

Supplementary material

A Exponential family

A generic exponential family distribution can be written

$$p(z|\eta) = h(z) \exp(\eta^\top s(z) - a(\eta)) \quad (13)$$

where h is the data dependent base measure, a is the log-partition function, and s is a sufficient statistic storing all the information about the natural parameters η contained in the data z . The conjugate prior for an exponential family distribution is also in the exponential family and is given by

$$\pi(\eta|\chi, \nu) = g(\chi, \nu) \exp(\eta^\top \chi - a(\eta)\nu) \quad (14)$$

where g is a normalizing factor and χ, ν are hyperparameters. The parameter posterior for a prior (14) and a likelihood (13) is given by

$$p(\eta|z, \chi, \nu) \propto p(z|\eta)p(\eta|\chi, \nu) = h(z)g(\chi, \nu) \exp(\eta^\top (\chi + s(z)) - a(\eta)(\nu + 1)) \quad (15)$$

where Bayes' rule was used in the first proportionality. If we compare the exponential factor in the posterior (15) with the prior (14) we note that the posterior indeed is of the same form as the prior, but with updated hyperparameters $\chi_{\text{new}} = \chi + s(z)$ and $\nu_{\text{new}} = \nu + 1$. Hence, we have conjugacy for distributions in the exponential family, and the parameter posterior is given by $p(\eta|z, \chi, \nu) = \pi(\eta|\chi + s(z), \nu + 1)$. We can obtain the likelihood of the data z by marginalizing out the natural parameters η from the joint distribution $p(z, \eta|\chi, \nu) = p(z|\eta)p(\eta|\chi, \nu)$ which gives

$$\begin{aligned} p(z|\chi, \nu) &= \int p(z, \eta|\chi, \nu) d\eta = h(z)g(\chi, \nu) \int \underbrace{\exp(\eta^\top (s(z) + \chi) - a(\eta)(1 + \nu))}_{\text{Unnormalized } \pi(\eta|\chi_{\text{new}}, \nu_{\text{new}})} d\eta \\ &= h(z) \frac{g(\chi, \nu)}{g(\chi_{\text{new}}, \nu_{\text{new}})} \end{aligned}$$

where (14) and (13) were inserted in the second equality. Hence, for members of the exponential family there is a closed form expression for the distribution of the data when the parameters have been marginalized out.

B Restricted exponential family

When working with SSMs we can run into problems when using the standard formulation for likelihoods and priors for the exponential family presented in Supplementary A. The reason for this is that we typically have a dependence on the previous state x_{t-1} in the transition density that can result in a log-partition function which depends on this previous state. Put in the same general notation as in the previous section we get the likelihood

$$p(z|\zeta, \eta) = h(z, \zeta) \exp(\eta^\top s(z, \zeta) - a(\eta, \zeta))$$

where the variable ζ is a known extra parameter (e.g the previous state x_{t-1}) and we note that the log-partition function a depends on this parameter. If we wish to formulate an exponential family prior for this likelihood we get

$$\pi(\eta|\chi, \nu, \zeta) = g(\chi, \nu) \exp(\eta^\top \chi - a(\eta, \zeta)\nu),$$

that is, we get a different prior depending on the value on ζ and, consequently, we cannot easily formulate a general prior distribution which is conjugate to the complete data likelihood. Instead we propose to use a restricted exponential family where the log-partition function is assumed to be separable into one parameter-dependent part and one ζ -dependent part. The likelihood is given by

$$p(z|\zeta, \eta) = h(z, \zeta) \exp(\eta^\top s(z, \zeta) - A^\top(\eta)r(\zeta))$$

where $A(\eta)$ is the restricted log-partition function and $r(\zeta)$ is some function which only depends on ζ . A conjugate prior for this likelihood is

$$\pi(\eta|\chi, \nu, \zeta) = g(\chi, \nu) \exp(\eta^\top \chi - A^\top(\eta)\nu).$$

The parameter posterior is

$$\pi(\eta|z, \zeta, \chi, \nu) \propto p(z|\zeta, \eta)p(\eta|\chi, \nu, \zeta) = h(z, \zeta)g(\chi, \nu) \exp(\eta^\top(\chi + s(z, \zeta)) - A^\top(\eta)(\nu + r(\zeta))).$$

We note that the posterior is of the same form as the prior but with updated hyperparameters $\chi_{\text{new}} = \chi + s(z, \zeta)$ and $\nu_{\text{new}} = \nu + r(\zeta)$. Comparing with the standard exponential family we note that the only difference is that the statistic now depends also on ζ and that we get a statistic, $r(\zeta)$, to update also for ν .

C Derivation of ancestor weights for the restricted exponential family

For SSMS with joint state and observation likelihood in the restricted exponential family the likelihood is given by (5) and the parameter prior is $\pi(\theta|\chi_0, \nu_0) = g(\chi_0, \nu_0) \exp(\theta^\top \chi_0 - A^\top(\theta)\nu_0)$. The ancestor weights are given by equation (9), which can be expanded to

$$\begin{aligned} \tilde{w}_{t-1|T}^i &\propto \bar{w}_{t-1}^i \frac{p(x'_{t:T}, y_{t:T}|x_{0:t-1}^i, y_{1:t-1})p(x_{0:t-1}^i, y_{1:t-1})}{p(x_{0:t-1}^i, y_{1:t-1})} = \bar{w}_{t-1}^i p(x'_{t:T}, y_{t:T}|x_{0:t-1}^i, y_{1:t-1}) \\ &= \bar{w}_{t-1}^i p(x'_t, y_t|x_{0:t-1}^i, y_{1:t-1}) \prod_{k=t+1}^T p(x'_k, y_k|x_{0:t-1}^i, y_{1:t-1}, x'_{t:k-1}, y_{t:k-1}) \\ &= \bar{w}_{t-1}^i \int p(x'_t, y_t|x_{t-1}^i, \theta)p(\theta|x_{0:t-1}^i, y_{1:t-1})d\theta \\ &\quad \prod_{k=t+1}^T \int p(x'_k, y_k|x'_{k-1}, \theta)p(\theta|x_{0:t-1}^i, x'_{t:k-1}, y_{1:k-1})d\theta. \end{aligned} \tag{16}$$

Now, to continue we need to compute the two integrals in the last equality for members in the restricted exponential family. First, note that

$$p(x_{0:t-1}^i, y_{1:t-1}|\theta) = \left(\prod_{j=1}^{t-1} h_j^i \right) \exp\left(\theta^\top \sum_{j=1}^{t-1} s_j^i - A^\top(\theta) \sum_{j=1}^{t-1} r_j^i\right)$$

and therefore

$$\begin{aligned} p(\theta|x_{0:t-1}^i, y_{1:t-1}) &\propto p(x_{0:t-1}^i, y_{1:t-1}|\theta)p(\theta|\chi, \nu) \\ &= g(\chi_0, \nu_0) \left(\prod_{j=1}^{t-1} h_j^i \right) \exp\left(\theta^\top \left(\chi_0 + \sum_{j=1}^{t-1} s_j^i\right) - A^\top(\theta) \left(\nu_0 + \sum_{j=1}^{t-1} r_j^i\right)\right) \\ &\propto \pi(\theta|\chi_{t-1}^i, \nu_{t-1}^i), \end{aligned}$$

where $\chi_{t-1}^i = \chi_0 + \sum_{j=1}^{t-1} s_j^i = \chi_0 + s_{1:t-1}^i$ and $\nu_{t-1}^i = \nu_0 + \sum_{j=1}^{t-1} r_j^i = \nu_0 + r_{1:t-1}^i$ and the last proportionality comes from comparison with the exponential factor in the prior. Similarly, by splitting in three terms (one for the ancestral part, one for the cross-over and one for the reference trajectory part) we have

$$\begin{aligned} p(x_{0:t-1}^i, x'_{t:k-1}, y_{1:k-1}|\theta) &= \left(\prod_{j=1}^{t-1} p(x_j^i, y_j|x_{j-1}^i, \theta) \right) p(x'_t, y_t|x_{t-1}^i, \theta) \left(\prod_{j=t+1}^{k-1} p(x'_j, y_j|x'_{j-1}, \theta) \right) \\ &= \underbrace{\left(\prod_{j=1}^{t-1} h_j^i \right) h_t^i \left(\prod_{j=t+1}^{k-1} h_j^i \right)}_H \exp\left(\theta^\top \underbrace{\left(s_{1:t-1}^i + s_t(x'_t, x_{t-1}^i, y_t) + s'_{t+1:k-1} \right)}_S\right) \\ &\quad - A^\top(\theta) \underbrace{\left(r_{1:t-1}^i + r_t(x_{t-1}^i) + r'_{t+1:k-1} \right)}_R \end{aligned}$$

and therefore, using the abbreviations H, S , and R above, the parameter posterior is

$$\begin{aligned} p(\theta|x_{0:t-1}^i, x'_{t:k-1}, y_{1:k-1}) &\propto g(\chi_0, \nu_0) H \exp\left(\theta^\top(\chi_0 + S) - A^\top(\theta)(\nu_0 + R)\right) \\ &\propto \pi(\theta|\chi_{k-1}^i, \nu_{k-1}^i) \end{aligned}$$

where

$$\begin{aligned}\chi_{k-1}^i &= \chi_0 + s_{1:t-1}^i + s_t(x'_t, x_{t-1}^i, y_t) + s'_{t+1:k-1} = \chi_{t-1}^i + s_t(x'_t, x_{t-1}^i, y_t) + s'_{t+1:k-1} \\ \nu_{k-1}^i &= \nu_0 + r_{1:t-1}^i + r_t(x_{t-1}^i) + r'_{t+1:k-1} = \nu_{t-1}^i + r_t(x_{t-1}^i) + r'_{t+1:k-1}.\end{aligned}$$

We can now compute the integrals in (16):

$$\begin{aligned}& \int p(x'_t, y_t | x_{t-1}^i, \theta) p(\theta | x_{0:t-1}^i, y_{1:t-1}) d\theta = \\ & \int h_t \exp(\theta^\top s'_t - A^\top(\theta) r'_t) g(\chi_{t-1}^i, \nu_{t-1}^i) \exp(\theta^\top \chi_{t-1}^i - A^\top(\theta) \nu_{t-1}^i) d\theta \\ & = h_t g(\chi_{t-1}^i, \nu_{t-1}^i) \int \underbrace{\exp(\theta^\top (s'_t + \chi_{t-1}^i) - A^\top(\theta) (r'_t + \nu_{t-1}^i))}_{\text{Unnormalized } \pi(\theta | \chi_t^i, \nu_t^i)} d\theta \\ & = h_t \frac{g(\chi_{t-1}^i, \nu_{t-1}^i)}{g(\chi_t^i, \nu_t^i)}.\end{aligned}\tag{17}$$

Note that $h_t = h_t(x'_t, x_{t-1}^i, y_t)$ depends on the ancestral path. In a similar fashion, for the second integral we obtain

$$\begin{aligned}& \int p(x'_k, y_k | x'_{k-1}, \theta) p(\theta | x_{0:t-1}^i, x'_{t:k-1}, y_{1:k-1}) d\theta = \\ & h'_k g(\chi_{k-1}^i, \nu_{k-1}^i) \int \exp(\theta^\top (s'_k + \chi_{k-1}^i) - A^\top(\theta) (r'_k + \nu_{k-1}^i)) d\theta \\ & = h'_k \frac{g(\chi_{k-1}^i, \nu_{k-1}^i)}{g(\chi_k^i, \nu_k^i)},\end{aligned}\tag{18}$$

where $h'_k = h'_k(x'_k, x'_{k-1}, y_k)$ only depends on the reference trajectory and observations. Now, by substituting (17) and (18) into (16), we obtain

$$\begin{aligned}\tilde{w}_{t-1:T}^i &\propto \bar{w}_{t-1}^i h_t \frac{g(\chi_{t-1}^i, \nu_{t-1}^i)}{g(\chi_t^i, \nu_t^i)} \prod_{k=t+1}^T h'_k \frac{g(\chi_{k-1}^i, \nu_{k-1}^i)}{g(\chi_k^i, \nu_k^i)} \\ &\propto \bar{w}_{t-1}^i h_t \frac{g(\chi_{t-1}^i, \nu_{t-1}^i)}{g(\chi_T^i, \nu_T^i)},\end{aligned}$$

where the (surprisingly) simple final expression is due to all terms except $g(\chi_{t-1}^i, \nu_{t-1}^i)$ and $g(\chi_T^i, \nu_T^i)$ canceling in the product. Note also that h'_k can be removed (in proportionality) since they are the same for all ancestor particles.

D Theoretical justification of the blocking scheme

To show that the blocking scheme in Algorithm 3 is correct we start by verifying that the underlying (hypothetical) partially collapsed Gibbs (PCG) sampler is correct. Following the notation in [41] we set $X_1 = x_{0:B}$, $X_2 = x_{B+1:B+L}$ (the overlap), $X_3 = x_{B+L+1:T}$, $Y = y_{1:T}$, and the superscript * indicates intermediate quantities. The joint distribution we wish to sample from is $p(X_1, X_2, X_3, \theta | Y)$. The underlying hypothetical PCG of the blocking scheme is

$$\begin{aligned}X_1, X_2^* &\sim p(X_1, X_2 | X_3, Y, \theta) \\ X_2, X_3 &\sim p(X_2, X_3 | X_1, Y) \\ \theta &\sim p(\theta | X_1, X_2, X_3, Y).\end{aligned}\tag{19}$$

We note that step 1 and 3 are correct hypothetical Gibbs steps, the only issue is the marginalized step 2. However, in step 2 it is possible to also sample θ from its full conditional, that is to sample $X_2, X_3, \theta^* \sim p(X_2, X_3 | X_1, Y) p(\theta | X_1, X_2, X_3, Y) = p(X_2, X_3, \theta | X_1, Y)$. This sampling of θ^* is redundant since the sampled value is never conditioned on in the following step and can be removed according to the reasoning in [41], concluding that the underlying hypothetical PCG of the blocking scheme is indeed a correct hypothetical PCG. If one or more of the conditionals in (19) are not

possible to sample from directly some MCMC scheme (such as PG/mPG) can be used. Describing the PMCMC steps in terms of kernels, K , Algorithm 3 can be formulated

$$\begin{aligned} X_1, X_2^* &\sim K(X_1, X_2 | X_1', X_2'; X_3, Y, \theta) \\ X_2, X_3 &\sim K(X_2, X_3 | X_2^*, X_3'; X_1, Y) \\ \theta &\sim p(\theta | X_1, X_2, X_3, Y) \end{aligned}$$

where $'$ indicates the (reference) trajectory from the previous iteration. Again, step 1 and 3 are correct (as long as K in step 1 is correct, e.g. a PG kernel), the main concern is step 2. However, as was the case for the hypothetical sampler, it is possible to add the sampling of θ in step 2. Sampling θ is, again, redundant and step 2 is also correct provided that the stationary distribution of K is the marginalized target.

E Additional results

In this section we provide some additional results from numerical simulations.

E.1 Toy-model in Section 3.2

Figure 6 shows the results for our implementation of (12) in Matlab (left) and Birch (right). We generated $T = 250$ observations from (12) with $\sigma_v^2 = 5.3$ and $\sigma_w^2 = 9$. We used hyperparameters $\alpha_v = \alpha_w = 2$ and $\beta_v = \beta_w = 10$, and all four methods were run for $M = 10000$ iterations. We initialized with $\sigma_v^2 = 4.3$ and $\sigma_w^2 = 9.4$ and used a bootstrap proposal for PG/PGAS and a marginalized bootstrap proposal for mPG/mPGAS. We note that for both implementations there is a clear improvement from marginalizing out the parameters (for both PG and PGAS), and that it is beneficial to use PGAS/mPGAS rather than PG/mPG for this model.

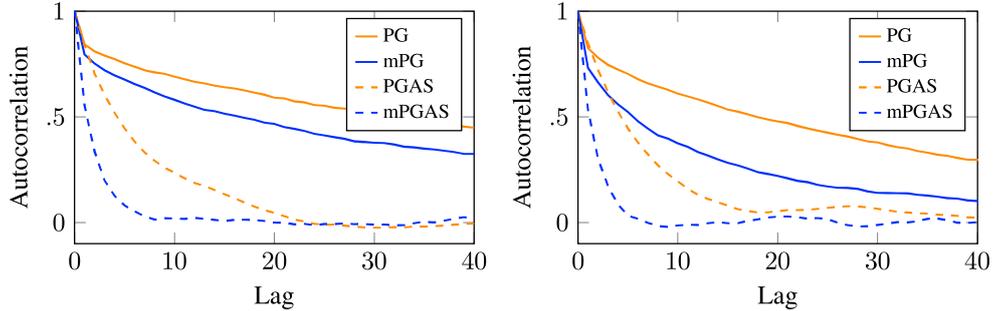


Figure 6: Results of the simulation of the model (12). *Left*: autocorrelation function for all methods, obtained from simulations in Matlab. *Right*: autocorrelation function for all methods, obtained from simulations in Birch.

E.2 Vector-borne disease model in Section 4.2

Figure 7 shows the results of a simulation of the vector-borne disease model in Birch of four different chains; for each of these 10000 samples were drawn using PGAS and mPGAS. The samplers used $N = 1024$ particles each. The autocorrelation of the samples is improved by marginalizing out the parameters. However, we note that there is no obvious improvement from using PGAS rather than PG for this model. One explanation for this could be that it is difficult to match the reference trajectory with an ancestor which is not the reference trajectory for this compartmental model. The reference trajectory would then be sampled in the ancestor sampling step, which in turn means that PGAS would reduce to PG.

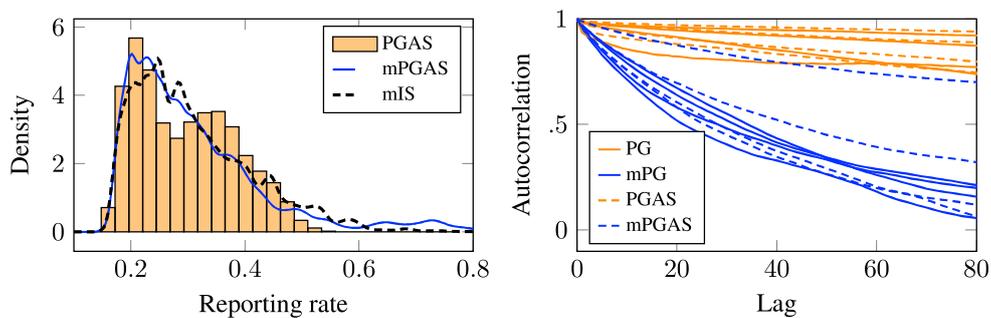


Figure 7: Results of the simulation of the vector-borne disease model. *Left:* estimated density of the reporting rate parameter for PGAS and mPGAS, mean of four chains. For comparison, a run of marginalized importance sampling (mIS) is also shown. *Right:* estimated autocorrelation function of the reporting rate parameter for all methods, four chains for each method.