

RAL-TR 2000045
R3 STOLE



CCLRC Library & Info Services




C4050835

Technical Report

RAL-TR-2000-045

Estimating Signal Amplitudes in Optimal Fingerprinting, Part 1: Theory

M R Allen and P A Stott



19th January 2001

© Council for the Central Laboratory of the Research Councils 2000

Enquiries about copyright, reproduction and requests for additional copies of this report should be addressed to:

The Central Laboratory of the Research Councils
Library and Information Services
Rutherford Appleton Laboratory
Chilton
Didcot
Oxfordshire
OX11 0QX

Tel: 01235 445384 Fax: 01235 446403

E-mail library@rl.ac.uk

ISSN 1358-6254

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.

Estimating signal amplitudes in optimal fingerprinting, Part I: Theory

Myles R. Allen
Space Science & Technology Department
Rutherford Appleton Laboratory
Chilton, Didcot, OX11 0QX, U.K.
m.r.allen@rl.ac.uk

&

Peter A. Stott
Hadley Centre for Climate Prediction and Research
The Met. Office
London Road, Bracknell, RG12 2SZ, U.K.

May 7, 2000

Abstract

There is increasingly clear evidence that human influence has played a role in the large-scale climatic changes that have occurred over the past few decades. Our attention now turns to the physical implications of the emerging anthropogenic signal. Of particular interest is the question of whether, and by how much, models may be over- or under-estimating the amplitude of the climate system's response to various external influences, including anthropogenic forcing. Evidence of a significant error in a model-simulated response amplitude would signal the existence of amplifying or damping mechanisms that are inadequately represented in the model. The range of uncertainty in the factor by which we can scale model-simulated changes while remaining consistent with observed change provides an estimate of uncertainty in model-based predictions. With any model that displays a realistic level of internal variability, the problem of estimating this factor is complicated by the fact that it represents a ratio between two incompletely known quantities: both observed and simulated responses are subject to sampling uncertainty, primarily due to internal chaotic variability. Sampling uncertainty in the simulated response can be reduced, but not eliminated, through ensemble simulations. Accurate estimation of these scaling factors requires a modification of the standard "optimal fingerprinting" algorithm for climate change detection, drawing on the conventional "Total Least Squares" approach discussed in the statistical literature.

This paper describes a variant of the regression-based technique of climate change detection and attribution that is generally known as “optimal fingerprinting” (see, for example, *Hasselmann, 1979; Bell, 1986; Hasselmann, 1993; North et al., 1995; Hasselmann, 1997; Leroy, 1998; Allen & Tett, 1999*). The fingerprinting approach is to define a pattern of response to external climate forcing using a climate model and then to estimate the amplitude of that pattern, or signal, in the observed climate record. If the hypothesis of zero pattern-amplitude can be rejected with confidence, then that signal is said to be detected.

The standard approach to optimal fingerprinting assumes that the model-simulated response-pattern is known exactly, that is, it is not subject to sampling uncertainty. By sampling uncertainty, we mean the variability in the model-simulated response which would be observed if the simulation (or ensemble of simulations) were repeated with an identical model and forcing and different initial conditions. A more general definition of sampling uncertainty would encompass how the response-pattern might vary were we to use a different but equally possible model or forcing series. The generalisation of the algorithm described here to encompass these wider sources of uncertainty (which are also much more difficult to quantify) represents work in progress.

Early applications using atmosphere-ocean general circulation models (A-OGCMs) in optimal fingerprinting ensured that the assumption of zero sampling uncertainty was satisfied by using response-patterns derived from simulations of mid-21st century climate change (e.g. *Hegerl et al., 1996; Hegerl et al., 1997*). By that time, the signal-to-noise is so high that sampling uncertainty in the response-pattern can be safely neglected. The main limitation of this approach is that it confines the analysis to spatial patterns of trends – 21st century information cannot be used to determine the temporal evolution of the climate response to external forcing over the past few decades. This excludes information on the differing time-histories of different components of anthropogenic climate change and, even more importantly, excludes direct consideration of naturally-forced signals (even if the information were available, the model response to 21st century solar forcing is not relevant to its 20th century response).

An alternative approach is to use a simpler model, such as an Energy Balance Model (EBM), that is not subject to internal variability, to obtain noise-free estimates of the signals in question (*North & Stevens, 1998*). The practical difficulty here is that processes not represented in the EBM may have an influence on observed climate change, so the EBM-defined signal may focus attention “in the wrong place”. A more fundamental problem is that EBMs are no longer the primary vehicles for climate change prediction, so the information that a particular EBM is over- or under-estimating the response to a particular forcing agent is of limited value.

Direct comparison of noisy A-OGCM simulations and the observed record can proceed through correlation-based approaches, as in *Barnett & Schlesinger, 1987; Santer et al., 1993; Santer et al., 1996; Tett et al., 1996; Folland et al., 1998*. The difficulty here, again, lies in physical interpretation. Even if the correlation between modelled and observed changes is very high, this does not provide any information on whether the model-simulated amplitude of the change is accurate.

The simplest results to interpret are those based on a direct comparison of “like with like”: A-OGCM simulations of 20th century climate change compared directly with

the corresponding period in the observed record, exploiting both pattern and amplitude information in the model-data comparison (*Allen & Tett, 1999; Tett et al., 1999; Stott et al., 2000a*). To date, such studies have used ensembles to reduce sampling uncertainty in model-simulated responses and not attempted to account for this uncertainty explicitly in their analyses. The problem, of course, is that even with the four-member ensembles used in the above studies, sampling uncertainty is still far from negligible, particularly in weak signal-to-noise situations such as the analysis of the response to solar forcing. With more advanced models (*Röckner et al., 1999*), even four-member ensembles may be unfeasible.

In this paper, we describe a revised approach to optimal fingerprinting that provides unbiased estimates of pattern-amplitudes and amplitude uncertainties when model-simulated response-patterns are subject to a finite level of noise. The algorithm is applicable even in the case where the model is subject to the same level of noise as the observations, as is the case (assuming model-simulated variability is realistic) in a single-member ensemble simulation. Uncertainties will typically be quite large in this situation, for obvious reasons: if the signal is poorly known in the first place, the uncertainty in its amplitude in a noisy observed record is even larger.

Accounting for this noise is particularly important if pattern-amplitudes are to be subject to a physical interpretation, such as addressing the question of whether the A-OGCM is over- or under-estimating the response to a particular forcing agent. Standard estimates are subject to a known bias towards zero and may therefore give a misleading impression of a weaker response in the real world than that simulated by the models, even if the models are perfectly accurate. The extent of this bias in typical climate change detection problems, and the extent to which we can eliminate it, is assessed in an accompanying paper: *Stott et al., 2000b*.

Estimated upper bounds are particularly severely affected. This should be intuitively clear if we recognise that, in estimating a pattern amplitude, we are estimating the ratio between the amplitude of the observed and the model-simulated response. The presence of noise in the denominator means that the underlying noise-free model-simulated response (that which we would obtain from a hypothetical infinite ensemble) could be smaller than the response simulated in this particular experiment. If the uncertainty is large enough that this unknown noise-free response could approach zero, then the ratio between the noise-free model-simulated response and the observed response could be very high.

This point is illustrated in figure 1, which shows hypothetical estimation results based on a single observable quantity (global mean temperature trend, for example) that is subject to unit variance due to Gaussian internal climate variability. We assume that the model-simulated variability in this quantity is realistic and a four-member ensemble is available to simulate the response to external forcing, so the variance of the simulation is one quarter that of the observations. Suppose we find an observed trend of two units and an ensemble-mean model-simulated trend of one unit. The distribution of noise-free, underlying trends in the real world consistent with this observation is shown as the solid line in the upper panel, while the distribution of noise-free model trends (those we would obtain given a hypothetical infinite ensemble) consistent with this ensemble mean is shown as the dashed line. The question is, given only an observed trend of two units

and an ensemble mean model trend of one unit, by how much can we conclude that the model is over- or under-estimating the underlying noise-free response, and with what range of uncertainty?

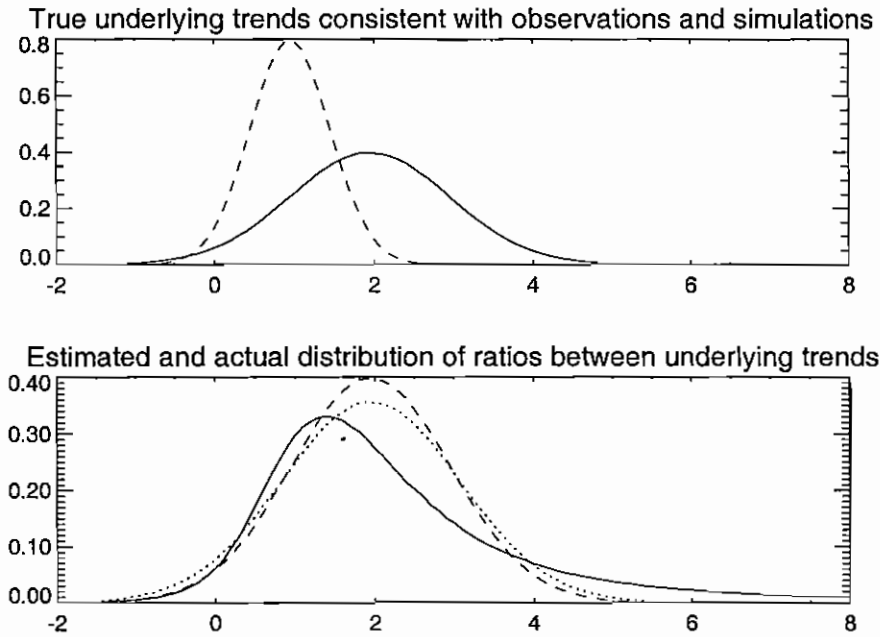


Figure 1: Upper panel: Distributions of underlying, noise-free trends in the real world (solid line) and a hypothetical climate model (dashed line) given an observed trend of two units, a four-member-ensemble mean trend of one unit and known unit variance in both modelled and observed trends. Lower panel, dashed line: Estimated distribution of ratios of real-world vs. model-simulated trends implied by conventional optimal fingerprinting (obtained simply by taking the model-simulated trend as given and accounting only for variance in the observations). Lower panel, dotted line: ditto, including the $1 + 1/\ell$ correction for finite ensemble size as used in *Allen & Tett, 1999; Tett et al., 1999*. Lower panel, solid line: true distribution of ratios obtained from random samples from the two distributions shown in the upper panel.

The answer provided by standard optimal fingerprinting, neglecting uncertainty in the model simulation altogether, is shown as the dashed line in the lower panel: this is obtained simply by dividing the distribution of trends consistent with that observed by the mean trend obtained in this particular ensemble. A simple correction on the variance of this estimate is to scale it by $1 + 1/\ell$, where ℓ is the ensemble size: this is the approach taken by *Allen & Tett, 1999; Tett et al., 1999; Stott et al., 2000a*, and is shown as the dotted line in the lower panel. It is only valid provided the signal-to-noise in the model-simulated response is high, which is not the case here. The correct answer for the distribution of ratios of observed/model-simulated responses consistent with these results is shown as the solid line in the lower panel, obtained simply by computing the distribution from samples drawn at random from the two distributions shown in the upper left panel (giving a Cauchy distribution). This is the distribution we aim to obtain (in the more general case of multiple observable quantities) using the algorithm provided

in this paper. Notice how the most important discrepancies arise in estimated upper bounds on this ratio: both standard and “corrected” variants of optimal fingerprinting suggest a very low probability of the ratio exceeding 4.5 units, whereas a significant fraction of the true distribution lies above this threshold.

Before proceeding to details, we would like to stress that, although this algorithm allows the biases due to sampling uncertainty in model-simulated responses to be quantified and largely eliminated, it is not a substitute for reducing this uncertainty directly through the use of larger ensembles. The size of ensemble required to pin down the response to an external forcing agent accurately depends, of course, on the signal-to-noise ratio. For strong signals, such as the response to greenhouse gases in the large-scale surface temperature record over the past few decades (*Tett et al., 1999; Stott et al., 2000a*), a 3- to 4-member ensemble may be enough, but for weaker signals such as the response to solar forcing, larger ensembles are likely to be required. In the drive for higher-resolution models, it should not be forgotten that single-member ensembles are of relatively little value in the analysis of observed climate change, so a compromise will always be required between model resolution and ensemble size.

1 Optimal fingerprinting as linear regression

1.1 Estimation procedure

We begin by summarising the standard linear regression algorithm as applied to the detection of climate change in order to introduce notation and make the link to the standard approach. The standard detection model assuming noise-free model-simulated response-patterns is as follows:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{v} \quad (1)$$

where \mathbf{y} is the rank- n vector of observations, the m columns of \mathbf{X} are the model-simulated response-patterns, the elements of $\boldsymbol{\beta}$ are the unknown pattern-amplitudes to be estimated and \mathbf{v} is the climate noise term. The climate noise covariance,

$$\mathbf{C}_N \equiv \mathcal{E}(\mathbf{v}\mathbf{v}^T), \quad (2)$$

is generally unknown and must be estimated from a control run of a climate model. Because the noise is generally far from white, or $\mathbf{C}_N \neq \sigma^2\mathbf{I}$, ordinary least squares regression gives highly inefficient estimators for $\boldsymbol{\beta}$ and strongly biased estimates of the errors in $\boldsymbol{\beta}$. The solution is to introduce a “pre-whitening” operator, \mathbf{P} , defined such that

$$\mathcal{E}(\mathbf{P}\mathbf{v}\mathbf{v}^T\mathbf{P}^T) = \mathbf{I}_\kappa, \quad (3)$$

where \mathbf{I}_κ is the rank- κ unit matrix. Note that κ , the rank of \mathbf{P} , may be much smaller than n , the rank of \mathbf{y} – in general, a full-rank pre-whitening operator (which would render the noise variance equal on all spatio-temporal scales) will not be available because small-scale noise variance is either unknown or poorly simulated by the model.

Because $\mathbf{P}\mathbf{v}$ is independent, identically distributed (i.i.d.) “white noise”, the best (lowest-variance) linear unbiased estimator of $\boldsymbol{\beta}$, $\bar{\boldsymbol{\beta}}$, is given by minimisation of the merit

function

$$r^2(\tilde{\boldsymbol{\beta}}) = (\mathbf{P}\mathbf{X}\tilde{\boldsymbol{\beta}} - \mathbf{P}\mathbf{y})^T(\mathbf{P}\mathbf{X}\tilde{\boldsymbol{\beta}} - \mathbf{P}\mathbf{y}) = \tilde{\mathbf{v}}^T\mathbf{P}^T\mathbf{P}\tilde{\mathbf{v}} \quad (4)$$

with respect to $\tilde{\boldsymbol{\beta}}$. At the minimum,

$$\tilde{\boldsymbol{\beta}} = (\mathbf{X}^T\mathbf{P}^T\mathbf{P}\mathbf{X})^{-1}\mathbf{X}^T\mathbf{P}^T\mathbf{P}\mathbf{y} \equiv \mathbf{F}^T\mathbf{y}, \quad (5)$$

where the rows of \mathbf{F} are the ‘‘distinguishing fingerprints’’ (Allen & Tett, 1999) and

$$r_{\min}^2 \sim \chi_{\kappa-m}^2. \quad (6)$$

The variance of the estimator $\tilde{\boldsymbol{\beta}}$ is given by the inverse of the Hessian or curvature matrix,

$$\mathbf{V}(\tilde{\boldsymbol{\beta}}) \equiv \mathcal{E}(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta})(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta})^T) \quad (7)$$

$$= \left[\frac{1}{2} \frac{\partial^2(r^2)}{\partial(\boldsymbol{\beta}')^2} \right]^{-1} \quad (8)$$

$$= (\mathbf{X}^T\mathbf{P}^T\mathbf{P}\mathbf{X})^{-1}, \quad (9)$$

as may be verified by expanding $\tilde{\boldsymbol{\beta}}$ on the RHS of (7) in terms of \mathbf{X} and \mathbf{v} .

Thus, if we ignore uncertainty in the estimated noise variance, discussed below, $\tilde{\boldsymbol{\beta}}$ is normally distributed with mean $\boldsymbol{\beta}$ and variance

$$\mathbf{V}(\tilde{\boldsymbol{\beta}}) = (\mathbf{X}^T\mathbf{P}^T\mathbf{P}\mathbf{X})^{-1}. \quad (10)$$

Equivalently,

$$(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta})^T(\mathbf{X}^T\mathbf{P}^T\mathbf{P}\mathbf{X})(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}) = \Delta r^2(\boldsymbol{\beta}) \equiv r^2(\boldsymbol{\beta}) - r_{\min}^2 \quad (11)$$

$$\sim \chi_m^2. \quad (12)$$

To map a desired confidence region, therefore, we determine the critical value of the appropriate distribution (e.g. $\chi_m^2(0.05)$) and plot the surface of values of $\boldsymbol{\beta}'_m$ for which $\Delta r^2(\boldsymbol{\beta}'_m)$, or the LHS of (12), is equal to that critical value. We can then expect (at this confidence level) the vector of true coefficient values, $\boldsymbol{\beta}$, to lie within this region.

Many applications require an estimate of uncertainty in individual signal amplitudes, or in specific combinations of signals, rather than the joint uncertainty in all components of $\tilde{\boldsymbol{\beta}}$. For example, to quantify uncertainty in the trend in global mean temperature over a particular period, we introduce a vector \mathbf{c} , where c_i is the corresponding trend in the model simulation which provides the i^{th} column of \mathbf{X} . The vector $\tilde{\boldsymbol{\beta}}$ represents the combination of model-simulated signals which best reproduces the observations, so the best-guess trend over this period is simply $\tilde{\phi} = \mathbf{c}^T\tilde{\boldsymbol{\beta}}$. The variance in $\tilde{\phi}$ is provided by $\mathbf{V}(\tilde{\phi}) = \mathbf{c}^T\mathbf{V}(\tilde{\boldsymbol{\beta}})\mathbf{c}$. To determine a confidence interval, we find the values of ϕ for which

$$\frac{(\phi - \tilde{\phi})^2}{\mathbf{c}^T\mathbf{V}(\tilde{\boldsymbol{\beta}})\mathbf{c}} = \chi_1^2(0.05). \quad (13)$$

If ϕ depends on one signal only (e.g. only one of the signals displays a trend), then the denominator of the LHS of (13) is simply equal to the corresponding diagonal element of $\mathbf{V}(\tilde{\boldsymbol{\beta}})$, being the familiar standard error of the estimate.

It will prove helpful to note that an equivalent method of computing confidence intervals in individual signals or signal-combinations is to map the surface of values β_1 for which $\Delta\hat{r}^2(\beta_1) = \chi_1^2(0.05)$. We then compute $\phi_1 = \mathbf{c}^T \beta_1$ for all points on this surface, and the limits on the confidence interval are given by the maximum and minimum values of the ϕ_1 . This is exactly equivalent to equation (13) when an estimate $\hat{V}(\tilde{\beta})$ is available (Press et al., 1992), but proves useful when it is not.

1.2 Sampling uncertainty in the estimated noise variance

Ignoring uncertainty in the estimated noise properties simplifies the analysis but can lead to “artificial skill” (systematic bias towards underestimation of uncertainties) in a climate change detection context (Bell, 1986). The simplest way to deal with this problem is to base the uncertainty analysis on a set of ν noise realisations, $\hat{\mathbf{Y}}_2$, which are statistically independent of any noise realisations used to estimate \mathbf{P} . The ν columns of $\hat{\mathbf{Y}}_2$ correspond to \mathbf{y} -like vectors of pure noise (“pseudo-observations”) drawn, for example, from ν statistically independent segments of a control integration of a climate model. If adjacent columns of $\hat{\mathbf{Y}}_2$ are not statistically independent, the degrees of freedom of covariance estimated based on this noise realisation, ν_2 , will be less than the number of columns, ν .

The standard fingerprinting algorithm is linear, so the variance of $\tilde{\beta}$ does not depend on the actual amplitude of the signal in the observations, \mathbf{y} . An estimate of $V(\tilde{\beta})$ can therefore be obtained by applying the same operator, \mathbf{F}^T , which was used to extract $\tilde{\beta}$ from \mathbf{y} to the columns of $\hat{\mathbf{Y}}_2$ and estimating the variance of the result:

$$\hat{V}(\tilde{\beta}) \equiv \frac{\mathbf{F}^T \hat{\mathbf{Y}}_2 \hat{\mathbf{Y}}_2^T \mathbf{F}}{\nu_2}. \quad (14)$$

Taking into account the sampling uncertainty in $\hat{\mathbf{Y}}_2$ gives

$$(\tilde{\beta} - \beta)^T \hat{V}(\tilde{\beta})^{-1} (\tilde{\beta} - \beta) \sim m F_{m, \nu_2}, \quad (15)$$

which can be used in place of equation (12) to provide confidence intervals on $\tilde{\beta}$.

We also have

$$\hat{r}_{\min}^2 \equiv \sum_{i=1}^{\kappa} \frac{(\mathbf{P} \tilde{\mathbf{v}} \tilde{\mathbf{v}}^T \mathbf{P}^T)_{i,i}}{\frac{1}{\nu_2} (\mathbf{P} \hat{\mathbf{Y}}_2 \hat{\mathbf{Y}}_2^T \mathbf{P}^T)_{i,i}} \sim (\kappa - m) F_{\kappa - m, \nu_2}. \quad (16)$$

This is equivalent to the F -test for residual consistency proposed in Allen & Tett, 1999, assuming that $\mathbf{C}_{N_2}^{-1}$ in their equation (20) refers to the pseudo-inverse in the detection space defined by \mathbf{P} – i.e. ignoring the off-diagonal elements of $\mathbf{P} \hat{\mathbf{Y}}_2 \hat{\mathbf{Y}}_2^T \mathbf{P}^T$. Since \mathbf{C}_{N_2} as defined in Allen & Tett, 1999, is non-invertible, there is potential ambiguity as to which pseudo-inverse should be used. The algorithm detailed here gives reasonably unbiased results and is the one used in Allen & Tett, 1999, Tett et al., 1999 and Stott et al., 2000a. Not all pseudo-inverses are similarly well-behaved (Tett, 1999, *pers. comm.*).

Uncertainty intervals in individual signals or signal-combinations are given by replacing equation (13) with

$$\frac{(\phi - \tilde{\phi})^2}{\mathbf{c}^T \hat{V}(\tilde{\beta}) \mathbf{c}} = F_{1, \nu_2}(0.05). \quad (17)$$

It will prove helpful to note that an equivalent method of computing confidence intervals in individual signals or signal-combinations is to map the surface of values β_1 for which $\Delta\hat{r}^2(\beta_1) = \chi_1^2(0.05)$. We then compute $\phi_1 = \mathbf{c}^T\beta_1$ for all points on this surface, and the limits on the confidence interval are given by the maximum and minimum values of the ϕ_1 . This is exactly equivalent to equation (13) when an estimate $\hat{V}(\tilde{\beta})$ is available (Press et al., 1992), but proves useful when it is not.

1.2 Sampling uncertainty in the estimated noise variance

Ignoring uncertainty in the estimated noise properties simplifies the analysis but can lead to “artificial skill” (systematic bias towards underestimation of uncertainties) in a climate change detection context (Bell, 1986). The simplest way to deal with this problem is to base the uncertainty analysis on a set of ν noise realisations, $\hat{\mathbf{Y}}_2$, which are statistically independent of any noise realisations used to estimate \mathbf{P} . The ν columns of $\hat{\mathbf{Y}}_2$ correspond to \mathbf{y} -like vectors of pure noise (“pseudo-observations”) drawn, for example, from ν statistically independent segments of a control integration of a climate model. If adjacent columns of $\hat{\mathbf{Y}}_2$ are not statistically independent, the degrees of freedom of covariance estimated based on this noise realisation, ν_2 , will be less than the number of columns, ν .

The standard fingerprinting algorithm is linear, so the variance of $\tilde{\beta}$ does not depend on the actual amplitude of the signal in the observations, \mathbf{y} . An estimate of $V(\tilde{\beta})$ can therefore be obtained by applying the same operator, \mathbf{F}^T , which was used to extract $\tilde{\beta}$ from \mathbf{y} to the columns of $\hat{\mathbf{Y}}_2$ and estimating the variance of the result:

$$\hat{V}(\tilde{\beta}) \equiv \frac{\mathbf{F}^T \hat{\mathbf{Y}}_2 \hat{\mathbf{Y}}_2^T \mathbf{F}}{\nu_2}. \quad (14)$$

Taking into account the sampling uncertainty in $\hat{\mathbf{Y}}_2$ gives

$$(\tilde{\beta} - \beta)^T \hat{V}(\tilde{\beta})^{-1} (\tilde{\beta} - \beta) \sim mF_{m,\nu_2}, \quad (15)$$

which can be used in place of equation (12) to provide confidence intervals on $\tilde{\beta}$.

We also have

$$\hat{r}_{\min}^2 \equiv \sum_{i=1}^{\kappa} \frac{(\mathbf{P}\tilde{\mathbf{v}}\tilde{\mathbf{v}}^T\mathbf{P}^T)_{i,i}}{\frac{1}{\nu_2}(\mathbf{P}\hat{\mathbf{Y}}_2\hat{\mathbf{Y}}_2^T\mathbf{P})_{i,i}} \sim (\kappa - m)F_{\kappa-m,\nu_2}. \quad (16)$$

This is equivalent to the F -test for residual consistency proposed in Allen & Tett, 1999, assuming that $\mathbf{C}_{N_2}^{-1}$ in their equation (20) refers to the pseudo-inverse in the detection space defined by \mathbf{P} – i.e. ignoring the off-diagonal elements of $\mathbf{P}\hat{\mathbf{Y}}_2\hat{\mathbf{Y}}_2^T\mathbf{P}^T$. Since \mathbf{C}_{N_2} as defined in Allen & Tett, 1999, is non-invertible, there is potential ambiguity as to which pseudo-inverse should be used. The algorithm detailed here gives reasonably unbiased results and is the one used in Allen & Tett, 1999, Tett et al., 1999 and Stott et al., 2000a. Not all pseudo-inverses are similarly well-behaved (Tett, 1999, *pers. comm.*).

Uncertainty intervals in individual signals or signal-combinations are given by replacing equation (13) with

$$\frac{(\phi - \tilde{\phi})^2}{\mathbf{c}^T \hat{V}(\tilde{\beta}) \mathbf{c}} = F_{1,\nu_2}(0.05). \quad (17)$$

Note that, because $t_{\nu_2}(0.025) = \sqrt{F_{1,\nu_2}(0.05)}$, this corresponds to the same critical value of the two-tailed t -distribution (or the one-tailed distribution at confidence level $P/2$).

2 Accounting for noise in model-simulated response

We now consider the impact of introducing noise in \mathbf{X} as well as in \mathbf{y} . The only source of noise we consider here is the sampling uncertainty due to the use of a small ensemble of simulations to obtain the response-pattern. A second source of noise in the response-patterns arises from uncertainty in model-formulation itself, generally referred to as “model error”. To the extent that model error can be represented by a covariance matrix, it could also be incorporated into the following framework, but there are practical difficulties: the most obvious method of estimating the model error covariance matrix (averaging second-order moments over a large ensemble of independent climate models) is likely to be unfeasible for the foreseeable future.

One way of avoiding the use of large ensembles while still relying on the standard optimal fingerprinting algorithm for model-data comparison would be to use a noise-free model, such as an Energy Balance Model, to simulate the response-patterns (e.g. *Stevens & North, 1996; Wigley et al., 1997; North & Stevens, 1998*). Unless, however, the spatio-temporal pattern of response simulated by the EBM can be assumed to be accurate (which seems implausible, since so many important processes are omitted from these models), the advantage is illusory. In a straightforward detection problem (assessing whether the amplitude of a particular pattern is significantly different from zero), inaccuracies in the model-simulated pattern are of secondary importance because although they may reduce the power of the algorithm, they are unlikely to lead to a false-positive result. If, on the other hand, pattern-amplitudes are to be interpreted physically (e.g. *Allen et al., 2000*), inaccuracies in model-simulated responses become very important. A low estimated amplitude of an incorrectly-specified greenhouse response pattern does not mean that the true greenhouse response in the real world is small: it simply means that the algorithm was looking in the wrong direction.

Having recognised that optimal fingerprinting is simply a variant of linear regression (e.g. *Leroy, 1998; Allen & Tett, 1999*), the solution to the problem of noise in both “independent” (predictor) and “dependent” (predictand) variables (being the response-patterns and observations respectively) is readily available (*Adcock, 1878; Deming, 1943; van Huffel & Vanderwaal, 1994*). Several approaches have been proposed, the differences between which are likely to be much less important than the impact of neglecting response-pattern noise altogether.

Press et al., 1992, present the “iterated weighted least squares” (IWLS) solution of *Jeffreys, 1980*, reviewed in *Lybanon, 1984*. *Ripley & Thompson, 1987*, discuss biases in the IWLS solution in the situation where the noise variance increases monotonically with signal amplitude. This might easily be the case in a climate change detection problem, despite the use of a pre-whitening operator which, in principle, assigns equal noise variance to all input variables. The pre-whitening operator is based on a control simulation of a climate model, and because of non-linearities in the real climate system which may not be adequately represented in the model (e.g. *Palmer, 1999*), model-

simulated variance may be underestimated in precisely the patterns in which the signal amplitude is large. We therefore follow *Ripley & Thompson, 1987's* approach, based on maximum likelihood fitting of a functional relationship (MLFR), which is closely related to *Adcock, 1878's* original solution. Their discussion relates to the case of a single independent variable (response-pattern) and noise which is uncorrelated between observations. We have been unable to locate a discussion of the multivariate problem (taking either the IWLS or MLFR approach) when the noise has a complex autocorrelation structure, such as is the case in climate change detection. While the treatment of this problem is a completely straightforward generalisation of these previous discussions, we feel it would be helpful to present it here, to make it available to the climate research community.

2.1 The Total Least Squares algorithm

Having confined ourselves to sampling uncertainty due to finite ensemble size, and assuming model-simulated variability is consistent with that in the real world, we can assume that the noise has the same autocorrelation structure in \mathbf{y} and in every column of \mathbf{X} . Under these circumstances, the same pre-whitening operator, \mathbf{P} , may be applied to all variables. Note that this assumption will only be valid if the dominant source of noise is internal climate variability which is correctly simulated by the model used to generate the columns of \mathbf{X} . We should not *a priori* expect noise due to model error or observation error to share the autocorrelation structure of climate variability, so if either of these is a significant contributor then a more complex treatment is called for which we will pursue elsewhere. If the expected noise variance in \mathbf{X} is different from that in \mathbf{y} – for example, if ensemble means have been used to reduce noise in model-simulated response-patterns – individual columns of \mathbf{X} can simply be scaled up to make the expected noise variance in each the same as that in \mathbf{y} , and the same scaling factor(s) applied to the final parameter estimates. For simplicity, we will ignore these scaling factors in the following, so this discussion applies directly to the case of single-member ensembles.

Given that there is the same noise in each column of \mathbf{X} as expected in \mathbf{y} , we have

$$\mathcal{E}(\mathbf{P}\mathbf{v}_{\mathbf{x}_i}\mathbf{v}_{\mathbf{x}_i}^T\mathbf{P}^T) = \mathbf{I}_\kappa \quad (18)$$

and

$$\mathcal{E}(\mathbf{P}\mathbf{v}_y\mathbf{v}_y^T\mathbf{P}^T) = \mathbf{I}_\kappa. \quad (19)$$

If $m' \equiv m + 1$, we define the $m' \times \kappa$ matrix

$$\mathbf{Z} \equiv [\mathbf{P}\mathbf{X}, \mathbf{P}\mathbf{y}] \quad (20)$$

as the observed (pre-whitened but still noise-contaminated) values of \mathbf{X} and \mathbf{y} .

Our basic linear model asserts that there exists a \mathbf{Z}^{true} whose columns are linearly related, that is

$$\mathbf{Z}^{\text{true}}\mathbf{v} = (\mathbf{Z} - \mathbf{\Upsilon})\mathbf{v} = \mathbf{0}, \quad (21)$$

where \mathbf{v} is a rank- m' vector of coefficients, and $\mathbf{\Upsilon}$ is an $m' \times \kappa$ matrix representing the true (pre-whitened) noise contamination in the m' variables. All the elements of $\mathbf{\Upsilon}$ are

normally distributed with unit variance, so the maximum likelihood estimator of \mathbf{v} , $\tilde{\mathbf{v}}$, is given by maximising

$$L = \text{constant} - \frac{1}{2} \text{tr}(\tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}}), \quad (22)$$

where $\tilde{\mathbf{Y}} = \mathbf{Z} - \tilde{\mathbf{Z}}$ and $\tilde{\mathbf{Z}}\tilde{\mathbf{v}} = \mathbf{0}$. The rows of $\tilde{\mathbf{Y}}$ are uncorrelated with $\tilde{\mathbf{v}}$, so maximising L is equivalent to minimising the merit function

$$r^2 = \tilde{\mathbf{v}}^T \tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}} \tilde{\mathbf{v}}. \quad (23)$$

We require a constraint to avoid the trivial solution $\tilde{\mathbf{v}} = \mathbf{0}$. Since all columns of \mathbf{Z} are subject to noise of equal amplitude, this constraint should not discriminate between them, so we use the standard normalisation, $\tilde{\mathbf{v}}^T \tilde{\mathbf{v}} = 1$. Imposing the constraint $\tilde{v}_{m'} = -1$ (i.e. only constraining the coefficient on \mathbf{y}) gives the standard regression model (1).

Incorporating this constraint into our merit function gives

$$r^2(\tilde{\mathbf{v}}) = \tilde{\mathbf{v}}^T \mathbf{Z}^T \mathbf{Z} \tilde{\mathbf{v}} + \lambda^2 (1 - \tilde{\mathbf{v}}^T \tilde{\mathbf{v}}) \quad (24)$$

where λ^2 is a Lagrange multiplier. In geometric terms, minimising r^2 is equivalent to finding the m -dimensional plane in an m' -dimensional space which minimises the sum squared perpendicular distance from the plane to the κ points defined by the rows of \mathbf{Z} – the *Adcock*, 1878, solution.

Differentiation of (24) with respect to $\tilde{\mathbf{v}}$ gives an eigen-equation defining the stationary points of r^2 at which

$$\frac{\partial(r^2)}{\partial(\tilde{\mathbf{v}})} = \mathbf{Z}^T \mathbf{Z} \tilde{\mathbf{v}} - \lambda^2 \tilde{\mathbf{v}} = \mathbf{0}, \quad (25)$$

and the curvature matrix

$$\frac{1}{2} \frac{\partial^2(r^2)}{\partial(\tilde{\mathbf{v}})^2} = \mathbf{Z}^T \mathbf{Z} - \lambda^2 \mathbf{I}. \quad (26)$$

The solution which minimises r^2 is $\lambda^2 = \lambda_{\min}^2$, the smallest eigenvalue of $\mathbf{Z}^T \mathbf{Z}$, and $\tilde{\mathbf{v}}$ being the corresponding eigenvector (the vector normal to the best-fit m -dimensional plane). In a practical implementation, we simply take the singular value decomposition (SVD) $\mathbf{Z} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}^T$ so, after sorting, $\tilde{\mathbf{v}} = \mathbf{v}_{m'}$.

The m^{th} element of the solution vector corresponds to the best-fit scaling parameter on the observations, \mathbf{y} . Since we are looking for a model to reproduce the observations themselves, not some scaled version thereof, we translate these coefficients into more familiar pattern-amplitudes by taking the ratios $\tilde{\beta}_i = \tilde{v}_i / \tilde{v}_{m'}$. At the minimum,

$$r_{\min}^2 = \lambda_{\min}^2 \sim \chi_{\kappa-m}^2, \quad (27)$$

provided $\kappa \gg m$ and neglecting, for now, uncertainty in the noise estimate. The properties of r_{\min}^2 when $\kappa \simeq m$ require a discussion in terms of Wishart matrices, which we will not attempt here. This provides a check on residual consistency directly analogous to that which we use in standard regression (*Allen & Tett*, 1999).

Uncertainty analysis of the $\tilde{\beta}$ is somewhat more complicated. If the diagonal matrix Λ^2 contains the ranked eigenvalues of $\mathbf{Z}^T\mathbf{Z}$, and the columns of \mathbf{V} contain the corresponding eigenvectors, \mathbf{v}_i , we can rewrite (26) as

$$\frac{1}{2} \frac{\partial^2(r^2)}{\partial(\tilde{\mathbf{v}})^2} = \mathbf{V}(\Lambda^2 - \lambda_{\min}^2)\mathbf{V}^T. \quad (28)$$

Note the relationship between equation (28) and the perturbation analysis of *North et al., 1982*: like them, we are analysing the stability of an eigen-decomposition. It is tempting to treat the pseudo-inverse of (28) as a standard covariance matrix on $\tilde{\mathbf{v}}$, and this is indeed a reasonable approximation in the limit of high signal-to-noise. Because of the non-linearity introduced by the normalisation constraint on $\tilde{\mathbf{v}}$, however, the merit function r^2 is not quadratic, so more realistic confidence intervals are obtained by explicitly mapping surfaces \mathbf{v} where $\Delta r^2(\mathbf{v}) = r^2(\mathbf{v}) - r_{\min}^2$ has some constant value. As in the standard model,

$$\Delta r^2(\mathbf{v}) = \mathbf{v}^T\mathbf{V}(\Lambda^2 - \lambda_{\min}^2\mathbf{I})\mathbf{V}^T\mathbf{v} \quad (29)$$

$$\sim \chi_m^2, \quad (30)$$

so having selected a confidence level, we compute the corresponding critical values of the χ_m^2 distribution and map the vectors \mathbf{v} for which $\Delta r^2(\mathbf{v})$ is equal to this critical value.

In a practical implementation, this mapping is achieved by first defining a set of points on an m -sphere of radius $\sqrt{\Delta r_{\text{crit}}^2}$, the critical value of the χ_m^2 or (next sub-section) F_{m,ν_2} distribution:

$$\sum_{i=1}^m a_i^2 = \Delta r_{\text{crit}}^2. \quad (31)$$

For each of these points, we compute

$$b_i = \frac{a_i}{\sqrt{\lambda_i^2 - \lambda_{\min}^2}}. \quad (32)$$

If the b_i provide the weights on eigenvectors 1 – m in \mathbf{V} used to generate \mathbf{v} , then equation (29) is automatically satisfied. The weight on $\mathbf{v}_{m'}$ is provided by the normalisation constraint,

$$\sum_{i=1}^{m'} b_i^2 = 1. \quad (33)$$

If Δr_{crit}^2 is too large, then $b_{m'}$ will be zero or imaginary, and the confidence region will be unbounded in this direction. This means that the \mathbf{v} are unconstrained (at this confidence level) to rotations through 360° in some plane.

In order to express these uncertainties in terms of familiar pattern-amplitudes, we need to take ratios $\beta_i = v_i/v_{m'}$, giving confidence regions which may differ markedly from ellipsoidal. To compute confidence intervals on individual signals or signal-combinations, we compute surfaces β_1 for which $\Delta r^2(\mathbf{v})$ is equal to the appropriate critical value of the χ_1^2 distribution, and simply take the maximum and minimum values of $\phi_1 = \mathbf{c}^T\beta_1$ as before. Note that certain signal-combinations can be well constrained by the observations even if the full m -dimensional confidence interval is open-ended.

2.2 Uncertainty in noise variance under Total Least Squares

If we take into account noise in all variables, the estimation algorithm becomes non-linear, so we cannot simply compute a series of $\tilde{\beta}$ -like estimates from the columns of $\hat{\mathbf{Y}}_2$ and use the resulting distribution to provide a confidence interval. Instead, we note that each λ_i^2 represents the signal-to-noise ratio in the corresponding pair of singular vectors of \mathbf{Z} :

$$\lambda_i^2 = \frac{\mathbf{u}_i^T \mathbf{Z} \mathbf{Z}^T \mathbf{u}_i}{\frac{1}{\mu} \mathbf{u}_i^T \mathbf{P} \hat{\mathbf{Y}}_1 \hat{\mathbf{Y}}_1^T \mathbf{P}^T \mathbf{u}_i}, \quad (34)$$

where the denominator is identical to unity if the μ columns of $\hat{\mathbf{Y}}_1$ have been used to derive the prewhitening operator, \mathbf{P} . Relying on these λ for the uncertainty analysis may be misleading if $\hat{\mathbf{Y}}_1 \hat{\mathbf{Y}}_1^T$ is rank-deficient, as will generally be the case. Poorly sampled state-space directions will automatically be given high weight by the pre-whitening operator, since by construction $\mathbf{P} \hat{\mathbf{Y}}_1 \hat{\mathbf{Y}}_1^T \mathbf{P}^T = \mu \mathbf{I}$. This artificially inflates the differences between eigenvalues and reduces estimated uncertainties.

The solution, as in the linear case, is simply to replace $\hat{\mathbf{Y}}_1$ with $\hat{\mathbf{Y}}_2$ in equation (34), giving

$$\hat{\lambda}_i^2 = \frac{\mathbf{u}_i^T \mathbf{Z} \mathbf{Z}^T \mathbf{u}_i}{\frac{1}{\nu_2} \mathbf{u}_i^T \mathbf{P} \hat{\mathbf{Y}}_2 \hat{\mathbf{Y}}_2^T \mathbf{P}^T \mathbf{u}_i}. \quad (35)$$

Because $\hat{\mathbf{Y}}_2$ is independent of \mathbf{P} , if \mathbf{P} is artificially inflating variance in a particular state-space direction, the same bias will apply to both numerator and denominator in equation (35), whereas in equation (34) it only applied to the numerator. The estimates, $\tilde{\mathbf{v}}$ and $\tilde{\beta}$ are unaffected, but r^2 is replaced by \hat{r}^2 in the uncertainty analysis. Exactly as in the linear case, the check for residual consistency, equation (27), becomes

$$\hat{r}_{\min}^2 = \hat{\lambda}_{\min}^2 \sim (\kappa - m) F_{(\kappa-m), \nu_2}, \quad (36)$$

and equation (30) becomes

$$\Delta \hat{r}^2(\mathbf{v}) = \mathbf{v}^T \mathbf{V} (\hat{\Lambda}^2 - \hat{\lambda}_{\min}^2 \mathbf{I}) \mathbf{V}^T \mathbf{v} \quad (37)$$

$$\sim m F_{m, \nu_2}. \quad (38)$$

2.3 The problem of open-ended confidence intervals

In interpreting coefficients derived from standard regression, we are accustomed to see these coefficients decline to zero as the amplitude of the signal in the observations goes to zero. With noise in both observations and model-simulated patterns, there is no reason for the vector $\tilde{\beta}$ to prefer one orientation over any other as the signal amplitude in both model and observations goes to zero. Its orientation thus becomes arbitrary, and the ratios $\tilde{\beta}_i = \tilde{\mathbf{v}}_i / \tilde{\mathbf{v}}_{m'}$ can take any value. The physical interpretation of a near-infinite ‘‘pattern-amplitude’’ requires some thought. What it means is that, because we are allowing for the presence of noise in the model-simulated patterns, it may be the case that the true response-pattern (the pattern which we would obtain if we were to run an

infinite ensemble) may be close to zero everywhere. Supposing that this pattern has a finite amplitude in the observations, this mean we would have to multiply this near-zero pattern by an arbitrarily large number to get a reasonable fit.

At higher signal-to-noise levels, the main impacts of adopting these revised estimators in place of the standard approach will generally be to increase the best-guess pattern amplitudes and also to increase, possibly substantially, the estimated upper bounds on parameter uncertainty ranges. Lower bounds (which are crucial for claims of detection) may either increase or decrease depending on the confidence level specified. The question will doubtless arise as to whether the additional precision of these revised estimators justifies the additional challenges of interpretation. In view of the fact that we know that any response-patterns obtained from small ensembles are subject to sampling uncertainty, there should be no question that we should satisfy ourselves what the impact of this uncertainty is within an unbiased estimation framework. Whenever these revised estimators give qualitatively different results from the traditional approach, we have two options: either we use the revised approach in policy advice, accepting the additional complexity that this entails; or we increase the size of the ensembles we use to estimate the response-patterns until differences between the two approaches are no longer substantial. While the latter option is clearly preferable in principle, circumstances may dictate otherwise.

2.4 Reconstructing noise-reduced observations and signals

Under OLS, reconstructing the noise-reduced observations and signals scaled by best-fit scaling parameters is straightforward: we simply compute $\tilde{\mathbf{y}} = \mathbf{X}\tilde{\boldsymbol{\beta}}$. Under TLS, the problem is only slightly more complicated. Assuming, as before, equal noise in observations and signals, we project both onto the plane orthogonal to $\tilde{\mathbf{v}}$, thus:

$$\tilde{\mathbf{Z}} = \mathbf{Z} - \mathbf{Z}\tilde{\mathbf{v}}\tilde{\mathbf{v}}^T. \quad (39)$$

This provides a “best-fit” reconstruction of both observations and signals: note that the presence of noise in \mathbf{X} now means that the true (noise-free) model-predicted response-patterns must be estimated along with the noise-free observations.

Noise-reduced observations and model-simulated signals associated with the \mathbf{v} on a particular confidence surface are computed identically to $\tilde{\mathbf{Z}}$. This gives a set of possible observation-signal combinations that are consistent with the statistical model at a given confidence level. If we desire an uncertainty range on a single quantity, such as the trend attributable to a particular signal, then we require the \mathbf{v} corresponding to univariate confidence intervals (i.e. those on the surface where $\Delta\hat{\tau}^2(\mathbf{v}) = F_{1,\nu_2}$). For the joint distribution of two trends, we require bivariate intervals and so on. Note that reconstructions of noise-reduced observations and model-simulated signals will, in general, be much better behaved than the corresponding regression coefficients: if a particular \mathbf{v} corresponds to the model-simulated signal having zero amplitude while the observed signal has a finite amplitude, the corresponding β_i will be infinite, but the reconstructed noise-reduced observed and model-simulated signals will all be finite.

References

- Adcock, R. J. A problem in Least Squares. *The Analyst (Des Moines, Iowa)*, **5**, 53. 1878.
- Allen, M. R., & Tett, S. F. B. Checking internal consistency in optimal fingerprinting. *Climate Dynamics*, **15**, 419. 1999.
- Allen, M. R., Stott, P. A., Mitchell, J. F. B., Schnur, R., & Delworth, T. Uncertainty in forecasts of anthropogenic climate change. *Nature*. to appear. 2000.
- Barnett, T. P., & Schlesinger, M. E. Detecting changes in global climate induced by greenhouse gases. *J. Geophys. Res.*, **92**, 14772–14780. 1987.
- Bell, T. L. Theory of optimal weighting to detect climate change. *J. Atmos. Sci.*, **43**, 1694–1710. 1986.
- Deming, W. E. *Statistical Adjustment of Data*. New York: Wiley. 1943.
- Folland, C. K., Sexton, D., Karoly, D., Johnson, E., Parker, D. E., & Rowell, D. P. Influences of anthropogenic and oceanic forcing on recent climate change. *Geophys. Res. Lett.*, **25**, 353–356. 1998.
- Hasselmann, K. On the signal-to-noise problem in atmospheric response studies. *Pages 251–259 of: Shawn (ed), Meteorology of Tropical Oceans*. Royal Meteorological Society. 1979.
- Hasselmann, K. Optimal fingerprints for the detection of time dependent climate change. *J. Climate*, **6**, 1957–1971. 1993.
- Hasselmann, K. On multifingerprint detection and attribution of anthropogenic climate change. *Climate Dynamics*, **13**, 601–611. 1997.
- Hegerl, G., Hasselmann, K., Cubasch, U., Mitchell, J. F. B., Roeckner, E., Voss, R., & Waszkewitz, J. On multi-fingerprint detection and attribution of greenhouse gas and aerosol forced climate change. *Climate Dynamics*, **13**, 613–634. 1997.
- Hegerl, G. C., von Storch, H., Hasselmann, K., Santer, B. D., Cubasch, U., & Jones, P. D. Detecting greenhouse gas-induced climate change with an optimal fingerprint method. *J. Climate*, **9**, 2281–2306. 1996.
- Jeffreys, W. H. On the method of least squares: 1. *Astronomical Journal*, **85**, 177–181. 1980.
- Leroy, S. Detecting climate signals, some Bayesian aspects. *J. Climate*, **11**, 640–651. 1998.
- Lybanon, M. A Better Least-Squares Method When Both Variables Have Uncertainties. *Am. J. Physics*, **52**, 22–26. 1984.
- North, G. R., & Stevens, M. J. Detecting climate signals in the surface temperature record. *J. Climate*, **11**, 563–577. 1998.
- North, G. R., Bell, T. L., Cahalan, R. F., & Moeng, F. J. Sampling errors in the estimation of Empirical Orthogonal Functions. *Mon. Wea. Rev.*, **110**, 699–706. 1982.
- North, G. R., Kim, K. Y., Shen, S. S. P., & Hardin, J. W. Detection of forced climate signals, 1: Filter theory. *J. Climate*, **8**, 401–408. 1995.
- Palmer, T. N. A non-linear perspective on climate change. *J. Climate*, **12**, 575–591. 1999.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. *Numerical Recipes in FORTRAN: the Art of Scientific Computing*. 2 edn. Cambridge Univ. Press. 1992.

- Ripley, B. D., & Thompson, M. Regression techniques for the detection of analytical bias. *Analyst*, **112**, 377–383. 1987.
- Röckner, E., Bengtsson, L., Feichter, J., Lelieveld, J., & Rodhe, H. Transient climate change simulations with a coupled atmosphere-ocean GCM including the tropospheric sulfur cycle. *J. Climate*, **12**, 3004–3032. 1999.
- Santer, B. D., Wigley, T. M. L., & Jones, P. D. Correlation methods in fingerprint detection studies. *Climate Dynamics*, **8**, 265–276. 1993.
- Santer, B.D., Taylor, K.E., Wigley, T.M.L., Johns, T.C., Jones, P.D., Karoly, D.J., Mitchell, J.F.B., Oort, A.H., Penner, J.E., Ramaswamy, V., Schwarzkopf, M.D., Stouffer, R.J., & Tett, S. A Search for Human Influences on the Thermal Structure of the Atmosphere. *Nature*, **382**, 39–46. 1996.
- Stevens, M. J., & North, G. R. Detection of the Climate Response to the Solar Cycle. *J. Atmos. Sci.*, **53**, 2594–2608. 1996.
- Stott, P. A., Tett, S. F. B., Jones, G. S., Allen, M. R., Ingram, W. J., & Mitchell, J. F. B. Attribution of twentieth century climate change to natural and anthropogenic causes. *Climate Dynamics*. submitted. 2000a.
- Stott, P. A., Allen, M. R., & Jones, G. S. Estimating signal amplitudes in optimal fingerprinting: Application to general circulation models. *Climate Dynamics*, in preparation. 2000b.
- Tett, S. F. B., Mitchell, J. F. B., Parker, D. E., & Allen, M. R. Human influence on the atmospheric vertical temperature structure: Detection and observations. *Science*, **247**, 1170–1173. 1996.
- Tett, S. F. B., Stott, P. A., Allen, M. R., Ingram, W. J., & Mitchell, J. F. B. Causes of twentieth century temperature change near the Earth's surface. *Nature*, **399**, 569–572. 1999.
- van Huffel, S., & Vanderwaal, J. *The Total Least Squares Problem: Computational Aspects and Analysis*. SIAM. 1994.
- Wigley, T. M. L., Jones, P. D., & Raper, S. C. B. The observed global warming record: What does it tell us? *Proc. Nat. Acad. Sci.*, **94**, 8314–8320. 1997.

