

A Dynamic Modelling Strategy for Bayesian Computer Model Emulation

Fei Liu
Univeristy of Missouri

Mike West
Duke University

Abstract. Computer model evaluation studies build statistical models of deterministic simulation-based predictions of field data to then assess and criticize the computer model and suggest refinements. Computer models are often expensive computationally: statistical models that adequately emulate their key features can be very much more efficient. Gaussian process models are often used as emulators, but the resulting computations lack the ability to scale to higher-dimensional, time-dependent or functional outputs. For some such problems, especially for contexts of time series outputs, building emulators using dynamic linear models provides a computationally attractive alternative as well as a flexible modelling approach capable of emulating a broad range of stochastic structures underlying the input-output simulations. We describe this here, combining Bayesian multivariate dynamic linear models with Gaussian process modelling in an effective manner, and illustrate the approach with data from a hydrological simulation model. The general strategy will be useful for other computer model evaluation studies with time series or functional outputs.

Keywords: Computer model emulation; Dynamic linear model; Forwarding filtering, backward sampling; Gaussian process; Markov chain Monte Carlo; Time-Varying Autoregression.

1 Introduction

1.1 Background and Motivation

Computationally intensive simulation models, or *simulators*, are increasingly used in research on complex systems in science and technology (Kennedy and O'Hagan 2001). Computer model evaluation studies emulate simulator-based input-output relationships using flexible statistical models. Often applied to deterministic simulators, statistical emulators are typically far less computationally expensive than the model simulators of interest. This enables their use as proxies for more extensive study of input-output relationships, as well as to assess, criticize and suggest statistical refinements to the simulator based on analyses of simulator predictions of field data (e.g. Higdon et al. 2004a,b; Goldstein and Rougier 2003; Fuentes et al. 2003; Goldstein and Rougier 2005; Rougier 2008b). Recent applications range from modelling traffic flows (Paulo et al. 2005) to analyzing climate sensitivity (Rougier et al. 2008) with a range of variants of statistical methods (Santner et al. 2003; Bayarri et al. 2005, 2007a,b).

Following Sacks et al. (1989), statistical emulators of computer model simulators are commonly based on flexible Gaussian process (GASP) methodology (e.g. Currin

et al. 1991; Welch et al. 1992; Morris et al. 1993; Craig et al. 2001; Kennedy et al. 2002; Rougier 2008a). GASP is justifiably popular in view of its provision of a broad, flexible nonparametric model framework capable of adapting to complex patterns of input-output relationships. However, GASP methodology does not scale easily and is simply computationally infeasible in problems with high-dimensional, time-dependent or functional outputs (Higdon et al. 2008a). This is particularly challenging in contexts with time series outputs, such as in our motivating application contexts and the example in Section 3 of the *logSPM* hydrological model (Kuczera et al. 2006).

With this interest in extending GASP modelling to problems with long time series outputs, we have combined Bayesian multivariate dynamic linear models (DLMs) with Gaussian processes. Specifically, we create multivariate time-varying autoregressive models – a special class of DLMs – in which the stochastic innovations are Gaussian processes over computer model input space. Methodologically, this provides flexibility in representing complex patterns of stochastic structure in outputs, relevant in our application as well as in other areas (e.g. Sanso et al. 2008; Tebaldi and Sanso 2008; Sanso et al. 2007). Computationally, the approach is tractable and scalable as posterior inference via Monte Carlo Markov Chain method efficiently scales based on the forward filtering, backward sampling algorithm (Carter and Kohn 1994; Frühwirth-Schnatter 1994); see also West and Harrison (1997, section 15.2.3).

Section 2 outlines GASP modelling and then describes the novel DLM-GASP synthesis and resulting new class of statistical emulators, with details of MCMC computations for model fitting. Section 3 describes an example of the *logSPM* hydrological simulator and problem context, and exemplifies the approach in an analysis of the *logSPM* data. Summary comments appear in Section 4 and supporting technical material in an appendix.

2 DLM-GASP Emulators

2.1 Gaussian Process Models

Begin with a computer model that runs with d -dimensional inputs $\mathbf{z} = (z_1, \dots, z_d)'$ and generates univariate outputs $y(\mathbf{z})$. GASP uses a Gaussian process prior,

$$y(\cdot) \sim \text{GP}(\mu, c(\cdot, \cdot)/\lambda), \quad (1)$$

where μ , λ and $c(\cdot, \cdot)$ are the mean, precision and correlation function respectively. The latter often takes the specific, separable form

$$c(\mathbf{z}, \mathbf{z}^*) = \exp\left(-\sum_{i=1}^d \beta_i |z_i - z_i^*|^{\alpha_i}\right) \quad (2)$$

for any two inputs \mathbf{z}, \mathbf{z}^* ; the *range* parameters β_i control decay of correlation with respect to distance in individual input dimensions, while the *roughness* parameters α_i control smoothness of the output response.

Write $\mathbf{y}_{1:n} = (y(\mathbf{z}_1), \dots, y(\mathbf{z}_n))'$ for the vector of outputs of the computer model runs at inputs $\mathbf{z}_{1:n} = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$. At any query input \mathbf{z} , due to the Gaussian process assumption, the joint distribution of $(\mathbf{y}_{1:n}, y(\mathbf{z}))$ is a multivariate normal distribution. Therefore, we can obtain the posterior predictive distribution for $y(\mathbf{z})$ as $(y(\mathbf{z}) \mid \mathbf{z}_{1:n}) \sim \mathcal{N}(\mu_{n+1}(\mathbf{z}), v_{n+1}(\mathbf{z}))$ where

$$\mu_{n+1}(\mathbf{z}) = \mu + \boldsymbol{\gamma}'_{n+1} \boldsymbol{\Gamma}_n^{-1} (\mathbf{y}_{1:n} - \mu) \quad \text{and} \quad v_{n+1}(\mathbf{z}) = (1 - \boldsymbol{\gamma}'_{n+1} \boldsymbol{\Gamma}_n^{-1} \boldsymbol{\gamma}_{n+1}) / \lambda$$

where $\boldsymbol{\gamma}'_{n+1} = (c(\mathbf{z}, \mathbf{z}_1), \dots, c(\mathbf{z}, \mathbf{z}_n))$ and $\boldsymbol{\Gamma}_n = (c(\mathbf{z}_i, \mathbf{z}_j))_{i,j=1,\dots,n}$, giving the GASP prediction for the output at any query input.

Dealing with outputs that are time series extends the above to series of $y(\mathbf{z})$ values over a time axis, for each input \mathbf{z} . This could, of course, be done with spatial or other axes too, but we use the time series anchor context here. If outputs are time series, we could simply extend the above to include time t as an input so that the computer model outputs are $y(\mathbf{z}, t)$. Viewing t as an extension of the \mathbf{z} input vector, this leads to the direct extension of the GASP model above with an additional term for the time dimension. However, this would massively increase the input dimension and computations under the extended GP prior become infeasible with other than trivial length of the time series. Our new DLM-GASP synthesis provides an alternative, tractable approach.

2.2 A DLM-GASP Statistical Model

By way of notation, define the following:

- the training data inputs $\mathbf{z}_{1:n} = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$, where \mathbf{z}_i is a d -dimensional input vector, $\mathbf{z}_i = (z_{i1}, \dots, z_{id})'$;
- at each time t , the outputs in the n -vector $\mathbf{y}_t = (y(\mathbf{z}_1, t), \dots, y(\mathbf{z}_n, t))'$;
- at any input \mathbf{z} , the output time series values up to any chosen time t is the t -vector $\mathbf{y}_{1:t}(\mathbf{z}) = (y(\mathbf{z}, 1), \dots, y(\mathbf{z}, t))'$;
- all outputs up to time t are denoted by $\mathbf{y}_{1:t}(\mathbf{z}_{1:n}) = \{\mathbf{y}_1, \dots, \mathbf{y}_t\}$.

Consider the n -vector time series \mathbf{y}_t , $t = 1, \dots, T$ over a period up to time T . A flexible model for what may be quite intricate inter-dependencies among the elements of \mathbf{y}_t over time, including non-stationary, time-varying stochastic structure, is provided by the class of *time-varying autoregressive models* (TVAR), a special subclass of multivariate DLMS (Prado and West 1997; West and Harrison 1997; Aguilar et al. 1999; Prado et al. 2001). To reflect the input space context, we now extend TVAR models to be driven by stochastic innovations that are related over input space via a Gaussian process model. The new DLM-GASP model is as follows. For each \mathbf{z} , we have a TVAR model

$$y(\mathbf{z}, t) = \sum_{j=1}^p \phi_{t,j} y(\mathbf{z}, t-j) + \epsilon_t(\mathbf{z}), \quad (3)$$

where the lag p is specified, the autoregressive parameters $\boldsymbol{\phi}_t = (\phi_{t,1}, \dots, \phi_{t,p})'$ may vary over time, and $\epsilon_t(\mathbf{z}) \sim N(0, v_t)$ independently over time. Across series, we link via a GASP model having

$$\text{Cov}(\epsilon_t(\mathbf{z}), \epsilon_t(\mathbf{z}^*)) = v_t c(\mathbf{z}, \mathbf{z}^*),$$

with $c(\mathbf{z}, \mathbf{z}^*)$ from equation (2), assuming the same correlation functions $c(\mathbf{z}, \mathbf{z}^*)$ for all t while allow the variances v_t to change. With $\boldsymbol{\epsilon}_t = (\epsilon_t(\mathbf{z}_1), \dots, \epsilon_t(\mathbf{z}_n))'$, we can write these models over all inputs as

$$\mathbf{y}_t = \mathbf{F}'_t \boldsymbol{\phi}_t + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim N(0, v_t \boldsymbol{\Sigma}) \quad (4)$$

where

$$\mathbf{F}'_t = \begin{pmatrix} y(\mathbf{z}_1, t-1) & \dots & y(\mathbf{z}_1, t-p) \\ \vdots & \ddots & \vdots \\ y(\mathbf{z}_n, t-1) & \dots & y(\mathbf{z}_n, t-p) \end{pmatrix}$$

and where $\boldsymbol{\Sigma}$ is the n by n correlation matrix with (i, j) element $c(\mathbf{z}_i, \mathbf{z}_j)$.

To reflect the fact that the computer model may describe some non-stationary process, we allow the auto-regressive coefficients, $\boldsymbol{\phi}_t = (\phi_{t,1}, \dots, \phi_{t,p})'$, to vary over time. In particular, we assume the AR coefficients follow a random walk through time,

$$\boldsymbol{\phi}_t = \boldsymbol{\phi}_{t-1} + \mathbf{w}_t, \quad \mathbf{w}_t \sim N(0, \mathbf{W}_t), \quad (5)$$

where $\boldsymbol{\phi}_0$, v_t and \mathbf{W}_t are parameters with subsidiary models and prior distributions to be specified. In our analysis, we fix roughness parameters at $\alpha_i = 2$ throughout, and treat the range parameters $\beta_{1:d}$ as uncertain and also to be estimated.

The model equations (4) and (5) define a multivariate DLM, and we then have access to the full theory of dynamic models (West and Harrison 1997) including the forward filtering analysis and the related backward sampling algorithm that are critical to efficient computation. Some key technical details are given in Appendix 2 here.

2.3 DLM-GASP as an Emulator of GASP

Before proceeding to analysis, we note as an aside that the DLM-GASP model can also be viewed as a flexible, non-stationary approximation to an alternative GASP model that incorporates time into the input variable. In the latter framework, suppose the correlation function to have the separable form $c((\mathbf{z}, t), (\mathbf{z}^*, t^*)) = c(\mathbf{z}, \mathbf{z}^*) c_{time}(|t - t^*|)$ where $c_{time}(\cdot)$ is an additional correlation function on the time dimension. Over times $t = 1, \dots, T$, write $\boldsymbol{\Sigma}_{time}$ for the $T \times T$ matrix with (t, s) element $c_{time}(|t - s|)$. Then the GASP model on the extended (\mathbf{z}, t) inputs has

$$\text{vec}(\mathbf{y}_{1:T}(\mathbf{z}_{1:n})) \sim N_{nT}(\boldsymbol{\mu} \times \mathbf{1}_{nT}, \boldsymbol{\Sigma} \otimes \boldsymbol{\Sigma}_{time}/\lambda). \quad (6)$$

Assuming that the correlation $c_{time}(\cdot)$ becomes negligible after p lags, then

$$p(\mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1 \mid \boldsymbol{\theta}) \approx p(\mathbf{y}_t \mid \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-p+1}, \boldsymbol{\theta})$$

for $t > p$. Using the properties of the Kronecker product listed in Appendix 1, we then obtain the approximations

$$\mathbb{E}(\mathbf{y}_t \mid \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-p}, \boldsymbol{\theta}) = \mathbf{F}'_t \boldsymbol{\phi} \quad \text{with} \quad \boldsymbol{\phi} = \tilde{\boldsymbol{\Sigma}}_p^{-1} \boldsymbol{\rho}_{t,t-1:t-p}, \quad (7)$$

and

$$\mathbb{V}(\mathbf{y}_t \mid \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-p}, \boldsymbol{\theta}) = v \boldsymbol{\Sigma} \quad \text{with} \quad v\lambda = 1 - \boldsymbol{\rho}'_{t,t-1:t-p} \tilde{\boldsymbol{\Sigma}}_p^{-1} \boldsymbol{\rho}_{t,t-1:t-p}, \quad (8)$$

where $\boldsymbol{\rho}_{t,r:s} = (c_{time}(|t-r|), \dots, c_{time}(|t-s|))'$ and $\tilde{\boldsymbol{\Sigma}}_p$ is a $p \times p$ matrix with (i, j) element $c_{time}(|i-j|)$. This is a multivariate normal DLM with constant AR coefficients and innovations variance. Hence we see the TVAR model can be interpreted as approximating this structure but allowing more flexibility through time-varying AR coefficients and innovations variance. Alternatively, we can just view the DLM-GASP as a richer model specification in its own right, allowing that flexibility and generating richer possibilities for correlation structures in the time dimension.

2.4 Model Completion and Prior Specification

Model completion requires specifications for the quantities $\boldsymbol{\phi}_0$, $\{v_t, \mathbf{W}_t, t = 1, \dots, T\}$ and $\beta_{1:d} = (\{\beta_i, i = 1, \dots, d\})$. This follows standard specifications for all DLM parameters with a conjugate normal/inverse-gamma prior on $(\boldsymbol{\phi}_0, v_0 | D_0)$ and the use of variance discounting to define models for the v_t sequence and specific values for the \mathbf{W}_t sequence; See West and Harrison (1997) and the Appendix 2 here.

A key component of the analysis is learning on the range parameters $\beta_{1:d}$ for which we use the half-range correlation idea (Higdon et al. 2008b). Suppose the design inputs have been standardized. Then, for each $i = 1, \dots, d$, the half-range $\rho_i = \exp(-\beta_i/4)$ is the correlation between $\epsilon_t(\mathbf{z})$ and $\epsilon_t(\mathbf{z}^*)$ when input \mathbf{z} and \mathbf{z}^* differ only in the i^{th} component with a difference of 0.5. Following that approach, we use independent $\text{Ga}(1, 0.1)$ priors on ρ_i , placing substantial mass near $\rho_i = 1$ to reflect the expectation that only a subset of the inputs are influential in the computer model. The resulting prior for $\beta_{1:d}$ is then

$$p(\beta_{1:d}) \propto \prod_{i=1}^d (1 - \exp(-\beta_i/4))^{0.9} \exp(-\beta_i/4). \quad (9)$$

2.5 DLM-GASP Prediction for Simulator Emulation

At any query input \mathbf{z} , we can predict the output of the computer model as follows based on assumed values of all parameters that will be generated from posterior simulation. Given the realized values of $\epsilon_t(\mathbf{z}_i) = y(\mathbf{z}_i, t) - \sum_{j=1}^p \phi_{t,j} y(\mathbf{z}_i, t-j)$ the predictive distribution for $y(\mathbf{z}, t)$ conditional on all model parameters, on the output training data $y_{1:T}(\mathbf{z}_{1:n})$, and on the “previous” time series values $\mathbf{y}_{t-1:t-p}(\mathbf{z})$ is normal with mean

$$\boldsymbol{\mu}_t(\mathbf{z}) = \sum_{j=1}^p \phi_{t,j} y(\mathbf{z}, t-j) + v_t^{-1} \boldsymbol{\gamma}'_{\mathbf{z}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\epsilon}_t$$

and variance $\sigma_t^2(\mathbf{z}) = v_t(1 - \boldsymbol{\gamma}_z^t \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}_z)$, where $\boldsymbol{\gamma}_z = (c_1(\mathbf{z}, \mathbf{z}_1), \dots, c_1(\mathbf{z}, \mathbf{z}_n))'$ and $\boldsymbol{\epsilon}_t = (\epsilon_t(\mathbf{z}_1), \dots, \epsilon_t(\mathbf{z}_n))'$. These distributions define the emulator that results from our DLM-GASP model. It is apparent that the emulator interpolates the data with $\mu_t(\mathbf{z}) = y_t(\mathbf{z})$ and $\sigma_t^2(\mathbf{z}) = 0$ at the realized points \mathbf{z}_i ; this is a typical feature of the emulators for deterministic computer models that carries over to this new DLM-GASP framework.

2.6 Posterior MCMC Sampling

We use MCMC for simulation of the joint posterior of the unknowns $\boldsymbol{\phi}_{1:T} = \{\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_T\}$, $\mathbf{v}_{1:T} = (v_1, \dots, v_T)$ and $\beta_{1:d}$. The joint posterior distribution can be written as

$$p(\boldsymbol{\phi}_{1:T}, \mathbf{v}_{1:T}, \beta_{1:d} \mid \mathbf{y}_{1:T}) \propto p(\mathbf{y}_{1:T} \mid \boldsymbol{\phi}_{1:T}, \mathbf{v}_{1:T}, \beta_{1:d}) p(\beta_{1:d}) \prod_{t=1}^T p(\boldsymbol{\phi}_t, v_t \mid \mathbf{y}_{1:t-1}, \boldsymbol{\beta})$$

with conditional likelihood function proportional to

$$p(\mathbf{y}_{1:T} \mid \boldsymbol{\phi}_{1:T}, \mathbf{v}_{1:T}, \beta_{1:d}) \propto |\boldsymbol{\Sigma}|^{-(T-p)/2} \exp(-Q/2) \prod_{t=p+1}^T v_t^{-n/2}$$

where

$$Q = \sum_{t=p+1}^T (\mathbf{y}_t - \mathbf{F}'_t \boldsymbol{\phi}_t)' \boldsymbol{\Sigma}^{-1} (\mathbf{y}_t - \mathbf{F}'_t \boldsymbol{\phi}_t) / v_t.$$

We use a block Gibbs sampler (Gelfand and Smith 1990) with a Metropolis-within-Gibbs step for $\beta_{1:d}$. The overall MCMC iterates between Gibbs draws from the complete conditional posterior distribution of $(\boldsymbol{\phi}_{1:T}, \mathbf{v}_{1:T})$ and Metropolis-Hastings draws whose target distributions are the conditionals for $\beta_{1:d}$. Conditional on $\beta_{1:d}$, the first step here involves running the forward filtering, backward sampling computations to efficiently sample the full sequence of AR parameters and innovation variances; see Appendix 2. A key novelty is the Metropolis-Hastings component for $\beta_{1:d}$ targeting the conditional $p(\beta_{1:d} \mid \mathbf{y}_{1:T}, \mathbf{v}, \boldsymbol{\phi}_{1:T})$ at each Gibbs iterate. This requires careful choice of the proposal distributions and extensive exploration as discussed in Appendix 3.

Given the draws of the parameters, at any query input \mathbf{z} we can then make draws from the emulator sequentially as described in Section 2.5. This generates repeat samples of outcomes $y_{1:T}(\mathbf{z})$ to summarize for posterior predictive inferences.

3 Example: LogSPM Hydrological Model

3.1 Background and Problem Context

The hydrological *saturated path models* (logSPMs) of Kuczera et al. (2006) are deterministic simulators of water balance dynamics at catchment scales; these models are

central to environmental models at these scales. The logSPM has 13 inputs: 2 time series inputs of external forces – rainfall and potential evapotranspiration; 3 initial conditions – initial water volume per unit area in the soil, in the ground water and in the river; and 8 coefficients describing water dynamics. The external forces and initial conditions are typically known and here we take them fixed so that we have an input dimension $d = 8$ for water dynamic coefficients \mathbf{z} . The logSPM output is a river flow time series. See Reichert et al. (2008) for description of the partial differential equations of this computer model and data source.

We use the first data set in Reichert et al. (2008) as our training data, which consists of $n = 60$ computer model runs with inputs selected according to a 60-point Latin hypercube design (LHD); a 2-dimensional projection of the design after normalization is given in Figure 1. These input points reasonably fill the design space as is important to exploring the global characteristics of the simulator. The corresponding 60 output time series over $T = 1827$ time points outputs are shown in Figure 2.

We analyse the output data on the log scale (right frame in Figure 2), fitting the DLM-GASP model to the $n = 60$ training data inputs and output time series over $t = 1, \dots, T$. An additional data set of another $n = 60$ runs is also available in Reichert et al. (2008). These runs are based on a 60-point LHD set of inputs that is independent from the designed inputs of the training data, i.e., the design allows different patterns from the training data. We use this data set as our validation data set.

3.2 DLM-GASP Emulation of LogSPM

Based on aspects of the known structure of the LogSPM model (Reichert et al. 2008) we specify $p = 2$ in the TVAR model component, with hyper-parameters defining the initial priors reflecting the known range and scale of outputs while being hedged with substantial uncertainty: $\mathbf{m}_0 = \mathbf{0}$, $n_0 = 1$, $d_0 = 1$ and $\mathbf{C}_0 = 10\mathbf{I}_2$. Discount factors for the analysis exemplified here are $\delta_v = \delta_W = 0.9$ which allow some variation over time in AR parameters and innovations variance; repeat analyses with varying discounts confirm that the global patterns of variation over time in ϕ_t and v_t are evident across models, and the overall conclusions here are not sensitive to choices of discount factors δ_v and δ_W . The MCMC analysis was run for a total of $H = 200,000$ iterations generating draws $\Theta^{(h)}$, ($h = 1, \dots, H$), of all uncertain DLM-GASP quantities $\Theta = \{\phi_{1:T}, v_{1:T}, \beta_{1:d}\}$.

Figure 3 shows approximate posteriors for the range parameters $\beta_{1:d}$. That for β_3 is centered around 0.066 which corresponds to a half-range correlation around $\exp(-0.066/4) \approx 0.98$; in contrast, a β_8 near 2.8 corresponds to a half-range ≈ 0.50 . The high correlation in the 3^{rd} dimension and low correlation in the 8^{th} dimension of input space suggest that a more efficient design should put fewer design points in the 3^{rd} than in the 8^{th} dimension. Furthermore, noticing that the smallest correlation in the 3^{rd} direction is still very large, about $\exp(-0.066) \approx 0.94$, one may want to remove the 3^{rd} input from the Gaussian process – that is, posterior inference is suggestive of model refinements and input variable *screening* and selection in a natural manner.

The model identifies time variation in the TVAR model parameters ϕ_{t1} , ϕ_{t2} and v_t

via estimated trajectories summarized in Figure 4. Though the TVAR model seems consistent with *local* stationarity as a general rule, with values of ϕ_t sitting mainly within the AR(2) stationary region for parameter values (West and Harrison 1997), there are substantial global non-stationarities evident: major stochastic changes in AR parameters and volatility changes in the innovations variance consistent with periods of apparent volatility and spiking behavior in the simulator outputs.

For any of the given validation input points \mathbf{z} , the DLM-GASP emulator generates draws from the posterior predictive distribution of future outputs from further model runs at \mathbf{z} . Applying the emulator to all the 60 inputs in the validation data set we obtain posterior predictive draws illustrated in Figure 6; this shows pointwise posterior predictive medians for all 60 inputs in the validation data. These are comparable, in terms of overall nature of patterns over time, with the realized validation outputs from the simulator, also displayed. The emulator is defining and predicting global trends, appropriate levels of variability across the set of inputs, and, in particular, is capturing both timing and magnitudes of the peaks and troughs including the marked periods of bursts in volatility and spiking of the output processes.

For each validation run, we can compare the prediction given by the emulator with the true validation data by the mean squared error (MSE), defined as

$$\text{MSE} = \frac{1}{T} \sum_{t=1}^T (\hat{y}(\mathbf{z}, t) - y(\mathbf{z}, t))^2,$$

where $\hat{y}(\mathbf{z}, t)$ is the posterior median of the posterior draws at time t , and $y(\mathbf{z}, t)$ the true validation data. The MSE can be potentially used as a measure to compare the performance of the present emulator with that of some other emulator, if both are applied to the same data set. For the 60 validation runs in the LogSPM data, the present emulator yields MSEs with mean and standard deviation 1.67 and 1.53, respectively.

We further investigate predictions of several individual runs, making comparisons of simulated predictions from the DLM-GASP emulator with the true outputs for validation sample inputs \mathbf{z}_1 , \mathbf{z}_{20} , \mathbf{z}_{40} and \mathbf{z}_{60} ; see Figure 7. This provides a further, good and subjective view of the ability of the emulator to predict global as well as local characteristics of the simulator-based hydrological process.

4 Further Comments

The DLM-GASP approach is particularly suitable as a candidate emulator for computer models generating long time series outputs, and when there is evidence of dynamic variation in underlying trends as well as in stochastic volatility in the outputs. TVAR models are well-established as flexible, tractable and sometimes interpretable empirical models of what may be quite complex and erratic processes, and we believe this synthesis with GASP will be of interest in computer model evaluation and emulation studies. DLM-GASP provides a computationally feasible approach to simulator emulation even with very long time series outputs, and will be of use in studies aiming to explore aspects

of variation in time and over input space to improve understanding of the underlying model structure as well as in studies where prediction at new design points is of central interest. The flexible model framework will also be relatively easily adaptable to other issues related to the computer experiments, such as design and calibration. We also note the potential for more customized models that build additional structural components into the multivariate DLM, including the opportunity to bring in dynamic regression terms as functions of inputs if desired; the overall methodology will easily extend to accommodate that since dynamic regressions are also special cases of DLMs and can be combined trivially with TVAR model components (West and Harrison 1997).

Supplementary Material: Code and data

Further details and complete information needed to recapitulate the analyses reported are available at www.stat.missouri.edu/~liufei/DlmComputerWeb.zip. This includes the data and the Matlab code to conduct the analysis and a brief readme on use of the code.

References

- Aguilar, O., Prado, R., Huerta, G., and West, M. (1999). “Bayesian inference on latent structure in time series (with discussion).” In Bernardo, J., Berger, J., Dawid, A., and Smith, A. (eds.), *Bayesian Statistics 6*, 3–26. Oxford University Press.
- Bayarri, M., Berger, J., Garcia-Donato, G., Liu, F., Palomo, J., Paulo, R., Sacks, J., Walsh, D., Cafeo, J., and Parthasarathy, R. (2007a). “Computer Model Validation with Functional Outputs.” *Annals of Statistics*, 35: 1874–1906.
- Bayarri, M. J., Berger, J. O., Kennedy, M., Kottas, A., Paulo, R., Sacks, J., Cafeo, J. A., Lin, C. H., and Tu, J. (2005). “Bayesian Validation of a Computer Model for Vehicle Crashworthiness.” Technical report, National Institute of Statistical Sciences, <http://www.niss.org/technicalreports.html>.
- Bayarri, M. J., Berger, J. O., Paulo, R., Sacks, J., Cafeo, J. A., Cavendish, J., Lin, C.-H., and Tu, J. (2007b). “A Framework for Validation of Computer Models.” *Technometrics*, 49(2): 138–154.
- Carter, C. K. and Kohn, R. (1994). “On Gibbs sampling for state-space models.” *Biometrika*, 81: 541–553.
- Craig, P. S., Goldstein, M., Rougier, J., and Seheult, A. (2001). “Bayesian forecasting for complex systems using computer simulators.” *Journal of the American Statistical Association*, 96(454): 717–729.
- Currin, C., Mitchell, T., Morris, M., and Ylvisaker, D. (1991). “Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments.” *Journal of the American Statistical Association*, 86: 953–963.
- Frühwirth-Schnatter, S. (1994). “Data augmentation and dynamic linear models.” *Journal of Time Series Analysis*, 15: 183–202.
- Fuentes, M., Guttorp, P., and Challenor, P. (2003). “Statistical assessment of numerical models.” *International Statistical Review*, 201–221.
- Gelfand, A. E. and Smith, A. F. M. (1990). “Sampling based approaches to calculating marginal densities.” *Journal of the American Statistics Association*, 85: 398–409.
- Goldstein, M. and Rougier, J. C. (2003). “Calibrated Bayesian forecasting using large computer simulators.” Technical report, University of Durham, <http://www.maths.dur.ac.uk/stats/physpred/papers/CalibratedBayesian.ps>.
- (2005). “Probabilistic formulations for transferring inferences from mathematical models to physical systems.” *SIAM Journal on Scientific Computing*, 26(2): 467–487.
- Higdon, D., Gattiker, J., Williams, B., and Rightley, M. (2008a). “Computer Model Calibration Using High-Dimensional Output.” *Journal of the American Statistical Association*, 103(482): 570–583.

- Higdon, D., Kennedy, M., Cavendish, J., Cafeo, J., and D., R. R. (2004a). “Combining field observations and simulations for calibration and prediction.” *SIAM Journal of Scientific Computing*, 26: 448–466.
- Higdon, D., Nakhleh, C., Gattiker, J., and Williams, B. (2008b). “A Bayesian calibration approach to the thermal Problem.” *Computer Methods in Applied Mechanics and Engineering (CMAME)*, 197: 2431–2441.
- Higdon, D., Williams, B., Moore, L., McKay, M., and Keller-McNulty, S. (2004b). “Uncertainty Quantification for Combining Experimental Data and Computer Simulations.” Technical Report LA-UR 04-6562, Los Alamos National Laboratories.
- Kennedy, M. C. and O’Hagan, A. (2001). “Bayesian calibration of computer models (with discussion).” *Journal of the Royal Statistical Society B*, 63: 425–464.
- Kennedy, M. C., O’Hagan, A., and Higgins, N. (2002). “Bayesian analysis of computer code outputs.” In *Quantitative Methods for Current Environmental Issues*. C. W. Anderson, V. Barnett, P. C. Chatwin, and A. H. El-Shaarawi (eds.), 227–243. Springer-Verlag: London.
- Kuczera, G., Kavetski, D., Franks, S., and Thyer, M. (2006). “Towards a Bayesian total error analysis of conceptual rainfall-runoff models: Characterising model error using storm-dependent parameters.” *Journal of Hydrology*, 331: 161–177.
- Morris, M. D., Mitchell, T. J., and Ylvisaker, D. (1993). “Bayesian design and analysis of computer experiments: Use of derivatives in surface prediction.” *Technometrics*, 35: 243–255.
- Paulo, R., Lin, J., Roupail, N., and Sacks, J. (2005). “Calibrating and Validating Deterministic Traffic Models: Application to the HCM Control Delay at Signalized Intersections.” *Transportation Research Record: Journal of the Transportation Research Board*, 1920: 95–105.
- Prado, R., Huerta, G., and West, M. (2001). “Bayesian time-varying autoregressions: Theory, methods and applications.” *Resenhas*, 4: 405–422.
- Prado, R. and West, M. (1997). “Exploratory modelling of multiple non-stationary time series: Latent process structure and decompositions.” In Gregoire, T., Brillinger, D., Diggle, P., Russek-Cohen, E., Warren, W., and Wolfinger, R. (eds.), *Modelling Longitudinal and Spatially Correlated Data*, 349–362. New York: Springer-Verlag.
- Reichert, P., White, G., Bayarri, M., Pitman, E., and Santer, T. (2008). “Mechanism-based Emulation of Dynamic Simulators: Concept and Application in Hydrology.” http://www.eawag.ch/kuerze/personen/homepages/reichert/index_EN.
- Rougier, J. (2008a). “Efficient Emulators for Multivariate Deterministic Functions.” *Journal of Computational and Graphical Statistics (In press)*.

- (2008b). “Formal Bayes Methods for Model Calibration with Uncertainty.” In Beven, K. and Hall, J. (eds.), *Applied Uncertainty Analysis for Flood Risk Management (In press)*. Imperial College Press/World Scientific. Draft version available at <http://www.maths.bris.ac.uk/~mazjcr/FRMbox2-4.pdf>.
- Rougier, J., Sexton, D., Murphy, J., and Stainforth, D. (2008). “Analysing the climate sensitivity of the HadSM3 climate model using ensembles from different but related experiments.” Technical report, University of Bristol. Available at <http://www.maths.bris.ac.uk/~mazjcr/qump1.pdf>.
- Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). “Design and analysis of computer experiments (C/R: p423-435).” *Statistical Science*, 4: 409–423.
- Sanso, B., Forest, C., and Zantedeschi, D. (2007). “Statistical Calibration of Climate System Properties.” Technical Report asm2007-06, Dept. of Applied Math & Statistics, University of California, Santa Cruz, <http://www.soe.ucsc.edu/research/report?ID=476>.
- (2008). “Inferring Climate System Properties Using a Computer Model.” *Bayesian Analysis*, 3(1): 1–38.
- Santner, T., Williams, B., and Notz, W. (2003). *The Design and Analysis of Computer Experiments*. New York: Springer-Verlag.
- Tebaldi, C. and Sanso, B. (2008). “Joint Projections of Temperature and Precipitation Change from Multiple Climate Models: A Hierarchical Bayes Approach.” *Journal of the Royal Statistical Society: Series A (In press)*.
- Welch, W. J., Buck, R. J., Sacks, J., Wynn, H. P., Mitchell, T. J., and Morris, M. D. (1992). “Screening, predicting, and computer experiments.” *Technometrics*, 34: 15–25.
- West, M. and Harrison, P. (1997). *Bayesian Forecasting and Dynamic Models*. New York: Springer-Verlag, 2 edition.

About the Authors

Fei Liu is Assistant professor, Department of Statistics, University of Missouri, Columbia, MO, 65211, USA. liufei@missouri.edu.

Mike West is The Arts & Sciences Professor of Statistical Science in the Department of Statistical Science, Duke University, Durham, NC 27708-0251, USA. mw@stat.duke.edu.

Acknowledgments

This research was supported in part by the Statistical and Applied Mathematical Sciences Institute (SAMSI) 2006-7 research program on *Development, Assessment and Utilization of Complex Computer Models* and by the University of Missouri-Columbia summer research fellowship and research board award. The computer model data used in this paper was kindly

provided by Peter Reichert. The first author would like to thank Dr. James Berger for his advisory during her Ph.D. study at Duke University and support through the National Science Foundation (GRANTS DMS-0103265, DMS-0112069). We also thank the editor and reviewers for comments on the original version of the paper. Research of Mike West was partially supported by National Science Foundation (DMS-0342172). Further, the final versions of this research were partly supported during the Sequential Monte Carlo program at SAMSI under the National Science Foundation Grant DMS-0635449. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Any opinions, findings and conclusions or recommendations expressed in this work are those of the authors and do not necessarily reflect the views of the NSF.

1 Kronecker product

The Kronecker product of two matrices $A = (a_{ij})_{i,j}$ and $B = (b_{ij})_{i,j}$ is defined as,

$$A \otimes B = \begin{bmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{bmatrix}.$$

It has the following properties.

1. $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$ if A and B are both invertible.
2. $|A \otimes B| = |A|^{d_2} |B|^{d_1}$, where d_1 and d_2 are the dimensions of A , B .
3. $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$ if the dimensions are matched.
If we assume $A = U_1 U_1'$, $B = U_2 U_2'$, then $(A \otimes B) = (U_1 \otimes U_2)(U_1 \otimes U_2)'$.

2 DLM and Forward Filtering, Backward Sampling

2.1 DLM Completion

Refer to the DLM representation of the TVAR model of equations (4) and (5) and let $D_t = \{\mathbf{y}_t, D_{t-1}\}$ be the information set at time t where D_0 is the initial information set. For this appendix material for notational clarity, we assume all other model parameters not mentioned, including the GASP parameters α_i, β_i , are conditioned upon and absorbed into D_0 . The initial prior is $(\phi_0 | v_0, D_0) \sim N(\mathbf{m}_0, v_0 \mathbf{C}_0)$ and $(v_0^{-1} | D_0) \sim \text{Ga}(n_0/2, d_0/2)$ with specified initial hyper-parameters. The models for the sequences v_t and \mathbf{W}_t are based on standard variance discounting using a variance discount factor δ_v for v_t and one δ_W for \mathbf{W}_t , each lying in the unit interval and typically taking values very close to 1 to represent relative stability over time in stochastic changes in the (v_t, \mathbf{W}_t) sequences. The nature and roles of the two variance discount factors (δ_v, δ_W) is clear in the resulting filtering equations summarized here.

2.2 Forward Filtering Computations

Initializing at $t = 1$, the forward filtering analysis sequences through $t = 1, 2, \dots$ updating the one-step forecast and posterior distributions as follows, valid for all $t > 0$:

- (a) The posterior for $(\phi_{t-1}, v_{t-1} | D_{t-1})$ is the conjugate normal/inverse-gamma form with margins $\phi_{t-1} \sim T_{n_{t-1}}(\mathbf{m}_{t-1}, \mathbf{C}_{t-1})$ and $v_{t-1}^{-1} \sim \text{Ga}(n_{t-1}/2, d_{t-1}/2)$; we note that $s_{t-1} = d_{t-1}/n_{t-1}$ is an implied point estimate of v_{t-1} .
- (b) The prior at for $(\phi_t, v_t | D_{t-1})$ has the same form with $\phi_t \sim T_{\delta_v n_{t-1}}(\mathbf{m}_{t-1}, \mathbf{R}_t)$ and $v_t^{-1} \sim \text{Ga}(\delta_v n_{t-1}/2, \delta_v d_{t-1}/2)$ where $\mathbf{R}_t = \mathbf{C}_{t-1} + \mathbf{W}_t$. The variance discounting model sets $\mathbf{W}_t = \mathbf{C}_{t-1}(1 - \delta_W)/\delta_W$ resulting in $\mathbf{R}_t = \mathbf{C}_{t-1}/\delta_W$. The effect of variance discounting is clear from these equations.

- (c) The one-step forecast distribution for $(\mathbf{y}_t \mid D_{t-1})$ is $\mathbf{y}_t \sim \text{T}_{\delta_v n_{t-1}}(\mathbf{f}_t, \mathbf{Q}_t)$ with $\mathbf{f}_t = \mathbf{F}'_t \mathbf{m}_{t-1}$ and $\mathbf{Q}_t = \mathbf{F}'_t \mathbf{R}_t \mathbf{F}_t + s_{t-1} \mathbf{\Sigma}$.
- (d) The posterior for $(\phi_t, v_t \mid D_t)$ is normal/inverse-gamma – as in part (a) following the update to time t . The margins are:
- $v_t^{-1} \sim \text{Ga}(n_t/2, d_t/2)$ with $n_t = \delta_v n_{t-1} + n$ and $d_t = \delta_v d_{t-1} + s_{t-1} \mathbf{e}'_t \mathbf{Q}_t^{-1} \mathbf{e}_t$ based on realized forecast error vector $\mathbf{e}_t = \mathbf{y}_t - \mathbf{F}'_t \mathbf{m}_{t-1}$. The updated point estimate of v_t is $s_t = d_t/n_t$.
 - $\phi_t \sim \text{T}_{n_t}(\mathbf{m}_t, \mathbf{C}_t)$ where $\mathbf{m}_t = \mathbf{m}_{t-1} + \mathbf{A}_t(\mathbf{y}_t - \mathbf{F}'_t \mathbf{m}_{t-1})$ and $\mathbf{C}_t = (\mathbf{W}_t - \mathbf{A}_t \mathbf{Q}_t \mathbf{A}'_t) s_t / s_{t-1}$ where $\mathbf{A}_t = \mathbf{R}_t \mathbf{F}_t \mathbf{Q}_t^{-1}$.

These results follow from the theory of general multivariate DLMs (West and Harrison 1997).

2.3 Backward Sampling Computations

Having completed the forward filtering computations over times $t = 1, \dots, T$, backward sampling generates a posterior draw for the full sequence of states and variances from the joint posterior $p(\phi_{1:T}, v_{1:T} \mid D_T)$. This operates in reverse time, using the general theory of DLMs that generate this computationally efficient algorithm for what can be a very high-dimensional full set of states when T is large. Backward sampling proceeds as follows.

- Initialize at $t = T$: draw v_T^{-1} from $\text{Ga}(n_T/2, d_T/2)$ followed by $\phi_T \mid v_T$ from $\text{N}(\mathbf{m}_T, \mathbf{C}_T v_T / s_T)$.
- Set $n_T^* = n_T$, $d_T^* = d_T$, $m_t^* = m_T$ and $C_T^* = C_T$. Then, for $t = T-1, T-2, \dots, 0$ in sequence, sample:
 - $v_t^{-1} \sim \text{Ga}(n_t^*/2, d_t^*/2)$ where $d_t^* = n_t^* s_t^*$ with $n_t^* = (1 - \delta_v) n_t + \delta_v n_{t+1}^*$ and $1/s_t^* = (1 - \delta_v)/s_t + \delta_v/s_{t+1}^*$;
 - $\phi_t \mid v_t \sim \text{N}(\mathbf{m}_t^*, \mathbf{C}_t^* v_t / s_t^*)$ where $\mathbf{m}_t^* = (1 - \delta_W) \mathbf{m}_t + \delta_W \mathbf{m}_{t+1}^*$ and $\mathbf{C}_t^* = (1 - \delta_W) \mathbf{C}_t / s_t + \delta_W^2 \mathbf{C}_{t+1}^* / s_{t+1}^*$

3 MCMC algorithm in DLM-GASP

At each Gibbs iteration h , the overall MCMC steps through the sequence of simulations as follows:

- (a) Sample $\beta_{1:d}^{(h)}$ from $p(\beta_{1:d} \mid \mathbf{y}_{1:T}, \phi_{1:T}^{(h-1)}, v^{(h-1)})$ by the Metropolis-Hastings algorithm detailed below.
- (b) Sample $(\phi_{1:T}, v_{1:T} \mid \mathbf{y}_{1:T}, \beta_{1:d}^{(h)})$ using the DLM forward filtering, backward sampling algorithm.

The Metropolis-Hastings algorithm to draw $\beta_{1:d}^{(h)}$ works as follows. We first set $\eta_i = \log(\beta_i)$ and $\boldsymbol{\eta} = (\eta_1, \dots, \eta_d)$ and propose a new value for η_i as

$$\eta_i^* = \eta_i^{(h-1)} + \text{N}(0, 0.1^2),$$

then setting $\beta_i^* = \eta_i^*$. The choice of the proposal distribution is customized to each application and requires experimentation with the step-width of the random walk proposal here, but the overall strategy is broadly relevant.

Denoting the correlation matrix at the proposal values by $\boldsymbol{\Sigma}_*$, calculate the likelihood function

$$|\boldsymbol{\Sigma}_*|^{-(T-p)/2} \exp(-Q^*/2) \quad \text{where} \quad Q^* = \sum_{t=p+1}^T \mathbf{e}_t' \boldsymbol{\Sigma}_*^{-1} \mathbf{e}_t / v_t.$$

With Jacobian

$$J(\beta_{1:d}) = \prod_{i=1}^d \frac{d\eta_i}{d\beta_i} = \prod_{i=1}^d \beta_i^{-1},$$

this gives the Metropolis-Hastings acceptance ratio

$$r = \max \left\{ \frac{L(\beta_{1:d}^*) \pi(\beta_{1:d}^*) J(\beta_{1:d}^{h-1})}{L(\beta_{1:d}^{(h-1)}) \pi(\beta_{1:d}^{(h-1)}) J(\beta_{1:d}^*)}, 1 \right\}.$$

Set $\beta_{1:d}^{(h)} = \beta_{1:d}^*$ with probability r ; otherwise, set $\beta_{1:d}^{(h)} = \beta_{1:d}^{(h-1)}$. In our example, the proposal distribution for $\beta_{1:d}$ leads to an acceptance rate of about 10%.




Figure 1: Pairwise scatter plots of normalized values of the $d = 8$ input parameters in the 60 runs of the logSPM model.




Figure 2: Outputs of the logSPM model at the $n = 60$ selected inputs in the training data set (left), and after log transformation and shifting to zero-mean (right).

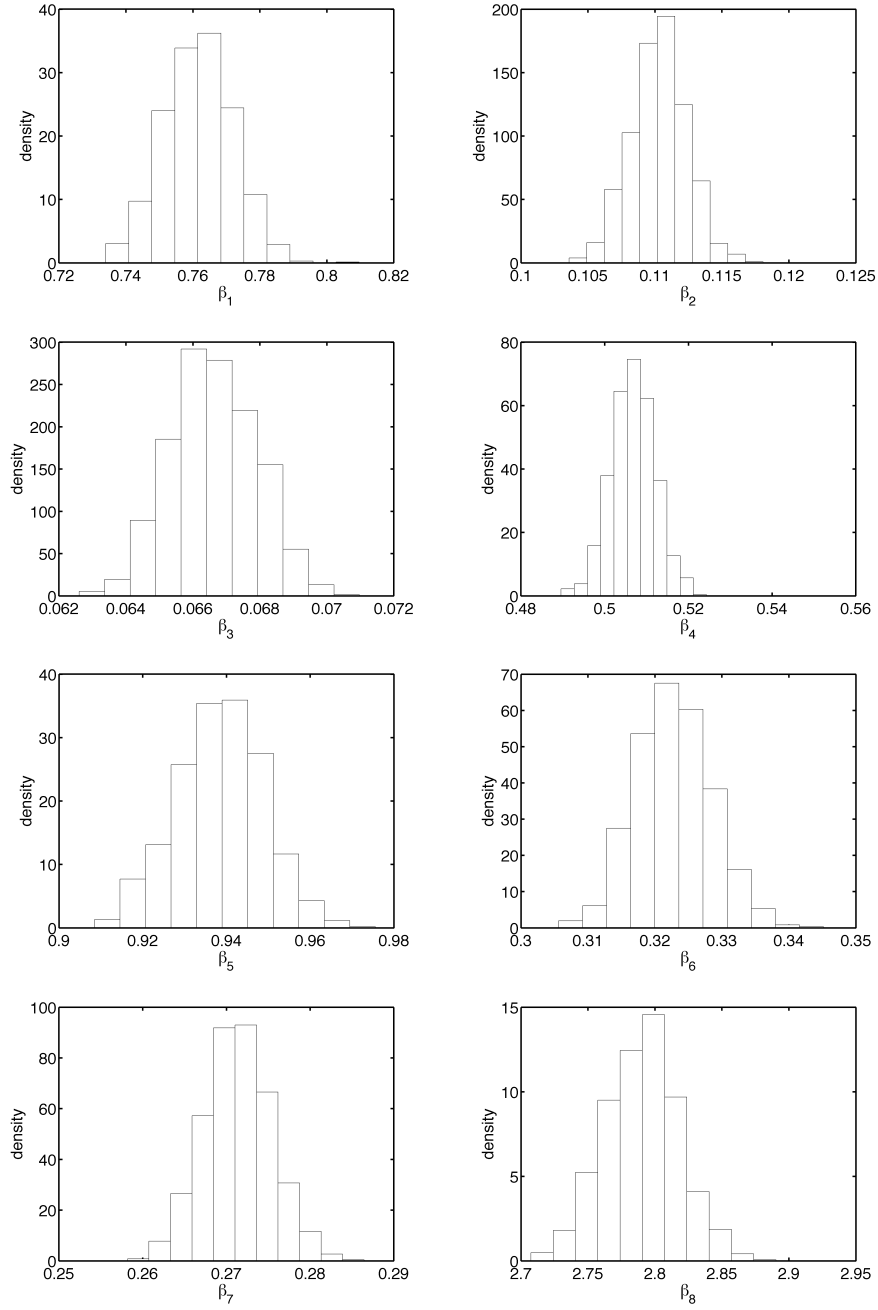


Figure 3: Histograms of posterior samples for the range parameters $\beta_{1:8}$.




Figure 4: Pointwise posterior medians and 95% credible interval for ϕ_{t1} and ϕ_{t2} defining estimated trajectories over $t = 1, \dots, T$.




Figure 5: Pointwise posterior medians and 95% credible interval for v_t defining estimated trajectories over $t = 1, \dots, T$.




Figure 6: Validation data output time series (left) and pointwise posterior predictive medians from the DLM-GASP emulator (right).

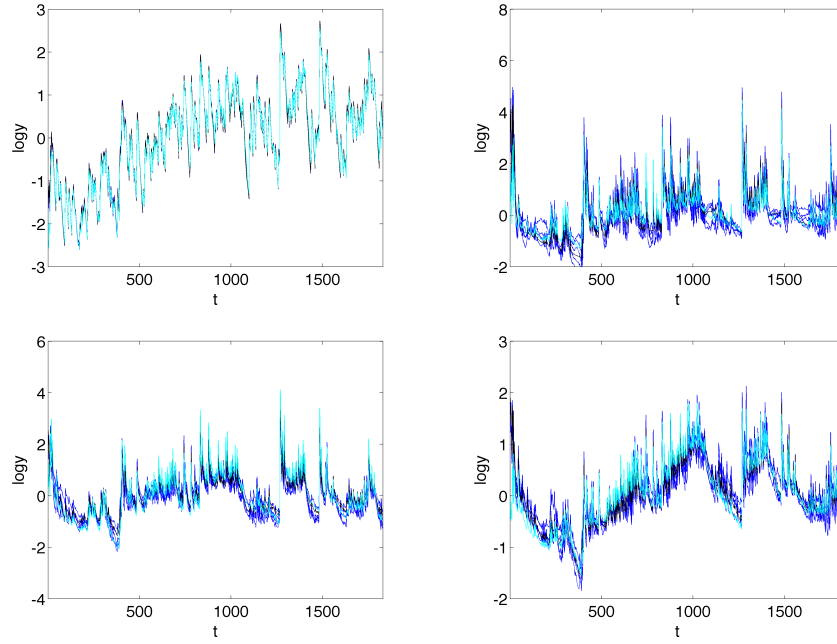


Figure 7: Simulated predictions of outputs at selected validation sample inputs z_1 , z_{20} , z_{40} and z_{60} . Light blue curves are the true validation output series and the darker blue curves are 5 randomly selected draws from the posterior predictive distribution of the DLM-GASP emulator. The black curves are the posterior medians.