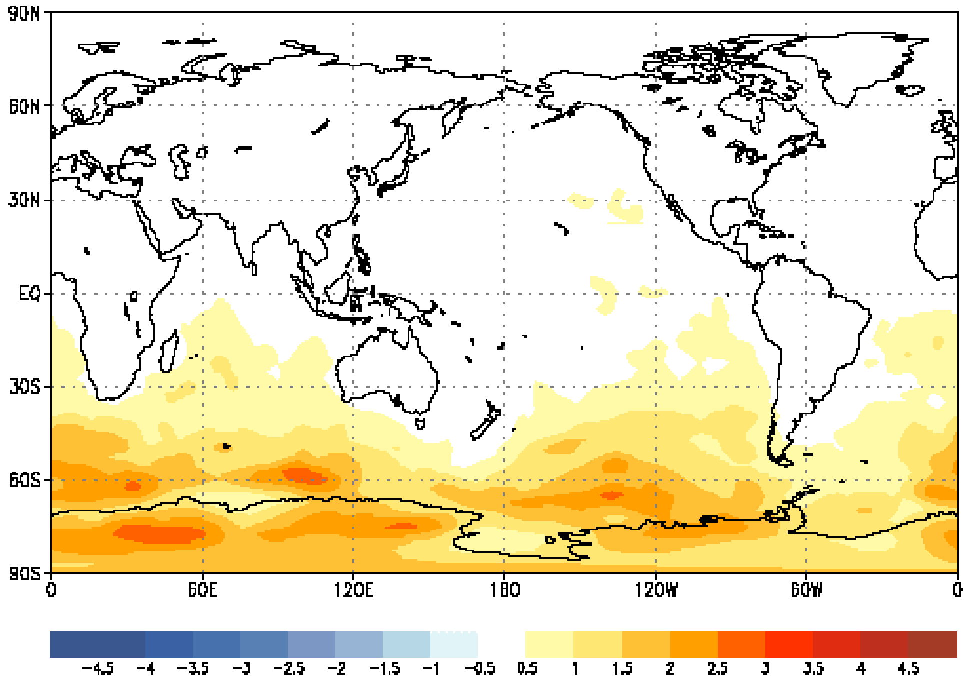

Figure 2. RMS difference in analyses of temperature (K) at 850 hPa between analyses with PAOBS in the correct and incorrect positions for July 1981 (from Kalnay, 1996).

from the European Centre for Medium-range Weather Forecasting (ECMWF). Overall the impact of the PAOBs problem on the reanalysis data used in this project is considered minimal.

Kidson (1994, 1997, 2000) has reported on the use of reanalysis datasets to describe the weather and climate patterns affecting the New Zealand region. His work forms an important part of the analysis of severe fire seasons conducted by Heydenrych et al. (2001). In his papers Kidson used principal component analysis to extract from the reanalysis data the principal components of either daily or monthly anomaly charts. He then grouped the charts into clusters based on the distance between each chart as measured along the principal components. By classifying the reanalysis data in this way, the complete range of different synoptic situations present in the reanalysis dataset was re-expressed as a much smaller number (10 to 13) of 'typical' situations.

In his 1997 paper, Kidson found that clusters based on patterns in geopotential height across multiple levels were almost the same as clusters derived from geopotential height at a single level. Upper level patterns in the multi-level clusters tended to be simply displaced to the west of lower level patterns. Theory predicts just such a relationship between geopotential heights for a barotropic (no horizontal temperature gradient) atmosphere, suggesting the clustering process had effectively averaged out temperature gradients. This result was disappointing because the same study showed that in clusters based on data from a single level, observed climatic variables did show variations that depended on geopotential heights at other levels. Kidson (2000) also mentions the possible use of PCA with ensemble prediction systems, although his focus is on climate models for the next few months rather than the 1-2 week period addressed here.

In an attempt to avoid some of these problems, this study differs in two ways from Kidson's work. The nature of eigenvalue analysis, which is at the heart of PCA, ensures that each principal component pattern is orthogonal to, or independent of, all the other patterns. Firstly clustering techniques were not used to identify a finite number of 'typical' weather maps from the principal component patterns. The pattern shown in an individual principal component does not necessarily represent a real weather situation, just some underlying aspect of many patterns, such as strength of north-westerly flow. By maintaining these underlying patterns, any particular situation can be described as part of a continuum rather than being binned with patterns that are somewhat similar. Unusual situations, which by their nature are interesting, are likely to be better described.

Secondly, the principal component analysis in this work uses data from different levels in the atmosphere, and also from different variables. In common with most synoptic typing studies, including some relating specifically to fire weather (Janz and Nimchuk, 1985), geopotential height at 500 hPa (H500) was the first field chosen for its ability to represent the overall synoptic situation. Mean sea level pressure (PMSL) was the next choice because its gradient gives an indication of both wind speed and direction. To minimise the effect of the barotropic relationship between PMSL and H500, and to better identify the location of frontal systems, temperature at 850 hPa (T850) was included. Finally relative humidity at 700 hPa (RH700) was chosen as a proxy for mid level cloud and its associated precipitation.

Gridded precipitation from the reanalysis was deliberately not included, despite the importance of observed precipitation to fire weather. A key feature of the reanalysis data is that exactly the same algorithms to assimilate the observation data are used from the start of the reanalysis project's data in 1954 to the present time, in order to create an internally consistent set of analyses. However the spatial resolution and the internal physics of operational models are regularly upgraded. The fields chosen in this study to describe the synoptic situation are relatively insensitive to such changes, but precipitation can be dramatically affected. Part of the reason 850 hPa temperatures were preferred to temperatures at the surface was that surface temperatures also vary greatly in space and are strongly dependent on the physics and boundary layer representations within forecast models.

2.1 Principal Component Analysis of NCAR/NCEP Reanalysis data

A training set containing 18 years (1985-2002 inclusive) of 0000 UTC data from the NCEP/NCAR reanalysis set, at each of 56 grid points over New Zealand and for the 4 fields chosen above (4×56=224 variables) was assembled. The seasonal trend was removed from each variable using an expression of the form

$$\overline{V_{\theta,x,y}} = a_{v,x,y} + b_{v,x,y}\sin\theta + c_{v,x,y}\cos\theta + d_{v,x,y}\sin 2\theta + e_{v,x,y}\cos 2\theta$$

Where V is the variable being processed, the subscripts *x, y* refer to each of the gridpoints, *a, b, c, d, e* are regression coefficients, and theta represents the day of the year.

The resulting anomalies were normalised by dividing through by their standard deviations, and then passed through the factor analysis function in the SYSTAT package. Since the variables were already normalised the analysis was performed on the covariance matrix, and the resulting eigenvectors were not rotated. The first four principal components are shown in Figure 3.

Care is required in the interpretation of these patterns, as they do not necessarily correspond to 'real' weather maps. Firstly, the analysis was performed on anomalies from the climatological values at each grid point. This removes the obvious north-south gradients in temperature and pressure. Secondly, the gradients within each pattern are essentially non-dimensional since equal weighting, in terms of number of standard deviations from the climatological value, was given to each of the weather elements. The gradients in each pattern represent the rate of change of the variable at a point relative to the other points in the domain. Also the negative of each pattern is equally valid. For example, the first principal

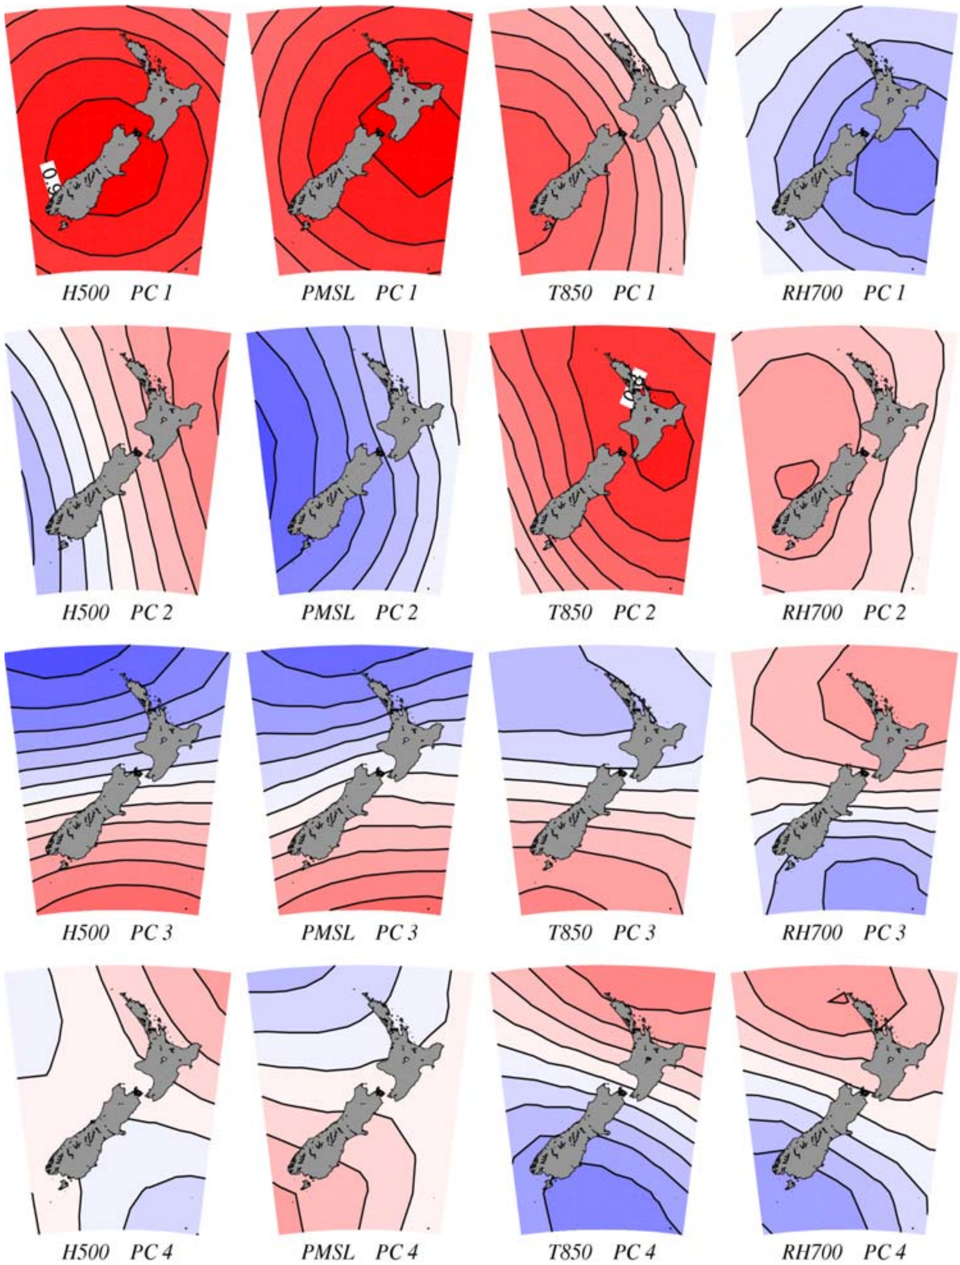

Figure 3. The first four principal components (eigenvectors) obtained from the principal component analysis of NCEP/NCAR reanalysis data from 1985-2002.

component of mean sea level pressure in Figure 3 could equally well represent a centre of high or low pressure. The shape of the patterns in H500 and PMSL are very similar to those found by Kidson (1997); however their order differs, indicating that the variability of the other fields moderates the importance of variations in H500 and PMSL.

Figure 4 shows the amount of variance, plotted on a logarithmic axis, explained by successive principal components. On such a diagram the magnitude of the components corresponding to un-correlated noise should decay exponentially (Wilks, 1995, p. 381), falling along a straight line to the right of the diagram. Based on this test the first 10-15 components of the present analysis contain most of the information, with 15 components explaining 89% of the variance. Examination of the remaining principal components showed that, while the contribution to overall variance from H500, PMSL, and T850 had dropped away by the 15th component,

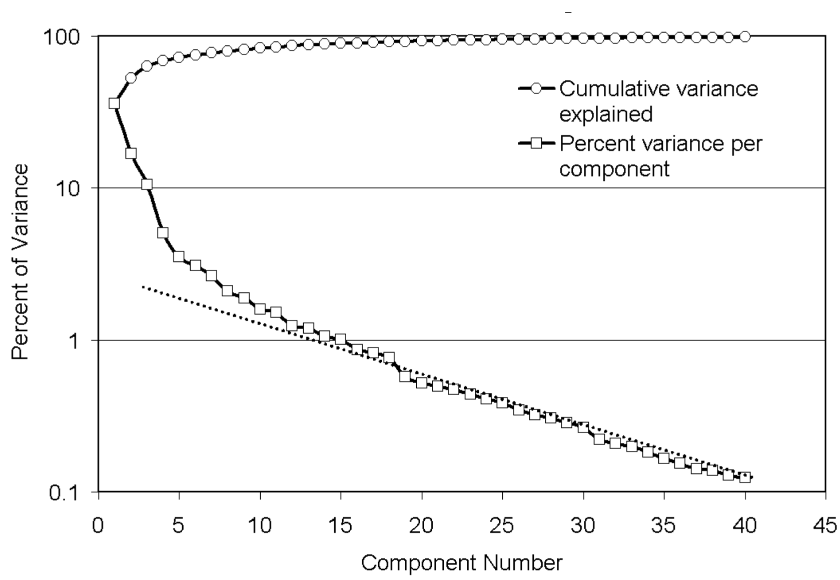

Figure 4. Amount of variance in the reanalysis training data explained by successive principal components. The straight dotted line indicates the region where the magnitude of the principal components is decreasing exponentially, which is indicative of un-correlated noise.

RH700 showed coherent structure as far as the 25th. This reflects the significantly more complex patterns contained in RH700 fields. Because the overall aim was to develop an independent, rather than minimal, set of predictors for the forecast scheme, it was decided to retain 30 principal components, which together explained 96% of the variance in the original 224 dimensions of the reanalysis data. When these 30 components were used to reconstruct independent reanalysis data from 2003-2004, the explained variance dropped only slightly to 95%, indicating that any noise introduced by higher order components was indeed un-correlated and did not introduce significant errors.

3. Observation data

The archive of observed weather used to develop the prediction scheme came from an updated version of the fire danger climatology prepared by Pearce et al. (2003). The observations come from over 150 stations, with two thirds owned and operated by rural fire authorities, and the remainder by MetService. The length of record for individual stations in the climatology varies from 3 to over 40 years. Data from July 2002 to June 2004 was reserved as an independent data set to be used when evaluating the skill of the forecast scheme.

Because the FWI depends on the values of component indices from the previous day, a continuous record of observations was required to compile the fire danger climatology of Pearce et al. (2003). To achieve this, Pearce replaced missing or obviously erroneous values in that climatology either with observations from the same station made within an hour of midday, or from the nearest appropriate neighbouring station.

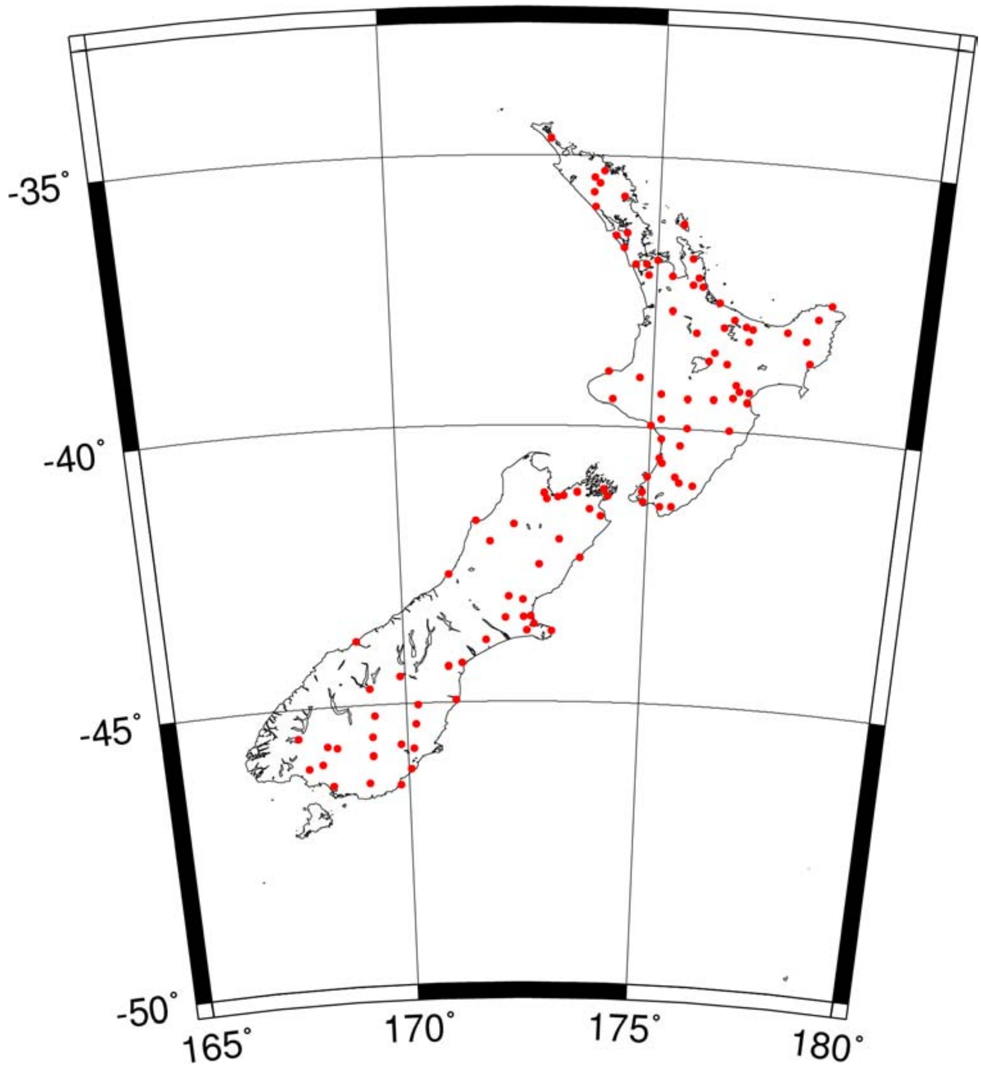

Figure 5. Area over which data on a 2.5 degree latitude/longitude grid was extracted from the NCAR/NCEP reanalysis data set, together with the location of fire weather stations used in the study. For a full list of the stations used see the appendix.

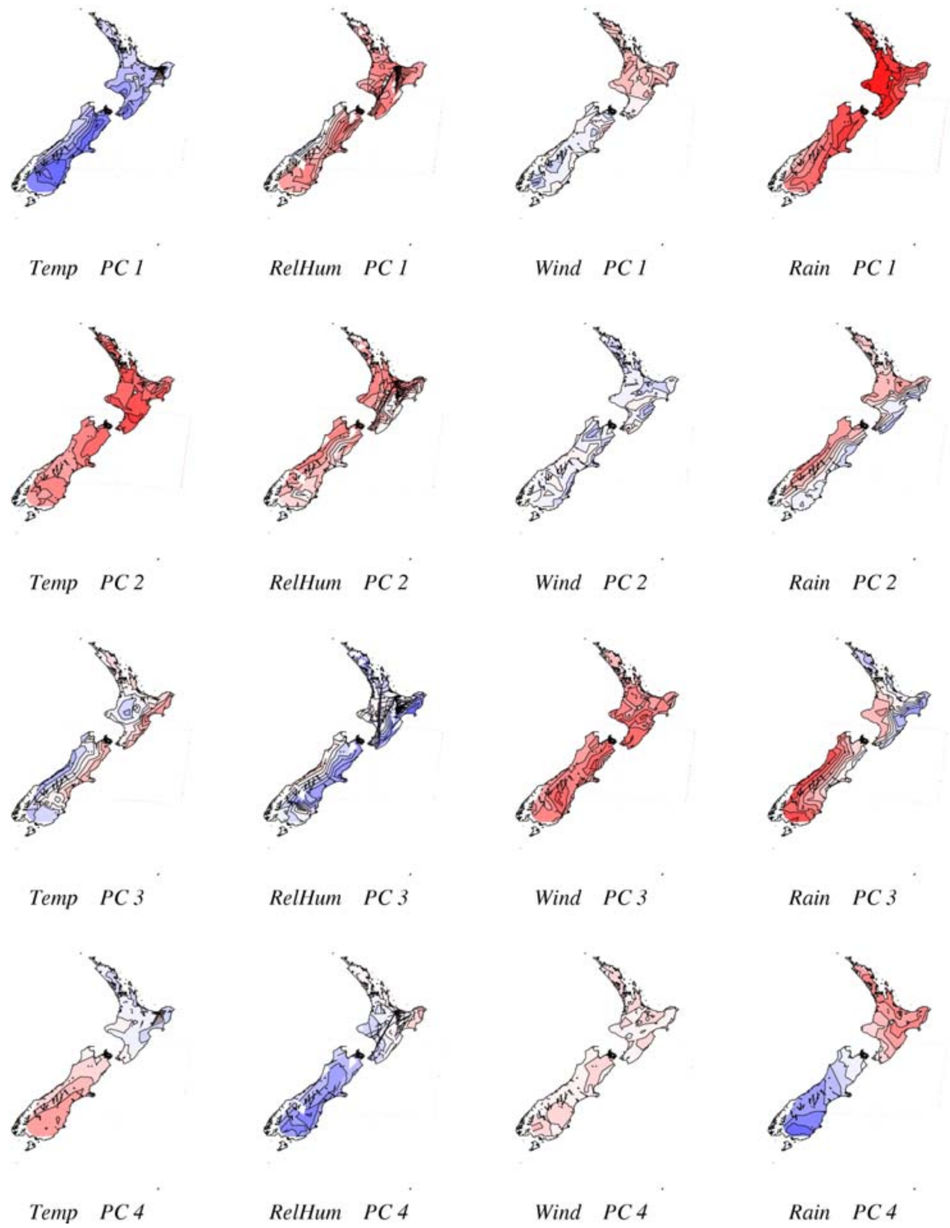

Figure 6 (a) The first four principal components (eigenvectors) obtained from the analysis of observations from the updated fire weather climatology of Pearce et al. (2003).

Despite these efforts there remained a small number of problems with the observed weather elements. Tait and Zheng (2005) have reported in more detail on issues with the data set and made a number of recommendations as to how the quality of the data may be improved. As an example of the problems found, observations at one or two stations would be inconsistent with those from nearby stations for a period, then return to normal. Where such problems were detected in this study, the values were simply marked as missing. In some cases the problems affected a longer period, such as the inclusion of wind directions instead of wind speeds at Dunedin Airport. Because the current study is an evaluation of a forecast scheme, the presence of any one particular station was not vital. This being the case, stations with significant problems were omitted from the analysis.

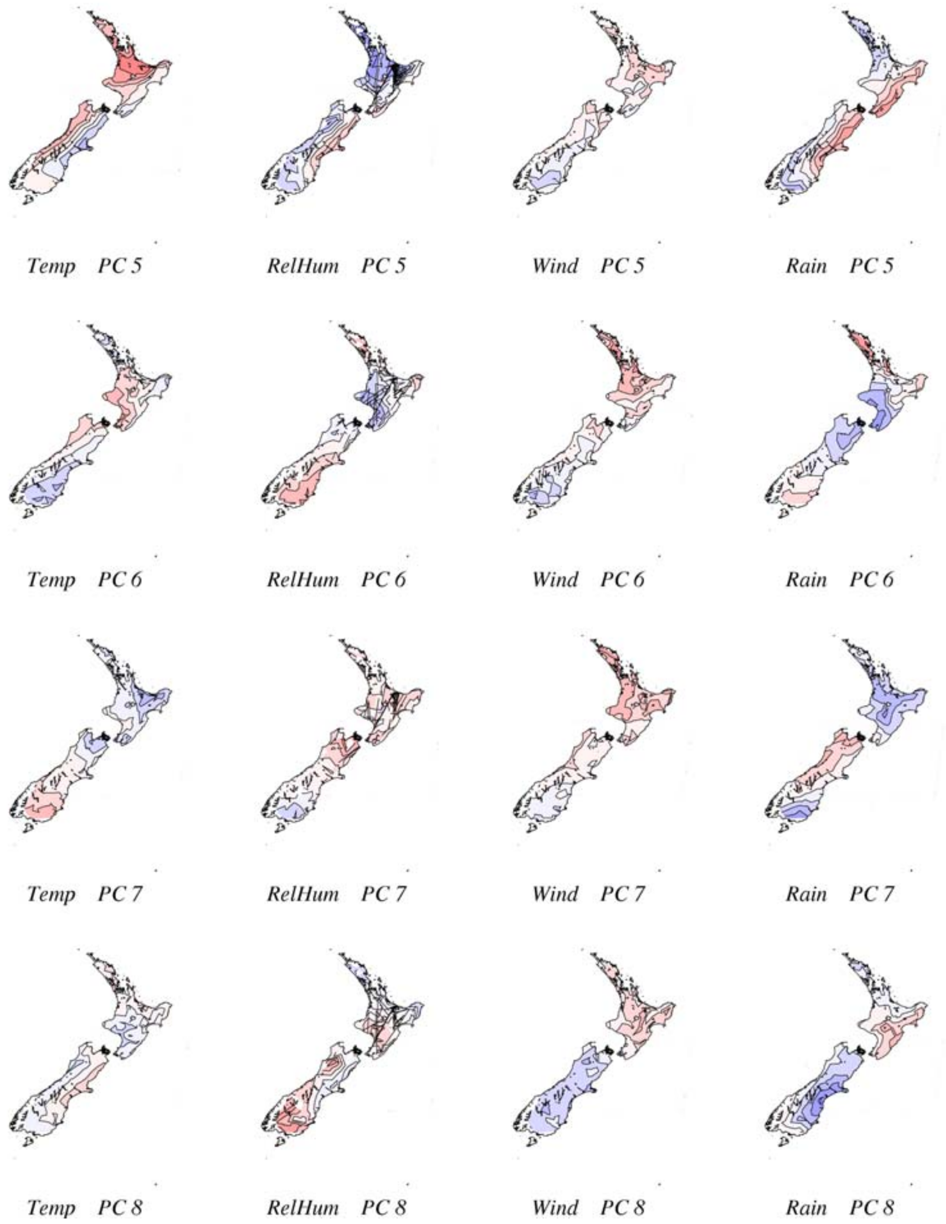

Figure 6 (b) The 5th to 8th principal components (eigenvectors) obtained from the analysis of observations from the updated fire weather climatology of Pearce et al. (2003).

To conduct a principal component analysis the covariance matrix of a data set is analysed, and, unlike computation of FWI, neither a continuous nor a complete record is required for each station. Pairwise deletion, where only combinations involving the missing elements are omitted, was used to calculate the covariance matrix. There were obvious patterns in where the missing data occurred, most notably due to the variation in length of record between stations. A check on the covariance calculated from the full 1985-2002 data set and one with observations from 1995-2002, showed only minor differences. What did become obvious was that some stations had ceased reporting before others had started, making the calculation of covariance between those stations impossible. In some cases, such as Ohakea where reports after 1996 were missing, it is possible more recent data could be found to fix the problem but this was not attempted.

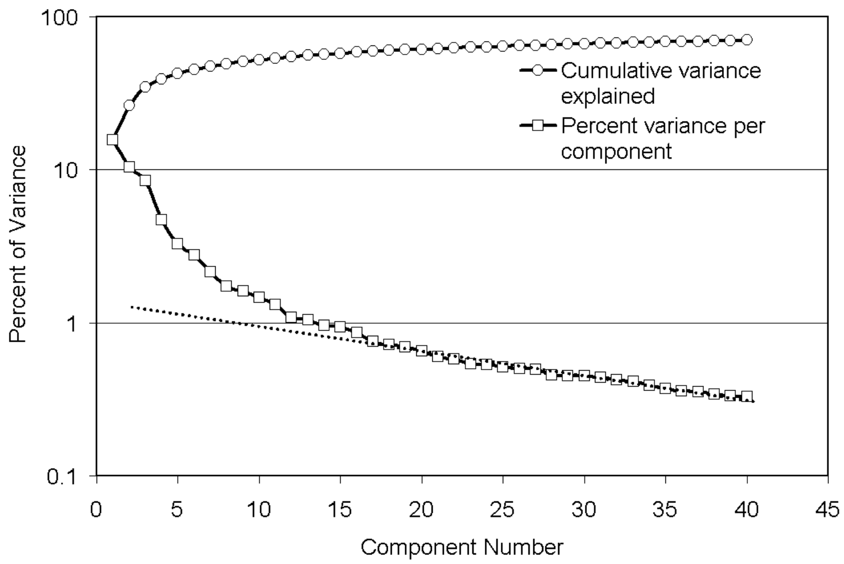

Figure 7. Amount of variance in the observation data set explained by successive principal components. The straight dotted line indicates the region where the magnitude of the principal components is decreasing exponentially which is indicative of un-correlated noise.

In the end, 137 stations (Figure 5) were selected on the basis of having at least 1000 days of data that overlapped with all other stations. In most cases there were more than 2000 days where observations for pairs of stations overlapped. The full set of stations used in the analysis appears in the appendix.

As with the gridded data, the covariance matrix was formed from standardised anomalies of the observed weather elements. Seasonality was removed only from the temperatures, since the annual variation was much smaller in the other elements and for some stations only three years of data was available from which to determine the seasonal trend.

The frequency distribution of precipitation amounts is decidedly non-normal, and dominated by a large number of days on which no rain falls. Fire weather indices are insensitive to rainfalls below 0.5 mm/day. The impact of extreme rainfall is also muted since a number of terms depend only on the log of rainfall, and some indices effectively become saturated after heavy rain. To avoid the variance due to the few extreme cases of rainfall dominating the values in the covariance matrix, and to better represent the days with no rain, precipitation was transformed by taking its cube root. The cube root of precipitation is sometimes used to normalise accumulated rainfall over months or seasons, where the likelihood of no rainfall is small. Its use here, on daily accumulations, is a more pragmatic choice that balances the need to differentiate between significant and nil or slight rainfall, against the desire to not have very large deviations from the transformed mean.

For similar reasons, the wind speed was also transformed by taking its square root.

3.1 Principal Component Analysis of Observation Data

Once the anomalies from climatology were determined, the principal components of the observation set were found by following a similar process to that used on the gridded reanalysis data. There were 548 variables available for analysis (137 stations × 4 weather elements) for the period 1985-2002 inclusive.

Table 1. Percent variance explained within the independent set (2002-2004) of observations (power transformed, and standardised anomalies) by 30 principal components with varying numbers of additional stations marked as missing. Estimated values for the missing data were obtained by three successive re-constructions of the data set and the procedure iterated for 100 different combinations of missing stations. The average variance explained for individual elements is shown together with that for the four elements combined.

Percent of extra stations omitted	All Elements Combined	Temp	RH	Wind	Rain
0.0	66.3	72.5	68.8	52.1	71.7
10.0	65.7±0.3	71.9±0.4	68.1±0.5	51.1±0.4	71.6±0.4
20.0	63.8±0.5	70.2±0.6	66.1±0.8	49.5±0.6	69.4±0.7
30.0	61.4±0.9	68.2±0.9	63.8±1.1	47.1±0.9	66.4±1.3
40.0	59.5±1.0	66.4±1.1	61.6±1.2	45.6±1.1	64.3±1.4
50.0	55.9±1.2	63.1±1.3	58.4±1.4	42.1±1.2	60.0±1.6
60.0	49.1±1.5	56.7±1.6	48.8±1.5	36.6±1.4	54.2±2.0

The first 8 principal component patterns retrieved are shown in Figure 6. The first pattern represents cold and wet (or warm and dry) conditions over the whole country, and the second warm and wet (or cold and dry). Further interpretation of these patterns is not easy for a number of reasons. Again the patterns are based upon the transformed and standardised anomalies, making it difficult to reconstruct the patterns mentally. Variations between neighbouring stations are also not as smooth as for the reanalysis data. For example, sites in and around the Southern Alps could be expected to vary markedly from those on the open Canterbury Plains. The contouring routine uses simple triangulation between the unevenly spaced data points which further hinders interpretation of the patterns.

Fortunately the lack of spatial continuity does not directly affect the derivation of the principal components, since the points could be assembled in any order without affecting the result. The 15 leading principal components could explain only 57% of variance in the observations, compared to 89% in the reanalysis data. Figure 7 shows that the point at which uncorrelated noise dominated the pattern occurred around component 15, compared to component 10 for the reanalysis data, indicating that there were approximately 5 extra degrees of freedom in the observation patterns. As with the reanalysis data, the inclusion of extra 'noisy' components was not going to cause problems in later analysis, so 30 principal components were retained which together explained 68% of the variance in the (normalised, and power transformed) observations. Note that this is considerably less than the 95% variance in the reanalysis data explained by 30 components. The reason for this is the smaller degree of spatial correlation between observations compared to the reanalysis fields that vary relatively smoothly between grid points.

3.2 Estimating Data for Missing Stations

Using independent data to test the principal components' robustness was more difficult than for the gridded data because 20 of the 137 stations did not have observations for the period after June 2002. In the first step of the test, where the contribution of each principal component to the whole pattern is found, approximately 15% of stations were missing. One approach was to simply set the anomalies for those missing stations to the average value of zero, and then test the fit of the derived principal components to the stations that had no missing data. When this was done, the explained variance was only 48% compared to 68% with the same 30 components for the dependent training dataset which had no missing values.

While this was a somewhat disappointing result, it should be noted that inverting the principal components to reconstruct the original data also automatically generates values for the stations that had no data. For example, on a day where most of the country was warmer than usual the reports from the stations that did have data would dominate the fewer, average, values at the missing sites. The principal component with the 'warm everywhere' signature would have a significant weighting. When the pattern was reconstructed this would result in warmer than average temperatures at the missing sites, balanced by not quite warm enough values at the sites that did have data.

This ability to fill in missing data points was used to significantly improve the test results when reconstructing the independent set of observations. The assumption that the missing data was average was refined with the results of the principal component patterns found from the non-missing stations. This process was iterated three times to eventually explain 66% of the variance in the non-missing stations, a result that closely agrees with the slight reduction in explained variance found with the reconstruction of the independent gridded data. Beckers and Rixen (2003) have reported on the use of a similar technique to fill in periods of missing data.

In an extension of this technique, a Monte Carlo test with 100 simulations was performed where an increasing additional percentage of stations were artificially set as missing. Each of the 100 simulations had a different set of missing stations. The results of the reconstructions were evaluated against the full, non-missing, data set. Table 1 shows that the decrease in explained variance is gradual, further supporting the potential of this technique as a means to fill in patches of missing data.

Table 1 also shows the contribution to total variance from each of the individual weather elements. The large explained variances for (transformed) precipitation and relative humidity may seem somewhat surprising. It is the larger amount of variance initially available for explanation that is believed to be responsible for this result. Another way of looking at this is to consider the temperatures at a site such as Auckland. Once the seasonal trend has been removed from Auckland midday temperatures, the root mean square difference from climatology is a little less than 2°C. A significant fraction of this residual variability will be due to instrumentation factors, and unresolvable microscale variations, leaving only a small amount that the principal components could reasonably be expected to explain.

4. Canonical Correlation Analysis

Canonical correlation analysis (CCA) is a statistical technique that identifies a sequence of pairs of patterns in two multivariate data sets (Wilks, 1995, p. 398). From each of the data sets, a new set of vectors is found such that the correlations between these projections of the original data are maximised. This procedure shares many similarities with PCA. Figure 9 shows how CCA relates the **y** vector that we wish to forecast (in this case the principal components of the observations) to the **x** vector of predictors (the principal components of the gridded data) using correlations between linear combinations of those original data sets. The correlations found can be used much like a traditional multiple regression equation, except that the multiple predictors are used to generate *multiple* predictands.

The results from CCA are sensitive to the presence of internal correlations between the input and output data sets. This is the reason that both the observations and gridded data were 'pre-filtered' by first changing them into their principal components. By definition each principal component is independent (orthogonal) to all the others, which improves the chances that the relationships found will perform well when applied to independent data. Decreasing amounts of variance are explained with successive correlations. This approach assumes that there is no significant reduction in information caused by the representation of the original data sets by a limited number of components. Because the dimensionality of the arrays involved is reduced, the method is also much less computationally expensive than retaining the original variables.

The SYSTAT package was used to perform CCA on the principal components extracted from the observational and gridded datasets. There were no problems with missing data in the reanalysis data so the period used for training the prediction scheme was determined by gaps

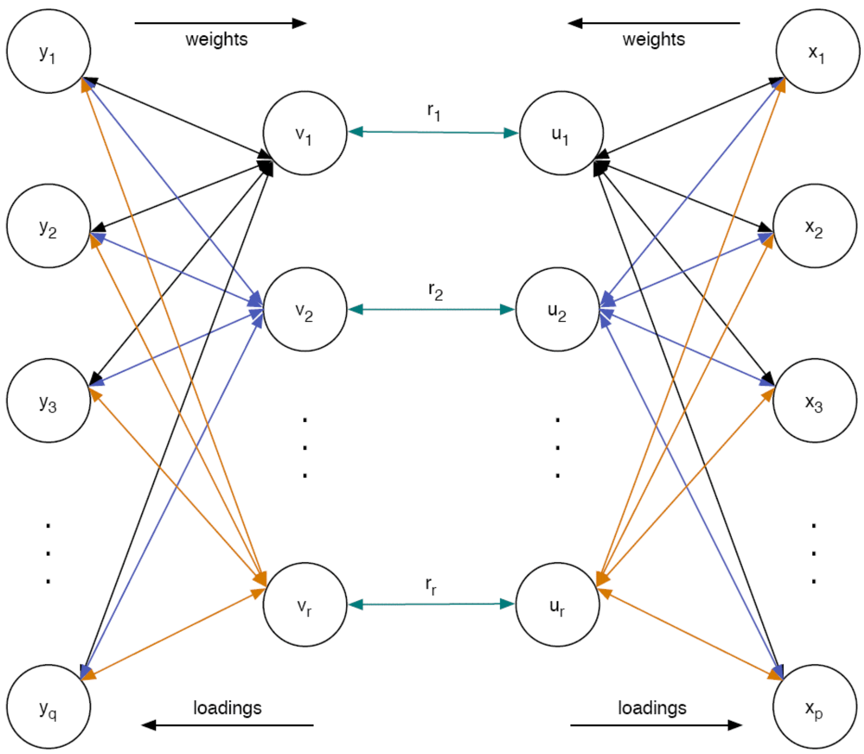

Figure 8. Schematic showing the relationship between variables in CCA. In this case **y** is the vector of principal components representing the elements we wish to forecast, and **x** is the vector of principal components extracted from reanalysis data. The CCA technique chooses the projections of **x** and **y** onto **u** and **v** that maximise the correlation between **u** and **v**.

in the observational data set. By choosing only days where not more than 10% of stations had missing data, 1394 days were identified as being suitable for analysis. For the stations that did have missing data, the missing values were estimated using three iterations of the PCA process before the final dataset was submitted for CCA.

The SYSTAT CCA module uses the correlation matrix in its analysis rather than the covariance matrix. In effect this is equivalent to normalising all the variables by their means and standard deviations before calculating the covariance. In order to retrieve predicted principal components of observations from the gridded data principal components, it was necessary to calculate and save the mean and standard deviations of both sets of weightings.

Initial results from the CCA were encouraging. Table 2 gives the canonical correlations, effectively R^2 values, between the vectors projected from the observational and reanalysis data sets (**u** and **v** in Figure 9). The high correlations indicate strong associations between the patterns in the two data sets. Because 30 orthogonal patterns from the reanalysis were used to explain 30 patterns in the observational data, the relationships found by CCA were, at least for the training set itself, able to completely describe the variance between the two sets. This would not have been the case if the number of principal components retained were not the same for both data sets. Of particular interest was the significance of the correlations found. The Bartlett test of residual correlations produced by the SYSTAT software was used determine the point beyond which the remaining correlations were likely to have been due to chance. Table 2 shows that this occurred at around the 23rd canonical correlation. As with the PCA results, the independence of the relationships means that including the full 30 correlations, rather than only those that are statistically significant, should only result in the addition of random noise to the forecasts and not affect their overall usefulness.

5. Validation of CCA Relationships

There were three stages to the validation of the CCA relationship between the reanalysis data and observations. Firstly the relationship was used to reconstruct observations from the reanalysis data itself. Next, to test whether the results derived from reanalysis data could be applied to forecast models, forecasts from a single ensemble member were used to generate forecasts of weather elements and fire weather indices. Finally the method was applied to the full 12 members of the NCEP ensemble.

5.1 Reconstructing observations with CCA

The correlation coefficients in Table 2 refer to the match between the principal components of two data sets when they are each projected on to an *intermediate* set of vectors. The 0.928 value of the first canonical correlation does not mean the relationship can explain 93% of the variation between the first two principal components, but that some *combination* of the 30 reanalysis principal components can explain 93% of the variation in a *combination* of the observation principal components.

The second column of Table 3 shows the percent variance in the first 10 principal components of observations that CCA could explain using the reanalysis principal components. The table also shows, in the third column, how much variance the observation principal components explained within the observations themselves. Multiplying together the two columns suggested that, when applied to independent data, the reanalysis should explain something over 30% of the variance in the observations.

Tests using independent reanalysis and observation data from the period May 2002 to June 2003 showed 32% of the variation between the normalised observation data and the values reconstructed from the reanalysis principal components was explained using the CCA relationships (Table 4). The effect of returning the anomalies from climatology to actual values can be seen clearly in the extra 30% variance explained between the normalised TT values and those that have the seasonal component added back in. There is very little difference in RH, while the variance explained in the normalised (square root of) wind and (cube root of) rain values changes somewhat when the power transformations are reversed.

Table 3 Estimate of variance explained in observations (transformed, standardised anomalies) using CCA with reanalysis data.

Observation Principal Component	% variance in each obs. PC explained by reanalysis PCs using CCA	% variance in obs. explained by obs. PCs	Estimate of % variance in obs. explained by reanalysis PCs using CCA
1	0.727	15.94	11.59
2	0.752	10.85	8.16
3	0.616	9.14	5.63
4	0.405	4.58	1.86
5	0.343	3.60	1.23
6	0.305	2.60	0.79
7	0.349	2.06	0.72
8	0.114	1.72	0.20
9	0.305	1.58	0.48
10	0.235	1.51	0.36
Sum		53.58	31.01

Table 4. Performance of CCA technique in reconstructing observations from reanalysis data for the period 2002-2004.

Weather Element	% Variance in *normalised* obs. explained by 30 obs. principal components	% Variance in *normalised* obs. explained by CCA on 30 reanalysis principal components	% Variance in *actual* obs. explained by CCA on 30 reanalysis principal components	RMSE between original and reconstructed observations.
TT	73.34	40.96	73.54	2.38°C
RH	69.80	32.45	32.49	12.9%
FF	53.18	16.16	18.57	8.7 km/h
RR	73.26	38.10	25.99	8 mm
Combined	67.39	31.91	37.65	

Looking more closely at the results indicated that, as expected when forecasting anomalies, the average value over the period was well modelled by the system. However the variance within the recreated data was approximately half of the original. Because the forecasts are initially expressed as deviations from the mean it would be relatively easy to artificially increase the variance by simply scaling up the forecast anomalies. This would of course incur a penalty in the form of significantly increased RMSE values, although the percent variance explained would not change.

If these were the only drawbacks to matching the variance of the forecast and observed data sets, it could make sense to trade off accuracy in the long term averages against producing values that were typical of observed extremes. However for relative humidity, wind and especially rainfall, increasing the variance led to large numbers of non-physical forecasts such as negative rain and wind speed. Also, bearing in mind that least squares minimisation lies at the heart of both PCA and CCA, there was little to be gained in tuning the variance at the expense of the RMSE score.

5.2 Treating Weather Elements Individually

In an effort to see if the amount of variance explained could be increased, the PCA of the observations was revisited. Instead of treating the midday observation as a combination of TT, RH, PP, and FF, principal component analysis was performed on each of the elements *independently* of the others. By doing this links between elements in a particular situation, such as 'wet and warm', would not be enforced and separate patterns for 'wet' and for 'warm' would eventuate. For example, by separating the elements it was hoped that for situations such as 'showery southwesterlies' the strong relationship between TT and T850 would not be weakened by the showers which may only fall 50% of the time.

As expected, the variance explained by principal components for individual weather elements did increase. Overall the variance explained rose from 67% when the principal components were derived from the four weather elements taken as a set, to 82% when each was considered individually. The principal components found are shown in Figure 9. Note that because they were derived for individual weather elements, a synoptic situation with a strong first relative humidity component, would not necessarily also be strong in the first temperature component. In fact, Figure 6a suggests that the sign for temperature is likely to be reversed.

Having obtained principal components for each element individually, the CCA procedure was repeated to establish a relationship between the reanalysis data and each individual weather element. From the results of the CCA an estimate, similar to that shown in Table 4, of the potential of the reanalysis data to forecast each element was made. Although the PCs for individual elements explained more variance within that element, the estimate of explained

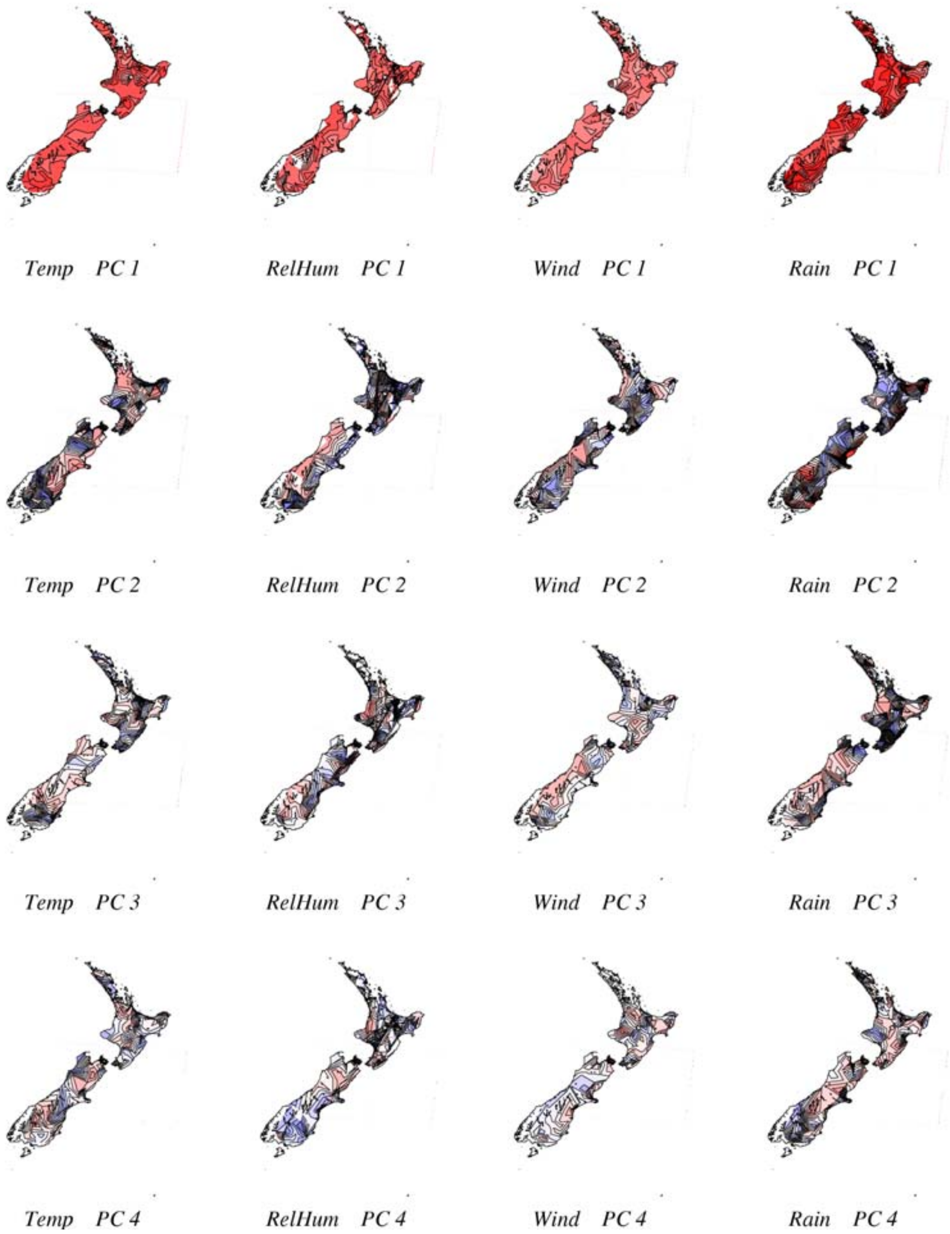

Figure 9. The first four principal components derived for individual weather elements from the climatology of Pearce et al. (2003).

variance explained by CCA rose only 2%. This reflects the substantial inter-correlation between the patterns of individual weather elements for a given weather situation which is exactly what the CCA was trying to find.

Forecasting the weather elements individually removes the inherent consistency between elements guaranteed by using a combined observation in the original PCA of observations. The method was also considerably more complex to implement. For these reasons, and in light of the very modest gains offered, it was decided to continue with the original approach where the four weather elements were treated together as a single observation.

6. CCA Based Forecasts of Fire Weather

Having established the extent to which it was possible to reconstruct or 'forecast' observations from the reanalysis data, the next step was to confirm that those relationships could also be applied to forecasts from the NCEP ensemble prediction system.

6.1 NCEP Ensemble Prediction System Forecasts

Ensemble forecast systems have been developed in an attempt to recognise the impossibility of exactly specifying the initial conditions for an atmospheric model. It is also known that the modelling of physical processes within numerical models is only approximate. The small differences between a model's start point and the true state of the atmosphere, and the approximations of reality in its governing equations, interact in a non-linear fashion. Initially the growth of differences between the ensemble members tends to be slow and linear. With time the differences begin to interact with each other and the growth becomes chaotic. By starting many separate versions of a model with a representative selection of possible initial conditions, and in some cases also stochastically varying the physical schemes within the model, it is possible to sample the range of possible outcomes. With information about a number of different scenarios, it is hoped to recognise those features that are only changing slowly, and are therefore more predictable. These features should be most visible in the average of all the ensemble members. It should also be possible to detect occasions when the atmosphere is delicately balanced, in which case chaotic behaviour would rapidly set in limiting the period for which forecasts are useful.

The National Centers for Environmental Prediction (NCEP) has been running an ensemble forecast system since 1994 (Toth and Kalnay 1997; Zhu et al. 2002), and MetService has been receiving data since early 2002. Each run of the system extends out to 16 days, with a total of 12 ensemble members. Ten of the members are formed from perturbations added to the operationally produced analysis of initial conditions. Two control members, one at high resolution and the other at the same, lower, resolution as the perturbed members, are started from the unperturbed analysis. The NCEP system does not alter the physics within any of the member models. In order to conserve computing resources the resolution of the models is reduced at longer lead-times.

NCEP uses a spectral model as the basis of their ensemble system. Spectral models solve the equations of motion using spherical harmonics as opposed to using rate of change at fixed grid points. This is computationally more efficient than a regular grid-point model for global models at the current horizontal resolutions. Spectral model resolution is typically expressed with the number of levels (e.g., T126L28 - where 'T' is the spectral resolution or maximum number of waves resolved around the circumference of the earth, and 'L' is the number of model levels). There are a number of ways of translating this to an effective horizontal resolution with half the smallest wavelength often being used. As computing power increases the resolution of models generally becomes finer and improvements are made to their internal modelling of the physical world.

During the period of this study the configuration of the NCEP ensemble system was relatively stable, having undergone changes to the way perturbations were calculated in April 2003. In March 2004 the model increased from running twice a day to four times a day, but still at the same resolution, with the same number of members. With the introduction of new hardware a significant increase in model resolution and number of vertical levels occurred in August 2005, together with slight changes to the way the initial perturbations are made. These improvements are expected to slightly increase accuracy in the first few days, and also in the

Table 5 Horizontal and vertical resolution of the members of the NCEP ensemble during the period of this study compared with the resolution after an increase in computing power in August 2005.

Horizontal Resolution			10 Members + low res. control		High res. control	
Spectral	km	Deg.	2001 to 2005	Sep 2005	2001 to 2005	Sep 2005
T382L64	50	0.5				To Day-7.5
T254L64	80	0.7			To Day-3.5	
T190L64	105	0.9				To Day-16
T170L42	120	1.0			To Day-7.5	
T126L28	160	1.5	To Day-7.5	To Day-16	To Day-16	
T62L28	320	3.0	To Day-16			

second week. Table 5 summarises the recent changes in the resolution of the NCEP ensemble system.

In this study the high-resolution control is included in the ensemble without giving it an increased weight compared to the other, lower resolution, members. The reason for this is that determining an appropriate weighting method is not straightforward. In any case, during the early part of the period the low resolution control, which starts from exactly the same initial conditions, will tend to follow and add weight to the solution found by its high resolution counterpart.

The ensemble forecasts that MetService receive are on a 2.5°×2.5° latitude/longitude grid, regardless of the resolution of the underlying model. Only forecasts from the 0000 UTC (midday local time) run were used in this study as they become available around midnight New Zealand time allowing new forecasts to be made before the start of the working day. The higher resolution at which the models run, compared to the resolution they are received at, results in better description of small-scale interactions which gradually change the evolution of the larger scale features. There is little information lost in receiving coarse resolution data, since beyond a day or so the absolute location and timing of small scale features is unlikely to be exactly correct, leading to the forecast being 'precisely wrong' on the fine scale. Upgrades in model resolution and physics schemes also can change the absolute values of quantities, particularly near the surface. The key point here is that although the forecast of a quantity such as temperature at a point may now be closer to the observation, the change may require the statistical relationships to be redeveloped. This would require a period of saved data that is typically longer than the interval between significant changes to the model. Recognising that resolution changes were inevitable played a significant part in the choice of variables used to describe the weather situation. The model variables selected were chosen because they are relatively insensitive to small-scale surface features. Because the variables are in a sense 'generic', it increases the likelihood that the technique could be successfully applied to forecasts from other models such as the ECMWF ensemble system.

6.2 Forecasts from a Single Ensemble Member

Initially, CCA forecasts from a single ensemble member were compared with observations reconstructed using CCA on the reanalysis data. The reanalysis data and ensemble data are received on the same grid, and the same variables at the same vertical levels are available.

The use of the perfect prognosis technique for the fire weather forecasts assumes that the statistically derived predictions from a given weather situation are independent of the lead-time of that forecast. In contrast, a model output statistics (MOS) approach would assume that as time goes on it is better to 'play it safe' and tend towards climatology as lead time increases. The test of transferability of the CCA relationship from the reanalysis dataset to the ensemble forecasts needs to avoid differences arising from the skill of the model, implying that only very short-term forecasts should be compared with the reconstructed

Table 6. Comparison of RMSE values for CCA based temperature forecasts using different individual members of the NCEP ensemble for the period June 2003 to May 2004 averaged across 137 stations. CH is the high resolution control, CL the low resolution control, and N1 is the first negatively perturbed ensemble member.

	NCEP	Day-0 Forecast			Day-1 Forecast		
	Reanal	CH	CL	N1	CH	CL	N1
Ave. temperature RMSE	2.357	2.351	2.347	2.376	2.417	2.417	2.471
95% confidence limit	0.060	0.060	0.060	0.060	0.061	0.061	0.061

observations created from the reanalysis. Areas where the model could still differ from the reanalysis include factors such as model resolution, interpolation from the native model space to the grid points that are disseminated and potentially slightly different representations of the variables themselves between the reanalysis and the forecasts.

The technique used by NCEP to generate perturbations to the initial conditions is optimised to produce differences that, on average, grow to match observed errors at a forecast time of 48 hours. Adding one positive and one negative perturbation to locations where the model is particularly sensitive creates a pair of ensemble members. For a number of reasons ensemble systems are generally under-dispersive, and growth of the differences is slower within the model than is observed in reality. In order to get variations of the size required by day-2, the initial perturbations are made larger, causing the early period perturbed forecasts to verify slightly less well than the control forecasts. For these reasons, the unperturbed control runs are the most appropriate ensemble members to compare with the reanalysis.

Forecasts and observations from April 2003 to June 2004 were run through the system and the average RMSE values from several individual members of the ensemble compared for all variables. The results for temperature are shown in Table 6, and indicate that the control runs very closely follow the reanalysis data. The small difference that does occur is greater between the reanalysis and the two control runs than between the control runs themselves. This may reflect the slight difference between the operational analysis scheme at NCEP, and the 1996 version that was held static and used for the full 40 years of historical data that went into the reanalysis project. At day-1 errors have increased slightly indicating the growth of error in the models. The perturbations added to the N1 member can be seen to be doing their job, with that member differing more from the reanalysis data than the two control members.

6.3 Comparison with Other Forecasts from Single Models

It was difficult to compare the results from the CCA based forecasts with those from other forecast methods used at MetService due to differences in the sets of stations used for verification, and locating verifications that relate only to midday forecasts. Table 7 summarises the closest comparable verifications that could be found. The MM5 (Fifth-Generation NCAR/Penn State Mesoscale Model) forecasts are from the 12 km resolution, mesoscale model run at MetService before and after statistical correction (MOS) of the raw

Table 7. Root mean square errors of CCA forecasts averaged over day-1 to day-3 based on individual members of the NCEP ensemble, compared to similar RMS errors from other forecast models available at MetService.

	CCA on CH Midday Fcsts 137 Stations	CCA on P1 Midday Fcsts 137 Stations	Raw MM5 Midday Fcsts 17 Stations	MOS on MM5 Midday Fcsts 17 Stations	Raw GFS Tmax Fcst 17 Stations
Temp. Jun03-May04	2.44°C	2.51°C	2.46°C	2.04°C	
Temp. Sep02-Aug03			2.43°C	2.02°C	3.25°
Wind Jun03-May04	9.02 km/h	9.10 km/h			
Wind Sep02-Aug03			9.34 km/h	7.44 km/h	

model output. The forecasts are for 17 sites, mainly airports, around New Zealand, most of which are also in the set of 137 stations forecast by the CCA method. The New Zealand Fire Service currently receives daily maps of Fire Weather Indices based directly on the raw output from the MM5 model, which it displays on its web site. The NCEP Global Forecast System (GFS) forecasts are for daily maxima interpolated from exactly the same model as the ensemble high-resolution control, but as received by MetService on a finer, 1°×1° grid. Because the exact timing of a daily extreme is not part of the forecast, verification scores of daily extremes are usually slightly better than those for temperature at a particular time of day. It is not known whether this is in fact the case with GFS daily maximum forecasts.

Overall, the CCA forecasts based on gridded data received at about 250 km resolution are shown to be similar in terms of RMSE to raw 12 km resolution forecasts from the mesoscale model. After statistical post processing the MM5 based forecasts are significantly improved. The CCA forecasts are better than simple interpolation from the GFS and, although not shown here, seem to be about the same as statistically corrected GFS forecasts. The slight increase in RMSE for the lower resolution P1 member forecast compared to the higher resolution CH forecast is also evident.

7. Ensemble Forecasts of Fire Weather

The CCA relationships derived using reanalysis data have been shown to also apply to individual members of the NCEP ensemble. The resulting forecasts are comparable in skill to the uncorrected output from a much higher resolution mesoscale model. It remains to be shown that considering the ensemble as a whole adds to skill, and also that the results for weather elements transfer to forecasts of fire weather indices.

The original methodology for calculating fire weather indices for Canadian forests is described in Van Wagner (1987). Building on that work, Alexander (1992, 1994) adjusted the indices to better reflect New Zealand conditions. The moisture indices depend in part on the previous day's value. This means, depending on the nature of the fuel that is drying, the effect of previous conditions takes some time to be erased by environmental forcing. The algorithms used in this project were checked for consistency against those in use by the New Zealand Fire Service. The following indices (Van Wagner and Pickett, 1985) were calculated:

FFMC Fine Fuel Moisture Code, which represents the moisture content of litter and other cured fine fuels. Has a time lag of around 1 day.

DMC Duff Moisture Code, which represents the moisture content of loosely compacted decomposing organic matter. Has a time lag of around 12 days.

DC Drought Code, which represents a deep layer of compact organic matter. Has a time lag of around 52 days.

BUI Build Up Index, a combination of DMC and DC that represents the total fuel available to a spreading fire.

ISI Initial Spread Index, a combination of wind and FFMC that represents rate of spread alone, without the influence of variable quantities of fuel.

FWI Fire Weather Index, a combination of BUI and ISI that represents the intensity of the spreading fire as energy output rate per unit length of the fire front.

DSR Daily Severity Index, a power transformed version of FWI that better represents the non-linear increase in work required to control a fire as FWI rises.

WSR Weekly Severity Index, an average of DSR over the previous seven days.

There were a number of days for which ensemble forecast data were missing. In each case the missing data affected all forecast periods from that run. This was not a significant issue because each forecast was initialised with actual values based on observations, which were present for all dates. As long as there was data for a run, forecasts could be made for all periods out to 16 days. In an operational setting there would be a number of ways to mitigate the impact should the ensemble data not be received by MetService. A relatively simple option would be to use data from the run starting six or twelve hours previous to the run that did not arrive.

The most straightforward forecast available from an ensemble system is the average of all its

Table 8. Percent variance in fire weather indices explained by CCA forecasts from the first positively perturbed member (P1), the high-resolution control (CH), and the average of the full 12 member NCEP ensemble (EN) for the period June 2003 to May 2004.

	D+03			D+06			D+09			D+12		
	P1	CH	EN	P1	CH	EN	P1	CH	EN	P1	CH	EN
BUI	79.4	81.6	81.4	61.8	66.2	67.6	44.4	50.1	54.1	35.9	38.8	46.9
ISI	23.4	27.2	25.9	12.1	16.6	19.6	5.5	7.8	12.2	2.2	2.4	6.7
FWI	41.8	46.5	44.4	27.5	33.1	37.3	16.4	20.5	25.7	11.5	12.4	20.9
DSR	32.7	37.0	35.5	21.0	26.2	29.7	12.6	16.0	20.7	7.6	9.8	15.2
WSR				63.3	66.4	67.1	40.9	47.3	49.4	29.0	33.8	40.0

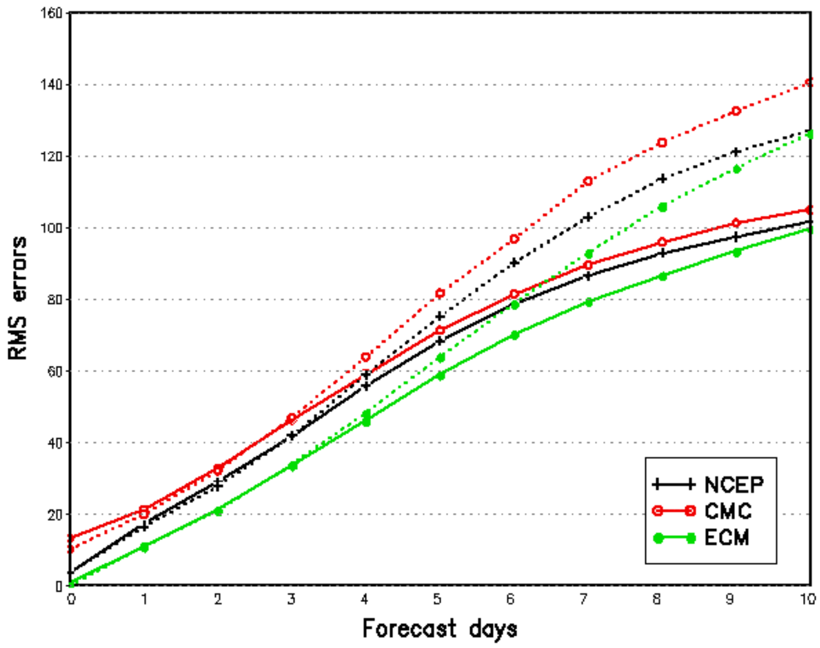

Figure 10. Verification scores for forecasts of Southern Hemisphere 500 hPa height for Jun-Aug 2005 from three centres running operational ensemble prediction systems: National Centres for Environmental Prediction (NCEP); Canadian Meteorological Centre (CMC); European Centre for Medium Range Weather Forecasting (ECM). Each ensemble limited to 10 members for this verification. Dotted lines are forecasts from a single control run. Solid lines are the ensemble average forecast. From Zhu (2005).

members. To assess the ensemble average forecasts of fire weather, and to quantify the difference that resolution makes to a forecast, CCA forecasts based on data from a single, low resolution, perturbed member of the ensemble (P1) were compared with the high resolution control (CH), and also with the average value from the complete 12 member ensemble (EN). In order to compare results from indices with different typical values, the verification measure chosen was the percent variance in observed indices explained by the forecast indices. Table 8 shows that the higher resolution control run consistently outperforms the lower resolution perturbed member. In fact out to day-3 it is a little better than the ensemble average for all indices. This is in large part due to the artificially increased perturbations that are applied by NCEP in an effort to get an appropriate degree of spread in later prognosis periods. By day-6 the ensemble average is performing better than the single control member, and this improving trend continues as the forecasts extend into the future. Figure 10 shows this result is common to other ensemble systems, and reflects the focus the designers of the forecast systems have on forecasts for days 5-10.

7.1 Impact of Persistence

The influence of the past drying conditions on the predictability of some fire indices shows up clearly in Table 8. The BUI, which depends only on the slowly changing DMC and DC, shows the most explained variance. The extent to which this predictability is due solely to the continued influence of initial conditions is examined in Table 9, where a persistence forecast of BUI, made by retaining the value observed on the day of the run, is compared with the ensemble average forecast. During the first half of the forecast period, the ensemble forecast of BUI steadily builds a lead in terms of variance explained over the persistence forecast. Later in the forecast period, where the skill from the individual weather elements has dropped away, the early advantage is simply maintained. It can also be seen from Table 9 that the explained variance for a persistence forecast at day-12 is the same as that in Table 8 for a

Table 9. Percent variance explained in BUI by a forecast of persistence, and by the CCA ensemble average forecast over the period June 2003 to May 2005.

	Forecast Lead-time														
	D+00	D+01	D+02	D+03	D+04	D+05	D+06	D+07	D+08	D+09	D+10	D+11	D+12	D+13	D+14
Persistence	100.0	92.1	82.2	73.1	65.3	58.7	53.3	48.7	45.0	42.0	39.6	37.4	35.2	32.6	29.9
Ens. Ave.	100.0	94.0	87.3	81.4	76.5	72.3	67.6	62.6	57.8	54.1	51.3	49.0	46.9	45.0	43.2
Difference	0.0	1.9	5.1	8.3	11.2	13.6	14.3	13.9	12.8	12.2	11.7	11.6	11.7	12.4	13.3

single ensemble member. This highlights the value of ensemble averaging – since the single ensemble member has no skill over persistence but the average does.

To assess the impact of persistence across all the variables, the lagged correlation between the observations of each element for 2003-04 was calculated. A high level of serial correlation effectively decreases the number of truly independent observations. Table 10b shows the results of this analysis and indicates the point for each variable where the variance explained by the forecast is no longer significant at the 5% level. The results are both interesting and somewhat counterintuitive. The correlation for temperature, which is a 'well-behaved' variable in that it is relatively continuous in both time and space, is still the most significant despite only 48 effectively independent verification pairs remaining after allowing for serial correlation. Wind speed, and the ISI which depends on it, benefit from the relatively small amount of association between values from day to day, with the variance explained (which is equal to r^2) remaining significant at the 5% level for most of the second week despite being less than 10%. Conversely the correlations for DC and DMC indices, which explicitly include persistence as a significant factor, loose significance at about the same lead-time as the wind speed despite having values of explained variance around 50%.

In terms of the absolute percent variance explained, rainfall and the weekly severity index benefit somewhat from being aggregated variables since the importance of getting the timing of changes exactly right is somewhat reduced. Rainfall is still very variable from day to day, however the serial correlation within the WSR is the highest of all the variables due to the explicit averaging of seven consecutive values. This means that despite having a comparatively high degree of correlation between forecast and observed values of WSR, the relationship is not significant at the 5% level much beyond the first week.

The spatial patterns in explained variance between forecasts from the high resolution control (Figure 11) and the whole ensemble (Figure 12), confirm the results of Table 10. The consistent but relatively small gain the ensemble shows over the high resolution control is evident, and it also indicates that a greater percent variance in most indices is explained in eastern regions, compared to the relatively even climates in the west. This last result is because typically there is more variance available to explain in the east compared to the west. In an absolute sense however, most stations in the west do show smaller RMS errors than eastern stations.

The care required in interpreting the percent variance results is also evident when they are compared with RMSE values, as is done in Table 10 a) and b). Results for temperature forecasts, boosted by the approximately 50% seasonal component in the variance, behave as one might expect them to; a slow increase in RMSE and a corresponding gradual decrease in percent variance explained. Both measures plateau after about 10 days indicating that this is effectively the limit of the non-seasonal predictability for temperature. Although less obvious, the pattern is much the same for relative humidity. For rainfall and wind speed however, the change in RMSE is significantly less compared to the change in variance. In part this is thought to be a result of the decision to train the CCA forecasts on power transformed versions of rainfall (cube root) and wind speed (square root). In each case the

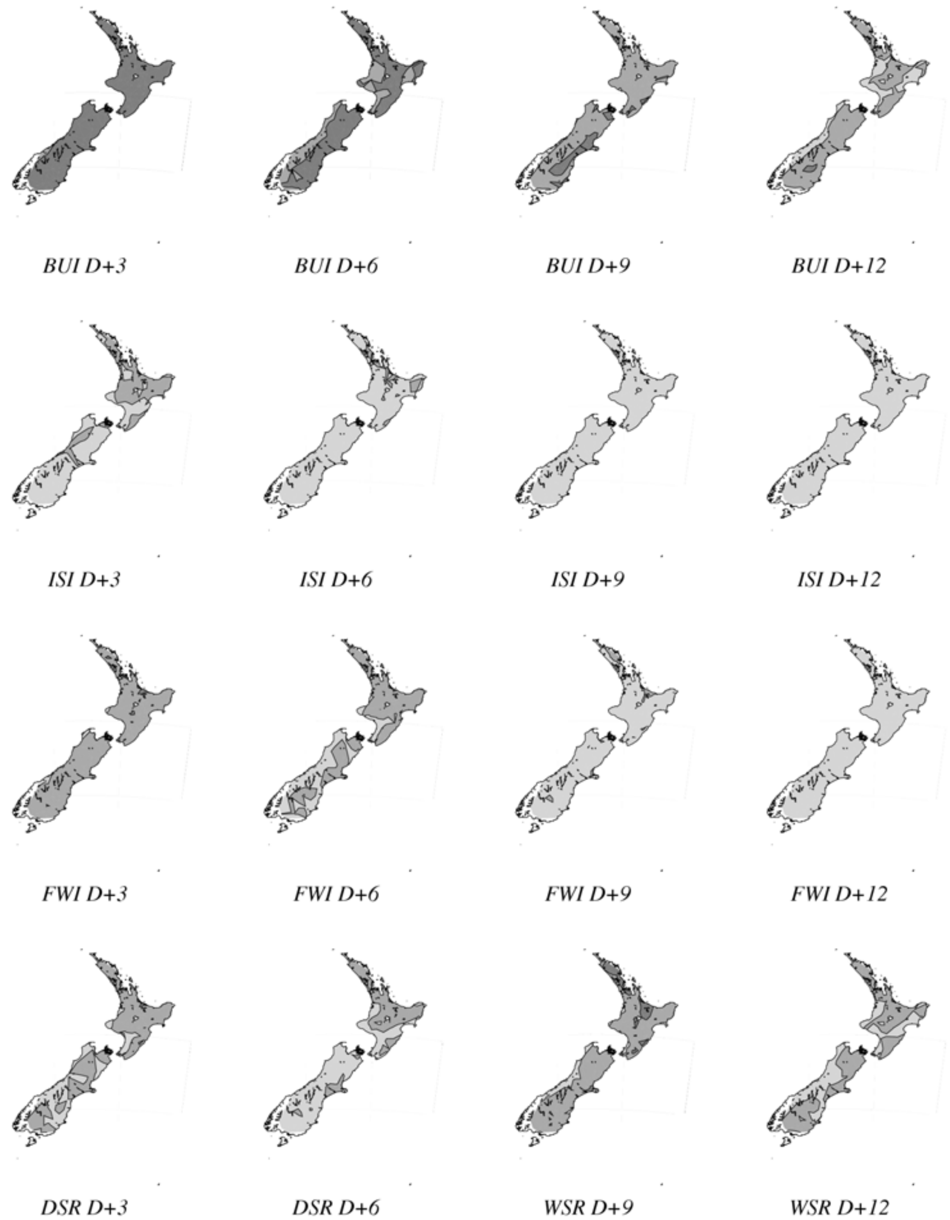

Figure 11. Percent variance explained in fire weather indices, averaged over the period Jun 2003 to May 2005, for CCA forecasts from the single high resolution control member from the NCEP ensemble. Shading is in 33% intervals. Note there are no forecast sites available in the Fiordland area, and DSR is replaced with WSR after day-6.

influence of extreme values was reduced in an effort to improve discrimination of values that were more typically observed. In the case of rainfall for instance, the influence of the occasional 100 mm/day event would overshadow the more common variation between a few millimetres of rain and none at all.. A comparison between the overall distributions of forecast rainfall amounts across all stations showed that the CCA technique predicted virtually no rain events above 30 mm/day, but did have a climatologically appropriate discrimination between events of greater and less than 1 mm/day.

For rainfall, where the difference between rain and no rain is important for fire weather, this approach seems to be appropriate. For wind speed the case is not so clear-cut. The ISI is heavily influenced by wind speed, and particularly so by high winds. In its turn, the ISI is a key driver of the FWI and the derived DSR and WSR. Extremes in these indices are usually

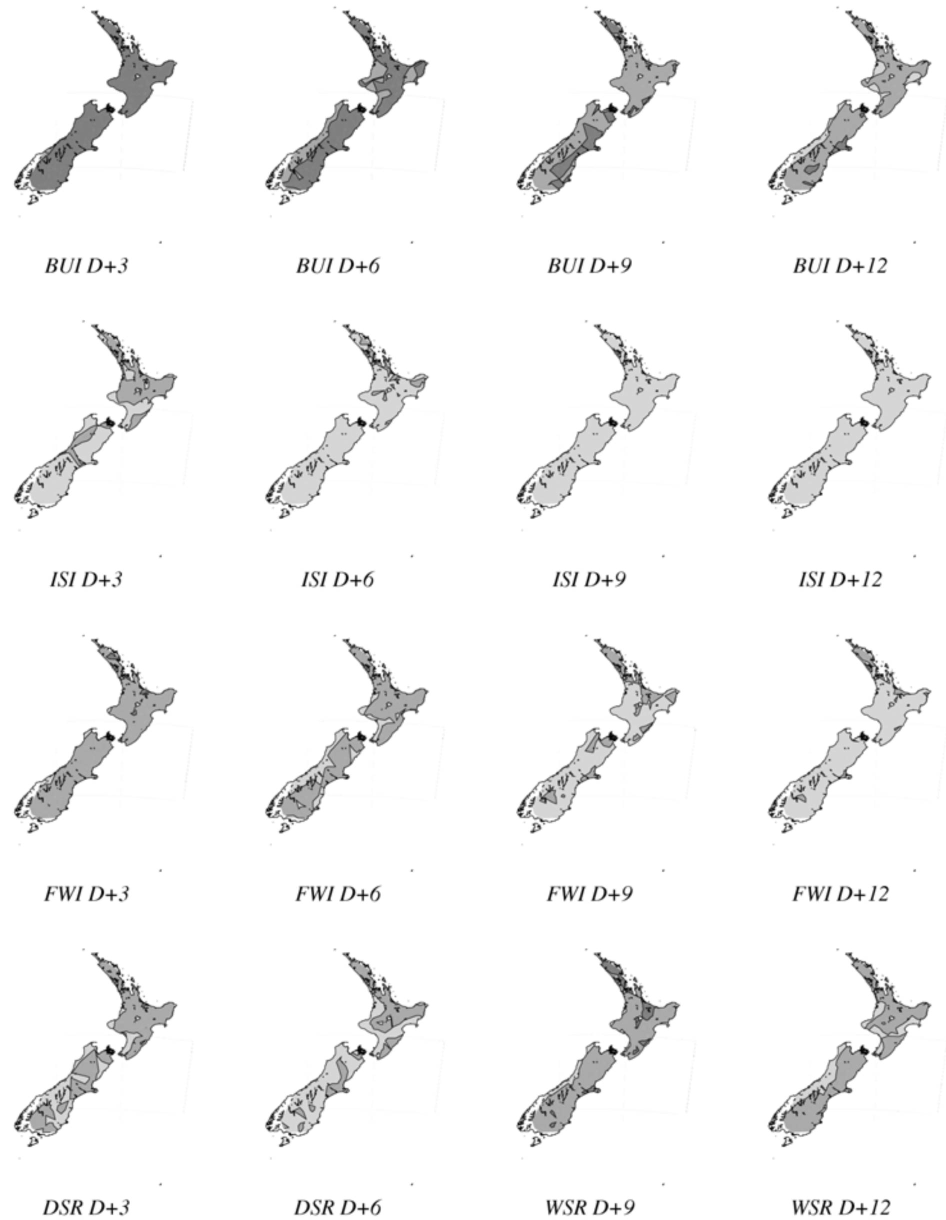

Figure 12. Percent variance explained in fire weather indices, averaged over the period Jun 2003 to May 2005, for CCA forecasts from the single high-resolution control member from the NCEP ensemble. Shading is in 33% intervals. Note there are no forecast sites available in the Fiordland area, and DSR is replaced with WSR after day-6.

due to high wind speeds on a day preceded by a period of dry weather. It has been noted before that forecasting wind speed is particularly difficult, and the results of using the CCA method on 2.5° gridded data are, in RMSE terms, comparable with the raw output from a much higher resolution numerical model. However, given the sensitivity of fire weather to high wind speeds, it may be more appropriate to focus on identifying extreme winds rather than looking to discriminate between more common ones. This could be done by either not transforming the wind speeds at all, or even raising them to a higher power, instead of working with the square root. Unfortunately time did not permit investigation of this during this project, but any future work should reconsider the need to power transform the wind speed when developing the CCA relationships.

7.2 Correction for Bias

Another artifact of the power transformation of wind speed and rainfall was that the average of forecast values did not equal the average of observed values. Insufficient rainfall particularly affected the moisture codes of the fire weather indices, which tended to dry out during the forecast period. Because high wind speeds were not forecast often enough, wind speeds were, on average, biased 3 km/h too light. A closer look at the average errors indicated that there was also a slight seasonal trend in the bias of relative humidity and wind speed. The impact of these errors was to shift the absolute value of the fire weather indices at the later forecast periods to unrealistically high values (Figure 13). This is of considerable practical significance because although the direction of change, as scored by the variance explained, was giving useful information, users of the forecasts would find it difficult to factor in the trend towards higher fire risk as the outlook period increased.

In order to address these issues a simple correction to inflate the absolute amount of rainfall forecast to match the total amount observed was introduced. At the same time, bias in a moving window over forecasts from the previous 14 days was removed from the other weather elements. The biggest effect of these measures was to align the average values of the forecast fire weather indices more closely with what was observed (Figure 13). There were also improvements in the RMS errors and variance explained for the fire weather elements (Table 11). The impact of the changes to the verification scores for the base weather elements was less noticeable. However it was still significant, with the short-range temperature forecasts improving in terms of RMSE almost to the same level as the statistically corrected forecasts from MetService's mesoscale model. The small improvements in variance explained for wind speed and relative humidity are likely to be due to the capturing of the slight seasonal character shown by these elements

It should be emphasised that the correction for bias does little to improve the information content of the forecasts, but the improvements do result in forecasts that are much easier to interpret. If the forecasts were adjusted by either a constant scaling factor, or by a constant bias correction, theory dictates the variance explained would remain steady although the RMSE could change. In the case of rainfall the variance explained did vary slightly between Table 10 and Table 11 due to truncation effects near zero rainfall.

After applying the corrections, there remains a slight tendency to drier forecasts at the longer lead-times. The adjustment to rainfall was calculated by simply making the total forecast rainfall over all sites for the year equal what was observed. It would be useful in future work to investigate whether results could be improved by applying a more tailored correction based either on the records at each site, or perhaps on the amount of rainfall forecast in a 24 hour period.

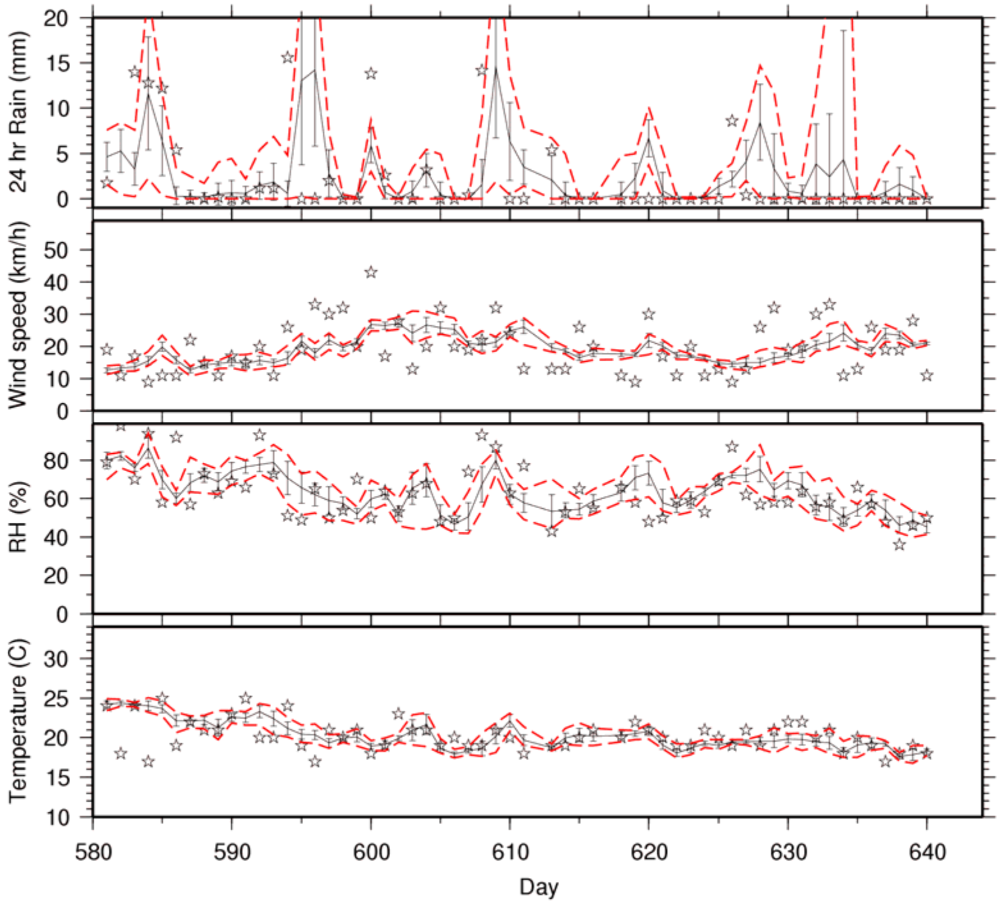

Figure 14a. Day-5 CCA forecast weather and observations for Tauranga between 1 Feb 2004 (day 580) and 31 March 2004. Solid line represents the ensemble average, bars the ensemble standard deviation about the average, and dashed lines are the bounding forecasts from the highest and lowest valued member forecasts. Verifying observations are marked by stars.

7.3 Forecasts for 2003-04 Fire Season

Having established the performance of the system over a 12-month period, attention was then turned to the 2003-04 fire season. The reason all months were used in the development of the initial CCA relationships was to maximise the number of weather situations in the training set.

The verification scores for the fire season forecasts (Table 12) are all slightly poorer that those for the full year. This was not unexpected, since during the summer months the amount of solar radiation received at the surface of the earth is more variable. During the winter the fire weather indices are typically lower valued and vary less. The daily severity index, which responds more to higher values of FWI, is particularly muted and more predictable in winter.

In summary, ensemble average forecasts created from applying the CCA relationships to members of the NCEP ensemble show an ability to explain significant variance in the day to day fire weather indices at least as far as day-10. Beyond day-10 the variance explained is due mostly to the seasonal variation of temperature, and to persistence within the moisture codes of previous drying conditions. Forecasts of the weather variables at all lead-times do not differ from their climatological mean values as much as would be expected. This is most

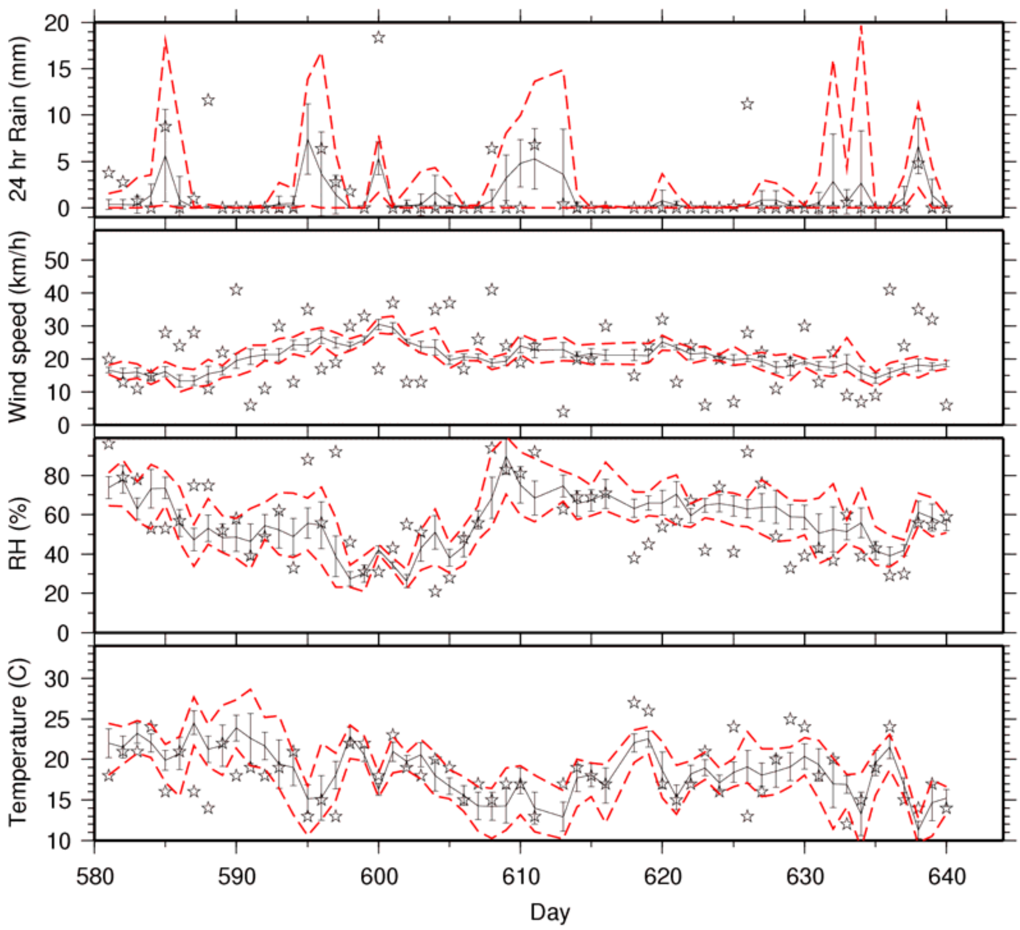

Figure 14b. Day-5 CCA forecast weather and observations for Christchurch between 1 Feb 2004 (day 580) and 31 March 2004. Solid line represents the ensemble average, bars the ensemble standard deviation about the average, and dashed lines are the bounding forecasts from the highest and lowest valued member forecasts. Verifying observations are marked by stars.

notable for rainfall where, although the discrimination between days with some rain and no rain is about right, there are very few forecasts of large amounts of rainfall. Correcting for bias helped reduce the tendency for fire weather indices to dry out at longer lead-times. For forecasts out to day-3 the RMS errors from the ensemble average forecasts were only slightly worse than statistically corrected forecasts from MetService's high resolution mesoscale model.

7.4 The Spread Skill Relationship

An ensemble system allows more than a simple single-valued forecast for a given element to be made. By considering the range of results from the ensemble, probabilistic forecasts of various kinds can be produced. Examples include the probability that a set value might be exceeded, an estimate of the likely standard deviation about the mean value, or identifying a scenario that might only be exceeded a certain percentage of the time. Figure 14 shows how some of these values might be displayed for the forecast at a single site, with the average and standard deviation of the ensemble forecast plotted together with the highest and lowest valued forecasts. To give an idea of the temporal and spatial variation in predictability, the values of these quantities for all stations could be displayed equally well using a separate map of New Zealand for each forecast day.

At the heart of the potential value available from the spread of values in an ensemble is the assumption that a narrow spread indicates increased certainty in a particular outcome, while the reverse is true if the ensemble members vary greatly between themselves. On average the degree of spread will increase with prognosis period until it matches the value of spread obtained by repeatedly sampling a climatology of forecasts. If the forecasts are modelling the growth of differences correctly this value should be the same as the climatological spread in observations. If there is a link between spread and skill it would be possible on some occasions to identify that there was as much confidence in the day-10 forecast as there usually was in a day-6 forecast.

To test this widely held assumption, the error in the day-5 CCA based ensemble average forecasts of temperature at all locations for the period June 2003 to May 2004 was compared with the ensemble spread. As can be seen from comparing Figure 14a with Figure 14b temperatures at places like Auckland or Tauranga typically vary only a little, and because of this have smaller errors on average than places such as Christchurch where temperatures may vary by 10°C between two days. To allow comparison between stations that had different levels of natural variability in temperature both the spread and error were normalised about their mean values at each station and then plotted in Figure 15. The distribution of spread is slightly positively skewed, reflecting the fact that spread is bounded at an absolute value of zero but has no upper bound, while the distribution of error is essentially Gaussian. What is of most interest, is that the line of best fit is almost indistinguishable from the horizontal line through the average error. The extremely low correlation coefficient confirms that, taking all stations together, there is no real link between spread and skill for temperature forecasts. An examination of other variables and individual stations confirmed the lack of correlation applied more generally.

Whitaker and Loughe (1998) offer insights as to how this disappointing result arises. They found that NCEP ensemble forecasts of 500 hPa geopotential height, a very smoothly varying and 'well behaved' quantity, also showed a relatively small correlation between spread and skill. The greatest correlation (r^2=0.08) occurred at day-5. Their study indicated that the value in the spread information came mostly where the values of spread were extreme. However they also found that the typical values of spread themselves varied with season and between years, making it difficult without a very long period of data to determine what 'normal' levels of spread were.

Any single valued forecast can be turned into a probabilistic forecast by using historical errors to estimate the uncertainty in a particular forecast value. This method could be applied to the NCEP ensemble average forecasts and would identify that day-3 forecasts were more reliable than day 10 forecasts, and also that forecasts for places like Auckland tended to be better than those at places like Christchurch. The key difference between these probabilistic forecasts and those that could be made if there was a relationship between spread and skill, is that they can not give information about whether the skill in the forecast varies from day to day.

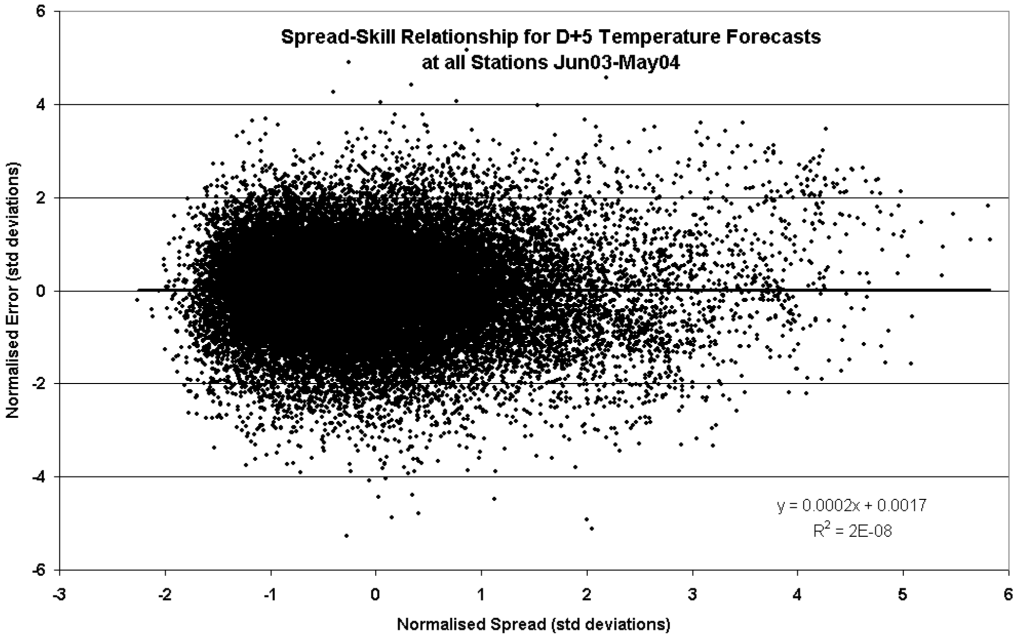

Figure 15. Error in CCA based ensemble average forecasts of day-5 temperature (normalised at each of 137 stations), compared with the normalised standard deviation between ensemble members.

There is some hope that useful ensemble based probability forecasts may become available in the future. The recent upgrade to the resolution of the NCEP system should increase the skill of the member forecasts and may strengthen the spread skill relationship. However one of the factors limiting the usefulness of the NCEP ensemble is its relatively small number of members. Slight gains in skill have been shown (Toth and Szunyogh, 2000) when members from a previous run have been included, but the drop in accuracy from including old forecasts is only made up later in the period. There are plans to combine the members from the Canadian and American ensemble systems before the end of 2006 which would more than double the number of members.

Another option is to use data from the European Centre for Medium Range Forecasting's (ECMWF's) 51-member ensemble system. Verifications (e.g., Figure 10) have consistently shown their ensemble to be the best in the world due to factors such as the high resolution of their model and the stochastic variation of physics schemes between ensemble members. MetService purchased access to this data at the end of 2004, but there was not enough archived data available for the ECMWF ensemble to be included in the current work. The CCA relationships developed for this report were designed with the possibility of using forecasts from other models in mind. Any small systematic differences between forecasts from NCEP and ECMWF should be largely resolved by the correction of bias based on the moving 14 day window.

8. Conclusions

Verifications using 12 months of independent data have shown the ensemble average CCA forecasts of fire weather indices have skill out to day-10. Beyond this lead-time the true skill of the forecasts is low, however they still contain useful information about typical seasonal variation of temperature, and the persistence of the initial drying conditions, out to day-15. The forecasts are produced for 137 sites around New Zealand from using output from the NCEP ensemble prediction system, which is received with a spatial resolution of about 250 km. Verifications show that for days one to three the skill of the ensemble CCA is comparable with the raw output from the 12 km resolution mesoscale model run by MetService. After applying a bias correction, the CCA forecasts are only marginally worse in terms of RMS errors than statistically corrected forecasts from the mesoscale model.

The forecasts show a tendency to dry out during the forecast period, due mostly to problems forecasting rainfall amounts. Because of this, fire managers using any forecasts based on this scheme should be aware that the day-10 forecasts of fire weather, on average, will indicate a higher fire risk than is actually observed. Applying a simple correction to the rainfall forecasts reduced this tendency somewhat, but it is likely a more targeted approach could reduce the difference further.

The ensemble CCA forecasts demonstrated a very poor relationship between the spread within the ensemble and the skill of the ensemble average. Unfortunately this meant it was not possible to generate probability forecasts that recognise differences in the predictability of weather situations. It had been hoped that some types of weather, perhaps such as large anticyclones, might have been more predictable than others, and that this information could have been passed on in the form of an expression of confidence. Part of the reason for the poor relationship between spread and skill is believed to be the relatively small number, 12, of members in the NCEP ensemble.

There are several ways in which the forecast information might be displayed to fire managers. Through the MetConnect web pages MetService already routinely produce charts showing the spatial variation around New Zealand of fire weather indices using data from their mesoscale model. The ensemble forecasts could be used to generate similar maps that extend coverage from day-3 out to day-10. The same data could easily be made available in table form. Forecasts of Weekly Severity Rating were found to explain nearly 25% of the variance in the observed index as far out as day-15. While these forecasts are not truly skilful in that they represent the seasonal variation of temperature and the persistence of conditions early in the period, fire managers may be interested in the overall view of fire conditions two weeks ahead. Until forecasts with a demonstrated ability to distinguish between weather situations that are more and less predictable are available, line charts showing the spread of the ensemble for an individual site are not considered useful. Similarly, maps of the best or worst case scenario cannot be justified.

An operational implementation of the prototype forecast system described in this report is likely to be slightly more skilful, and may even show a useful relationship between spread and skill. The decision to use the square root of wind speed in the analysis needs to be revisited, and a better way of correcting for the under forecasting of rainfall amount found. The NCEP ensemble system was upgraded recently which should result in slight improvements in accuracy. However the biggest improvement may come from applying the system to the 51 members of the ECMWF ensemble.

Acknowledgements

This work was carried out under a grant from the New Zealand Fire Service Commission's Contestable Research Fund. Thanks are due to Grant Pearce at Ensis for an updated version of his 2003 fire weather climatology, help with the fire weather index algorithms, and many helpful comments. At MetService, Ross Marsden helped organise the data, and Neil Gordon gave freely of his considerable experience in analysis techniques.

References

Alexander, M.E., 1992: Standard specifications for Fire Weather Index System calculations. Paper prepared for discussion at the 3rd meeting of the Advisory Committee on Forest and Rural Fire Research, 21 October 1992, NZ Fire Service National Headquarters, Wellington. 3 p + attachments.

Alexander, M.E., 1994: Proposed revision of fire danger class criteria for forest and rural fire areas in New Zealand. National Rural Fire Authority, Wellington. Circular 1994/2. 73 p.

Beckers, J. M., Rixen, M., 2003: EOF Calculations and Data Filling from Incomplete Oceanographic Datasets. Journal of Atmospheric and Oceanic Technology (2003) 20: 1839-1856.

Bertrand, D. and R. Verret, 2004: Toward a new Canadian medium-range Perfect-Prog temperature forecast system. 17th Conference on Probability and Statistics in Atmospheric Sciences. The 84th American Meteorological Society Annual Meeting, Seattle, WA. Viewed June 2005 at http://ams.confex.com/ams/84Annual/techprogram/paper 67887.htm.

Fogarty, L.G., H.G. Pearce, W.R. Catchpole and M.E. Alexander, 1998: Adoption vs. adaptation: lessons from applying the Canadian Forest Fire Danger Rating System in New Zealand. In Viegas, D.X. (editor). Proceedings, 3rd International Conference on Forest Fire Research and 14th Fire and Forest Meteorology Conference, Luso, Coimbra, Portugal, 16-20 November, 1998. pp 1011-1028.

Heydenrych, C., M.J. Salinger and J. Renwick, 2001: Climate and Severe Fire seasons. NIWA Client Report AK00125, National Institute of Water and Atmospheric Research, Auckland. 117 p.

Janz, B.; Nimchuk, N. 1985. The 500mb anomaly chart - a useful fire management tool. In: Proceedings of the Eighth Conference on Fire and Forest Meteorology, Detroit, Michigan, April 29- May 2, 1985. Society of American Foresters, Bethesda, Maryland. SAF Publication 85-04. pp 233-238.

Kalnay, E., et al., 1996: The NCEP/NCAR 40-year reanalysis project. Bull Amer. Meteor. Soc., 77, 437-471.

Kalnay, E., 1996: Image retrieved 4 November 2005 from http://www.cpc.ncep.noaa.gov/products/wesley/paobs/paobs.html

Kidson, J.W., 1994: Relationship of New Zealand daily and monthly weather patterns to synoptic weather patterns. Intl J. Climatol., 14, 723-737.

Kidson, J.W., 1997: The utility of surface and upper air data in synoptic climatology specification of surface climatic variables. Intl J. Climatol. 17, 399-413.

Kidson, J.W., 2000: An analysis of New Zealand synoptic types and their use in defining weather regimes. Int. J. Climatol. 20, 299-316.

Pearce, H. G., 2001: A New Zealand scrubland fire danger model: scientific rigour versus operational need. In Proceedings, Bushfire 2001 Conference, 3-6 July 2001, Christchurch, New Zealand. pp 53.

Pearce, H.G., K.L. Douglas and J.R. Moore, 2003: A fire danger climatology for New Zealand. New Zealand Fire Service Commission Research Report No. 39. 289 p.

Reifsnyder, W. E. and B. Albers, 1994: Systems for evaluating and predicting the effects of weather and climate on wild-land fires. Special Environment Report No. 11, World Meteorological Organisation, Geneva. 34 p.

Stocks, B.J., M.E. Alexander, R.S. McAlpine, B.D. Lawson and C.E. Van Wagner, 1987: Canadian Forest Fire Danger Rating System: users guide. Government of Canada, Canadian Forestry Service, Petawawa, Ontario. Loosebound.

Tait, A. and X. Zheng, 2005: Final report: Optimal mapping and interpretation of fire weather information. NIWA Client Report: WLG2005-1, National Institute of Water and Atmospheric Research, Wellington. 62 p.

Toth, Z and I Szunyogh, 2000: Results from an increase in horizontal resolution and membership of the NCEP global ensemble. Retrieved 8 November 2005 from http://wwwt.emc.ncep.noaa.gov/gmb/ens/verif000830/verifresults.html

Van Wagner, C.E. and T.L. Pickett, 1985: Equations and FORTRAN program for the Canadian Forest Fire Weather Index System. Forestry Technical Report 33, Government of Canada, Canadian Forestry Service, Ottawa, Ontario. 18 p.

Van Wagner, C.E., 1987: Development and structure of the Canadian forest fire weather index system. Forestry Technical Report 35, Government of Canada, Canadian Forestry Service, Ottawa, Ontario. 36 p.

Wilks, D. S., 1995: Statistical methods in the atmospheric sciences. Academic Press, California. 467 p.

Whitaker, J.S., and A.F. Loughe, 1998: The relationship between ensemble spread and ensemble mean skill. Mon. Wea. Rev., 126, 3292-3302.

Zhu, Y., Z. Toth, R. Wobus, D. Richardson, and K. Mylne, 2002: The economic value of ensemble based weather forecasts. Bulletin of the American Meteorological Society, 83, 73-83.

Zhu, Y., 2005: Image retrieved 6 November 2005 from http://wwwt.emc.ncep.noaa.gov/gmb/yzhu/html/opr/Z500_PAC_RMS_mens.html

Appendix: Details of Stations Used

The period of data for each station refers to the data used train the forecast system. Most stations also had data up to the end of 2004, which was used for forecast validation.

	Station Name	ID	Source	Lat	Lon	Elev (m)	Start	End	n Days
AKL	Auckland Aero	181	Met	-37.00	174.80	10	01/01/85	31/12/02	6574
APA	Taupo Aero	128	Met	-38.73	176.06	407	01/01/85	31/12/02	6574
APP	Aupouri Peninsula	24	CR10	-34.69	173.02	40	21/11/94	31/12/02	2963
ASH	Ashburton Plains	38	CR10	-43.89	171.74	103	21/09/94	31/12/02	3024
ASY	Ashley	9	CR10	-43.17	172.51	280	19/12/93	28/05/02	3083
ATH	Athol	81	CR10	-38.23	175.80	365	03/11/92	31/12/02	3711
AWV	Awatere Valley	37	cr10	-41.64	174.07	80	02/09/94	31/12/02	3043
BEL	Belmont	78	CR10	-41.18	174.88	260	24/11/97	31/12/02	1864
BML	Balmoral	10	CR10	-42.86	172.75	205	06/11/94	31/12/02	2978
BMT	Blackmount	3	CR10	-45.76	167.67	275	04/11/93	28/05/02	3128
BPO	Big Pokororo	149	CR10	-41.24	172.94	630	07/12/98	28/05/02	1269
BTL	Bottle Lake	8	CR10	-43.47	172.68	5	28/10/93	31/12/02	3352
BUR	Burnham	165	CR10	-43.61	172.75	64	16/09/96	28/05/02	2081
CAN	Cannington	5	cr10	-44.35	170.93	180	24/02/94	31/12/02	3233
CDT	Cornwallis Depot	18	CR10	-37.00	174.60	20	18/11/93	28/05/02	3114
CHA	Christchurch Aero	96	Met	-43.48	172.53	30	01/01/85	31/12/02	6574
CLV	Clevedon Coast	26	CR10	-36.92	175.00	100	05/10/94	31/12/02	3010
CLY	Clyde	162	CR10	-45.20	169.31	169	19/09/97	31/12/02	1930
CPX	Castle Point	94	Met	-40.90	176.20	120	01/10/91	31/12/02	4110
CRK	Cricklewood	74	CR10	-38.96	177.15	440	26/08/96	31/12/02	2319
CYB	Glenledi	41	CR10	-46.17	170.06	100	24/08/94	31/12/02	3052
DAR	Dargaville	75	cr10	-35.96	173.84	30	01/01/85	31/12/02	6574
DNP	Dansey Pass	39	CR10	-45.03	170.26	495	20/08/94	31/12/02	3056
DOV	Dovedale	2	CR10	-41.34	172.99	320	29/12/93	31/12/02	3290
DPS	Deep Stream	87	CR10	-45.73	169.85	700	04/02/98	31/12/02	1792
FPL	Darfield	7	CR10	-43.49	172.15	190	28/10/93	31/12/02	3352
GAL	Galatea	160	CR10	-38.33	176.79	160	06/03/96	31/12/01	2127
GBI	Great Barrier Islan	20	cr10	-36.24	175.46	8	06/07/94	31/12/02	3101
GCE	Gore	141	Met	-46.10	168.88	123	01/10/91	31/12/02	4110
GDE	Goudies	159	CR10	-38.54	176.51	238	18/10/93	31/12/02	3362
GSA	Gisborne Aero	100	Met	-38.65	177.98	5	01/01/85	31/12/02	6574
HAN	Hanmer	47	cr10	-42.53	172.85	350	16/07/96	31/12/02	2360
HAU	Haurangi	57	CR10	-41.44	175.25	200	01/10/95	31/12/02	2649
HIR	Hira	1	CR10	-41.28	173.33	180	07/12/93	31/12/02	3312
HIX	Hicks Bay	145	Met	-37.55	178.30	15	18/09/94	31/12/02	3027
HKA	Hokitika Aero	104	Met	-42.70	170.98	45	01/01/85	31/12/02	6574
HNA	Hamilton Aero	102	Met	-37.85	175.33	52	02/10/91	31/12/02	4109
HNE	Hunua East	21	CR10	-37.20	175.29	25	24/01/95	28/05/02	2682
HNW	Hunua West	22	CR10	-37.07	175.06	100	16/07/96	28/05/02	2143
HOK	Hokianga	63	CR10	-35.48	173.37	80	29/11/96	31/12/02	2225
HTX	Haast	101	Met	-43.85	169.00	3	09/04/93	31/12/02	3554
HWT	Holdsworth Station	35	CR10	-40.89	175.52	240	22/08/94	31/12/02	3054
KAI	Kaipara	82	wr62a	-36.48	174.23	120	28/07/96	28/05/02	2131
KAW	Kawerau	66	CR10	-38.06	176.73	20	26/08/96	31/12/02	2283
KHD	Keneperu Head	12	CR10	-41.16	174.12	20	25/10/93	31/12/02	3355
KIX	Kaikoura	108	Met	-42.41	173.68	105	01/01/85	31/12/02	6574
KOE	Kaikohe	107	Met	-35.41	173.81	204	30/11/94	31/12/02	2954
KWK	Kaiwaka	32	CR10	-39.27	176.87	400	11/08/94	31/12/02	3065

	Station Name	ID	Source	Lat	Lon	Elev (m)	Start	End	n Days
KX	Kaitaia Observatory	109	Met	-35.13	173.24	17	01/01/85	30/05/02	6359
LAE	Lauder	164	CR10	-45.03	169.68	370	01/10/91	31/12/02	3161
LBX	Le Bons Bay	110	Met	-43.73	173.11	237	16/11/94	31/12/02	2968
LEV	Lees Valley	147	fws1a	-43.11	172.22	480	08/11/97	31/12/02	1880
LIS	Lismore	65	CR10	-39.83	175.20	292	05/09/96	31/12/02	2309
LNX	Levin	111	Met	-40.65	175.26	45	01/10/91	31/12/02	4110
LUX	Lumsden	112	Met	-45.75	168.45	193	01/10/91	28/05/02	3893
MAH	Mahurangi	60	CR10	-36.42	174.43	300	23/10/95	31/12/02	2627
MAT	Matawaia	58	cr10	-35.51	173.91	170	24/10/95	22/03/01	1977
MGF	Mangatu Forest	46	CR10	-38.24	177.88	475	03/12/94	28/05/02	2734
MHX	Mahia	113	Met	-39.11	177.95	136	14/10/94	31/12/02	3001
MIN	Minginui	151	FWS1A	-38.62	176.68	569	01/07/98	31/12/02	1645
MLX	Molesworth	69	CR10	-42.08	173.25	881	01/10/92	28/05/02	2538
MOA	Manapouri Aero	114	Met	-45.53	167.63	209	01/10/91	31/12/02	4110
MOS	Barn Hill	70	CR10	-45.71	168.25	400	09/10/96	28/05/02	2058
MSX	East Taratahi	99	Met	-41.00	175.61	91	01/10/91	31/12/02	4110
MTB	Mount Benger	71	CR10	-45.58	169.25	1167	07/12/98	31/12/02	1486
MTE	Matea	130	FWS1A	-38.77	176.41	682	18/10/93	31/12/02	3362
MTK	Motukarara	152	FWS1A	-43.72	172.59	30	15/08/99	31/12/02	1235
MUR	Murchison	77	CR10	-41.80	172.33	160	30/03/98	31/12/02	1738
NAT	National Park	51	cr10	-39.16	175.42	825	28/04/96	31/12/02	2439
NGA	Ngapaenga	84	FWS1A	-38.35	174.91	38	04/12/97	31/12/02	1854
NGX	Nugget Point	146	Met	-46.45	169.81	129	01/09/99	31/12/02	1218
NMU	Ngaumu	14	cr10	-41.04	175.88	50	04/11/93	31/12/02	3345
NOE	Normanby	120	Met	-39.50	174.25	122	30/09/91	31/12/02	4111
NPA	New Plymouth Aero	118	Met	-39.00	174.16	30	01/01/85	31/12/02	6574
NRA	Napier Aero	116	Met	-39.45	176.85	2	30/09/91	31/12/02	4111
NSA	Nelson Aero	117	Met	-41.30	173.21	5	01/01/85	31/12/02	6574
NTA	Ngamatea	15	CR10	-39.44	176.19	980	22/10/93	31/12/02	3358
NVA	Invercargill Aero	105	Met	-46.41	168.33	7	01/01/85	31/12/02	6574
NWX	Ngawihi	119	Met	-41.58	175.23	6	16/11/94	31/12/02	2968
OKT	Okato	64	CR10	-39.25	173.88	90	15/12/96	31/12/02	2208
OMT	Omataroa	27	CR10	-38.10	176.85	205	22/11/94	08/12/02	2939
OPO	Opouteke	59	CR10	-35.69	173.81	110	13/10/95	31/12/02	2637
OSN	Opua Bay	11	cr10	-41.27	174.21	5	10/12/93	31/12/02	3309
OUA	Oamaru Aero	121	Met	-44.96	171.08	30	01/10/91	31/12/02	4110
PAX	Paeroa	122	Met	-37.35	175.68	17	01/10/91	31/12/02	4110
PKE	Pukekohe	168	CR10	-37.20	174.85	82	30/09/91	30/05/02	3896
PMA	Palmerston North Ae	123	Met	-40.31	175.60	45	10/07/96	31/12/02	2366
PPA	Paraparaumu	144	Met	-40.90	174.98	6	01/01/85	31/12/02	6574
PTU	Pouto	19	CR10	-36.25	174.05	125	01/12/93	31/12/02	3318
QNA	Queenstown Aero	125	Met	-45.01	168.73	357	01/01/85	31/12/02	6574
RAI	Rai Valley	68	CR10	-41.21	173.59	89	10/10/96	31/12/02	2274
RAU	Raumai	83	CR10	-40.20	175.22	18	24/12/95	31/12/02	2565
REF	Reefton	155	CR10	-42.11	171.85	181	25/07/99	31/12/02	1256
RFP	Rimutaka Forest Par	36	cr10	-41.35	174.91	40	03/08/94	31/12/02	3073
RHU	Rotoehu	158	CR10	-37.94	176.50	160	06/03/96	31/12/02	2492
RIP	Ruatoria	30	CR10	-37.82	178.07	724	29/11/94	28/05/02	2738
RNP	Rock and Pillar	48	CR10	-45.38	170.20	270	22/02/96	31/12/02	2505
ROA	Rotorua Aero	126	Met	-38.10	176.31	285	01/01/85	31/12/02	6574
ROT	Rotoaira	23	CR10	-38.85	175.60	630	30/07/98	31/12/02	1616
RTF	Porapora	45	CR10	-37.80	178.27	470	02/12/94	31/12/02	2952
RUX	Waiouru Aero	132	Met	-39.46	175.70	821	01/10/91	30/06/02	3926
SLP	Slopedown	44	CR10	-46.39	169.13	140	05/10/94	31/12/02	3010
STO	Stony Creek	56	cr10	-41.42	175.48	130	19/09/95	31/12/02	2661

	Station Name	ID	Source	Lat	Lon	Elev (m)	Start	End	n Days
TAH	Tahorakuri	28	cr10	-38.58	176.16	440	01/01/95	31/12/02	2922
TEP	Te Pohue	53	CR10	-39.26	176.68	370	19/09/95	31/12/02	2661
TGA	Tauranga Aero	129	Met	-37.66	176.20	4	01/01/85	31/12/02	6574
THA	Te Haroto	67	CR10	-39.15	176.61	554	27/08/96	31/12/02	2318
THE	Tara Hills	127	Met	-44.51	169.90	488	01/10/91	31/12/02	4110
TNI	Totaranui	13	CR10	-40.82	173.00	2	05/11/93	31/12/02	3344
TPE	Te Puke	169	cr10	-37.81	176.31	91	01/10/91	30/06/02	3926
TPN	Tapanui	40	CR10	-45.91	169.23	200	21/08/94	31/12/02	3055
TPU	Tapuae	89	FWS1A	-39.99	175.72	585	10/09/96	31/12/02	2304
TRQ	Traquair	4	cr10	-45.81	170.13	425	04/11/93	31/12/02	3345
TTA	Toatoa	17	CR10	-38.11	177.51	700	12/11/93	31/12/02	3337
TUA	Timaru Aero	131	Met	-44.30	171.23	27	01/01/92	31/12/02	4018
TUT	Tuatapere	43	CR10	-46.08	167.81	85	15/08/94	31/12/02	3061
WAF	Waimarino Forest	50	CR10	-39.39	175.18	625	27/04/96	31/12/02	2440
WAH	Waihau	52	CR10	-39.39	176.56	350	02/11/95	31/12/02	2617
WAO	Waione	55	CR10	-40.46	176.30	100	09/10/95	31/12/02	2641
WAV	Waverly	54	CR10	-39.78	174.60	80	30/10/95	31/12/02	2620
WAX	Chatham Island	95	Met	-43.95	-176.71	46	28/11/95	31/12/02	2591
WBA	Woodbourne Aero	140	Met	-41.51	173.85	33	01/01/92	31/12/02	4018
WDH	Woodhill	62	CR10	-36.70	174.38	220	17/07/96	31/12/02	2359
WFA	Wanaka	133	Met	-44.71	169.23	348	11/03/94	31/12/02	3218
WGF	Waitangi Forest	25	CR10	-35.28	173.98	60	13/12/94	31/12/02	2941
WGM	Whangamata	61	cr10	-37.21	175.78	220	22/03/96	31/12/02	2476
WGO	Waihi Gold	73	CR10	-37.38	175.87	115	31/05/98	31/12/02	1676
WHG	Marco	29	cr10	-39.10	174.76	160	31/07/94	31/12/02	3076
WKA	Whakatane Aero	137	Met	-37.91	176.91	6	01/01/92	30/06/02	3834
WNA	Wellington Aero	135	Met	-41.33	174.81	6	01/01/85	31/12/02	6574
WPK	Waipukurau	33	CR10	-39.99	176.53	143	17/08/94	31/12/02	3059
WRA	Whangarei Aero	138	Met	-35.76	174.36	37	01/01/92	31/12/02	4018
WRY	Wreys Bush	42	cr10	-46.02	168.11	110	16/08/94	28/05/02	2843
WSA	Westport	136	Met	-41.73	171.56	4	01/01/85	31/12/02	6574
WTA	Whitianga Aero	139	Met	-36.86	175.66	4	01/01/92	31/12/02	4018
WTF	Waitarere Forest	34	CR10	-40.55	175.20	1	28/07/94	31/12/02	3079
WUA	Wanganui Aero	134	Met	-39.96	175.01	8	01/01/85	31/12/02	6574