The use of predictors in the modal control of an AO system.

R. M. Brockie[1], M. Wells[2], P. Gallant[3] and G. J. M. Aitken[3]

[1] *Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh, EH9 3HJ.*
[2] *United Kingdom Astronomy Technology Centre, Royal Observatory, Edinburgh, EH9 3HJ.*
[3] *Department of Electrical and Computer Engineering, Queen's University, Kingston, Ontario, Canada.*

Abstract

We present results of our investigation into improving AO system performance by incorporating an extra predictive algorithm into the control loop of a "best is last" system. A number of innovations are employed in this scheme; a) the use of a deformable mirror with zero hysteresis which allows prediction of the atmosphere rather than wavefront sensor errors, b) many parallel predictors which each predict a coefficient of an orthogonal modal representation of the atmosphere, and c) switches which, based on immediate previous system performance, choose between "best is last" or an alternative prediction each iteration of the control loop. Simulations of a number of conditions have established that there is a clear benefit, which will most probably be realised as increased sky coverage for low to medium Strehl ratios.

1 Introduction

Experiments on measurements of wavefront tilt data taken with the Multiple Mirror Telescope (MMT), Mt. Hopkins, Arizona in *open loop* (*i.e.* no adaptive optics correction) have shown that the tilts are predictable using simple neural networks of linear predictors (Jorgenson and Aitken, 1992). Using data from the COME–ON system at the 3.6m ESO telescope at La Silla, again in open loop, Aitken and McGaughey have confirmed these findings and conclude that short–term prediction could overcome the inherent time delays in an AO system (Aitken and McGaughey, 1995). Their results suggest that in good seeing, neural network and linear predictors are comparable, whereas for poor seeing, neural networks outperform linear predictors by a factor of 2 in reducing the root mean square error signal.

Aitken and McGaughey have gone on to characterise the atmospheric statistics as being well described by fractional Brownian motion. They find that, in principle, the atmospheric distortions are not predictable, and that the predictability comes from the spatial filtering of the distortions by the Shack–Hartmann subaperture pupils (McGaughey and Aitken, 1997).

The question investigated here is the following. Given that the measurements of atmospheric turbulence are predictable, can this knowledge be included in a practical (*closed loop*) implementation of an AO system to boost the performance? If so, quantify the potential benefits.

One proposition is to optimise the control matrix of the AO system servo–loop given prior knowledge of the statistics of the atmospheric effects. This has been explored for the Kolmogorov case (Wild, 1996). The question arises of what to do when the statistics are non–Kolmogorov. A more general approach would be beneficial.

The work by Aitken and others shows that the atmosphere is predictable through measuring local wavefront slopes in open loop. In the case of an AO system working in closed loop, however, the measurements made by the the WFS are of the error signal: the difference between the shape of the DM and the incoming atmospheric distortion.

The wavefront error signals may not be predictable for the following reason. Assume that the position adopted by the DM is the result of a prediction made by some predictive algorithm, and the next measurement of the wavefront error position is then performed. The error signal is therefore a difference between the atmosphere and the prediction of the predictor. For the next prediction to be accurate, the predictor has to understand how both it and the atmosphere are behaving. Whether this is predictable or not will depend on the type of predictor, and is not easily known. What is certain however, is that the predictor is not concentrating specifically on the process which is known to be predictable — the measurements of atmospheric distortions.

2 Recovering *open loop* data whilst the loop is closed

When in operation, the wavefront sensor measures the following,

$$E(\underline{r}) \; = A(\underline{r}) \; + \; DM(\underline{r}), \tag{1}$$

where, $E(\underline{r})$ is the error signal used to determine the correction to the shape of the deformable mirror, $DM(\underline{r})$, and $A(\underline{r})$ is the atmospheric distortion. Equation 1 can be rearranged to give

$$A(\underline{r}) = E(\underline{r}) - DM(\underline{r}). \tag{2}$$

The behaviour of the atmosphere can therefore be recovered if the position of the DM is known accurately. In the case of piezoelectric stacked actuator mirrors, accurate positional estimates cannot be guaranteed as hysteretic effects can cause a discrepancy between the expected shape given the applied voltages and that actually attained by the mirror.

The Electra mirror (Buscher et al., 1995), to be used in Naomi (Gilmore et al., 1995) in the UK AO programme, has strain gauges on each segment of the mirror creating a positional servo–loop for each segment to ensure that the overall mirror figure is that which is desired. This results in a mirror with zero hysteresis, and allows equation 2 to be used to accurately reconstruct the behaviour of the atmosphere. The open loop case can therefore be recovered from the closed loop measurements and this is the approach used in this paper.

A second way of recovering the atmosphere from wavefront sensor error measurements is to use knowledge of the workings of the AO loop. This has been investigated in recently published concurrent work by a group at ONERA (Dessenne et al., 1997; Dessenne et al., 1998). Their method is closely related to the optimisation of the gain of individual modes of correction (Gendron and Léna, 1994). Through determination of the transfer function of an AO system, the closed loop WFS measurements can be converted into what would have been obtained in the open loop case.

3 Implementation of Predictive Algorithm

The aim is to make the addition of a predictor to the Naomi system as straightforward as possible. As much work as possible is done by the AO system, with the predictor only being employed to tackle the very well–defined problem of predicting the next atmospheric phase screen. It is envisaged as a "bolt on" addition to the existing control loop with as little interference as possible in the operation of the AO system, *e.g.* no characterisation of the AO system is required with this algorithm.

A novel feature of this implementation scheme is that the training set used to train the predictor is continually updated with the latest complete example of the behaviour of the atmosphere. The values recorded in the training set are the coefficients of the orthogonal mirror modes. The length of the training set is fixed to keep the computation required in each AO system iteration constant. The oldest data example is discarded to make way for the newest, and the predictor is re–trained at each AO system iteration so that the most recent behaviour of the atmosphere is characterised. The latest measurement of the atmosphere can the be presented to the predictor and a prediction obtained for the next AO system iteration.

A useful measure of how well the predictor is trained is to determine how well the predictor reproduces the training set and compare it to the performance of the *best is last*, standard AO system over the training set. If there is significant improvement, we describe the predictor as "well trained".

The AO system now has two predictions for what the atmosphere will do next: that of the unaided AO system and that of the predictor. The crucial element of this scheme is that the system is able to test which prediction has been performing better in the most recent past and select appropriately. A schematic diagram of this implementation is shown in figure 1.

The test is performed by comparing the variance of the error signal measured by the wavefront sensor with that that would have been obtained using the other predicted mirror shape. This is possible as the difference in the predicted

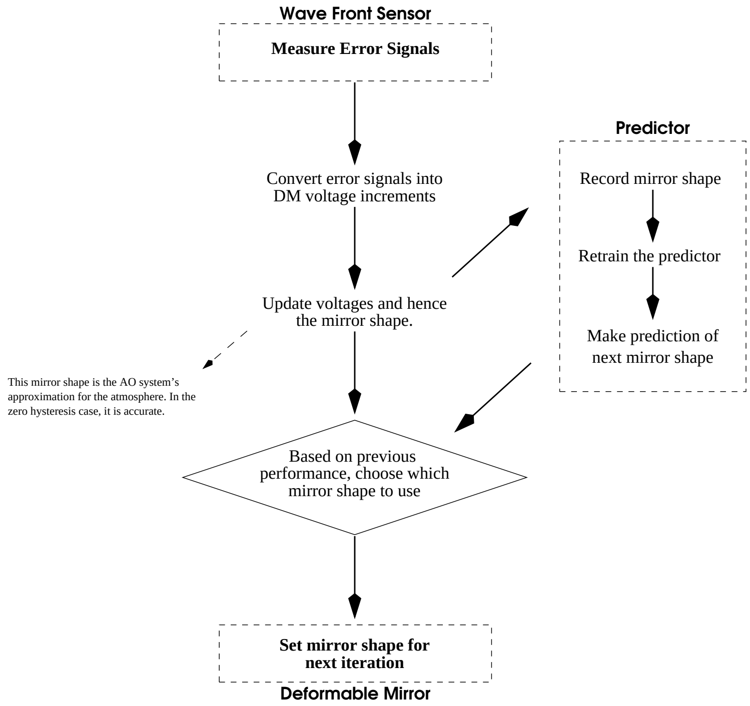

Figure 1: Schematic of the predictor implementation. The vertical path minus the rhombus is the standard AO system control loop. The addition is the predictor path and choice between the two predictions.

mirror shapes is known. Note that it would not be possible to perform this test if the mirror suffered hysteretic effects.

The default decision is to use the *best is last* AO system shape and there are three conditions which must be satisfied before the prediction is used. Most important first:

1. the prediction in the previous AO system iteration must have performed better than the unaided AO system,
2. the predictor must be performing reasonably well over the training set and
3. the correction error signal $E(\underline{r})$ is small.

This decision (switching) stage is essential as the atmosphere can evolve on a short timescale causing the predictor to become out of date. The predictor does not always outperform the *best is last* prediction from the AO system, so it is essential that the control loop can determine which estimate to use.

Unfortunately, there is still a lag in the system. It is now a lag in determining whether or not to use the predicted mirror shape. It is conceivable that there are many single, isolated examples of the predictor outperforming the unaided AO system. However, because the choice of mirror shape to use is based on the performance in the previous iteration, it is possible that the overall performance of the system will suffer. An overall gain in system performance requires that there are clumps of AO system iterations for which the predicted mirror shape outperforms the unaided AO system.

3.1 Predictors

Two predictors are investigated in this work. Linear predictors, as described in *Numerical Recipies* (Press et al., 1988), where the prediction is a linear combination of the previous values and neural networks (Müller and Reinhardt,

m_R	$t(ms)$	$\mathcal{S}_{AO}$	$\mathcal{S}_{NN}$	$\mathcal{S}_{LP}$
5.0	2.5	0.882 ± 0.001	–	–
10.0	2.5	0.872 ± 0.001	–	–
12.5	2.5	0.648 ± 0.005	0.662 ± 0.003	0.507 ± 0.010
13.5	10	0.367 ± 0.014	0.435 ± 0.010	0.323 ± 0.013
14.5	15	0.184 ± 0.014	0.226 ± 0.014	0.196 ± 0.012
15.5	35	0.117 ± 0.007	0.126 ± 0.009	0.131 ± 0.009
16.5	55	0.093 ± 0.006	0.096 ± 0.006	0.098 ± 0.006
17.5	75	0.084 ± 0.006	0.081 ± 0.005	0.082 ± 0.005

Table 1: Parameters and performance for each guide star magnitude.

1991). The linear predictors used here are more specifically defined than the general linear controller used by the ONERA group. The predictive algorithm was implemented into an existing AO simulation kindly supplied by François Rigaut (Rigaut et al., 1997).

4 Results

The goal of this work was to obtain a graph of the performance of the simulated NAOMI system with guide star magnitude, similar to that reported by Rigaut for his model of the proposed Gemini AO systems (Rigaut et al., 1997). The graph would compare unaided AO, AO with neural network predictors and AO with linear predictors.

As stated earlier, the implementation of the predictive scheme is to be as straightforward as possible. Therefore, the case against which the prediction scheme should be compared is an optimised AO system. For this reason, good seeing of $r_0(0.55\mu m) = 30cm$ was selected to make the prediction task as difficult as possible. The predictor can be thought of as working in the background, observing what is going on, and making suggestions which may or may not be accepted by the switch which has been introduced to the control loop.

A suitable integration time for the unaided AO system for each of the chosen magnitudes was determined, with the chosen magnitudes tabulated in table 1 along with the selected WFS integration time.

For the performance graph, the total system time was $10s$ with performance data being collected for the last $9s$. The length of the learning set was $0.5s$, with the servo warm–up being $0.25s$, again giving $0.25s$ of system time for convergence of the predictor to occur.

Experience has shown that predicting the lowest order modes has greatest effect on AO system performance. Each mode has it's own individual predictor, making the implementation of the predictors an inherently parallel problem. The prediction cases in this work only predicted the 3 lowest mirror mides to keep run–time on a single processor machine within reasonable limits. The AO system was correcting ~ 70 modes.

The results of this investigation are listed in table 1 and are graphed in figure 2. The time taken to obtain this data was about 3 weeks of real time or about 1 week of CPU time using a DEC AXP AlphaStation 500/400.

5 Conclusions

The use of a zero hysteresis mirror in an AO system allows the addition of the "bolt–on" predictive algorithm described in this paper. This approach avoids interference with the normal operation of the AO system as far as possible as the predictor can be on a parallel processor. The predictive element makes suggestions to the control loop, the only additional task in the control loop is to choose between the values suggested by the predictor and the normal operation of the system.

It is clear from figure 2 that the predictors of choice in this implementation scheme are the parallel neural networks. The simple linear predictor case investigated here is, generally, poorer performing than the neural network case, and

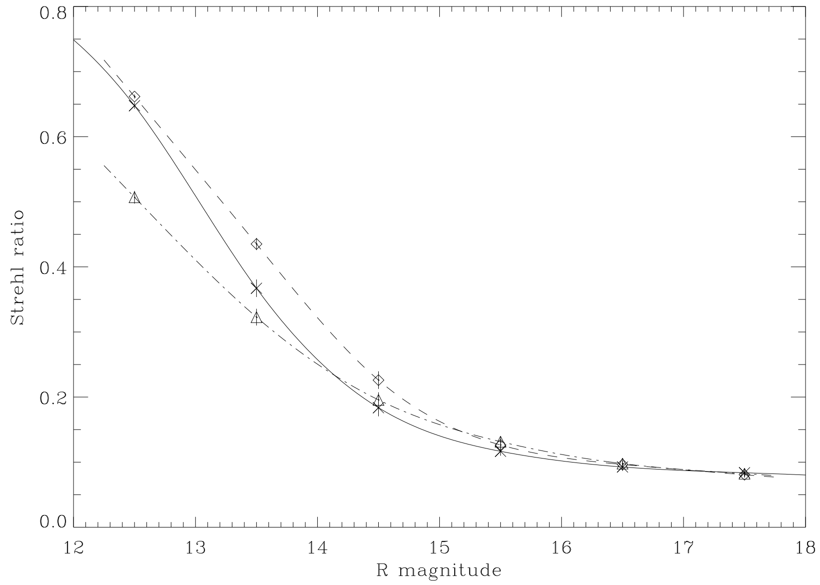

Figure 2: Comparison of system performance with magnitude:
—×— no prediction,
–◇– neural networks and
-·△- linear predictors.

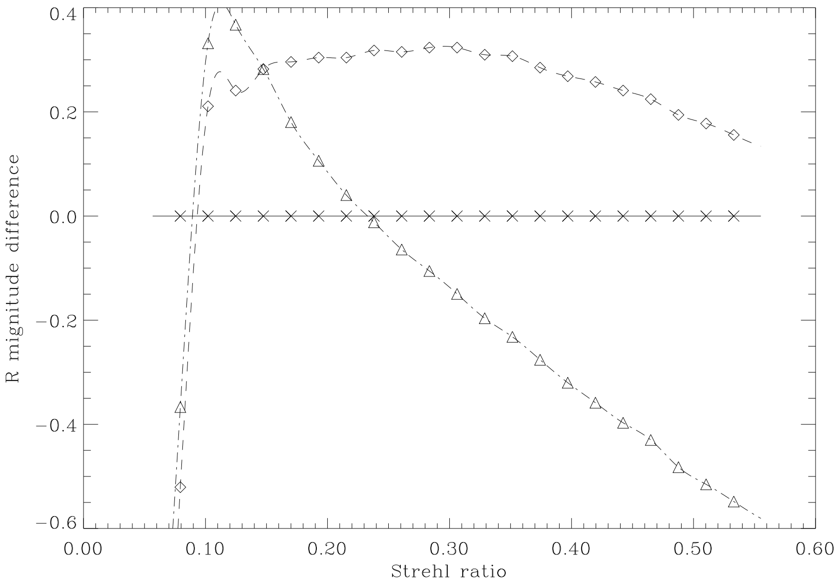

Figure 3: Increase in guide star magnitude for given Strehl:
—×— no prediction,
–◇– neural networks and
-·△- linear predictors.

for bright magnitudes, can make matters worse than the unaided AO system case — see section 3. It is possible that a more general linear predictor such would be better performing.

An interesting result can be established if the data in figure 2 is plotted somewhat differently. For a range of Strehl ratios, the difference in the magnitude required to obtain that performance is calculated with the unaided AO system providing the baseline. This is graphed in figure 3 which shows that for a wide range of Strehl, the NN's allow a guide star of ~ 0.25 magnitudes fainter to be used to obtain the same performance of the AO system.

A fuller account of the work summarised here will be presented in a paper which is currently in preparation.

References

Aitken, G. J. M. and McGaughey, D.: 1995, in (Cullum, 1995), pp 89–94

Buscher, D. F., Doel, A. P., Andrews, N., Dunlop, C., Morris, P. W., Myers, R. M., Sharples, R. M., Vick, A. J. A., Zadrozny, A., Haniff, C. A., and Wilson, R. W.: 1995, in (Cullum, 1995), pp 63–68

Cullum, M. (ed.): 1995, *Adaptive Optics*, Vol. 54 of *Conference and Workshop Proceedings*, ESO

Dessenne, C., Madec, P.-Y., and Rousset, G.: 1997, *Optics Letters* **22(20)**, 1535

Dessenne, C., Madec, P.-Y., and Rousset, G.: 1998, *Applied Optics* **37(21)**, 4623

Gendron, E. and Léna, P.: 1994, *Astron. Astrophys.* **291**, 337

Gilmore, G., Longmore, A., Myers, R., and Wilson, R.: 1995, in (Cullum, 1995), pp 9–15

Jorgenson, M. B. and Aitken, G. J. M.: 1992, *Optics Letters* **17(7)**, 466

McGaughey, D. R. and Aitken, G. J. M.: 1997, *J. Opt. Soc. Am. A* **14(8)**, 1967

Müller, B. and Reinhardt, J.: 1991, *Neural Networks: An Introduction*, Physics of Neural Networks, Springer–Verlag

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T.: 1988, *Numerical Recipies in C: The Art of Scientific Computing*, CUP

Rigaut, F., Ellerbroek, B. L., and Northcott, M. J.: 1997, *Applied Optics* **36(13)**, 2856

Wild, W. J.: 1996, *Optics Letters* **21(18)**, 1433