BARCAST and the Kalman Smoother

Martin Tingley (tingley@fas.harvard.edu)

January 12, 2010

1 Introduction

There are a number of similarities between the Kalman smoother [4], and BARCAST [7]. Both methods seek to infer the values of a state vector through time and an associated uncertainty estimate, given noisy, transformed, and/or incomplete observations, and both methods assume that the process and the observational error structures are Gaussian. Under certain simplifying assumptions, the estimates of the field and the associated uncertainty from BARCAST are equivalent to those from the Kalman smoother. If the state vector is augmented to include both the values of the climate field, and the scalar parameters inferred by BARCAST, then generalizations of the Kalman filter that apply to non-linear systems, such as the Extended Kalman Filter (EKF) or Unscented Kalman FIlter (UKF), could be used to infer the unknown quantities. These modifications of the Kalman filter necessarily involve approximations.

The Bayesian approach has a number of advantages over the various Kalman smoothing approaches. In particular, BARCAST is not based on an approximation, gives simultaneous estimates of the field and the scalar parameters that define the model, can easily be generalized to include proxy observations assumed to integrate the climate field, and results in not only point estimates of the field, with associated uncertainty estimates, at each year, but also an ensemble of realizations of the space-time field consistent with the data and model assumptions.

2 The Kalman filter and smoother

2.1 The standard filtering and smoothing algorithms

This development follows that of [10]. The Kalman filter and smoother are designed to infer the realized values of a Markov chain from noisy, incomplete, and/or transformed measurements of the state vector.

Assume that the Markov chain evolves according to,

$$\vec{X}_t = \mathbf{M}_t \vec{x}_{t-1} + \vec{\eta}_t, \quad \eta_t \sim N(\vec{0}, \mathbf{Q}_t), \tag{1}$$

where $\vec{X}_t$ indicates the state variable at time t, and $\vec{x}_{t-1}$ indicates the realized value of the process at time $t-1$. The values of the linear propagators $\mathbf{M}_t$ are assumed known at each time point, as are the values of the noise covariance matrices $\mathbf{Q}_t$.

The observations $\vec{W}_t$ of the system at each time point are modeled as,

$$\vec{W}_t = \mathbf{H}_t \vec{x}_t + \vec{\epsilon}_t, \quad \epsilon_t \sim N(\vec{0}, \mathbf{R}_t), \tag{2}$$

where the observation operator $\mathbf{H}_t$ and measurement error covariance matrix $\mathbf{R}_t$ are assumed known at each time point.

Define the conditional expectation of the process at time t, given observations up to $t-1$ and t, respectively, as $\vec{x}_{t|t} \equiv E\left[X_t | w_{1:t}\right]$, and $\vec{x}_{t|t-1} \equiv E\left[X_t | w_{1:t-1}\right]$. Define the corresponding error covariances matrices as

$$\mathbf{P}_{t|t} = E\left[\left(\vec{X}_t - \vec{x}_{t|t}\right)\left(\vec{X}_t - \vec{x}_{t|t}\right)' | w_{1:t}\right], \tag{3}$$

$$\mathbf{P}_{t|t-1} = E\left[\left(\vec{X}_t - \vec{x}_{t|t-1}\right)\left(\vec{X}_t - \vec{x}_{t|t-1}\right)' | w_{1:t-1}\right] \tag{4}$$

The Kalman *filter* provides a set of recursive formulas to infer first the values of $\vec{x}_{t|t-1}$ and $\mathbf{P}_{t|t-1}$ in term of $\vec{x}_{t-1|t-1}$, $\mathbf{P}_{t-1|t-1}$, $\mathbf{M}_t$, and $\mathbf{Q}_t$, and then the values of $\vec{x}_{t|t}$ and $\mathbf{P}_{t|t}$ in terms of $\vec{x}_{t|t-1}$, $\mathbf{P}_{t|t-1}$, y_t, $\mathbf{R}_t$, $\mathbf{H}_t$. Given an initial condition $\vec{x}_{0|0}$ and $\mathbf{P}_{0|0}$, these formulas can be iterated to obtain sequential estimates of the underlying state and the associated uncertainty. Conceptually, the Kalman filter involves propagating the first two moments through time, and can be thought of as a predict-then-correct algorithm: $\vec{x}_{t|t-1}$ and $\mathbf{P}_{t|t-1}$ predict the first two moments given the previous information, and then $\vec{x}_{t|t}$ and $\mathbf{P}_{t|t}$ correct these predictions using the new observation.

The Kalman *smoother* seeks estimates of $\vec{x}_{t|T}$ and $\mathbf{P}_{t|T}$, for $t = 0 \ldots T$. In other words, the smoother seeks estimates of the state vector and associated uncertainty at each time t, conditional on all observation, both before and after t. The smoother provides a set of recursive formulas to infer $\vec{x}_{t|T}$ and $\mathbf{P}_{t|T}$ from $\vec{x}_{t+1|T}$, $\mathbf{P}_{t+1|T}$, the results of the filter, and the sequence of $\mathbf{M}_t$.

Schematically, the Kalman filter and smoother infer the values of the hidden state for the model illustrated in Figure 1, provided that the time propagation of the hidden variable is specified, along with the observation operator and error covariance matrix at each time step.

2.2 The Kalman filter and non-linear systems

Many systems do not conform to the basic assumptions of the the Kalman fitler/smoother that both the observation and propagation equations (Eqs. 1 & 2) are linear. Such systems can be written, using

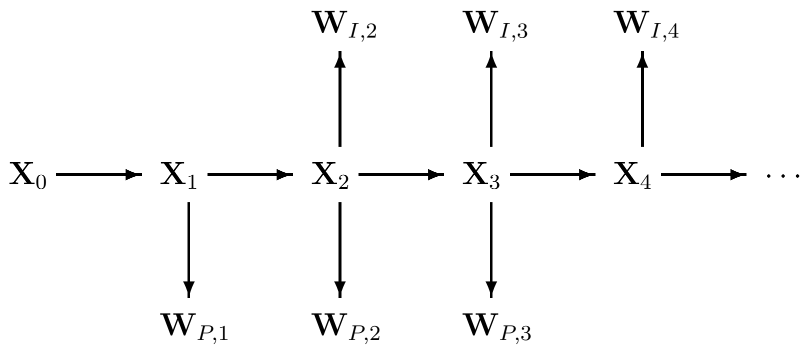

Figure 1: Schematic of a simple model used to assimilate different climate field data types. In this case, both the proxy ($\mathbf{W}_P$) and instrumental ($\mathbf{W}_I$) observations at a given year depend only on the corresponding field values ($\mathbf{X}$) for that year.

the notation above, as,

$$\vec{X}_t = f\left(\vec{x}_{t-1}\right) + \vec{\eta}_t, \tag{5}$$

$$\vec{W}_t = g\left(\vec{x}_t\right) + \vec{\epsilon}_t, \tag{6}$$

where f and g are arbitrary functions. As the transformations are non linear, the process is no longer Gaussian, despite the assumption that all noise terms are Gaussian. Several variants of the Kalman filter have been developed to solve the filtering and smoothing problems for non-linear systems.

- The **Extended Kalman Filter (EKF)** is based on linearizing the functions f and g at each time step, and then employing the basic machinery of the Kalman filter (see, for example, [9]). As the EKF is based on linearizations, the results are necessarily approximate. If the process is strongly nonlinear on the time scale of the spacing of the observations, the performance of the EKF can be highly unstable [2, 3].

- The **Unscented Kalman Filter (UKF)** is "founded on the intuition that it is easier to approximate a Gaussian distribution than it is to approximate an arbitrary non-linear function" [2]. The UKF involves propagating a deterministically chosen set of points (i.e. not a random sample) that follow a $N\left(\vec{x}_{t|t}, \mathbf{P}_{t|t}\right)$ distribution through the non linear propagator, and then calculating the sample mean and covaraince of these transformed points [2, 3]. This approach offers a number of advantages over the EKF [2, 3], but like the EKF is based on an approximation.

3 BARCAST and the Kalman Smoother

Under certain simplifying assumptions, the basic implementation of BARCAST, designed to infer the unknown field values for a scenario like that in Figure 1, gives results that are in some sense equivalent to the Kalman smoother. In particular, running BARCAST with all scalar parameters fixed at their correct values results in point estimates of the field, and associated uncertainty estimates, equivalent to the estimates $\vec{x}_{t|T}$ and $\mathbf{P}_{t|T}$ from the Kalman smoother. The way in which the two methods arrive at these estimates is different: the Kalman smoother is a recursive algorithm that provides estimates of the mean and uncertainty covariance at each time point, whereas BARCAST produces an ensemble of draws of the space-time field, consistent with the model assumptions and the data, from which percentiles can be calculated.

Both BARCAST and the Kalman smoother involve prior information about the underlying field — BARCAST requires a prior distribution on the field values for the year before the first observations (assumed to be normal), whereas $\vec{x}_{0|0}$ and $\mathbf{P}_{0|0}$ are specified for the Kalman smoother. If the prior information is in disagreement with the data (i.e. if the field values at time $t = 0$ are many prior standard deviations from the prior mean), the field estimates from both the Kalman smoother and BARCAST are inaccurate for the first number of time steps.

In practice, the various scalar parameters estimated by BARCAST could be estimated beforehand (or in the case of the instrumental observational error variance, taken from the literature [1]) and then the Kalman smoother used to estimate the missing field values. However, accurately estimating these parameters is non trivial. For example, there is a spatial misalignment issue in determining the relationship between the proxies and the instrumental records, as the proxy records are not in general co-located with an instrumental record.

If the Kalman smoother is used with incorrect parameter values, then the point estimates of the field, as well as the associated uncertainties, can be biased. This is true in particular of the range parameter in the spatial covariance matrix, which specifies that correlation decays exponentially with separation [7]. Even if the parameters are correctly specified, the uncertainties from the Kalman smoother are in some sense overly narrow, as they do not consider the uncertainties in the estimates of the scalar parameters. The uncertainty estimates from BARCAST, in contrast, integrate out the uncertainty in the scalar parameters so in some sense are more trustworthy. In the context of the climate field reconstruction problem, there is great interest in obtaining accurate uncertainty estimates.

A method that simultaneously estimates both the field values and the scalar parameters is therefore preferred to one that first estimates the scalar parameters, and then inputs point estimates of these parameters into a second stage of the analysis.

The EKF or UKF provide one approach to the simultaneous estimation of the parameters and the field values, as the state vector in the climate field reconstruction problem can be augmented to include the scalar parameters inferred by BARCAST. It should then be possible to set up a non linear filtering/smoothing problem (as in Eqns. 5&6) that can then be solved using the EKF or UKF.

However, BARCAST is still preferred to this approach, for a number of reasons:

- The mathematical development required by the EKF or UKF quickly become far less tractable than that of BARCAST. In his seminal paper [4], R.E. Kalman proposed the filter which now bears his name as a new and improved solution to the Weinar problem. In reference to existing solutions, he wrote, "The mathematics of the derivations are not transparent. Fundamental assumptions and their consequences tend to be obscured." The assumptions made by BARCAST are simpler than those required by the variants of the Kalman fitler, and the impacts these assumptions have on the results simpler to determine in the case of BARCAST.

- BARCAST provides a cohesive framework for estimating both the field and the scalar parameters that define the model, that can be expanded to deal with data types which have more complicated relationships with the underlying field. For example, a tree ring width measurement at a given year likely reflects temperatures integrated over several years, so the analysis model is like that illustrated in Figure 2. These assumptions can easily be incorporated into BARCAST, as the conditional posterior distribution of the field at a given year will then depend on tree ring width measurements over several years. In contrast, incorporating such assumptions into a Kalman smoother requires (I believe) augmenting the state vector to include the field values from the number of years reflected in the tree ring width measurements. In other words, incorporating low frequency proxies is easier in the BARCAST framework.

- There is a general interest in quantities smoothed through time, both time series at individual locations and time series of spatial means. BARCAST allows for a very simple estimation of the associated uncertainty; Kalman methods do not. This is true in general — BARCAST can provide both a point estimate and an associated uncertainty for any function of the space-time field, simply by applying the function to each member of the ensemble of posterior draws, and then calculating the percentiles of the resulting distribution.

- BARCAST can easily weigh in on something like Michael Mann's claim that 1998 was likely the warmest year in the last millennium [5], by determining the percentage of ensemble members for which each year was the warmest. Kalman methods do not provide a simple solution to such a problem.

- There is a general utility in having an ensemble of realizations of the climate variable consistent with the data and assumptions. Provided the model assumptions are correct, the ensemble members have, on average, the same temporal variance as the true field values [8]. In contrast, the point estimates from BARCAST, the Kalman smoother, and RegEM [6, 8], which minimize the mean squared error, have reduced temporal variance, with the reduction a function of the amount and quality of available data. In general, the ensemble of posterior draws produced by

BARCAST allows for a richer, more interesting analysis of the data, and is likely useful in ways not considered here.

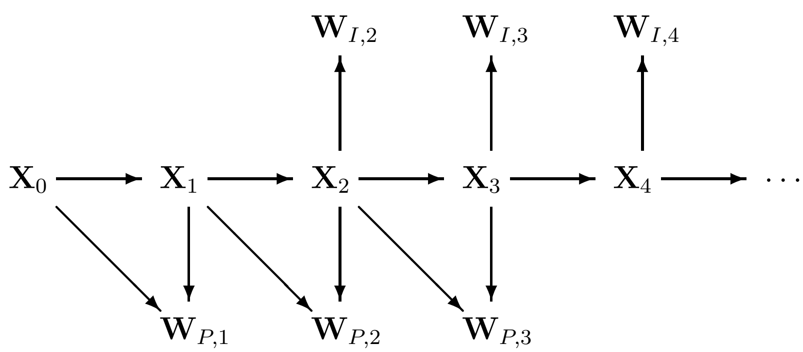

Figure 2: Schematic of a more complex model used to assimilate different climate field data types. In this case, the proxy observations ($\mathbf{W}_P$) at a given year depend on the corresponding field values ($\mathbf{X}$) for that year and the previous year.

4 Summary

There is a general correspondence between Kalman filtering/smoothing and recursive Bayesian estimation, such as that performed by BARCAST. The Kalman smoother gives, for each time point, a best estimate of the field, and the covariance structure of the associated uncertainty, conditioned on all available data, the model, the specified scalar parameters, and the initial values of the state vector and uncertainty. BARCAST, in contrast, produces an ensemble of draws from the posterior distribution of the space time field and scalar quantities, conditioned on the data, model assumptions, and priors. As discussed above, this ensemble allows for a richer analysis than the two moment analysis of the Kalman approach. Any data analysis performed with a standard Kalman filter can equivalently be carried out via recursive Bayesian estimation, but many analyses that are theoretically very straightforward using recursive Bayesian estimation can only be performed using variants of the Kalman filter/smoother, such as the EKF or UKF, which rely on additional assumptions or are based on simplifying approximations. In the case of the climate field reconstruction problem, BARCAST results in simultaneous estimation of both the field values through time and the scalar parameters, is based on a simple set of assumptions, and does not make recourse to approximations.

References

[1] P. Brohan, J. J. Kennedy, I. Harris, S. F. B. Tett, and P. D. Jones. Uncertainty estimates in regional and global observed temperature changes: A new data set from 1850. *Journal of Geophysical Research*, 2:99–113, 2006.

[2] S.J. Julier and J.K. Uhlmann. A new extension of the Kalman filter to nonlinear systems. In *Int. Symp. Aerospace/Defense Sensing, Simul. and Controls*, volume 3, 1997.

[3] S.J. Julier, J.K. Uhlmann, and H.F. Durrant-Whyte. A new approach for filtering nonlinear systems. In *American Control Conference, 1995. Proceedings of the*, volume 3, 1995.

[4] R.E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

[5] M.E. Mann, R.S. Bradley, and M.K. Hughes. Northern hemisphere temperatures during the past millennium: inferences, uncertainties, and limitations. *Geophysical Research Letters*, 26(6):759–762, 1999.

[6] T. Schneider. Analysis of Incomplete Climate Data: Estimation of Mean Values and Covariance Matrices and Imputation of Missing Values. *Journal of Climate*, 14(5):853–871, 2001.

[7] M.P. Tingley and P. Huybers. A Bayesian Algorithm for Reconstructing Climate Anomalies in Space and Time. Part 1: Development and applications to paleoclimate reconstruction problems. *Journal of Climate*, 2010.

[8] M.P. Tingley and P. Huybers. A Bayesian Algorithm for Reconstructing Climate Anomalies in Space and Time. Part 2: Comparison with the Regularized Expectation-Maximization Algorithm. *Journal of Climate*, 2010.

[9] G. Welch and G. Bishop. An introduction to the Kalman filter. *University of North Carolina at Chapel Hill, Chapel Hill, NC*, 1995.

[10] Christoper K. Wikle and L. Mark Berliner. A Bayesian tutorial for data assimilation. *Physica D*, 230:1–16, 2006.