

# Efficient estimation and ensemble generation in climate modelling

BY J.D. ANNAN AND J. C. HARGREAVES

*FRCGC/JAMSTEC*  
*3173-25 Showa-machi Kanazawa-ku*  
*Yokohama, Kanagawa 236-0001, Japan*

In this paper we review progress towards efficiently estimating parameters in climate models. Since the general problem is inherently intractable, a range of approximations and heuristic methods have been proposed. Simple Monte Carlo sampling methods, although easy to implement and very flexible, are rather inefficient, making implementation only possible in the very simplest models. More sophisticated methods based on random walks and gradient descent methods can provide more efficient solutions, but it is often unclear how to extract probabilistic information from such methods and the computational costs are still generally too high for their application to state of the art GCMs. The ensemble Kalman filter is an efficient Monte Carlo approximation which is optimal for linear problems, but we show here how its accuracy can degrade in nonlinear applications. Methods based on particle filtering may provide a solution to this problem but have yet to be studied in any detail in the realm of climate models. Statistical emulators show great promise for future research and their computational speed would eliminate much of the need for efficient sampling techniques. However, emulation of a full GCM has yet to be achieved and the construction of such represents a substantial computational task in itself.

**Keywords:** Probabilistic climate prediction, ensemble Kalman filter

## 1. Introduction

Model projections of anthropogenically forced climate change are sensitive to many aspects of model construction, in particular the values of parameters which are not well determined by direct observations or theoretical arguments but which directly impact on the future trajectory of model simulations. Although short-term numerical weather prediction is essentially an initial-value problem, the behaviour of a coupled atmosphere-ocean model in response to a forcing scenario on climatological (multidecadal and upward) time scales is determined much more by the details of its parameterisations rather than the initial state. Therefore, much attention has recently been given to the parameter estimation problem in this context [eg Forest et al., 2002, Knutti et al., 2002, Gregory et al., 2002, Hargreaves et al., 2004]. One initial aim may be to simply try to improve the realism of the model, in order to give the most accurate forecast possible. This model tuning is often seen as part of the model-building process, and has generally been performed in the rather ad-hoc and inefficient manner of “hand-tuning” by the model developers. More generally, we would also like to estimate the uncertainty surrounding this central forecast,

typically through an ensemble of model integrations. The methods initially used for estimation and ensemble generation were generally rather simple applications of Monte Carlo sampling (discussed in Section 2 below), and thus were computationally inefficient and could only be applied to the the cheapest models. Here we discuss progress in this field.

The computational problem amounts to an application of Bayes' Theorem: given prior belief  $f(x)$  over any uncertain parameters  $x$ , and observational data  $o$  which gives rise to a likelihood function  $f(o|x)$ , how can we efficiently estimate the posterior pdf  $f(x|o) = f(o|x)f(x)/f(o)$ ? In fact we are unlikely to be directly interested in the full posterior pdf, but will in general hope to calculate statistics of interest such as the expected value of some function  $g$  of the model output over the posterior pdf via  $E(g(x)) = \int (g(x)f(x|o))$ , or marginal and cumulative distributions thereof. Typical examples might be to estimate probabilistic percentiles for climate sensitivity, or calculate a probability for global temperature rising by more than a certain threshold under a specific emissions scenario.

Researchers have used a range of different observational data, including observations of recent and long-term climate change in conjunction with computationally cheap energy-balance earth system models [Andronova and Schlesinger, 2001, Forest et al., 2002, Knutti et al., 2002], or the use of the quasi steady-state climates in conjunction with more state of the art GCMs [Murphy et al., 2004, Annan et al., 2005a] to estimate equilibrium climate changes. In idealised situations where we can assume the model is perfect (that is, there are correct parameter values to estimate, for which the model behaviour perfectly mirrors the real world), the formulation of the problem is straightforward, and this is a natural environment in which to develop and test methods.

However, the perfect model assumption is rarely tenable in practice. No matter how well tuned, all climate models are readily distinguished from reality, with systematic biases in climate variables, and limited representation of climatic phenomena. The parameters may have no readily-identifiable counterparts in reality and this, combined with the need to compensate for missing processes and inevitable biases means that concept of a "correct" parameter value does not apply. Thus, the concept of the "best input" [Rougier, 2007] is introduced, which refers to the input set for which the model provides the "best" overall simulation of past and future climate. If the model was perfect, the correct parameter values would of course fulfill this role. Estimating this best input is complicated by the fact that the residual between the model and the data is not merely due to observational uncertainty, but also our uncertainty concerning model inadequacy (the discrepancy between the "best model" and the real climate). This uncertain discrepancy term adds to the observational uncertainty in broadening the posterior distribution. These issues are discussed more fully in Rougier [2007]. In the applications presented here, a very simple approach to model inadequacy has been adopted, based on the variance of the data fields. More detailed specification of this term would help to provide authoritative and credible predictions of climate change.

It is worth mentioning at the outset that there is no reasonable possibility of finding a general solution to the estimation problem that will work efficiently in all applications. The general parameter estimation problem can easily be seen to contain generic optimisation problems which are at least nondeterministic polynomial-time hard [Garey and Johnson, 1979]. For example, an efficient and general param-

eter estimation method would also be capable of solving quadratic programming problems, which is widely believed to be computationally intractable. Thus, we do not expect to find such a general solution, and our focus is on developing methods which aim to achieve an acceptable level of precision, by making reasonable approximations and taking advantage of any structure that exists in the climate prediction problem.

## 2. Numerical Integration and Monte Carlo sampling

In very low dimensional problems, the posterior pdf can be approximated by deterministic numerical integration, for example via Gaussian quadrature. This requires a multifactorial sampling of all adjustable parameters. Clearly, this type of approach is only feasible when the number of dimensions is rather low and/or the model is computationally cheap, since the number of integrations required increases exponentially with the number of dimensions. For example, Forest et al. [2000] used 61 integrations to sample two parameters at several levels in an energy balance model. In a more recent example, Stainforth et al. [2005] undertook a multifactorial sampling of 6 parameters of an AGCM at only 3 values each, and this experiment required massive computational resources. Using the same approach to sample even 10 parameters at 5 levels each would require resources some  $10^5$  times greater still.

Monte Carlo methods [Hammersley and Handscomb, 1964] provide an alternative randomised approach which has been widely used: instead of attempting to exhaustively integrate over the whole parameter space, we can simply draw a set of samples  $x_i$  from the prior (note that we are implicitly assuming that this itself is a straightforward operation), calculate their likelihood-based weights  $w_i \equiv f(o|x_i)$  and then approximate the required integrals in the obvious manner, eg

$$E(g(x)) = \int g(x)f(x|o) \simeq \sum (w_i g(x_i)) / \sum w_i$$

This approach is trivial to implement and very flexible, and has been widely used in climate science [eg Knutti et al., 2002, Schneider von Deimling et al., 2006]. A minor variant (rejection sampling) is to reject or accept each sample in a probabilistic manner based on the weight, which produces a smaller posterior ensemble (of accepted samples) for further calculation, and eliminates the need to keep track of a weight for each member. The main drawback of these approaches is that if the weights are low across a large proportion of the prior space, then it may take a large number of samples to populate the posterior and achieve reasonable convergence. In particular, the vast majority of samples may be given negligible weights if the prior is substantially more vague than the posterior, and this situation is likely to arise in high dimensional problems, especially when (as is common in climate science) the prior is deliberately chosen to be on the vague side so as to minimise the risk of double-counting of data. Latin hypercube sampling [McKay et al., 1979] is generally a rather more efficient approach, but only by a modest factor which is unlikely to completely overcome this disadvantage. For example, in some work which we describe below, a 1000 member Latin Hypercube experiment failed to find a single sample in the support of the posterior distribution defined by the problem (in other words, rejection sampling an ensemble of this size would have almost certainly have generated the empty set). Importance Sampling, in which samples are drawn not

from the prior but from some “proposal distribution” which is thought to more closely approximate the posterior, could potentially lead to further improvements, but it is not clear how such a proposal distribution could be found.

### 3. Random walk and gradient-descent methods

There are a wide variety of methods of varying generality, sophistication, and efficiency based on random walks or gradient descent in parameter space. Perhaps the simplest is the Monte Carlo Markov Chain [Metropolis et al., 1953, Hargreaves and Annan, 2002, Tarasov and Peltier, 2005]. This method undertakes a directed random walk in parameter space, in which trial steps are generated, and they are accepted or rejected probabilistically depending on how the likelihood of the new parameter values compares to that of the last-visited point. The distribution of points visited in parameter space converges to the posterior pdf. This method does not suffer from the curse of dimensionality, and is also highly flexible and simple to implement, requiring no special assumptions or approximations with the problem. However, since the method is inherently sequential and generally requires (at a minimum) thousands of integrations to converge, it again seems to be limited to computationally cheap models which can be integrated in a matter of seconds.

A range of more sophisticated heuristic approaches such as very fast simulated annealing [Jackson et al., 2004], genetic algorithms [Price et al., 2006] and Proximal-ACCPM [Beltran et al., 2005] have also been applied to optimisation and estimation with intermediate complexity climate models. These methods have significantly lower computational cost which allows application to at least intermediate complexity models (with integration times of around one hour). Perhaps the most efficient at present is the Proximal-ACCPM method, which can achieve good results with as few as 20 full model integrations, but it is unclear as to how probabilistic information can be extracted from this method which directly searches for the optimal (maximum likelihood) parameter set.

The use of an adjoint (which efficiently calculates the gradient of the cost function with respect to all of the control variables), can provide a highly efficient multivariate parameter estimation method in some cases [eg Hall et al., 1982, van Scheltinga and Dijkstra, 2005]. Moreover, it also generates a Jacobian of sensitivities which (at least in the near-linear case) can be used to provide probabilistic information. Although the construction of the adjoint code may be a rather complex task (even similar in magnitude to the creation of the model itself), automatic tools are available to help with this, and in some cases the adjoint may already exist if a version of the climate model is also used for numerical weather prediction. However, the application of this method to climatological estimation in complex models is problematic, due to the chaotic dynamics of the models [Lea et al., 2000, 2002]. In general, for a chaotic model integrated towards a climatological steady state, the gradient of the cost function is undefined everywhere, due to the fact that infinitesimal initial perturbations will grow exponentially over a long integration time. This remains the case even when the climatological response of the model to finite perturbations is close to linear. Therefore, various approximations and simplifications have to be developed for particular applications in order to generate usable results. These may include, for example, averaging over an ensemble of shorter integrations, or even direct modifications to the adjoint code through an artificially large

diffusion coefficient or the use of a low-resolution non-chaotic approximation to the full model. Furthermore, the Jacobian of sensitivities which an adjoint generates are only locally valid, so in the nonlinear case there is again no clear probabilistic interpretation of this output. No general solutions are known for these problems, and these obstacles appear to have inhibited research in this direction.

#### 4. Ensemble Kalman Filter

The ensemble Kalman filter (EnKF) was introduced by Evensen [1994] as an efficient Monte Carlo approximation to the Kalman filter [Kalman, 1960], which is the optimal filter for linear systems. The method is now in increasingly widespread use for state estimation problems in initial-value problems such as numerical weather prediction and ocean forecasting. A thorough description of the theory and basic methodology together with a survey of recent applications is provided in Evensen [2003]. Although the EnKF has generally been used for sequential initial state estimation, parameter estimation can readily be included in the same framework, by the means of state space augmentation [Derber, 1989, Anderson, 2001]. The principle here is that the parameters can be considered to be part of the model state alongside the conventional variables, and then the covariances sampled by the ensemble members can be used directly to update parameters in exactly the same manner as for the state variables.

The Kalman filter equations are optimal for linear systems, and this method does not suffer from the “curse of dimensionality”. Later, we investigate the performance in some simple non-linear cases. The method is trivially parallelisable, at least under the generally realistic assumption that the dominant cost is in the integration of  $O(100)$  models.

This method has recently been used to solve the parameter estimation problem in a range of models such as the Lorenz model [Annan and Hargreaves, 2004], a simple Earth system model [Hargreaves et al., 2004], a simplified primitive equation atmospheric GCM [Annan et al., 2005b] and even a modern state-of-the-art AGCM with slab ocean, albeit at low resolution [Annan et al., 2005a]. Figure 1 illustrates some results obtained with a simplified T21L7 atmospheric AGCM (a version of the IGCM of Forster et al. [2000]). The seasonally-averaged differences between the control run and the target of NCEP seasonal surface air temperature fields is shown on at the top. Below, the ensemble mean shows a marked improvement after 5 parameters were tuned. Precipitation fields were also used as a target in this experiment, but tuning was unable to significantly improve the model fields (not shown), thus demonstrating that there were substantial structural problems and motivating further model development.

Figure 2 illustrates an application of the EnKF to the simultaneous estimation of 12 parameters in an intermediate complexity EMIC in order to investigate the stability and predictability of the long-term response of its meridional overturning circulation to various emissions scenarios [Hargreaves and Annan, 2005]. Despite sampling a wide range of parameter values, the model climatologies were all of a reasonably high quality in comparison to the observations used, and the root mean square error of the ensemble mean was also very similar to the results obtained by other optimisation methods which have been applied to this model [Beltran et al., 2005, Price et al., 2006]. Minor differences in model version and problem

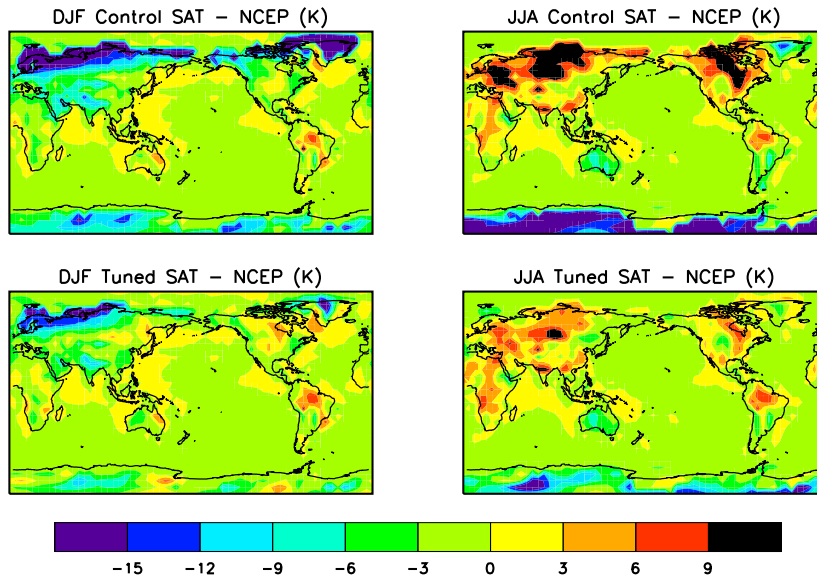


Figure 1. Results obtained by tuning a low-resolution AGCM to NCEP data, allowing 5 parameters to vary simultaneously. Upper plots show errors in the control run, lower plots show greatly reduced errors in posterior ensemble mean. Values are plotted in degrees Kelvin. From Annan et al. [2005b].

specification (such as uniform versus Gaussian parameter priors) prevent a direct comparison, but clearly all methods were working well when applied to this problem. In contrast, a 1000 member Latin Hypercube experiment [Edwards and Marsh, 2005] with the same model failed to find a single solution comparable to any of these methods, illustrating the limitations of more naive methods. That is, it failed to produce any samples from the posterior defined by our problem specification.

The EnKF approach is only formally correct in the linear case. Of course, the linearity here refers not to the model itself (which in general will exhibit chaotic dynamics) but to the climate of the model, as a function of the control variables. Despite this theoretical limitation, it appears to produce surprisingly good solutions to problems where the exact solution is known to be strongly non-Gaussian [eg Zhou et al., 2006].

It is hard to validate the method in complex applications where the correct solution is not known, but we can readily investigate the effect of nonlinearity in idealised cases. For example, Annan et al. [2005b] used a simple single-variable problem with a low degree of nonlinearity to demonstrate the benefit of the iterative procedure we have developed. Here we extend this idea to investigate the degradation of the technique with increasing nonlinearity. Our toy model here takes a single input parameter  $x$  for which we have a prior estimate  $x = 10 \pm 3$  (the indicated uncertainties are all one standard deviation, Gaussian, and assumed independent) and generates the output  $y$  via the equation

$$y = 10 \times \left(\frac{x}{10}\right)^k \quad (4.1)$$






Figure 2. Ensemble response of meridional overturning under a GHG forcing scenario. Each line represents one of the 54 ensemble members. The probability of a long-term shut-down of the overturning depends strongly on several factors which are poorly understood at present. From Hargreaves and Annan [2005].

where  $k$  will vary between experiments.

We have a single observation  $y_o = 10 \pm 3$  with which we aim to improve our estimate of  $x$ . For all  $k$ , the posterior maximum likelihood value is given by  $(x, y) = (10, 10)$ , but for  $k > 1$ , the model is nonlinear and the posterior is skewed. The results, shown in Figure 3, are the average of 10 experiments with 10000 ensemble members each, so as to focus solely on the influence of nonlinearity. For the case  $k = 1$ , the problem is linear and trivial, and the EnKF generates an accurate solution. As  $k$  increases, the posterior pdf becomes more strongly skewed and the accuracy of the EnKF solution (which is more symmetrical in shape) decreases. For the case  $k = 16$ , the bias in the pdf is comparable in magnitude (and of course additional to) to the  $2\sigma$  uncertainty that would also arise due to sampling error with use of a typical practical ensemble size of 100 members, as indicated by the standard formula for the sampling error of an event of probability  $p$ :  $\sigma^2 = p(1 - p)/N$ . As was noted by Zhou et al. [2006], the EnKF generally gives fairly good estimates for the mean even for nonlinear problems, with the accuracy of higher order moments being more noticeably degraded.

We can evaluate the nonlinearity in the posterior automatically, by comparing the analysis output to the model output at the subsequent iteration. The EnKF automatically generates an ensemble of “analysis” values  $(x_i^a, y_i^a)$  from linear combinations of the previous iteration’s “forecast”  $(x_i^f, y_i^f)$ . If the model is linear, the  $(x_i^a, y_i^a)$  pairs should be in perfect balance. Therefore, we examine the RMS error function




Figure 3. Illustrating the influence of nonlinearity on the EnKF performance. The top left plot shows the function  $y = 10 \times \left(\frac{x}{10}\right)^k$  for three values of  $k$ . The top right plot shows the probability plot of EnKF output versus correct solution (the red line overlies the exact  $x = y$  line, shown in black). The lower plot shows the correct posterior pdfs (solid lines) with EnKF estimates (dashed lines). Colour coding indicates  $k$  (as shown in the first plot). Vertical lines in first and last plots indicate the 5% and 95% thresholds for the posterior pdfs.

$$R = \sqrt{\frac{\sum_i (y_i^a(t) - y_i^f(t+1))^2}{n}}$$

where  $y_i^a(t)$  is the analysed output for  $y$  from the  $t$ th iteration, and  $y_i^f(t+1)$  is the forecast of  $y$  for the subsequent iteration. For the linear case  $k = 1$ ,  $R$  is essentially the machine precision. For  $k = 4$ ,  $R$  is about 0.30, and for  $k = 16$  it grows to around 0.45. The bias of the analysis versus forecast is negligible in all cases. These results suggest that it might be possible to diagnose the presence of nonlinearity and determine to what extent the results from particular experiments can be considered reliable. However, practical considerations may limit the integration time to the



point that the model's internal variability and lack of true convergence contaminates the calculation.

For strongly bimodal pdfs, such as would be generated by the function  $y = x^2$  with an observation of  $y_o = 10$  and a prior that assigns significant probability to both high probability regions  $x = \pm\sqrt{10}$ , the EnKF method fails to converge, with the main mass of the ensemble instead oscillating between the two modes of the posterior. Such behaviour is easy to demonstrate in artificial situations, but has not yet been observed in any of the practical experiments we have undertaken with realistic models.

## 5. Particle filters

Particle filtering is a sequential Monte Carlo method in which the pdf is, as with the EnKF, sampled by an ensemble of models (particles). Unlike the EnKF, the particles are not directly repositioned in light of observational data, but instead likelihood-based weights are assigned to them. To compensate for the inevitable concentration of weights on very few members, regular resampling is generally required. For linear problems, the unequal weights result in the method being generally less accurate than an EnKF of similar size, but it has the advantage of being able in principle to handle nonlinear problems with no loss of precision. Zhou et al. [2006] illustrate the superior accuracy of a particle filter method over the EnKF for a highly nonlinear problem when the ensemble size is not a limiting factor. Particle filtering methods are not ideally suited for parameter estimation, since there is no direct means by which replicated particles (generated in the resampling step) become separated. However, there are techniques such as jittered resampling Liu and West [2001] which address this problem. The hybrid kernel resampling method of Anderson and Anderson [1999] offers another related avenue by which nonlinearity can be handled within the particle/Kalman filter framework, without the need for such a large ensemble as the particle filter typically requires. Such methods have not been widely tested.

## 6. Emulators

A statistical emulator is a computationally efficient approximation to a complex computer code. A simple approach based on a linear interpolation was presented by Murphy et al. [2004], and a more sophisticated version is described by Rougier and Sexton in this issue. Due to the computational efficiency of the emulator, it is possible to evaluate its pdf with high precision, even using simple Monte Carlo sampling. This approach appears to have great potential for more accurate evaluation of probabilities, although the construction of the emulator is itself a complex task. Success in this research would have strong implications for the need to find efficient optimisation techniques, but the ability to emulate the full multivariate output (ie fields of climatological variables) of a complex GCM, which would be needed for calibration and prediction of spatially-resolved climate change, has not yet been achieved.

## 7. Conclusions

Most methods of uncertainty analysis in common use in climate science rely on large numbers of model evaluations sampled directly from prior distributions on the input parameters, and thus can only be applied to low-dimensional problems and simplified models. Even within this setting, significant improvements in efficiency can often be achieved by careful attention to sampling strategy such as the use of a Latin Hypercube.

More sophisticated estimation and data assimilation methods based on gradient descent and other heuristics attempt to sample more effectively from the posterior (defined by some measure of fit to observational data). The methods are generally flexible, optimal (in the sense of correctly solving an arbitrary problem, given adequate computational resources) and simple to implement, but are still usually too expensive for application to state of the art models. The ensemble Kalman filter is highly parallelisable and requires relatively few model integrations, but it is only approximate in the nonlinear case, and the validity and accuracy of its approximation in real applications remains rather uncertain. However, it has proven very effective in at least finding well-calibrated solutions across a wide range of applications, and the more general particle filter or variants such as kernel resampling may have the potential to address its main limitations. The Bayesian emulator approach demonstrated by Sexton and Rougier (this issue) appears to be a highly promising alternative for the future.

## References

- J. L. Anderson. An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, 129:2884–2902, 2001.
- J. L. Anderson and S. L. Anderson. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127:2741–2758, 1999.
- N. G. Andronova and M. E. Schlesinger. Objective estimation of the probability density function for climate sensitivity. *Journal of Geophysical Research*, 108 (D8):22,605–22,611, 2001.
- J. D. Annan and J. C. Hargreaves. Efficient parameter estimation for a highly chaotic system. *Tellus*, 56A:520–526, 2004.
- J. D. Annan, J. C. Hargreaves, R. Ohgaito, A. Abe-Ouchi, and S. Emori. Efficiently constraining climate sensitivity with paleoclimate simulations. *SOLA*, 1:181–184, 2005a.
- J. D. Annan, D. J. Lunt, J. C. Hargreaves, and P. J. Valdes. Parameter estimation in an atmospheric GCM. *Nonlinear processes in geophysics*, 12:363–371, 2005b.
- C. Beltran, N. R. Edwards, A. Haurie, J.-P. Vial, and D. S. Zachary. Oracle-based optimization applied to climate model calibration. *Environmental Modelling and Assessment*, 11(1):31–43, 2005.

- J. Derber. A variational continuous assimilation scheme. *Monthly Weather Review*, 117:2437–2446, 1989.
- N. R. Edwards and R. Marsh. Uncertainties due to transport-parameter sensitivity in an efficient 3-D ocean-climate model. *Climate Dynamics*, 2005. In press.
- G. Evensen. The ensemble kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53:343–367, 2003.
- G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):10143–10162, 1994.
- C. E. Forest, M. R. Allen, P. H. Stone, and A. P. Sokolov. Constraining uncertainties in climate models using climate change detection techniques. *Geophysical Research Letters*, 27:569–572, 2000.
- C. E. Forest, P. H. Stone, A. P. Sokolov, M. R. Allen, and M. D. Webster. Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science*, 295(5552):113–117, 2002.
- P. M. Forster, M. Blackburn, R. Glover, and K. P. Shine. An examination of climate sensitivity for idealised climate change experiments in an intermediate general circulation model. *Climate Dynamics*, 16:833–849, 2000.
- M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, 1979.
- J. M. Gregory, R. J. Stouffer, S. C. B. Raper, P. A. Stott, and N. A. Rayner. An observationally based estimate of the climate sensitivity. *Journal of Climate*, 15(22):3117–3121, 2002.
- M. C. G. Hall, D. G. Cacuci, and M. E. Schlesinger. Sensitivity analysis of a radiative-convective model by the adjoint method. *Journal of the Atmospheric Sciences*, 39:2038–2050, 1982.
- J. M. Hammersley and D. C. Handscomb. *Monte Carlo methods*. Methuen & Co Ltd, London, 1964.
- J. C. Hargreaves and J. D. Annan. Assimilation of paleo-data in a simple Earth system model. *Climate Dynamics*, 19(5–6):371–381, 2002.
- J. C. Hargreaves and J. D. Annan. Using ensemble prediction methods to examine regional climate variation under global warming scenarios. *Ocean Modelling*, 11(1–2):174–192, 2005.
- J. C. Hargreaves, J. D. Annan, N. R. Edwards, and R. Marsh. Climate forecasting using an intermediate complexity Earth System Model and the ensemble Kalman filter. *Climate Dynamics*, 23(7–8):745–760, 2004.
- C. Jackspon, M. K. Sen, and P. L. Stoffa. An efficient stochastic Bayesian approach to optimal parameter and uncertainty estimation for climate model predictions. *Journal of Climate*, 17:2828–2841, 2004.

- R. E. Kalman. A new approach to linear filtering and prediction problems. *J. Basic Engineering*, 82D:33–45, 1960.
- R. Knutti, T. F. Stocker, F. Joos, and G.-K. Plattner. Constraints on radiative forcing and future climate change from observations and climate model ensembles. *Nature*, 416:719–723, 2002.
- D. J. Lea, M. R. Allen, and T. W. N. Haine. Sensitivity analysis of the climate of a chaotic system. *Tellus*, 52A:523–532, 2000.
- D. J. Lea, T. W. N. Haine, M. R. Allen, and J. A. Hansen. Sensitivity analysis of the climate of a chaotic ocean circulation model. *Q. J. R. Meteorological Soc.*, 128(586):2587–2605, 2002.
- J. Liu and M. West. Combined parameter and state estimation in simulation-based filtering. In A. Doucet et al, editor, *Sequential Monte Carlo Methods in Practice*, Springer, 2001.
- M.D. McKay, R.J. Beckman, and W.J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.
- N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *J. Chem. Phys.*, 21: 1087–1091, 1953.
- J. M. Murphy, D. M. H. Sexton, D. N. Barnett, G. S. Jones, M. J. Webb, M. Collins, and D. A. Stainforth. Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature*, 430:768–772, 2004.
- A. R. Price, I. I. Voutchkov, G. E. Pound, N. R. Edwards, T. M. Lenton, S. J. Cox, and the GENIE team. Multiobjective tuning of grid-enabled earth system models using a non-dominated sorting genetic algorithm (nsga-ii). In *2nd International Conference on eScience and Grid Computing*, 2006. (Submitted).
- J. C. Rougier. Probabilistic inference for future climate using an ensemble of simulator evaluations. *Climatic Change*, 2007. doi: 10.1007/s10584-006-9156-9.
- T. Schneider von Deimling, H. Held, A. Ganopolski, and S. Rahmstorf. Climate sensitivity estimated from ensemble simulations of glacial climate. *Climate Dynamics*, 27(2–3):149–163, 2006.
- D. A. Stainforth, T. Aina, C. Christensen, M. Collins, N. Faull, D. J. Frame, J. A. Kettleborough, S. Knight, A. Martin, J. M. Murphy, C. Piani, D. Sexton, L. A. Smith, R. A. Spicer, A. J. Thorpe, and M. R. Allen. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature*, 433:403–406, 2005.
- L. Tarasov and W. R. Peltier. Arctic freshwater forcing of the Younger-Dryas cold reversal. *Nature*, 435:662–665, 2005.
- A. D. Terwisscha van Scheltinga and H. A. Dijkstra. Nonlinear data-assimilation using implicit models. *Nonlinear Processes in Geophysics*, 12:515–525, 2005.

Y. Zhou, D. McLaughlin, and D. Entekhabi. Assessing the performance of the ensemble Kalman filter for land surface data assimilation. *Monthly Weather Review*, 134:2128–2142, 2006.