Nowcasting Astronomical Seeing: Towards an Operational Approach

Fionn Murtagh

Space Telescope – European Coordinating Facility,
European Southern Observatory, Karl-Schwarzschild-Str. 2,
D-85748 Garching, Germany
Email: fmurtagh@eso.org
(Affiliated to Astrophysics Division, Space Science Department,
European Space Agency)

Alex Aussem

European Southern Observatory, Karl-Schwarzschild-Str. 2,
D-85748 Garching, Germany
Email: aaussem@eso.org
and
Université René Descartes, UFR de Mathématiques et Informatique,
45, rue des Saintes-Pères, F-75006 Paris, France

Marc Sarazin

European Southern Observatory, Karl-Schwarzschild-Str. 2,
D-85748 Garching, Germany
Email: msarazin@eso.org

Running Head: Nowcasting Astronomical Seeing

Send correspondence to: Fionn Murtagh (fmurtagh@eso.org)

ABSTRACT

In Murtagh and Sarazin (1993), the ability to regress seeing on a set of near-ground meteorological variables was assessed. The present article proceeds towards the following operational framework: we obtain median seeing estimates for the night, given knowledge of ambient ground meteorological conditions. The data used are extensive meteorological and seeing measurements made at European Southern Observatory sites in Chile over the years 1989–1993. In addition we used radiosonde data, to monitor effects on seeing at higher altitudes, but did not thereby obtain improved results. Seeing nowcasts are based on a novel, fuzzy analysis approach. Our tests show that we successfully nowcast good seeing in around 70%-80% of the cases in which this is possible.

Key words: forecasting, prediction, nowcasting, astronomical seeing, fuzzy calculus, correspondence analysis, meteorological environment, telescope operations.

1. INTRODUCTION

1.1 Data and Overall Goal

An automated meteorological station (of Vaisala manufacture), microthermal sensors, and differential image motion monitors (Sarazin 1990) are used to study environmental conditions at current or future observing sites of the European Southern Observatory (ESO) in Chile. Summaries are produced as quarterly reports. These reports provide information on the instrumentation used, the data and its precision, and plots showing variability. Recently, this data has become available on-line through the STARCAT database interface (now in turn available on the World-Wide Web) which is also used for other ESO and Hubble Space Telescope archived data.

The overall goal of this work is the Astronomical Weather Station site monitor (ASM), a future interface between the observer and the terrestrial observatory environment. The ASM is a component part of the Very Large Telescope (VLT) observatory which is currently under construction. Advance knowledge of fulfillment of conditions for excellent seeing (e.g. 0.3 arcsec full width at half-maximum at Paranal over a period of a few hours) will allow the telescope to be set in the most requiring instrumental mode.

Murtagh and Sarazin (1993) examined seeing nowcasting at short time-scales, i.e. one-hour separation between cases, without considering physical models or time-related trends. To improve on the flexibility of telescope scheduling, we are now looking at nowcasting on larger time-scales, i.e. per night. We characterize the night's seeing as its median seeing value. Thus the night is good or bad if the median seeing value is low or high. Other definitions of a night's seeing are of course possible, but any interesting measure would probably be site-dependent (we found this to be so). The median is a general measure, and is distributionally robust.

Similar methods are used as in Murtagh and Sarazin (1993), but now on a night basis. The central task remains the same: to identify meteorological parameters, such that good association between seeing and these parameters can be obtained. In addition to the meteorological parameters used in Murtagh and Sarazin (1993), we also report here on our findings in regard to usage of radiosonde data from higher altitudes.

1.2 Nowcasting

Murtagh and Sarazin (1993) pursued the goal of nowcasting seeing. This related seeing measurements with meteorological measurements at the same timepoint, disregarding the time of the night. A given night gave rise to a variable number of cases.

In the terminology of statistics, this is multiple regression of seeing on a set of meteorological variables. A nonlinear solution which was preferred in Murtagh and Sarazin (1993) was a nearest neighbor method. This method is used below in this article also. It is additionally complemented by another nonparametric method for obtaining nonlinear class discrimination boundaries, the feedforward connectionist network.

The nearest neighbor method (for discriminant analysis, regression, or prediction) is simply explained as follows. Given a set of meteorological variables, we seek the associated seeing value. We search for the set of meteorological variables, among all those stored in our database, which is closest to our set. We use the seeing associated with the nearest neighbor vector as the estimate of the desired seeing. This method is powerful, in that it leads to nonlinear decision boundaries between different classes of seeing (e.g. good and bad). It and related methods have a long tradition in pattern recognition and statistics. It also concords well with common sense practices for carrying out predictions.

Multidimensional "closeness" may be measured using the Euclidean distance. However pressure, measured in millibars, has a value around 750 at ground level, while temperature at ground level may have a value, say, around 10 if measured in degrees centigrade. Clearly, the question of scale of measurement variables must be addressed. One way of doing so is to suitably recode the measurement variables and this will be looked at in the next section.

The feedforward network is the most widely-used neural network method. A "network" here represents a set of recursive, simple functions (a sigmoid of the sum of weighted inputs) which maps the input nonlinearly onto an output value. The free parameters in the network, the weights, are estimated using a training set. Further details of the feedforward network implementation used are given in Section 5 below, and in Aussem et al. (1995).

1.3 Data Coding

Murtagh and Sarazin (1993) categorized each variable into three categories: low-valued ($= 0$), a linearly varying moderate-value, and high-valued ($= 1$). The boundaries between these categories were set at the 33rd and 67th percentiles. If the percentiles are defined with respect to all data for a particular season, then the data are seasonally-referenced. This is a robust approach: extreme values are not allowed to play a disproportionate role.

When we recode seeing in this way, we can refer to low values as good seeing, and high values as bad seeing, again referenced to the season or other subset of the data on which we base our definitions of the 33rd and 67th percentiles. The Euclidean distance may be more easily used between vectors which have been recoded in this way. We can go further, however, making use of the notions of "high" and "low" which are used in this recoding approach. In order to be consistent with these notions, we have taken them as fuzzy values. When dealing with such data, we therefore use fuzzy generalizations of traditional notions such as conditional probability. The fuzzy version of this is termed conditional possibility. For definitions, see Murtagh and Sarazin (1993).

1.4 Meteorological Variables Used

The meteorological variables used were: wind direction and velocity, at 10m and 2.5m above ground; seeing at 5m; temperature, humidity and pressure at 2m; and ground temperature at -0.1m. Very high correlation among some of these variables led to us retaining the following: wind velocity in ms^{-1} at 10m above ground; standard deviation of the windspeed over 20 min. using 2-second samples; relative humidity, in percent; air temperature, in degrees at 2m above ground; pressure in mB; and seeing, measured at 5m above ground, and defined as the full width at half-maximum of a stellar image with a perfect large telescope at 0.5 μm wavelength and at zenith.

Unlike in Murtagh and Sarazin (1993), nightly median values were obtained for these variables. This was done with the ultimate aim of deriving predictions one time unit in advance, where the unit time-step was the observing night. In this work, we report on nowcasting with the unit interval being the 24-hour period between two nights.

1.5 Higher Altitude Data

Following Bély (1987), to complement the near-ground variables, upper air meteorological data was obtained. This resulted from daily radiosonde balloon ascents, with launch taking place at Antofagasta (23 26S, 70 26W; Paranal is 24 37S, 70 24W).

As an example of the data obtained in such an ascent, we can consider an ascent to 25000m, with about 30 readings on: pressure, height, temperature, relative humidity, wind direction and wind speed. Heights were of course irregularly sampled, and information on changing position downwind from the location of launch is not available. Launch time was 12:00 UT daily. Balloon ascents before the observing night, and after the night, were both used to check for relatedness. Only the first has immediate practical (i.e. predictive) usefulness. An ascent could last up to about 3 hours, reaching a height of 20–30 km. The irregular sampling (often separated by hundreds of meters in height) led to the use of this data set to model an air column (supposed to be relevant for the Paranal area, approximately 120 km south of the balloon launch site).

A preliminary analysis did not show correlation between seeing and wind shear or wind velocity. Rather than classical models using wind only (Hufnagel 1978) or wind and temperature profiles (VanZandt 1978), we applied a seeing model proposed by Coulman (1988) and used by Benkhaldoun (1994) with some success on monthly medians. The desire was to capture information related to thin atmospheric temperature-gradient layers (Dalaudier et al. 1994) which could be correlated either (1) with ground-based observed seeing; or (2) together with the near-ground meteorological variables, all (i.e. high atmosphere, and ground- or boundary-layer) would correlate well with observed seeing. Model seeing values were on average about 0.15 arcsec lower than observed seeing, thus giving credence to the possibility that the high atmosphere turbulence, and the near-ground turbulence, could form an additive combination.

Daily modeled seeing values were cross-correlated with daily meteorological and observed seeing values. In view of the need to transform variables, the fuzzy coding described above was employed. Analyses took place in seasons, which implied the seasonal-referencing of all values prior to analysis. For exploratory purposes, correspondence analysis (further discussed below) was used to generate plots indicating interrelationships.

No clear relationships emerged, however, between model seeing values and observed seeing values. No clear relationship emerged either between local meteorological variables and a "corrected" local seeing, defined as observed seeing with the high altitude model seeing factored out.

In agreement with Bély (1987), but in absence of dome seeing, it was concluded that the high altitude data sampling was grossly insufficient to capture thin horizontal atmospheric temperature gradient sheets. Data in considerably greater vertical resolution would be needed in order to attempt to verify the physical hypotheses used in the modeling, and to provide subsequent correlations with observed seeing.

2. SEASONAL NOWCASTING RESULTS

Correspondence analysis is a dimensionality-reduction method for the analysis and display of data. It is approximate, insofar as we often consider only two-dimensional projections of a multivariate space. In presenting projections in such a planar display, we may look for proximities (indicating high correlatedness, or perhaps even causal relationships), oppositions (for example, between good and bad seeing), and various other interrelationships. Input to correspondence analysis was, in all cases, based on the fuzzily (also termed piecewise linear) coded data, as described above in Section 1.3. For further aspects of the use of correspondence analysis, reference may be made to Murtagh and Sarazin (1993).

Table 1 lists the fuzzily coded variables used for each case. "High" and "low" were specified with respect to the 33rd and 67th percentile values within a season. Values of these percentiles can be seen in Table 2. Each case referred to a median of night-time recorded values, and was therefore associated with an observing night. In Table 2, it can be seen how many such observing nights were retained in each season. The maximum of about 90 nights in each season was not arrived at because of missing data. The possibilities of good seeing, conditional on (say) observed low windspeed and low windspeed standard deviation are used; together with the possibility of all the variables simultaneously happening in the particular season.

TABLE 1 ABOUT HERE

TABLE 2 ABOUT HERE

TABLE 3 ABOUT HERE

The data for all seasons between Spring 1989 and Autumn 1991 were from Cerro Paranal. Construction work then started at this site. The data for the four final seasons, Summer 1992 to Spring 1992, came from the nearby Cerro NTT.

Nowcasting interrelationships as in Murtagh and Sarazin (1993) were provided by correspondence analyses. Using the heuristic value of these plots, we looked at various relationships which are potentially of interest. These are overviewed in Table 3. Two fuzzy criteria are used for each such relationship, which involve a trade-off: how good is the relationship; and even if very good, how relevant in practice is the relationship?

Some of the relationships summarized in Table 3 are very good indeed. One sees for instance that in Autumn 1992, the meteorological variables `s1L`, `s1dL` and `pH` (respectively, low wind speed, low wind speed standard deviation, and high pressure) had a 90% association with good (i.e. low seeing); and that such a "constellation" of meteorological variables arose 31% of the time.

By seasonally-referencing our data, Table 3 presents an appealing list of useful findings. Table 3 shows that variables `s1L` and `s1dL` often co-occur with good seeing. But there are nonetheless differences between the same season in different years, which limits the usefulness of this finding.

3. SEASONALLY-REFERENCED DATA CODING

Within a season, it appears from our analysis in the previous section that there is information available which would allow good forecasts. The most relevant meteorological variables are difficult to specify, though, because we find that they differ in different seasons. Let us therefore consider all variables at the same time; and let us check out whether the versatile nearest neighbor algorithm provides adequate nowcasts (as a preliminary to seeking forecasts). In Section 1.2 above we have motivated the use of the nearest neighbor algorithm.

Following the seasonally-referenced recoding of all values (Section 1.3 above), each case is considered in turn. For it, we seek the best-matching set of meteorological variables. The corresponding seeing value is the estimate returned for seeing. We compare the estimate with the known seeing. This experimental strategy is a "leaving-one-out" one.

The best match is based on a fuzzy Jaccard measure (Miyamoto 1990). The fact that the measure is fuzzy is so that it is consistent with the coding used. The Jaccard measure is well-known in information retrieval, and in some tests which we carried out, it outperformed a fuzzy Euclidean distance. The definition, for vectors of meteorological variables $\{A_i\}$ and $\{B_i\}$ is:

$$\bigcap_i (A_i, B_i) / \bigcup_i (A_i, B_i) = \sum_i \min(A_i, B_i) / \sum_i \max(A_i, B_i)$$

When equi-valued similarities were found, we averaged the seeing values in an ad-hoc way, using the ordinary definition of mean.

Four different cross-tabulation measures can be used to quantify the outcome of all "leaving-one-out" seeing nowcasts as they relate to the good and bad seeing categories:

Possibility of good seeing (A), given that good seeing has been estimated (B): $\sum \min(A, B) / \sum B$ (where the summation is over all "leaving-one-out" trials).

The other situations relate to the possibility of bad seeing, given that good seeing is estimated; the possibility of good, given bad; and the possibility of bad, given bad.

We found a 63% possibility of good seeing, given that good seeing was estimated; a 60.6% possibility of bad seeing, given that bad seeing was predicted; and a 43% possibility of bad seeing, given that good seeing was predicted.

A criticism of this analysis relates to the fact that all meteorological variables, and seeing, are seasonally referenced. Therefore what is potentially future information is being used to define the data values used in this analysis. We would prefer to use only past values to carry out the fuzzy recoding. A way to do this is proposed in the next section.

4. FUZZY NOWCASTS OF SEEING

A season, for us, is a set of about 90 nights which are batched together in time. In fact, around 50 nights remain when missing values are removed. The question arises as to whether we really need to consider the real seasons, or whether instead the locality in time is the most important aspect of this. We therefore consider only the observations made on 50 successive nights. It is not important that these nights be consecutive: for us, all that is important is that they capture or express the time-local meteorological conditions.

The recoding of every variable is carried out as before. But the "seasonal"-referencing which requires a relevant set of 33rd and 67th percentiles will now be provided by the immediately preceding "season" of 49 cases.

The "leaving-one-out" algorithm was as follows. The meteorological variables (but not seeing) for the case corresponding to a given time point were obtained. The most similar set of meteorological variables, in the entire database, was sought. The seeing value associated with the latter provided the estimate. This estimate was compared with the actual seeing at the initial time-point. Note that we are using the meteorological variables (pressure, temperature, etc.) here, without trying to forecast them. In a practical scenario, high quality estimates of such variables (as "high", "low" or "moderate" category values) are needed, with respect to the previous set of 49 values. ESO is currently consulting with meteorological institutes for the development of a LAM model (Limited Area Model) spanning the Chilean observatories.

As before, a fuzzy Jaccard criterion was used to assess multivariate similarity. Equi-valued similarities led to an ad-hoc averaging of seeing values. The results obtained are cross-tabulated in Table 4. One may note that the combined good–good, moderate–moderate and bad–bad cases amount to 40.2%; and the number of unacceptable outcomes (good–bad, and bad–good) amount to 9.6% of the cases.

TABLE 4 ABOUT HERE

First we look at the "hit rate". The possibility of good seeing, given that good seeing was predicted, is 68/(68+42+31) = 48.2%. The prediction obtained with an uninformed policy based on the systematic prediction of good seeing (i.e. related to the 33rd percentile) is 33.3% of cases. We remark that a "carbon copy" of the previous night's value, which would be an alternative baseline for our results, would have an even smaller success rate. The uninformed policy, again by definition, would lead to the possibility of good or moderate seeing of 66.7%. With our prediction approach, the possibility of good or moderate seeing, given that good seeing was predicted is (68+42)/(68+42+31) = 78.0%.

Finally, looking at "failure rate", with our prediction approach, the possibility of bad seeing given that good seeing was predicted is 31/(68+42+31) = 22.0%, which is more favorable than the 33.3% with the formerly defined uninformed policy.

5. FUZZY/NEURAL NOWCASTING

A feedforward connectionist network was used on the same data, to assess the potential for further improvement over the nearest neighbor approach. This feedforward network was a static implementation of a very general dynamical recurrent network. A fuller description of the latter, and its application to various seeing- and temperature-related datasets, can be found in Aussem et al. (1995).

The static network used for the seeing experiment is a four-layer feedforward network with $5 \times 14 \times 14 \times 1$ nodes without any delayed connection. A recurrent network did not help improve the results. The network had, in all, 618 links including bias links. Overfitting was prevented by a weight elimination method (Aussem 1995; Weigend et al. 1990).

Since the neural network output $o(t)$, at any time t, fell inside the open interval (0,1), two thresholds (s_1, s_2) were required to interpret the output value. The prediction was assumed to be low whenever $o(t) < s_1$, a moderate value when $s_1 \leq o(t) \leq s_2$ and a high value when $o(t) > s_2$. These thresholds were optimally estimated from the performance achieved on the training set and then applied to the prediction set to appraise the predictive performance of the model. The thresholds can be adjusted on-line in a practical implementation as time progresses, among the chronologically most recent tuples. Consequently, the results presented are lower bounds for expected performance. An upper bound of the expected prediction performance is also provided by optimally adjusting the thresholds on the prediction set in a suitable way, according to a given cost function (see Table 5).

TABLE 5 ABOUT HERE

The idea of a cost function, or policy, is interesting. It allows us to automatically determine the output thresholds, on the basis of an a priori judgement of what is important and what should be prioritized in regard to predicted versus actual scenarios.

Denoting good (low-valued) seeing by `seeL`, and similarly for bad (high-valued) and moderate (respectively, `seeH` and `seeM`), then possibility, a scalar bounded by 0 and 1, is a notion like probability but subject to different composition laws. The possibility of good seeing will be denoted by `poss(seeL)`. The possibility of good seeing, given that bad seeing was predicted is `poss(seeL|seeH)`. A measure of relevance is provided by `poss(seeL)`; and a measure of prediction quality, subject to a relevant ambiance, is given by `poss(seeL|seeL)`.

Various performance measures, e.g. cost policies, may be proposed so as to meet the observers' requirements. With a view to presenting the performance of the connectionist model with the above objective in mind, three pairs of thresholds were used, related to the following performance measures of interest:

Policy 1: maximum hit rate, when actual and estimated seeing are in the same category: `poss(seeL|seeL) + poss(seeM|seeM) + poss(seeH|seeH)`.

Policy 2: maximum conditional possibility to have good seeing when good seeing was predicted: `poss(seeL|seeL)`.

Policy 3: minimum of the cost function given by Table 5.

Table 5 expresses user-oriented utility, or the wish to diminish the number of "false alarms" (e.g. bad seeing, when good is expected). Positive and negative values express failure and success, respectively. Observing time is lost when an observation requiring good seeing is scheduled during an actual bad or moderate seeing night. Conversely, observing time is inefficiently used when an observation with no specific seeing requirements is scheduled during a good seeing night.

Beyond the quest for reliable forecasts, the use of cost tables can allow optimization of scheduling. The observation time assignment procedure would be based on an optimal thresholding, based on the judgements of what is important, as expressed by the cost policy table.

The results shown in Tables 6 and 7 are typical for two-season experiments. We can also attempt to minimize incorrect results, setting weights and threshold values in the connectionist network in order to seek 0% incorrect results, but this is very much at the cost of the relevance of such a result. Further discussion of such results can be obtained in Aussem et al. (1995).

The results confirm that Paranal is less predictable than Cerro NTT, a lower summit more subject to ground effects.

6. CONCLUSION

We tested several methods with a view towards nowcasting astronomical seeing on a night average (more strictly: median) basis.

The results obtained from correspondence analysis, the exploratory methodology used in Murtagh and Sarazin (1993), were improved upon by use of the connectionist network approach. The latter was, however, based on relatively small samples (two seasons at a time, in Tables 6 and 7).

Even when applied to an outstanding astronomical site, i.e. with little turbulence at ground level, like Paranal, fuzzy nowcasts of seeing from ground meteorological data allowed a success rate of 70%-80% (connectionist network results). A baseline, “uninformed” policy based on our global definition of good seeing is 33.3% (although alternative baseline models could use the time series information).

This is possible because our technique uses the tenuous statistical link between ground conditions and larger-scale meteorological patterns. We were not able to improve the prediction by using available high atmosphere data from classical meteorological soundings.

The requirements for operational use imply: choosing a policy; fuzzifying the meteorological variables; and having a reliable meteorological forecast for the observatory.

ACKNOWLEDGEMENTS

We are grateful for helpful comments by the referee on an earlier version of this paper.

REFERENCES

Aussem, A. 1995, Dynamical and recurrent neural networks towards time series prediction, PhD thesis, in preparation

Aussem, A., Murtagh, F. and Sarazin, M. 1994, Vistas in Astronomy 38, 357

Bély, P.-Y. 1987, PASP 99, 560

Benkhaldoun, Z. 1994, “Un Observatoire Astronomique au Maroc: Choix et Qualification Diurne du Site de l’Oukaimeden. Validation d’un Modèle de Prévision du Seeing Nocturne”, PhD Thesis (Marrakesh, Université Cadi Ayyad)

Coulman, C.E., Vernin, J., Coquengniot, Y. and Caccia, J.L. 1988, App. Opt. 27, 155.

Dalaudier, F., Sidi, C., Crochet, M. and Vernin, J. 1994, J. Atmos. Sci. 51, 237

Hufnagel, R.E. 1978, “Propagation Through Atmospheric Turbulence”, The Infrared Handbook (Washington, DC: US Government Printing Office), p. 6-1

Miyamoto, S. 1990, Fuzzy Sets in Information Retrieval and Cluster Analysis (Dordrecht, Kluwer)

Murtagh, F. and Sarazin, M. 1993, PASP 105, 932

Sarazin, M. and Roddier, F. 1990, A&A 227, 294

VanZandt, T.E., Green, J.L., Gage, K.S. and Clark, W.L. 1978, Radio Sci. 5, 819

Weigend, A.S., Rumelhart, D.E. and Huberman, B.A. 1990, Journal of Neural Systems, 1, 193

Table 1
Fuzzy variables used

s1H	windspeed – High
s1L	windspeed – Low
s1dH	windspeed standard deviation – High
s1dL	windspeed standard deviation – Low
rhH	relative humidity – High
rhL	relative humidity – Low
t1H	temperature – High
t1L	temperature – Low
pH	pressure – High
pL	pressure – Low
seeH	seeing – High
seeL	seeing – Low

Table 2
33rd and 67th percentile values, by season,
for meteorological variables and seeing
(Paranal to Aug91; thereafter Cerro NTT)

	s1		s1d		rh		t1		p		see	
Season	33rd	67th	33rd	67th	33rd	67th	33rd	67th	33rd	67th	33rd	67th
Spr89	4.5	7.9	0.1	0.2	4	8	10.6	11.9	741.6	742.7	0.57	0.68
Sum90	4.5	8.7	0.1	0.2	9	17	12.3	13.3	741.8	742.8	0.61	0.77
Aut90	6.0	8.8	0.1	0.2	5	9	10.3	13.3	741.6	743.8	0.59	0.78
Win90	3.6	9.3	0.1	0.2	4	8	9.4	11.9	741.7	743.5	0.51	0.66
Spr90	4.2	7.3	0.3	0.5	6	13	11.1	12.9	742.2	743.4	0.46	0.66
Sum91	3.9	7.2	0.1	0.2	15	23	12.1	13.2	742.6	743.6	0.54	0.68
Aut91	5.2	8.6	0.1	0.2	9	15	10.9	12.7	742.6	743.7	0.47	0.64
Sum92	3.1	6.1	0.1	0.3	12	20	11.6	14.4	751.4	752.8	0.53	0.72
Aut92	4.4	7.1	0.2	0.5	12	19	9.8	12.7	751.4	752.6	0.69	0.93
Win92	4.4	8.7	0.2	0.3	5	10	8.9	10.9	750.8	752.1	0.64	0.95
Spr92	3.7	6.5	0.1	0.2	6	13	11.4	12.9	750.6	752.3	0.59	0.77

Table 3
Nowcasting results
(Paranal to Aug91, thereafter Cerro NTT)

Season	No. of cases	Good seeing given ...	Conditional possibility of good seeing	Possibility of variables happening
Spr89	81	s1L	0.68	0.35
		s1dL	0.61	0.27
		s1L and s1dL	0.62	0.20
Sum89	77	pH	0.64	0.32
		s1dL	0.69	0.38
		s1L	0.68	0.35
		pH, s1dL and s1L	0.75	0.24
Aut90	48	rhL	0.76	0.35
Win90	69	rhL	0.63	0.32
		t1H	0.63	0.31
		pH	0.63	0.33
Spr90	47	s1L	0.58	0.29
		s1dL	0.67	0.36
		s1L and s1dL	0.67	0.27
Sum90	38	rhL	0.52	0.27
		t1H	0.62	0.31
		rhL and t1H	0.60	0.21
Aut91	61	s1L	0.61	0.29
		s1dL	0.70	0.41
		s1L and s1dL	0.60	0.25
Sum92	55	s1L	0.76	0.40
		s1dL	0.77	0.38
		pH	0.68	0.35
		s1L, s1dL and pH	0.88	0.24
Aut92	66	s1L	0.78	0.40
		s1dL	0.80	0.43
		pH	0.67	0.37
		s1L, s1dL and pH	0.90	0.31
Win92	44	s1L	0.89	0.46
		s1dL	0.87	0.42
		rhL	0.72	0.38
		s1L, s1dL and rhL	0.96	0.27
Spr92	69	s1L	0.74	0.38
		s1dL	0.68	0.33
		s1L and s1dL	0.72	0.29

Table 4
Vertically displayed: output of our forecast. Horizontally displayed: actual situation.

	bad	moderate	good
bad	58	49	27
moderate	104	117	109
good	31	42	68

Table 5
Costs associated with each possibility. Positive and negative values express failure and success respectively.

estimated \ actual seeing	bad	moderate	good
bad	$-1/4$	$+1/4$	$+1/2$
moderate	$+1/4$	0	$+1/2$
good	$+1$	$+1$	-1

Table 6
Conditional possibilities, expressed as percentages, for Winter 1992 - Spring 1992 (Cerro NTT)

Policy	`poss(seeL|seeL)`	`poss(seeH|seeL)`	`poss(seeL)`
1	80.0	6.7	28.6
2	87.0	4.3	21.9
3	85.7	3.6	26.7

Table 7
Conditional possibilities, expressed as percentages, for Spring 1989 - Summer 1989 (Paranal)

Policy	`poss(seeL|seeL)`	`poss(seeH|seeL)`	`poss(seeL)`
1	66.7	20.0	28.6
2	72.7	18.2	10.5
3	70.4	14.8	25.7