

Linear regression downscaling of seasonal predictions for the UK climate districts

Hadley Centre technical note 61

Anca Brookshaw and Mike Davey

April 2005



Linear regression downscaling of seasonal predictions for the UK climate districts

Anca Brookshaw and Mike Davey

Met Office, FitzRoy Road, Exeter, EX1 3PB

Abstract

This document describes a method for using non-local information to improve the skill of forecasts for a given geographic region. The example is the prediction of monthly-mean temperatures for the ten UK climate districts, using forecast data from the GloSea long-range forecast system. It is demonstrated that by using linear regression and nearby (rather than nearest) grid-points the forecast skill can be improved significantly. This work also provides the foundation for further developments and applications.

1 Introduction

The raw long-range forecast data produced from a general circulation model are provided on a spatial grid at the resolution of the model. Applications may require the forecast in a different form: e.g. information for a specific geographic location or region that does not correspond to the model grid. In this case, some algorithm is required to convert the model data to the required form.

The simplest procedure, often used in practice, is to use the nearest model grid-point (or some average if the region contains more than one grid-point) to represent the required location. However, this may not be the best choice, for various reasons. For example, the nearest point may be over the ocean, and have different characteristics to a desired land location. The model forecast (particularly at long range) may contain biases such that a non-local point is more representative of the target location than the nearest point.

There are many possible procedures for improving the 'nearest point' forecast, that fall into the general category of calibration and downscaling techniques (see

e.g. Barnston and Smith 1996, Goddard et al. 2001, Landman and Goddard 2002, Stephenson et al. 2005), but none have yet been applied specifically to UK regions. For demonstration purposes we use linear regression, as described in section 2. The particular application, as presented in section 3, is the deterministic prediction of monthly mean temperatures for the UK climate districts, for ranges up to six months ahead. Results based on 30 years of data demonstrate that standard cross-validated skill measures are improved substantially by incorporating non-local data.

The linear regression scheme has been written such that it can easily be adapted for application to other geographic regions, as required. It can be adapted for use with probabilistic as well as deterministic forecast formats. It also provides a basis for the construction of more advanced downscaling/calibration schemes.

2 Method

The standard linear regression method and the selection of predictors are briefly described here.

Given a predictand y , the problem is to determine it from known values of k predictors x_1, x_2, \dots, x_k , using information about the relationship between predictand and predictors over a training period.

Let the linear prediction equation be

$$y = b_0 + b_1 x_1 + \dots + b_k x_k, \quad (1)$$

with b_0, \dots, b_k unknown constants.

These $k + 1$ constants are determined from predictand values \hat{y}_n , known at time points $n = 1, \dots, N$, and corresponding known predictor values $\underline{x}_n = (x_{1n}, x_{2n}, \dots, x_{kn})$. (These values are referred to as training data.) The values of the predictand estimated by equation (1) at time point n are

$$y_n = b_0 + \sum_{i=1}^k b_i x_{in}, \quad n = 1, \dots, N. \quad (2)$$

The differences between these and the corresponding observed values \hat{y}_n measure the error of the forecast. A cost function is defined as the sum of the squares of these differences:

$$C = \frac{1}{N} \sum_{n=1}^N (\hat{y}_n - y_n)^2 = \frac{1}{N} \sum_{n=1}^N (\hat{y}_n - b_0 - \sum_{i=1}^k b_i x_{in})^2. \quad (3)$$

The coefficients b_0, \dots, b_k are then determined from the condition that they are chosen to minimise the mean squared error C .

With the regression coefficients thus determined, equation (1) can be used to estimate the predictand at any time, given values of the predictors.

To apply this method in practice the issue of selecting the predictors for the forecast equation must be addressed first. The number of predictors available, and potentially relevant for the forecast of a given meteorological variable, can be very large. Choosing those useful for a prediction equation is difficult. Too many, and the data are ‘overfit’ — that is, the equation provides a very good fit for the training data, but provides a poor estimate for the independent verification data. Too few, and it may happen that some of the predictive ‘opportunities’ are missed.

There are several ways of selecting useful predictors from a large number available. One obvious way is the exhaustive search: try all possible combinations, of all possible numbers of predictors, and choose that which optimises the skill measure used — e.g. minimum mean squared error (MSE). While this method ensures the selection of the best set, it is impractical, as it is very expensive in terms of computing resources.

For this study we first selected a set of M likely predictors (in this case data at grid-points located near the predictand: see section 3). From this group we then selected subsets by applying a widely used method, known as forward selection or stepwise regression. As a first step, each of the predictors from the pool of M available is used to write a trial 1-predictor equation. The cost function C (equation (3)) is calculated for each of these M equations, and that which achieves the lowest value for the cost function (i.e. which fits the data best) is chosen as x_1 . At the next step, the $M - 1$ remaining predictors are, each in turn, used in trial 2-predictor regression equations, which this time also contain the variable chosen at the previous step as x_1 . The best, in the sense described above, is chosen as the second predictor x_2 . The process is then repeated, at each step adding one predictor from those not yet used, until some stopping criterion is satisfied. In general, when the regression equations are recomputed, at each such step, the coefficients will change. This is because the predictors are usually not independent, so that the information about the predictand is ‘spread around’ among the predictors differently as more predictors are added to the equation.

For a given set of training data, the MSE cost function decreases as more pre-

dictors are added. While a stopping criterion could be based on the rate of decrease of MSE, such a calculation is not very robust. A better alternative is to withhold some of the data and use a stopping criterion based on predictions of the withheld data, using cross-validation.

Briefly, cross-validation is a resampling technique in which the available data are repeatedly divided into training and verification subsets. As is common practice, we use the ‘leave one out’ approach, in which a training data set with $n - 1$ time points is selected, with one data point reserved for verification. (‘Leave one out’ is appropriate for timeseries in which successive points are only weakly related, if at all, which is the case for our seasonal forecast application. This method does not eliminate the influence of slow — e.g. decadal — variability on skill measures however. The present choice is a commonly-used compromise between accuracy of validation and length of available timeseries.) There are n such partitions of the data, and a regression equation is calculated for each of these partitions, which provides a prediction for the predictand at the independent omitted time point. This yields n predictions, and the associated MSE can be computed.

This MSE calculated by cross-validation does not necessarily decrease as predictors are added in the stepwise regression procedure, and it is this MSE that we use as a basis for a stopping criterion. However, a real-time forecast equation would be calculated using all the available data, after a suitable set of predictors has been identified using cross-validation.

The choice of a stopping rule remains subjective, to some degree. The cross-validated MSE may not have a unique minimum as predictors are added. Moreover, even with a judicious choice of a stopping rule, this stepwise method of selection does not necessarily produce the best sequence of predictors, but the gain in speed of computation, relative to an exhaustive search, makes it much more usable than the latter. For these reasons, stepwise regression will be used here only as a guide for selecting suitable sets of predictors.

3 Application

The method described above is used to predict surface temperature for the ten UK climate districts, using output from a seasonal prediction model. The predictand is the monthly-averaged, district-averaged 2m mean temperature, for a given UK




Figure 1: UK climate districts and GloSea atmospheric grid

climate district. The procedure is repeated for all ten districts shown in figure 1 (district 0 — NW Scotland, district 1 — NE Scotland, district 2 — NE England, district 3 — E England, district 4 — Midlands, district 5 — SE England, district 6 — SW Scotland, district 7 — NW England, district 8 — SW England, district 9 — N Ireland).

The predictor data are forecast data provided at the forecast model grid-points. The districts are irregular in shape, and we wish to determine the combinations of model grid-points that best provide forecasts for the districts, taking into account that the most effective predictors may not be the nearest points. Thus we initially selected as predictors the temperatures in an area surrounding the UK (between 15° west and 2° east, in longitude, and 50° north and 62° north, in latitude). Each grid-point represents conditions in a surrounding box, as indicated in figure 1. With reference to the layout in figure 1, the boxes are numbered as follows

1	2	3	4	5
6	7	8	9 *	10
11	12	13	14 *	15
16	17	18 *	19 *	20 *
21	22	23	24	25 *

with the asterisks denoting model land boxes in the chosen domain. The forecast data are monthly-mean 2m temperatures that are the ensemble mean values from a 9-member ensemble generated with the Met Office global seasonal (GloSea) coupled model. The data used are from the DEMETER project, for which retrospective 6-month range forecasts were calculated four times a year (at the beginning of each February, May, August and November) in the period 1959-2001. (Details of the DEMETER project and the GloSea model can be found in Palmer et al. 2004 and Graham et al. 2005). Although long-range forecasts are usually provided in probability form, based on the forecast ensemble, in this study we concentrate on the ensemble-mean values as the initial objective is to examine the influence of the data from a range of nearby grid boxes on the forecast skill. The observational data are likewise monthly-mean 2m temperature values for each district, that are based on station data within each district, from the National Climate Information Centre (for details see Colman et al, 2002). Timeseries for a sample of observed data are illustrated in figure 2.

Observed and forecast data from the period 1971-2000 have been used. As we are interested in predicting departures from ‘normal’ climate, anomalies were calculated by subtracting the observed and forecast monthly climatologies respectively. (Note: as the forecast climatologies depend on the time elapsed from the forecast start — the ‘lead time’ —, there are four sets of forecast climatologies associated with the four start times.)

As measure of forecast error, the standardised mean squared error (SMSE) — the mean squared error divided by the variance of the corresponding observations — is chosen, as it allows easier comparison of results from different districts and easy comparison with other forecast strategies. (For a strategy in which the forecast is climatology, i.e. ‘zero anomaly’, SMSE is 1; for a strategy in which the forecast is randomly chosen from past observations the SMSE is 2.)

For each forecast start time, each lead time, and each district the best choice of

predictors is established using stepwise linear regression, at each step evaluating the SMSE in cross-validation mode.

3.1 Stepwise regression results

To illustrate how SMSE varies as the number of predictor grid boxes is increased in the stepwise procedure, SMSE values have been plotted as functions of the number of predictors in figures 3, 4, 5 and 6.

Each figure shows results for one of the start times. The six panels in each figure correspond to the six months after that start time, and the different coloured lines correspond to the ten districts. For each of the panels, the naming convention is chosen to indicate the month-of-year of the forecast and the number of months after the start time: i.e. may4 indicates forecasts for May, started 4 months earlier, at the beginning of February. (In this example the lead time is $4 - 1 = 3$ months.) The maximum number of predictors allowed in each regression equation is 10. The dashed line in each plot is the line $SMSE = 1$, the error of a climatology forecast.

The results show that for most start and lead times the forecasts obtained with stepwise regression from the GloSea model grid-point predictions have substantially lower SMSE than climatology forecasts, for all UK districts. Most exceptions occur in late months of the forecasts (nov4, dec5, may4, jun5) and only for some districts. There is one case (dec2) when at 1-month lead the forecast for some districts is no better than climatology, and two cases when this happens at 2-months lead (apr3 and jan3). The apparent skill for UK districts for monthly timescales at lead times beyond one month is in itself a new and significant result.

There are some months and some districts for which SMSE is lowest with only 1 or 2 grid-box predictors (e.g. feb1). In other cases, the error keeps improving gradually with the addition of further predictors (e.g. sep2). In others still, the number of predictors used seems to make very little difference to the skill of the forecast (e.g. may1). No conclusion can be easily drawn about the optimal number of predictors in the regression equation, especially since, as postulated before, the SMSE does not vary monotonically with the number of predictors.

Examining the identity of the predictors for a given forecast reveals little geographical relation between the locations of the districts for which the forecast is sought and their most useful predictors. Generally the nearest grid box(es) are not the best predictors, which vindicates the strategy of searching further afield to improve the

forecast skill (see below for further evidence).

3.2 A ‘common predictor’ model

Some predictors are selected by the stepwise algorithm for more than one district. This is not surprising as the UK districts tend to have similar monthly-average temperature anomalies at any one time (see for example figure 2). This suggests the possibility of using common predictor grid boxes for all districts, at the respective start and lead time.




Figure 2: UK district observations

To test this method, for each lead time and start time three sets of ‘common’ predictors were chosen, based subjectively on the information provided by the stepwise regression algorithm results, and then used to derive new forecast equations. Of these three sets, the combination that yielded the lowest total SMSE for the ten districts was selected. Figures 7– 10 show how SMSE for these ‘common’ forecasts compare with the ‘stepwise’ selected ones. In these figures, the colour-coded dashed lines

(horizontal as they are independent of the number of ‘stepwise’ selected predictors) indicate ‘common’ SMSE for the various districts.

The skill (as measured here) of the ‘common’ forecast is in all cases comparable with that of the stepwise regression forecast, and in a few cases it is actually better (lower SMSE). (Recall that stepwise regression is just one methodical way of testing plausible predictor combinations, that does not necessarily include the ‘common’ combination.) The ‘common’ combination, involving cross-validation with all districts, is less likely to have provided low SMSE by chance and is thus also a more robust prediction model overall than that provided by the individual stepwise models.

3.3 Comparison with ‘nearest grid-box’ predictors

A common method for using gridded forecast information to build a forecast for a chosen location is to use some form of linear interpolation involving boxes that overlap that location. The grid boxes that overlap each UK district are listed in the table below. (Some districts are contained within one grid box — see figure 1).

district	0	1	2	3	4	5	6	7	8	9
grid boxes	9, 8	9, 14	14, 15, 20	20	19, 20	20, 25	14	14, 19	19, 24	13

Cross-validated linear regression prediction equations were derived in the usual way using these combinations of predictors. The SMSE associated with these ‘nearest box’ forecasts are plotted in figure 11, as functions of the lead time of the forecast — each colour represents a district.

For the majority of districts and most start times, it is only in the first month of the forecast that the SMSE of the ‘nearest box’ predictions is less than 1, the value of the error for a climatology forecast. This is in contrast with the forecasts for which the choice of predictors is guided by the stepwise selection algorithm.

A direct comparison of the SMSE aggregated over all the districts, for forecasts obtained with the two methods, shows a clear disadvantage for the ‘nearest box’ choice, as shown in figure 12. While for almost all start times and all lead times the sum of the errors of the ‘common’ forecasts’ is lower than that of a climatology forecast, the reverse is true for the ‘nearest box’ forecasts. For all cases, the former

are better (in the SMSE metric) forecast equations than the latter.

Based just on the usual nearest-box approach, one would deduce that the GloSea model has no advantage over climatology for predicting monthly-average UK district temperatures. The results here demonstrate that substantial gains in forecast skill can be obtained by considering a wider range of nearby grid boxes.

4 Conclusions and implications

The purpose of this work was to test the hypothesis that, for long-range forecasts of conditions in a geographic location, model forecast information from a wide surrounding area may be more useful than data at nearest points.

Standard cross-validated linear regression forecast methods have been used to assess this idea quantitatively. Specifically, predictions of monthly averages of 2m temperature, for the ten climatic regions of the UK, were constructed from grid-point values of monthly averages of 2m temperature calculated with the Met Office GloSea dynamical seasonal prediction model. The data sets for both predictors and predictands cover the 30-year period 1971-2000, with the 6-month range dynamical hindcasts taken from the DEMETER dataset. Ensemble-mean hindcast values were used for the predictor data.

The ‘nearest point’ linear regression method used as predictors data from the grid boxes that overlapped each UK district. The ‘stepwise’ method used a stepwise algorithm to select predictors for each district from the wider area indicated in figure 1. From the ‘stepwise’ results, a set of ‘common’ predictors was subjectively selected for application to all ten districts. The forecast skill in each case is measured by standardised mean squared error, with cross-validation using the usual ‘leave one year out in turn’ strategy.

For all forecast times and lead times, the stepwise selection of predictors from a wider area surrounding the UK outperforms the use of the nearest grid-points as predictors. For most start and lead times the forecasts using non-local information are more skilful than equivalent forecasts based simply on climatological observed values. By contrast, the forecasts obtained from only nearest grid-points are less skilful than climatology at all lead times longer than 1 month. Use of the ‘common’ predictor set likewise proved to be more skilful than the nearest set for all UK districts, and more skilful than climatology.

This result may seem counterintuitive: surely the nearest points should be the most effective? One explanation is that the dynamical forecast model contains drifts and biases, in that as a forecast proceeds from observed initial conditions the model tends to move towards its own intrinsic climatology. Such shifts include changes in spatial characteristics, such that a predicted temperature anomaly (even if correct in size) may be displaced in geographic location from the ‘true’ location. Another factor is the distribution of land and ocean grid boxes in the model. The relatively coarse horizontal resolution of the forecast model means that land and ocean temperature changes will be distributed differently in the model relative to ‘true’ conditions. Further, land conditions are generally more prone to error in the long-range forecast model. In fact, the results here show that ocean points are often more effective predictors than land points among the selected predictors (e.g. they are the first to be selected in the stepwise procedure).

The implication of these results is profound. The current ‘default’ method of producing forecasts is to use the model grid or, for specific regions, to use nearest grid-points. While this strategy is effective in many parts of the globe, it does not make best use of the information in the gridded forecast.

The demonstration work presented here has used a deterministic approach with basic linear regression. The method can be adapted readily to produce probabilities for forecast categories, which is the preferred and intrinsically more valuable format for long-range forecasts. Further, the linear regression approach can be advanced by making use of empirical orthogonal functions and techniques such as canonical correlation analysis, by incorporating more extensive data with Bayesian methods, and by using multi-model data from the European multi-model system that incorporates the Met Office dynamical model.

Another important aspect of these results is the demonstration of potentially useful long-range forecast skill for the UK districts on monthly timescales. Such skill is not apparent in previous assessments based on individual grid-points. The implication is that it will be worth investing further effort in developing monthly-scale forecast products. Although such products have less skill than 3-month averages at seasonal ranges, in meteorological terms, they may have more value for applications that require monthly information.

References

1. Barnston A, Smith TM (1996) Specification and prediction of global surface temperature and precipitation from global SST using CCA. *J Clim*, **9**, 2660-2697
2. Colman A, Graham RJ, Davey M, Becker B (2002) A new UK district climatology for the Monthly Outlook. unpublished paper
3. Goddard L, Mason SJ, Zebiak SE, Ropelewski CF, Basher R, Cane MA (2001) Current approaches to seasonal-to-interannual climate predictions. *Int J Climatol*, **21**, 1111-1152
4. Graham RJ, Gordon M, McLean PJ, Ineson S, Huddleston MR, Davey MK, Brookshaw A, Barnes RTH (2005) A performance comparison of coupled and uncoupled versions of the Met Office seasonal prediction General Circulation Model. *Tellus*, to appear.
5. Landman A, Goddard L (2002) Statistical recalibration of GCM forecasts over Southern Africa using model output statistics, *J Clim* **15**, 2038-2055
6. Palmer TN, Alessandri A, Andersen U, Cantelaube P, Davey M, Delecluse P, Deque M, Diez E, Doblas-Reyes FJ, Feddersen H, Graham R, Gualdi S, Gueremy JF, Hagedorn R, Hoshen M, Keenlyside N, Latif M, Lazar A, Maisonnave E, Marletto V, Morse AP, Orfila B, Rogel P, Terres JM, Thomson MC (2004) Development of a European multimodel ensemble system for seasonal-to-interannual prediction (DEMETER). *Bull.Am.Meteorol.Soc*, **85** (6), 853-872
7. Stephenson DB, Coelho CAS, Doblas-Reyes FJ, Balmaseda M (2005) Forecast assimilation: a unified framework for the combination of multi-model weather and climate predictions. *Tellus*, to appear.
8. Wilks DS, Statistical methods in the atmospheric sciences: an introduction (1995), Academic Press.




Figure 3: Variation of forecast error with the number of predictors used in the regression equation (November start, all lead times, all districts)




Figure 4: Variation of forecast error with the number of predictors used in the regression equation (February start, all lead times, all districts)




Figure 5: Variation of forecast error with the number of predictors used in the regression equation (May start, all lead times, all districts)

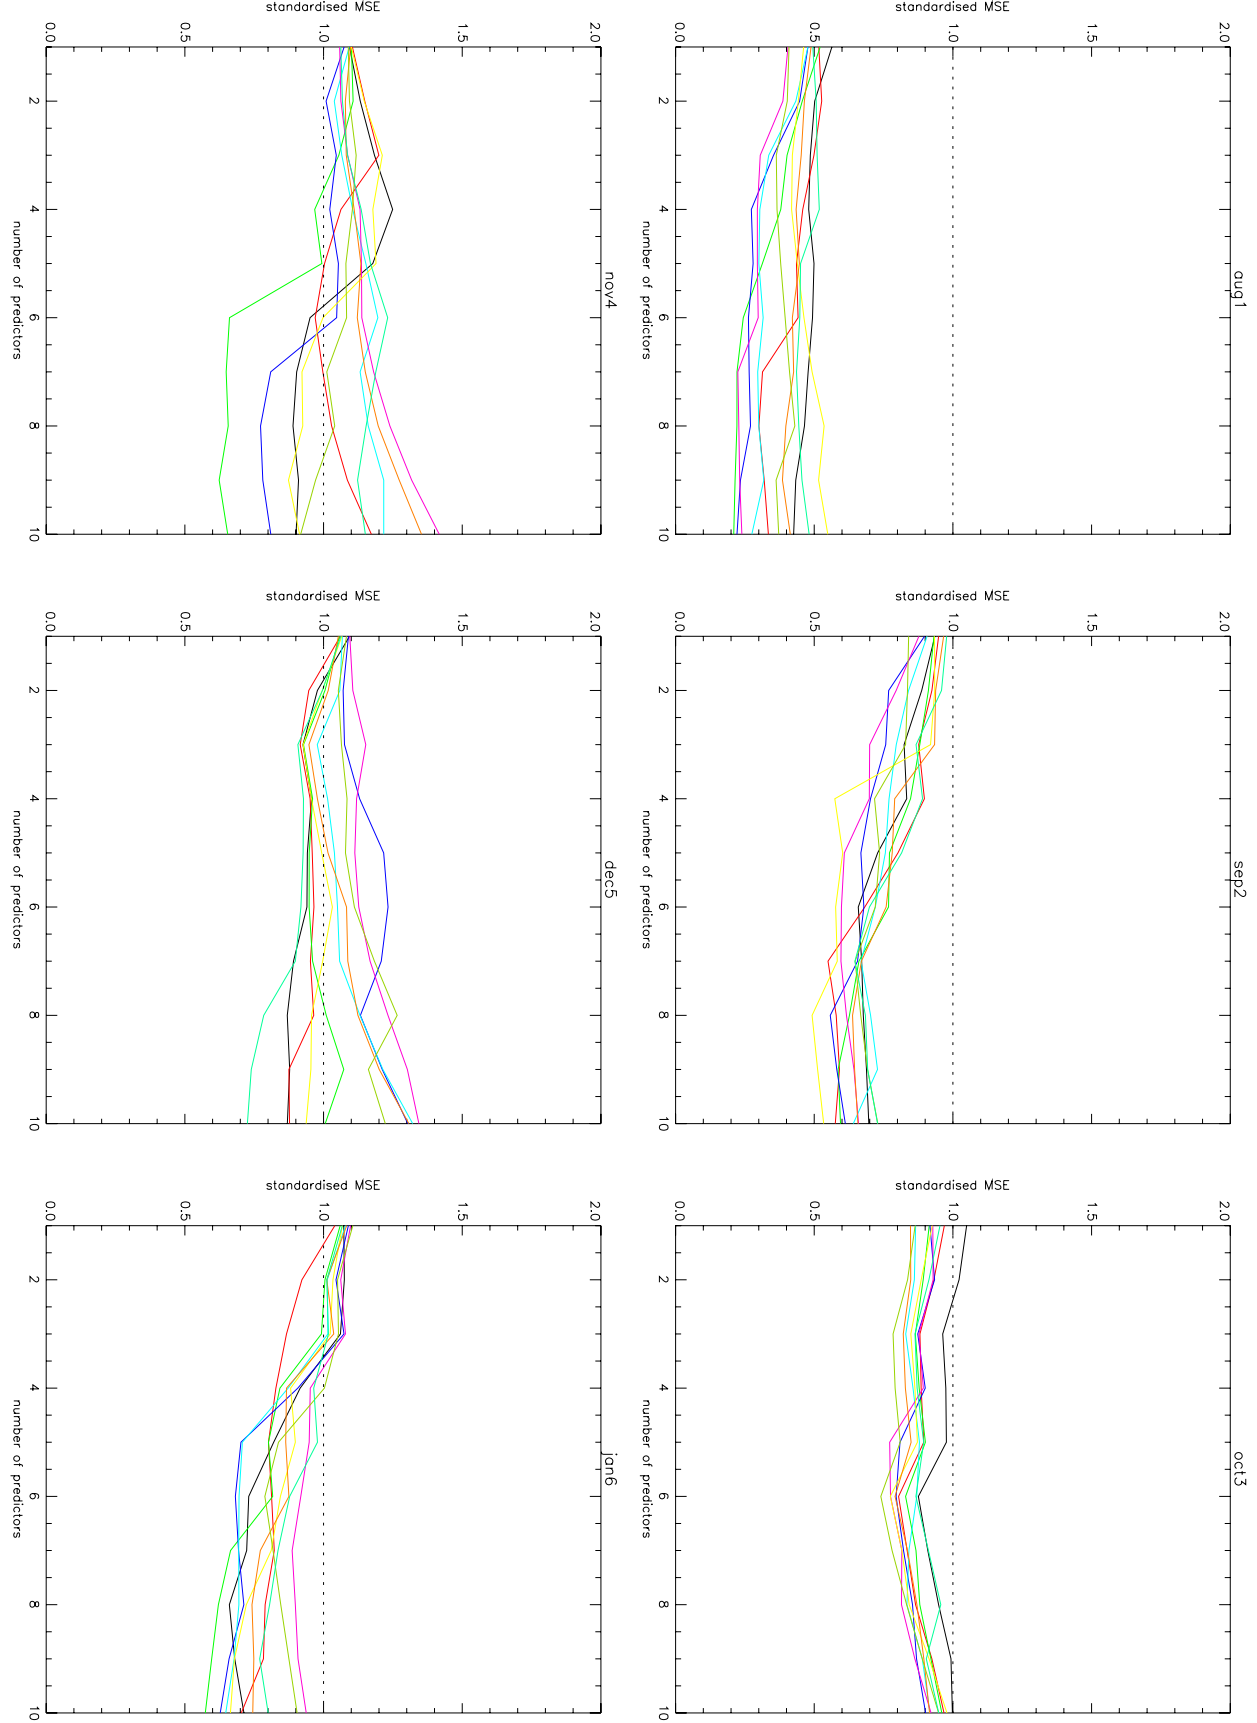


Figure 6: Variation of forecast error with the number of predictors used in the regression equation (August start, all lead times, all districts)




Figure 7: Comparison of ‘common’ (dashed lines) and ‘stepwise’ (solid lines) forecasts (November start, all lead times, all districts) — each colour (solid and dashed lines) represents one district

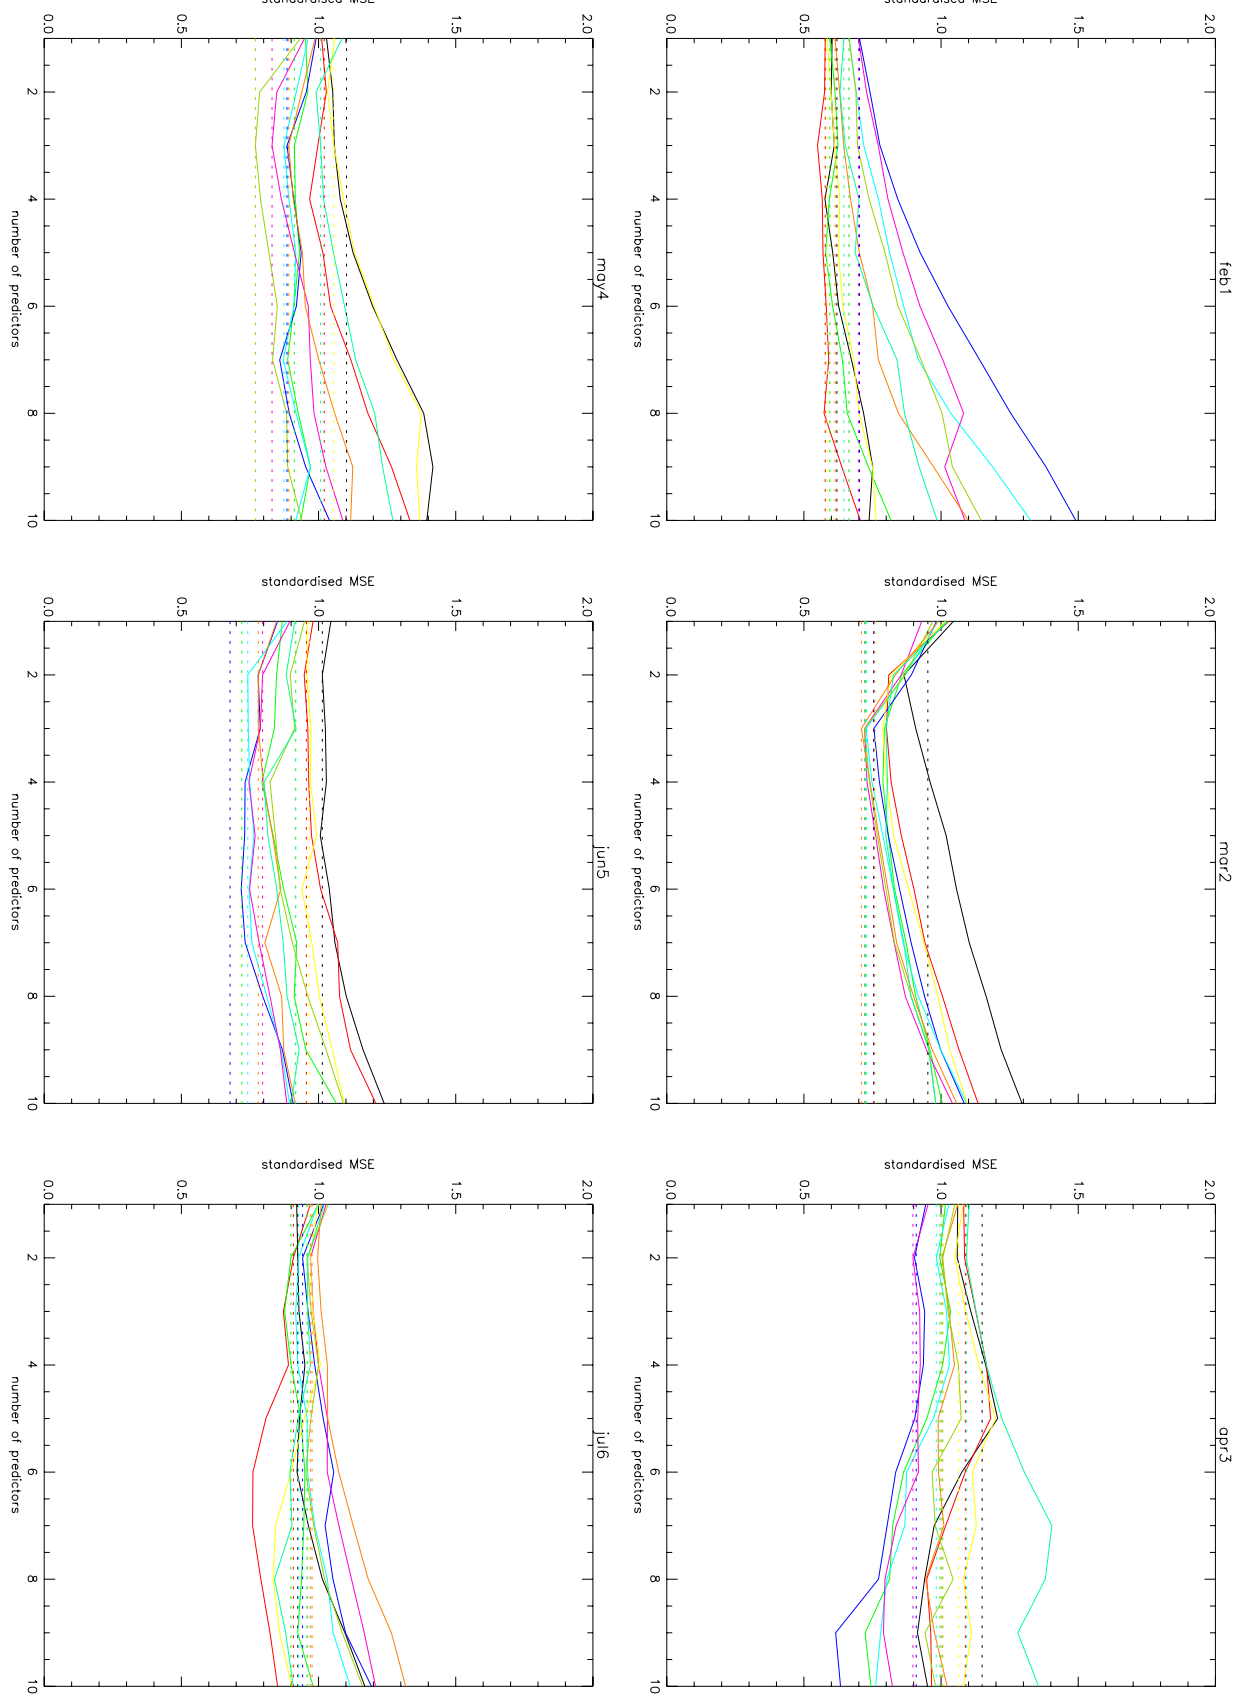


Figure 8: Comparison of ‘common’ (dashed lines) and ‘stepwise’ (solid lines) forecasts (February start, all lead times, all districts) — each colour (solid and dashed lines) represents one district




Figure 9: Comparison of ‘common’ (dashed lines) and ‘stepwise’ (solid lines) forecasts (May start, all lead times, all districts) — each colour (solid and dashed lines) represents one district

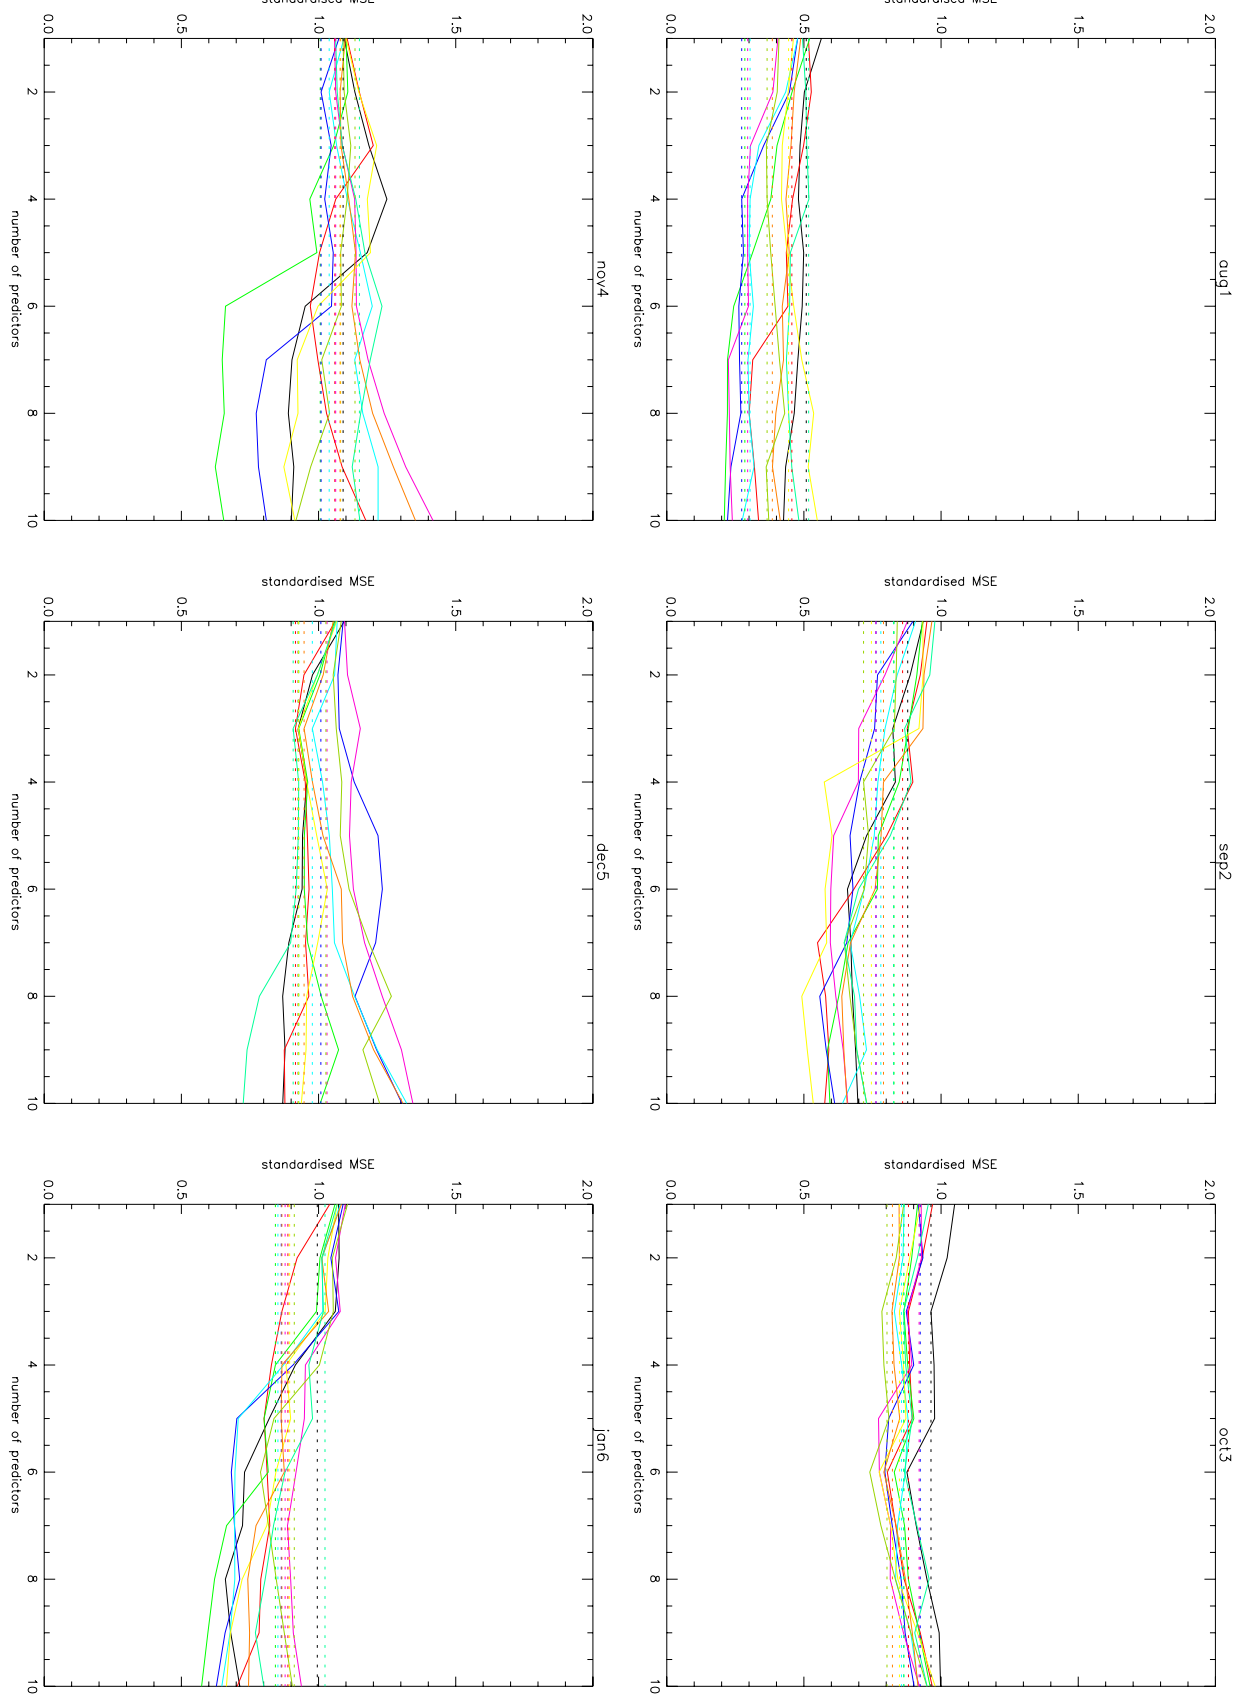


Figure 10: Comparison of ‘common’ (dashed lines) and ‘stepwise’ (solid lines) forecasts (August start, all lead times, all districts) — each colour (solid and dashed lines) represents one district




Figure 11: Variation of forecast error with the lead time, for 'nearest box' forecasts (all start times, all districts) — each colour represents one district

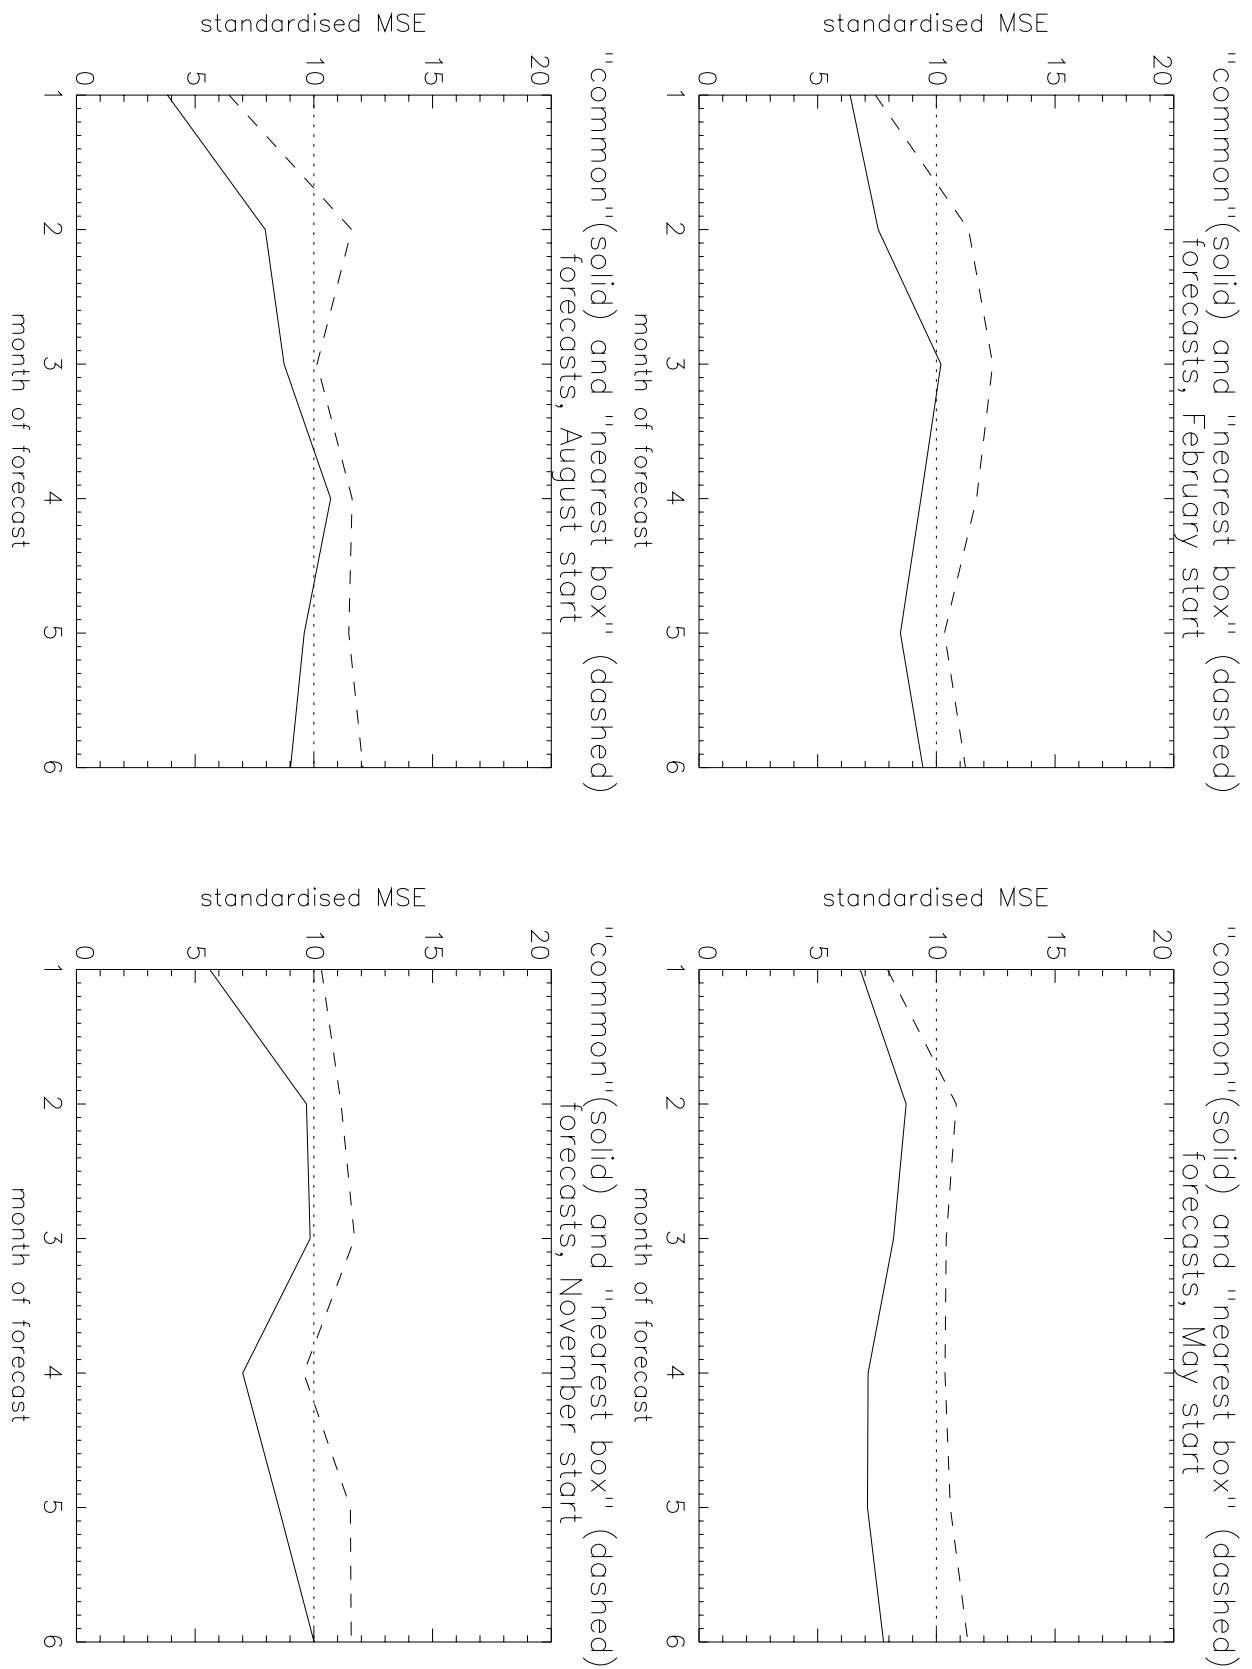


Figure 12: Comparison of 'common' (solid lines) and 'nearest box' (dashed lines) forecasts - total error for the 10 districts

A Example data

Forecasts with stepwise linear regression

SMSE of forecast and sequence of predictors chosen by stepwise regression — nov1

district 0

.7693	.7454	.5404	.5401	.5340	.5400	.5728	.6241	.6834	.7575
3	5	4	25	9	15	14	13	1	10

district 1

.8255	.8277	.4908	.4867	.5083	.5348	.5779	.6357	.6694	.7037
3	5	4	9	20	10	8	18	19	13

district 2

.8244	.5272	.4556	.4268	.4376	.4448	.4697	.4959	.5329	.5098
4	5	3	9	10	21	18	13	19	20

district 3

.7354	.5307	.4796	.4637	.4696	.4893	.5237	.5335	.5406	.5321
4	5	7	20	11	25	15	10	12	13

district 4

.7951	.5737	.5566	.5328	.5441	.5635	.5923	.6244	.6635	.7040
4	5	3	9	21	20	10	15	12	16

district 5

.7343	.5721	.5455	.5098	.5244	.5404	.5824	.6078	.6206	.6465
4	5	25	7	11	20	15	12	17	13

district 6

.8384	.8104	.5405	.5499	.5629	.5847	.6063	.6343	.6598	.6945
3	5	4	25	13	24	14	18	10	1

district 7

.8274	.6196	.6146	.5859	.6124	.6375	.5293	.5249	.5576	.5807
4	5	20	15	3	18	19	13	21	24

district 8

.7508	.6314	.5576	.5266	.5197	.5182	.4672	.4640	.4260	.4186
4	5	14	10	18	19	20	21	23	13

district 9

.8887	.9333	.9644	.8288	.8573	.9175	.9137	.9175	.9534	1.0093
11	7	5	4	20	12	18	21	14	6

for each district,

first line - standardised MSE

second line - sequence of predictors

Forecasts with same set of predictors for all districts

nov1

district 0	error .5248 - predictors: 4 5	district 1	error .5220 - predictors: 4 5
	error .5505 - predictors: 3 4 5 9		error .4867 - predictors: 3 4 5 9
	error .5401 - predictors: 3 4 5 25		error .5183 - predictors: 3 4 5 25
district 2	error .5272 - predictors: 4 5	district 3	error .5307 - predictors: 4 5
	error .4268 - predictors: 3 4 5 9		error .4396 - predictors: 3 4 5 9
	error .4910 - predictors: 3 4 5 25		error .5291 - predictors: 3 4 5 25
district 4	error .5737 - predictors: 4 5	district 5	error .5721 - predictors: 4 5
	error .5328 - predictors: 3 4 5 9		error .5437 - predictors: 3 4 5 9
	error .5799 - predictors: 3 4 5 25		error .5625 - predictors: 3 4 5 25
district 6	error .5315 - predictors: 4 5	district 7	error .6196 - predictors: 4 5
	error .5521 - predictors: 3 4 5 9		error .6227 - predictors: 3 4 5 9
	error .5499 - predictors: 3 4 5 25		error .6445 - predictors: 3 4 5 25
district 8	error .6314 - predictors: 4 5	district 9	error .7446 - predictors: 4 5
	error .6599 - predictors: 3 4 5 9		error .7995 - predictors: 3 4 5 9
	error .5995 - predictors: 3 4 5 25		error .7966 - predictors: 3 4 5 25
total	error 5.7778 - predictors: 4 5		
	error 5.6144 - predictors: 3 4 5 9		
	error 5.8114 - predictors: 3 4 5 25		

Nearest gridpoint forecasts - SMSE

nov1

district 0	error	.9461
district 1	error	1.0798
district 2	error	1.1584
district 3	error	1.0206
district 4	error	1.1109
district 5	error	1.0036
district 6	error	1.0407
district 7	error	.9823
district 8	error	.8956
district 9	error	1.0819