Journal of Hydrology 560 (2018) 127–140

Contents lists available at ScienceDirect

Journal of Hydrology

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

Research papers

Constraining the ensemble Kalman filter for improved streamflow forecasting

Deborah H. Maxwell *, Bethanna M. Jackson, James McGregor

School of Geography, Environment and Earth Sciences, Victoria University of Wellington, PO Box 600, Wellington, New Zealand

ARTICLE INFO

Article history:
Received 24 August 2017
Received in revised form 2 March 2018
Accepted 5 March 2018
Available online 8 March 2018
This manuscript was handled by Marco borga, Editor-in-Chief, with the assistance of Venkat Lakshmi, Associate Editor

Keywords:
Constrained Ensemble Kalman Filter
Data assimilation
Observation error
Hydrology

ABSTRACT

Data assimilation techniques such as the Ensemble Kalman Filter (EnKF) are often applied to hydrological models with minimal state volume/capacity constraints enforced during ensemble generation. Flux constraints are rarely, if ever, applied. Consequently, model states can be adjusted beyond physically reasonable limits, compromising the integrity of model output. In this paper, we investigate the effect of constraining the EnKF on forecast performance. A "free run" in which no assimilation is applied is compared to a completely unconstrained EnKF implementation, a 'typical' hydrological implementation (in which mass constraints are enforced to ensure non-negativity and capacity thresholds of model states are not exceeded), and then to a more tightly constrained implementation where flux as well as mass constraints are imposed to force the rate of water movement to/from ensemble states to be within physically consistent boundaries. A three year period (2008–2010) was selected from the available data record (1976–2010). This was specifically chosen as it had no significant data gaps and represented well the range of flows observed in the longer dataset. Over this period, the standard implementation of the EnKF (no constraints) contained eight hydrological events where (multiple) physically inconsistent state adjustments were made. All were selected for analysis. Mass constraints alone did little to improve forecast performance; in fact, several were significantly degraded compared to the free run. In contrast, the combined use of mass and flux constraints significantly improved forecast performance in six events relative to all other implementations, while the remaining two events showed no significant difference in performance. Placing flux as well as mass constraints on the data assimilation framework encourages physically consistent state estimation and results in more accurate and reliable forward predictions of streamflow for robust decision-making. We also experiment with the observation error, which has a profound effect on filter performance. We note an interesting tension exists between specifying an error which reflects known uncertainties and errors in the measurement versus an error that allows "optimal" filter updating.

1. Introduction and scope

The use of data assimilation through state estimation is increasingly recognised as an essential part of any hydrological forecasting application (Liu et al., 2012). Without it, relatively large errors can accumulate in the model output, even over the short term, i.e. hours to days (Collischonn et al., 2007). State estimation methods use observations to adjust model states by taking account of errors in both the observation(s) and the model (Clark et al., 2008; Reichle, 2008; Salamon and Feyen, 2009; Vrugt et al., 2005), with the aim to reduce errors associated with the data and model structure and improve the physical realism of the model output.

One of the most common state estimation methods used in hydrology is the Ensemble Kalman Filter (EnKF). The EnKF utilises Monte Carlo methods to generate an ensemble of model trajectories, consistent with the main sources of uncertainty in the given problem. Although some calibration or fine-tuning of relevant error specification is necessary for optimal filter performance (Noh et al., 2014), and most applications would ensure some physical consistency (e.g. non-negativity of sub-surface states), application of the "standard" EnKF does not necessarily require the ensemble of model states to be within physically realistic limits. Consequently, states can be perturbed to unrealistic and inconsistent values to obtain a closer match to the observations, which can lead to erratic streamflow simulations (Clark et al., 2008). Unrealistic perturbations can be particularly severe when modelled streamflow is significantly different to observed streamflow (although a high observation error term will limit the severity of

* Corresponding author.
E-mail addresses: Deborah.maxwell@vuw.ac.nz (D.H. Maxwell), Bethanna.jackson@vuw.ac.nz (B.M. Jackson), Jim.mcgregor@vuw.ac.nz (J. McGregor).

https://doi.org/10.1016/j.jhydrol.2018.03.015
0022-1694/

fluctuations, as we note later). Weerts and El Serafy (2006) contend that error specification should be given due consideration. While error settings which allow large perturbations can increase the possibility of successful forecasting, we argue that it is not realistic to allow unjustified perturbations which can subsequently compromise the accuracy and reliability of forward predictions.

To address this issue, constraints can be applied to ensure physical limits are not exceeded. Constraints have been widely applied in oceanic and atmospheric applications (Janjic et al., 2014; Simon and Simon, 2006; Simon and Tien Li, 2002; Thacker, 2007). In hydrologic modelling, states generally represent some component of water storage, for example groundwater and soil water storage; these should be non-negative and, where possible, have some maximum capacity constraint. Constraints can be imposed in a number of ways. Wang et al. (2009) compared the naive, projection and accept/reject methods for incorporating constraints into the EnKF. Non-negativity constraints were applied to five state variables (three quick flow states, one baseflow state and soil water content) and a capacity constraint to ensure the soil state does not exceed maximum soil water content. In dual state-parameter estimation approaches (Moradkhani et al., 2005; Shi et al., 2014; Wang et al., 2009), predefined ranges are used to bound the ensemble generation to ensure non-negative volumes. Pan and Wood (2006) and Li et al. (2012) add mass conservation constraints to ensure closure of the water balance for land surface modelling. Checking for physical realism or 'consistent model behaviour' is often recommended (Clark et al., 2008; Xie et al., 2014), although little specific guidance is provided on how to avoid model behaviour violating physical laws.

Few studies constrain the perturbations themselves. The danger of forecasting streamflow using unconstrained flux perturbations as the only innovations to states has been demonstrated (Lee et al., 2011; Samuel et al., 2014; Seo et al., 2009). However, rather than searching for ways to further constrain the state estimation procedure, these studies focussed on showing how assimilating soil moisture as well as streamflow could improve analysis and forecast performance. It is often the case, however, that soil moisture (and other state information) is not readily available, so streamflow may be the only obtainable variable that can be closely related to model states. Moradkhani et al. (2005) and Abaza et al. (2014) apply proportionality factors to ensure the generated ensemble spread is within a meaningful range. Clark et al. (2008), and subsequently McMillan et al. (2013) and Xie et al. (2014), impose fractional factors to allow larger model errors when model fluxes are large and smaller model errors when fluxes are small. While their approach does not account for periods where the model inadequately simulates large model fluxes or where these fluxes are missed (Clark et al., 2008), it does address the temporal variability of the error which is more representative of the nature of model errors and leads to improved forecast performance.

A large source of uncertainty lies in the forcing data (e.g. precipitation and evapotranspiration) due to the difficulty in fully representing the spatial and temporal variability of these inputs in the model. The ability to provide accurate estimates of forcing uncertainty would help to understand and identify other sources of error and uncertainty in the model (McMillan et al., 2011), and can lead to more reliable and successful data assimilation as shown in Noh et al. (2014), Rakovec et al. (2012) and Weerts and El Serafy (2006), amongst others. A number of studies have specifically investigated the effect of perturbing forcing, with few studies specifically considering constraining model states in the assimilation process. While the focus of this paper is on the latter, an important next step is to combine methodologies for perturbing forcing with constraints on state innovations, examining both the question of whether the combination improves model performance, and the question of whether inconsistencies in predictive outputs and state innovations can point to ways to improve specification of both input error definition and state constraints.

In summary, although the importance of plausible state updating has been acknowledged in the literature, the implications of inadequate constraint of the EnKF on the reliability of forward streamflow predictions have not been demonstrated in any detail. In addition, although the sensitivity of the EnKF to the specified measurement error has been noted, strategies to overcome this have focussed on bringing in supplementary data, rather than exploring methods to allow more robust state updating where streamflow observations remain the only obtainable measure. While the use of basic non-negativity and/or capacity constraints is not uncommon, this is the first study to specifically investigate the effect of also imposing flux constraints (controlling the rate of water movement to or from ensemble model states) on the accuracy and reliability of forecast streamflow. We compare three approaches: 1) a "naïve" EnKF approach in which no constraints are applied to model output; 2) an arguably more typical approach in hydrology, where mass constraints are enforced to ensure non-negativity and capacity thresholds (where possible) of model states are not exceeded; and 3) our new, further constrained approach where flux constraints are imposed in conjunction with the more typical mass constraints so that state perturbations as well as state extremes are kept within physically realistic ranges. The three approaches are demonstrated using a lumped conceptual hydrological model (Maxwell, 2013) in the Tauranga-Taupo catchment, New Zealand. Streamflow observations are used to update four (soil storage, throughflow, interflow and baseflow) model states. It is shown that mass constraints alone are not sufficient to significantly improve forecast performance in the majority of events analysed. The combination of mass and flux constraints results in more reliable forecasting of streamflow compared to the standard and mass constrained EnKF implementations. The implications for the accuracy and reliability of model predictions are demonstrated.

We also experiment with the observation error term in the calculation of the Kalman Gain. Often paid little attention in the literature, this term has a profound effect on filter performance. As further discussed in Crow and Van Loon (2006) and Reichle (2008), very high confidence in the observations relative to the model draws the filter rapidly toward the observations by allowing larger perturbations to model states. On the other hand, specifying a large measurement error will give greater weight to the model, preventing significant perturbations and drawing the filter toward an outcome similar to the model output from the free run. Despite these previous authors having noted this tension, most studies simply note this term should reflect the accuracy of the measurement (instrument, processing and representativeness) (Evensen, 2003; McMillan et al., 2013; Xie and Zhang, 2010), with little consideration given to its subsequent impact on filter performance. We found interesting conflicts between a "realistic" specification of the error and its impact on the filter's ability to constrain state adjustments.

The hydrological model, data used, catchment area, calibration procedure, implementation of the constrained EnKF and experimental set-up is described in Section 2. The results are discussed in Section 3 and include a comparative analysis of model performance over the eight events between 2008 and 2010. We conclude with a summary of the main findings and suggest areas for future work.

2. Material and methods

2.1. Study area and data

The Tauranga-Taupo catchment (Fig. 1) covers 197 km^2 and drains an area of impermeable sedimentary (basement greywacke)