28*th* IUGG Conference on Mathematical Geophysics, June 7-11, 2010, Pisa, Italy
Session 8: Data assimilation and model validation

Conditioning of the variational data assimilation problem

Stephen A. Haben[1], Amos S. Lawless[1], Nancy K. Nichols[1]

[1] *University of Reading, Reading, UK*

Keywords: *condition number; variational data assimilation*

Numerical weather prediction (NWP) centres use numerical models of the atmospheric flow to forecast the future weather states from an estimate of the current state of the atmosphere. Data assimilation methods combine observations with the forecast model in order to produce a 'best' estimate of the current state, called the analysis. Variational data assimilation (VAR) is popularly used in operational NWP and involves minimising a weighted non-linear least-squares objective function which measures the error between the model forecast and the available observations. The solution of the minimisation is the analysis and is found using an iterative optimisation algorithm. In practice an incremental version of VAR is implemented in many operational centres. This method solves a sequence of linear approximations to the nonlinear least-squares problem. Each approximate linearised least-squares problem is solved using an 'inner' gradient iteration method, such as the conjugate gradient method, and the linearization state is then updated in an 'outer' iteration loop. The rate of convergence of the inner loop of the VAR iteration scheme and the sensitivity of the analysis to perturbations are proportional to the condition number, that is, the ratio of the largest to the smallest eigenvalue of the Hessian of the linear least-squares objective function. The Hessian is generally assumed to be ill-conditioned and hence operationally the system is preconditioned by transforming the state variables to new variables where the errors are assumed to be approximately uncorrelated. Experimental comparisons have demonstrated that the preconditioning improves the speed of the assimilation scheme. In this work we examine the conditioning of the variational assimilation method theoretically. We derive bounds on the condition number of the Hessian in the case of a single, periodic, spatially-distributed system parameter. We show that the conditioning is sensitive to the length-scale in the correlation structures and confirm that the condition number is generally improved by preconditioning. The theoretical bounds are used to understand how the density and accuracy of the observations affect the conditioning of the preconditioned system. Finally, we examine the conditioning of the Met Office operational VAR method experimentally. We show, using both pseudo and real observations, that our experimental results using this system are consistent with our theoretical results.