Importance Densities for Particle Filtering using Iterated Conditional Expectations

Supplementary Material

This Supplementary Material provides additional simulation details and results. Note that in both examples, the integrals over the conditional expectations have closed form solutions and thus, the one-step OID approximation ($L = 1$) coincides with the method in [12]. References are as in the article.

I. Univariate Nonlinear Growth Model

The univariate nonlinear growth model [9], [35] is given by

$$x_n = \frac{x_{n-1}}{2} + \frac{25x_{n-1}}{1+x_{n-1}^2} + 8\cos(1.2n) + q_n,$$
$$y_n = 0.05x_n^2 + r_n,$$

with $x_0 \sim \mathcal{N}(0,5)$ and $q_n \sim \mathcal{N}(0,10)$, which is the standard parametrization used in the literature. However, we assume $r_n \sim \mathcal{N}(0, 1\times 10^{-2})$, which corresponds to two magnitudes lower measurement noise variance than commonly considered.

In addition to the results in the article, Fig. S.1 shows the optimal importance density (OID) approximation for particle $j = 1$ and time step $n = 1$. Here, the bootstrap proposal as well as the slightly adapted one-step OID approximation are both far from the true OID whereas the approximation based on the proposed method converges to the larger OID mode. This is because of the iterations being initialized by the bootstrap proposal, which is closer to the larger OID mode in this case. Furthermore, since a Gaussian density is used to approximate the OID, the smaller mode can not be accounted for.

Fig. S.2 (left) shows the mean effective sample size (ESS) together with the Monte Carlo standard deviation as a function of the number of particles J. Here, the ESS is highest for the proposed method, followed by the particle flow particle filter (PFPF), Gaussian flow OID approximation (GFPF), the one-step OID approximation, and the bootstrap filter (BPF). The one-step OID approximation only performs slightly better compared to the BPF, which is due to the highly concentrated likelihood, which illustrates the benefit of the iterative re-linearization. The GFPF performs better than both the BPF and the one-step OID approximation, but not as good as the proposed method and the PFPF. When comparing the computational time (Fig. S.2, right), it can be seen that the BPF achieves the lowest computational time per particle, mainly due to the possibility of exploiting the model structure. The proposed method is slower than the BPF but faster than the flow-based filters.

II. Multivariate Ricker Population Model

For the multivariate Ricker population model, the dynamic model of the logarithm of a species' ith population out of a total of I interacting populations is given by [36]

$$z_{i,n} = (1-\alpha)z_{i,n-1} + \alpha \sum_{j\in\mathcal{I}\setminus\{i\}} e^{x_{j,n-1}} \frac{e^{-c_{ij}d_{ij}}}{\sum_{k\in\mathcal{I}\setminus\{j\}} e^{-c_{jk}d_{jk}}},$$
$$x_{i,n} = \log(z_{i,n}) + \beta\left(1 - \frac{z_{i,n}}{C_i}\right) + u_{i,n} + v_n q_n,$$

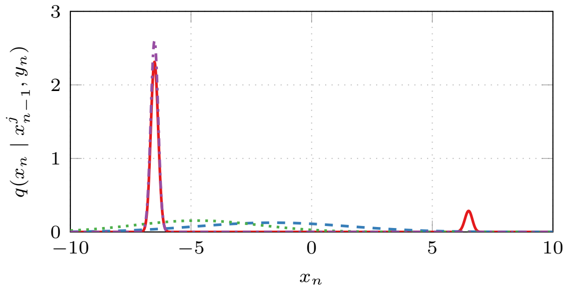

Fig. S.1. Illustration of the OID (——), the bootstrap proposal (– – –), the one-step Gaussian OID approximation (······), and the OID approximation using the proposed method (–·–·) for particle $j = 1$ at $n = 1$.

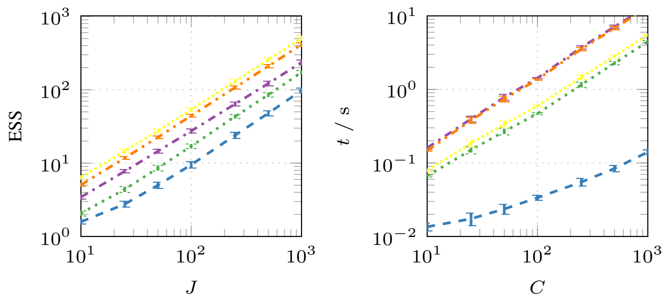

Fig. S.2. ESS (left) and computational time (right) as a function of the number of particles J for the BPF (– – –), one-step Gaussian OID approximation (······), GFPF (–·–·), the PFPF (–·–·), and the proposed method (······).

where $x_{i,n}$ is the logarithm of the ith population at time n, $z_{i,n}$ is the deterministic change in population due to migration, deaths, and births, and $\mathcal{I} = \{1,\dots,I\}$. Furthermore $\alpha = 0.1$ is the migration coefficient, $c_{ij} = 1$ the spatial reach of migration, d_{ij} the Euclidean distance between populations, $\beta = 1$ the growth rate of the population, $C_i = 20$ the carrying capacity for the corresponding population. The stochastic effects enter the model through $u_{i,n}$ and q_n. In particular, $u_{i,n} \sim \mathcal{N}(0,1)$ is the stochastic variation of the ith population, which is independent of all other populations, and $q_n \sim \mathcal{N}(0, 0.3^2)$ the stochastic variation affecting the whole species simultaneously. The latter only occurs sporadically and enters the model through $v_n \sim \mathcal{B}(0.05)$ where $\mathcal{B}(\rho)$ denotes the Bernoulli probability mass function with parameter ρ.

As a likelihood, we use a generalized Poisson model [37], [38] with the population as its mean, that is,

$$y_{i,n} \sim \mathcal{GP}(0.5\exp(x_{i,n}), 0.5).$$

Here,

$$\mathcal{GP}(x;\theta,\lambda) = \frac{\theta(\theta+\lambda x)^{x-1}e^{-\theta-\lambda x}}{x!}$$

is the generalized Poisson distribution [37] which is left-skewed and has a heavier right tail compared to the ordinary Poisson distribution and has mean $\mathrm{E}\{x\} = \theta(1-\lambda)^{-1}$ and variance $\mathrm{var}\{x\} = \theta(1-\lambda)^{-3}$. Furthermore, for each population and time step, a measurement is only obtained with probability 0.5.