Journal of Hydrology 533 (2016) 114–127

Contents lists available at ScienceDirect

Journal of Hydrology

journal homepage: www.elsevier.com/locate/jhydrol

Gradient-based model calibration with proxy-model assistance

Wesley Burrows [a,b,*], John Doherty [c,b]

[a] Flinders University, GPO Box 2100, Adelaide, SA 5001, Australia
[b] National Centre for Groundwater Research and Training, Adelaide, SA 5001, Australia
[c] Watermark Numerical Computing, 336 Cliveden Avenue, Corinda, QLD 4075, Australia

ARTICLE INFO

Article history:
Received 21 July 2015
Received in revised form 20 November 2015
Accepted 24 November 2015
Available online 28 November 2015
This manuscript was handled by Peter K. Kitanidis, Editor-in-Chief, with the assistance of Niklas Linde, Associate Editor

Keywords:
Proxy model
Calibration
Uncertainty analysis

SUMMARY

Use of a proxy model in gradient-based calibration and uncertainty analysis of a complex groundwater model with large run times and problematic numerical behaviour is described. The methodology is general, and can be used with models of all types. The proxy model is based on a series of analytical functions that link all model outputs used in the calibration process to all parameters requiring estimation. In enforcing history-matching constraints during the calibration and post-calibration uncertainty analysis processes, the proxy model is run for the purposes of populating the Jacobian matrix, while the original model is run when testing parameter upgrades; the latter process is readily parallelized. Use of a proxy model in this fashion dramatically reduces the computational burden of complex model calibration and uncertainty analysis. At the same time, the effect of model numerical misbehaviour on calculation of local gradients is mitigated, this allowing access to the benefits of gradient-based analysis where lack of integrity in finite-difference derivatives calculation would otherwise have impeded such access. Construction of a proxy model, and its subsequent use in calibration of a complex model, and in analysing the uncertainties of predictions made by that model, is implemented in the PEST suite.

1. Introduction

Environmental models that simulate the details of complex physical and chemical processes over domains wherein the properties which govern those processes are spatially and temporally heterogeneous are often characterized by long run-times and a propensity for problematic solver convergence. Furthermore, it is not uncommon for models of these types to exhibit good numerical behaviour when provided with one set of parameters, but suffer serious degradation of numerical performance when supplied with another set of parameters. Where this occurs, calibration and uncertainty analysis become very difficult undertakings. This can erode the use of such models in environmental decision-support.

In the present paper we focus on those aspects of a model's performance which compromise the ability of a model-independent inversion package such as PEST (Doherty, 2015a) to calculate derivatives of model outputs with respect to the parameters which require adjustment during calibration, and calibration-constrained uncertainty analysis. In PEST, derivatives are calculated using a finite-difference methodology based on a two, three or five point stencil. Model outputs are computed based on values of a particular parameter which are varied incrementally in accordance with the selected stencil; differences in these outputs form the basis for approximation of local partial derivatives with respect to that parameter. These derivatives are housed in a so-called Jacobian matrix. The Jacobian matrix is then employed in calculation of an improved set of parameters. Jacobian matrix and parameter upgrade calculations are undertaken repeatedly in an iterative process whose outcome is a set of parameter values that produce an acceptable level of fit between model outcomes and field observations of system state. Where a model is being calibrated, a set of parameters which constitute a minimum error variance solution to the inverse problem is sought through this process. Where calibration-constrained uncertainty analysis is being undertaken, multiple sets of parameters are sought, all of which are considered to be reasonable expressions of system properties, and all of which fit field measurements to within limits that reflect the noise content of those measurements.

A variety of numerical methods have been developed to expedite calibration and calibration-constrained uncertainty analysis. Many of these methods do not, in fact, require calculation of a Jacobian matrix. However, use of so-called "gradient methods" which do make use of partial derivatives of model outputs with respect to adjustable parameters to perform the above tasks accrues certain benefits. A major benefit that gradient methods have over other methods is their speed; see, for example Keating

* Corresponding author at: Flinders University, GPO Box 2100, Adelaide, SA 5001, Australia.
E-mail address: wesley.burrows@bigpond.com (W. Burrows).

http://dx.doi.org/10.1016/j.jhydrol.2015.11.033
0022-1694/

et al. (2010). Another benefit is that gradient-based inversion algorithms are easily extended to include mathematical regularisation schemes that readily accommodate parameter nonuniqueness (Aster et al., 2013; Menke, 1989). A further benefit is that, once a Jacobian matrix has been filled, it can be used in calculation of post-calibration statistics such as parameter identifiability (Doherty and Hunt, 2009), parameter and predictive uncertainty (Gallagher and Doherty, 2007a, 2007b; James et al., 2009), and the worth of existing and yet-to-be acquired data in terms of its ability to reduce the uncertainties of parameter and predictions of interest (Dausman et al., 2010).

Use of gradient methods is not without its problems, however. Their performance may be hampered where the relationship between model outputs and parameters is highly non-linear (Duan et al., 1992). Even worse, it may not be possible to use these methods at all where model performance is such that elements of the Jacobian matrix lose their integrity. Where these elements are calculated using finite parameter differences, loss of integrity can occur when incremental changes in model outputs employed in finite-difference derivatives calculation reflect more than simply incremental changes in parameter values. This is not an uncommon situation, particularly where the complex nonlinear environmental processes simulated by a model challenges its solver. While strategies such as adaptive time stepping that alter the solution procedure when convergence becomes problematic, may mitigate these problems as far as the model is concerned, they may exacerbate them as far as calculation of finite-difference derivatives is concerned, for model outputs may then become somewhat dependent on solution path.

Examples of modelling contexts in which calculation of finite-difference parameter derivatives may be compromised are not hard to find. The handling of "dry cells" in MODFLOW (Harbaugh et al., 2000) is a common example. Kavetski et al. (2006) discuss how algorithmic design of models that simulate surface water movement can lead to similar problems in these kinds of models. Other contexts in which model numerical behaviour can compromise finite-difference derived gradients include:

- simulation of the effects of mining and tunnelling operations on groundwater systems;
- interaction of ground and surface waters near streams and wetlands;
- high temperature geothermal reservoir simulation where water phase is a discontinuous function of temperature and pressure; and
- chemical reactions in mobile contaminant plumes.

In the difficult numerical circumstances that these modelling contexts present, so-called "global methods" which do not rely on calculation of derivatives of model outputs with respect to adjustable parameters, provide an alternative option for software-controlled history-matching. Examples of non-gradient based calibration and calibration-constrained uncertainty analysis algorithms include (among many others) particle swarm optimization (Kennedy and Mendes, 2002), shuffled complex evolution (Duan et al., 1992), genetic and evolutionary programming (Vrugt and Robinson, 2007), and covariance matrix adaption algorithms (Hansen et al., 2003). All of these replace the need to calculate derivatives with respect to adjustable parameters with intelligent random sampling of parameter values. While delivering robustness in the face of problematical numerical behaviour, the cost of this robustness is the requirement for a greater number of model runs than that required by gradient methods. This differential between model run requirements of the two different approaches tends to grow with the number of parameters that require estimation or adjustment.

To ease the computational burden of applying global methods to the problems of model calibration and calibration-constrained uncertainty analysis, increasing use is being made of fast-running model surrogates. In recognition of the fact that the complex simulator is the most accurate replicator of reality available, in many applications the surrogate model does not completely replace the original simulator. Rather it is strategically substituted for the simulator on many occasions that a model run is required. The greater is the ratio of surrogate to simulator runs, the greater is the efficiency of the overall process. The surrogate may be a simulator that runs much faster than that which it replaces because of its simpler algorithmic design. For example, the SWI package for MODFLOW (Bakker et al., 2013), which replaces mass conservative governing equations with equations based on continuity of flow, avoiding the need for fine-scale vertical discretisation. Another example is MODFLOW-USG, described by Panday et al. (2013), which is able to represent flow in grids with highly irregular spatial discretisation thereby reducing the number of simultaneous equations required in solution. More sophisticated model reduction strategies may be employed as are used by Efendiev et al. (2005, 2009) and Mondal et al. (2010) whereby a coarse-gridded simulator whose parameterization is based on single-phase upscaling procedures, surrogates for a fine scale, dual-phase reservoir model.

Alternatively the surrogate may undertake data-driven reproduction of simulator outputs, or interpolate between samples of simulator outputs to non-sampled parts of parameter space using devices such as radial basis functions, kriging or artificial neural networks; see for example, Regis and Shoemaker (2004), Bliznyuk et al. (2007), and Alam et al. (2004) respectively. More recently, Laloy et al. (2013) and Elsheikh et al. (2014) deploy polynomial chaos expansion theory to develop interpolators of simulator outputs. Also the statistical characteristics of simulator outputs can be modelled using Gaussian process theory; see Johnson et al. (2011) and Conti et al. (2009) as examples. Data-driven surrogates such as these are commonly known as model emulators or proxy models.

One analysis scheme that can benefit enormously from simulator run reductions through strategic use of surrogate models is Markov Chain Monte Carlo (MCMC). Various adaptations of the so-called "two-stage MCMC" approach have been documented (see Efendiev et al., 2005, 2009; Mondal et al., 2010; Cui et al., 2011 for examples) that seek to reduce unnecessary expensive simulator runs in assessment of low-probability proposal parameter fields. In these example studies calibration and/or uncertainty assessment of complex reservoir simulators is undertaken wherein a surrogate model is used in "stage one" of the process as a pre-screening mechanism. The goal is to increase the acceptance rate of proposed parameter fields in "stage two" where acceptance/rejection of the proposal is determined on the basis of the simulator. The studies just mentioned use surrogate models based on simplified algorithms, as has already been mentioned. The studies also cited earlier in relation to polynomial chaos expansion theory (that is Laloy et al., 2013; Elsheikh et al., 2014), also deployed their model emulators within the two-stage MCMC framework. Two-stage MCMC consistently demonstrates several fold savings in computational costs over full/direct MCMC, effected primarily through inexpensive pre-screening of proposals.

Of course use of a surrogate model, either as a direct substitute or as a companion to a more accurate simulator, will undoubtedly incur some cost on the analysis undertaken. Put simply, a simplified model cannot be expected to replicate the same level of accuracy at all spatial and temporal locations of a modelled domain as can a simulator. It is readily acknowledged in the literature of two-stage MCMC cited above, that there exists potential for rejection of parameter proposals in the pre-screening stage of the process when assessed by the surrogate, that would otherwise find support