
Measuring the non-asymptotic convergence of sequential Monte Carlo samplers using probabilistic programming

Marco F. Cusumano-Towner

Computer Science & Artificial Intelligence Laboratory
Massachusetts Institute of Technology
marcoct@mit.edu

Vikash K. Mansinghka

Department of Brain & Cognitive Sciences
Massachusetts Institute of Technology
vkm@mit.edu

1 Introduction

A key limitation of sampling algorithms for approximate inference is that it is difficult to quantify their approximation error. Widely used sampling schemes, such as sequential importance sampling with resampling and Metropolis-Hastings, produce output samples drawn from a distribution that may be far from the target posterior distribution. This paper shows how to upper-bound the symmetric KL divergence between the output distribution of a broad class of sequential Monte Carlo (SMC) samplers and their target posterior distributions, subject to assumptions about the accuracy of a separate gold-standard sampler. The proposed method applies to samplers that combine multiple particles, multinomial resampling, and rejuvenation kernels. The experiments show the technique being used to estimate bounds on the divergence of SMC samplers for posterior inference in a Bayesian linear regression model and a Dirichlet process mixture model.

This paper builds on a growing body of work begun by [1] and [2] into estimating upper bounds on KL divergences between a sampler’s output distribution and the posterior. In variational inference, the KL divergence of the variational approximation is the gap between the variational lower bound and the log-evidence. [1] and [2] recognized that certain stochastic inference Markov chains including annealed importance sampling (AIS) and single-particle SMC can be treated as variational approximations over an extended space that includes auxiliary random choices of the sampler. A similar insight was introduced independently in [3]. [1] and [2] also showed how to estimate upper bounds on the log-evidence for datasets simulated from the model using generalizations of the harmonic mean estimator, and introduced the bidirectional Monte Carlo (BDMC) technique for ‘sandwiching’ the log-evidence between these upper bounds and variational lower bounds. A related approach for sandwiching the partition function was previously used in the statistical physics literature [4]. Finally, [1] and [2] recognized that the gap between the bounds serves as an upper bound on the KL divergence of the sampler, allowing BDMC to be used for measuring sampler accuracy on simulated datasets.

Two independent papers [5] and [6] built on [1] and [2] to develop the technique further in different ways. Our previous paper [5] took a probabilistic programming perspective, and showed how to estimate the KL divergence bound described in [1] and [2] for general samplers using a ‘meta-inference’ sampler that generates sampler execution histories. [5] also provided meta-inference samplers for sampling importance resampling (SIR) and particle filtering without MCMC rejuvenation kernels. [5] also introduced an upper bound on the symmetric KL divergence between the sampler output and the posterior, analyzed optional use of approximate ‘reference’ samples as surrogates for exact posterior samples (prompting the label ‘subjective divergence’), and related the tightness of the bounds to the accuracy of the meta-inference sampler. A closely related but independent work [6] introduced Bounding Divergences with REverse Annealing (BREAD), which uses the same upper bound on the symmetric KL divergence given in [5], and showed how to evaluate AIS and single-particle SMC approximate inference quality using this bound. BREAD also includes a heuristic scheme, applicable to hierarchical Bayesian statistical models, for generating simulated datasets whose divergence profiles are used as proxies for divergence profiles on real-world datasets. [6] also integrated their technique into existing probabilistic programming platforms.

The main contribution of the current work is a meta-inference construction for generic SMC samplers [7] that is related to conditional SMC [8] and generalizes the existing meta-inference constructions for AIS, single-particle SMC, SIR, and particle filtering. By handling a broad class of samplers, the construction increases relevance for real world problems. The construction allows analysis of samplers that rely on MCMC rejuvenation kernels for good inference quality, while permitting use of multiple particles (instead of custom model-specific annealing schemes) to tighten the KL divergence bounds.

2 Background on subjective divergence

We first review the subjective divergence procedure of [5]. Let p denote an approximate inference sampling program that samples output $z \sim p(z)$ for $z \in \mathcal{Z}$. Suppose p also comes endowed with a side-procedure that evaluates the log probability $\log p(z)$ that the sampler produces any given output z . Let $\pi(z)$ denote the posterior distribution, and let $\tilde{\pi}(z) = \pi(z)Z_{\tilde{\pi}}$ denote an unnormalized posterior distribution. Suppose that we have access to samples from $\pi(z)$. Then the following is an unbiased Monte Carlo estimate of the symmetric KL divergence between $p(z)$ and $\pi(z)$:

$$\frac{1}{N} \sum_{i=1}^N \log \frac{\tilde{\pi}(z_1^i)}{p(z_1^i)} - \frac{1}{M} \sum_{j=1}^M \log \frac{\tilde{\pi}(z_2^j)}{p(z_2^j)} \quad \text{for} \quad \begin{array}{l} z_1^i \sim \pi(z) \quad i = 1 \dots N \\ z_2^j \sim p(z) \quad j = 1 \dots M \end{array} \quad (1)$$

Unfortunately, it is often not possible to efficiently evaluate $\log p(z)$ for sampling programs that sample auxiliary random choices during their execution, including MCMC and SMC sampling algorithms for approximate Bayesian inference. We denote the joint distribution over auxiliary random choices u and output z by $p(u, z)$. It is intractable to marginalize out the auxiliary random choices u because there is an exponentially large number of terms in the sum $p(z) = \sum_u p(u, z)$. Therefore, we instead compute the following unbiased estimate of an upper bound on the symmetric KL divergence, using a ‘meta-inference’ sampler program $u|z \sim q(u; z)$ which samples execution histories of the sampler p (assignments to the auxiliary variables u) given the output z :

$$\frac{1}{N} \sum_{i=1}^N \log \frac{\tilde{\pi}(z_1^i)q(u_1^i; z_1^i)}{p(u_1^i, z_1^i)} - \frac{1}{M} \sum_{j=1}^M \log \frac{\tilde{\pi}(z_2^j)q(u_2^j; z_2^j)}{p(u_2^j, z_2^j)} \quad \text{for} \quad \begin{array}{l} u_1^i, z_1^i \sim \pi(z)q(u; z) \quad i = 1 \dots N \\ u_2^j, z_2^j \sim p(u, z) \quad j = 1 \dots M \end{array} \quad (2)$$

The upper bound estimated is the symmetric KL divergence on an extended space that includes the auxiliary variables u of the sampler. As shown in [5], the tightness of the bound is governed by how well $q(u; z)$ approximates $p(u|z)$ on average for $z \sim p(z)$ and $z \sim \pi(z)$. When samples from a gold-standard approximate inference ‘reference sampler’ are used in place of posterior samples, the validity of the bound is subject to the accuracy of the reference sampler [5].

3 A probabilistic programming interface for subjective divergence

We now clarify the procedures associated with a sampler that are needed for subjective divergence estimation. In particular, we introduce the following probabilistic programming interface, which consists of two stochastic procedures, denoted (p, q) .SIMULATE and (p, q) .REGENERATE for some distributions $p(u, z)$ and $q(u; z)$:

$$\begin{aligned} (z, \log(p(u, z)/q(u; z))) &\leftarrow (p, q).\text{SIMULATE}() \quad \text{for } u, z \sim p(u, z) \\ \log(p(u, z)/q(u; z)) &\leftarrow (p, q).\text{REGENERATE}(z) \quad \text{for } u|z \sim q(u; z) \end{aligned} \quad (3)$$

The SIMULATE procedure runs a sampler with joint distribution $p(u, z)$ over execution histories u and output z , and returns z . The REGENERATE procedure takes a potential sampler output z as its input, and runs a ‘regeneration’ sampler that samples an execution history u of the original sampler. Both procedures also return a log-weight. The log-weight returned by SIMULATE can be interpreted as a log harmonic mean estimate of $p(z)$ and the log-weight returned by REGENERATE can be interpreted as a log importance sampling estimate of $p(z)$. When the sampler is an inference sampler, we call the regeneration sampler a ‘meta-inference’ sampler. As will be seen, the relationship between the original sampler and the regeneration sampler is analogous to the relationship between SMC and conditional SMC [8].

Note that the auxiliary random variables u are not exposed through the interface. Also note that a sampler with a tractable marginal output probability $p(z)$ trivially implements the interface because $\log(p(u, z)/q(u; z))$ reduces to the log output probability when there are no auxiliary variables u . Algorithm 1 shows a procedure that computes Equation (2) using the above interface.

Algorithm 1 Subjective divergence estimation using SIMULATE and REGENERATE

Require: Sampler package (p, q) implementing SIMULATE and REGENERATE; posterior sampler $z \sim \pi(z)$ or reference sampler $z \sim r(z)$; unnormalized posterior probability function $\tilde{\pi}(z)$.

procedure ESTIMATE-KL-BOUND((p, q) , π , $\tilde{\pi}$)

for $i \leftarrow 1 \dots N$ **do**

$z_1^i \sim \pi(z)$ \triangleright Replace with sample from reference sampler $z_1^i \sim r(z)$ if exact posterior sampler unavailable

$\ell_1^i \leftarrow (p, q).\text{REGENERATE}(z_1^i)$

end for

for $j \leftarrow 1 \dots M$ **do**

$(z_2^j, \ell_2^j) \leftarrow (p, q).\text{SIMULATE}()$

end for

return $\frac{1}{N} \sum_{i=1}^N (\log \tilde{\pi}(z_1^i) - \ell_1^i) - \frac{1}{M} \sum_{j=1}^M (\log \tilde{\pi}(z_2^j) - \ell_2^j)$

end procedure

4 Implementing SIMULATE and REGENERATE for sequential Monte Carlo

Algorithm 2 below shows how to implement SIMULATE and REGENERATE for the generic SMC sampler template introduced in [7], with independent resampling. The SMC sampler template (the SIMULATE procedure of Algorithm 2), permits use of MCMC kernels (within the k_t), provided that corresponding ‘backward kernels’ ℓ_t are defined such that the weights can be computed. Note that SIMULATE does not sample from the backward kernels. Building on the analysis of SMC used in [8], the auxiliary variables u for the SMC sampler are the random choices made during its execution: the resampling choices $a_t^i \in \{1 \dots N\}$ for $(i, t) \in \{1 \dots N\} \times \{1 \dots T - 1\}$ and $I_T \in \{1 \dots N\}$ and the values of all intermediate particles $x_t^i \in \mathcal{X}_t$ for $(i, t) \in \{1 \dots N\} \times \{1 \dots T\}$. The output of the SMC sampler is denoted $z \in \mathcal{X}_T$. The SMC stochastic regeneration template (the REGENERATE procedure of Algorithm 2), is given an output $z \in \mathcal{X}_T$, and samples an execution history u of the SMC sampler by first choosing the ancestral particle indices that led to the output (denoted I_t for $t \in \{1 \dots T\}$), then *sampling* from the backward kernels ℓ_t in reverse order to define the ancestral particle values $x_t^{I_t}$ for $t \in \{1 \dots T\}$ that led to the output, and finally running SMC forward, with the ancestral indices I_t and values $x_t^{I_t}$ for $t \in \{1 \dots T\}$ fixed. This is related to the conditional SMC update of [8], but differs in that only an output particle and not a full particle trajectory is required as input. The log-weight for this sampler and regeneration pair simplify to (see Appendix A for derivation):

$$\log \frac{p(u, z)}{q(u; z)} = -\log \left(w_{T+1}^1 \prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j \right) \quad (4)$$

Algorithm 2 SIMULATE and REGENERATE for SMC samplers with independent resampling

Require: Number of steps T ; hypothesis spaces \mathcal{X}_t (not necessarily related) and unnormalized target distributions \tilde{p}_t defined on \mathcal{X}_t where $\tilde{p}_t(x_t) > 0$ for $x_t \in \mathcal{X}_t$ for $t \in \{1 \dots T\}$; sampler for initialization kernel k_1 defined on \mathcal{X}_1 with $k_1(x_1) > 0$ for $x_1 \in \mathcal{X}_1$; samplers for kernels k_t indexed by \mathcal{X}_{t-1} and defined on \mathcal{X}_t for $t \in \{2 \dots T\}$; sampler for kernel k_{T+1} indexed by \mathcal{X}_T and defined on \mathcal{X}_T ; samplers for kernels ℓ_t indexed by \mathcal{X}_t and defined on \mathcal{X}_{t-1} such that $k_t(x_t; x_{t-1}) > 0 \iff \ell_t(x_{t-1}; x_t) > 0$ for $x_{t-1} \in \mathcal{X}_{t-1}, x_t \in \mathcal{X}_t$ for $t \in \{2 \dots T\}$; sampler for kernel ℓ_{T+1} such that $k_{T+1}(z'; z) > 0 \iff \ell_{T+1}(z; z') > 0$ for $z, z' \in \mathcal{X}_T$; evaluator procedures for weight functions $w_1(x_1) := \frac{\tilde{p}_1(x_1)}{k_1(x_1)}$, $w_t(x_{t-1}, x_t) := \frac{\tilde{p}_t(x_t)\ell_t(x_{t-1}; x_t)}{\tilde{p}_{t-1}(x_{t-1})k_t(x_t; x_{t-1})}$ for $t \in \{2, \dots, T\}$ and $w_{T+1}(x_T, x'_T) := \frac{\ell_{T+1}(x_T; x'_T)}{\tilde{p}_T(x_T)k_{T+1}(x'_T; x_T)}$ for $x_T, x'_T \in \mathcal{X}_T$; number of particles N

procedure SIMULATE()

for $i \leftarrow 1 \dots N$ **do**

$x_1^i \sim k_1(\cdot)$

$w_1^i \leftarrow w_1(x_1^i)$

end for

for $t \leftarrow 2 \dots T$ **do**

for $i \leftarrow 1 \dots N$ **do**

$a_{t-1}^i \sim \text{Categorical}(\text{NORMALIZE}(\mathbf{w}_{t-1}))$

$x_t^i \sim k_t(\cdot; x_{t-1}^{a_{t-1}^i})$

$w_t^i \leftarrow w_t(x_{t-1}^{a_{t-1}^i}, x_t^i)$

end for

end for

$I_T \sim \text{Categorical}(\text{NORMALIZE}(\mathbf{w}_T))$

$z \sim k_{T+1}(\cdot; x_T^{I_T})$

$w_{T+1}^1 \leftarrow w_{T+1}(x_T^{I_T}, z)$

return $\left(z, -\log \left(w_{T+1}^1 \prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j \right) \right)$

end procedure

procedure RAND-ANCESTRY(N, T)

for $t \leftarrow 1 \dots T$ **do**

$I_t \sim \text{Uniform}(1, \dots, N)$

end for

return (I_1, \dots, I_T)

end procedure

procedure REGENERATE(z)

$(I_1, \dots, I_T) \sim \text{RAND-ANCESTRY}(N, T)$

$x_T^{I_T} \sim \ell_{T+1}(\cdot; z)$

for $t \leftarrow T - 1 \dots 1$ **do**

$x_t^{I_t} \sim \ell_{t+1}(\cdot; x_{t+1}^{I_{t+1}})$

end for

for $i \leftarrow 1 \dots N$ **do**

if $i \neq I_1$ **then**

$x_1^i \sim k_1(\cdot)$

end if

$w_1^i \leftarrow w_1(x_1^i)$

end for

for $t \leftarrow 2 \dots T$ **do**

for $i \leftarrow 1 \dots N$ **do**

if $i = I_t$ **then**

$a_{t-1}^i \leftarrow I_{t-1}$

else

$a_{t-1}^i \sim \text{Categorical}(\text{NORMALIZE}(\mathbf{w}_{t-1}))$

$x_t^i \sim k_t(\cdot; x_{t-1}^{a_{t-1}^i})$

end if

$w_t^i \leftarrow w_t(x_{t-1}^{a_{t-1}^i}, x_t^i)$

end for

end for

$w_{T+1}^1 \leftarrow w_{T+1}(x_T^{I_T}, z)$

return $-\log \left(w_{T+1}^1 \prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j \right)$

end procedure

Having specified how to implement SIMULATE and REGENERATE for this generic variant of SMC, we can now estimate subjective divergences for SMC. We illustrate the use of Algorithm 1 and Algorithm 2 to estimate subjective bounds on symmetric KL divergences of SMC samplers and black box variational approximations to the posterior in Figure 1. Note that we optimized the performance of variational inference and SMC implementations separately, and the relative runtimes of the two approaches are not meant to be informative.




Figure 1: (a) and (b) show estimated ELBO lower bounds and estimated upper bounds on KL divergence to the posterior, respectively, for SMC samplers and two black box variational inference (BBVI) programs, in a Bayesian linear regression inference problem. SMC IMH samplers use single-site independent Metropolis-Hastings (MH) rejuvenation kernels, and SMC RW use single-site random-walk MH rejuvenation kernels. BBVI 1 and BBVI 2 optimize over different variational families. (c) and (d) show estimated ELBO lower bounds and subjective upper bounds on KL divergence, for SMC with single-site Gibbs rejuvenation kernels over cluster assignments in a Dirichlet process mixture model problem with collapsed cluster parameters. The SMC samplers in both problems use sequential observation to define the sequence of target distributions, and were parameterized by the number of particles (N , represented by color) and the number of applications of the MCMC rejuvenation kernels between target distribution updates (distinct estimates of the same color). Particles were initialized from the prior. Exact posterior reference samples were used for KL bound estimation in (a) and (b), and samples from a gold-standard approximate MCMC reference sampler were used in lieu of posterior samples for KL bound estimation in (c) and (d). In (a) and (b), the random-walk MH kernels appear more effective than the independent MH kernels. Increasing the number of particles tightens the KL divergence bound when the effect of rejuvenation kernels has already been saturated (compare SMC RW $N=1$ to SMC RW $N=40$).

References

- [1] Roger Baker Grosse. *Model selection in compositional spaces*. PhD thesis, Massachusetts Institute of Technology, 2014.
- [2] Roger B Grosse, Zoubin Ghahramani, and Ryan P Adams. Sandwiching the marginal likelihood using bidirectional Monte Carlo. *arXiv preprint arXiv:1511.02543*, 2015.

- [3] Tim Salimans, Diederik P. Kingma, and Max Welling. Markov chain monte carlo and variational inference: Bridging the gap. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 1218–1226, 2015.
- [4] John E Hunter III, William P Reinhardt, and Thomas F Davis. A finite-time variational method for determining optimal paths and obtaining bounds on free energy changes from computer simulations. *The Journal of chemical physics*, 99(9):6856–6864, 1993.
- [5] Marco F Cusumano-Towner and Vikash K Mansinghka. Quantifying the probable approximation error of probabilistic inference programs. *arXiv preprint arXiv:1606.00068*, 2016.
- [6] Roger B Grosse, Siddharth Ancha, and Daniel M Roy. Measuring the reliability of MCMC inference with bidirectional Monte Carlo. *arXiv preprint arXiv:1606.02275*, 2016.
- [7] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.
- [8] Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342, 2010.

Appendix A: Derivation of log weight for SMC

Recall that the set of auxiliary random choices u for the SMC sampler of Algorithm 2 is the set of all resampling choices $a_t^i \in \{1 \dots N\}$ for $(i, t) \in \{1 \dots N\} \times \{1 \dots T-1\}$ and $I_T \in \{1 \dots N\}$ and the values of all intermediate particles $x_t^i \in \mathcal{X}_t$ for $(i, t) \in \{1 \dots N\} \times \{1 \dots T\}$. The joint probability over auxiliary random choices u and output z for an execution of SMC's SIMULATE is:

$$p(u, z) := \left[\prod_{i=1}^N k_1(x_1^i) \right] \left[\prod_{t=2}^T \prod_{i=1}^N \frac{w_{t-1}^{a_t^i}}{\sum_{j=1}^N w_{t-1}^j} k_t(x_t^i; x_{t-1}^{a_t^i}) \right] \left[\frac{w_T^{I_T}}{\sum_{j=1}^N w_T^j} k_{T+1}(z; x_T^{I_T}) \right]$$

The joint probability over auxiliary random choices u for an execution of SMC's REGENERATE is:

$$q(u; z) := \left[\frac{1}{N^T} \right] \left[\ell_{T+1}(x_T^{I_T}; z) \prod_{t=2}^T \ell_t(x_{t-1}^{I_{t-1}}; x_t^{I_t}) \right] \left[\prod_{\substack{i=1 \\ i \neq I_1}}^N k_1(x_1^i) \right] \left[\prod_{t=2}^T \prod_{\substack{i=1 \\ i \neq I_t}}^N \frac{w_{t-1}^{a_t^i}}{\sum_{j=1}^N w_{t-1}^j} k_t(x_t^i; x_{t-1}^{a_t^i}) \right]$$

where the first N^{-T} factor is due to RAND-ANCESTRY. First, note that $w_t^i > 0$ for all $(i, t) \in \{1 \dots N\} \times \{1 \dots T\}$ for $u, z \sim p(u, z)$ and for $u|z \sim q(u; z)$. This is true for $t = 1$ by the requirements $\tilde{p}_1(x_1) > 0$ and $k_1(x_1) > 0$ for all $x_1 \in \mathcal{X}_1$. For $t \in \{2 \dots T\}$, $w_t^i = \frac{\tilde{p}_t(x_t^i) \ell_t(x_{t-1}^{a_t^i}; x_t^i)}{\tilde{p}_{t-1}(x_{t-1}^{a_t^i}) k_t(x_t^i; x_{t-1}^{a_t^i})}$. Either $x_t^i \sim k_t(\cdot; x_{t-1}^{a_t^i})$ or $x_{t-1}^{a_t^i} \sim \ell_t(\cdot; x_t^i)$. Using the requirements $\tilde{p}_t(x_t) > 0$ for all $x_t \in \mathcal{X}_t$ for all $t \in \{1 \dots T\}$ and $k_t(x_t; x_{t-1}) > 0 \iff \ell_t(x_{t-1}; x_t) > 0$ for all $x_{t-1} \in \mathcal{X}_{t-1}, x_t \in \mathcal{X}_t$ for all $t \in \{2 \dots T\}$, gives $w_t^i > 0$ for all $(i, t) \in \{1 \dots N\} \times \{2 \dots T\}$.

To see that $p(u, z) > 0 \iff q(u; z) > 0$, first consider some u, z such that $p(u, z) > 0$. Since $p(u, z) > 0$ we have $k_1(x_1^i) > 0$ for $i \in \{1 \dots N\}$. We also have $k_t(x_t^i; x_{t-1}^{a_t^i}) > 0$ which implies $\ell_t(x_{t-1}^{a_t^i}; x_t^i) > 0$ for $(i, t) \in \{1 \dots N\} \times \{2 \dots T\}$, and $k_{T+1}(z; x_T^{I_T}) > 0$, which implies $\ell_{T+1}(x_T^{I_T}; z) > 0$. Combined with $w_t^i > 0$ for all $(i, t) \in \{1 \dots N\} \times \{1 \dots T\}$ these ensure $q(u; z)$ is defined for output z and $q(u; z) > 0$.

Next, assume $q(u; z)$ is defined for output z and $q(u; z) > 0$. Then we have $\ell_{T+1}(x_T^{I_T}; z) > 0$ which implies $k_{T+1}(z; x_T^{I_T}) > 0$. We also have $\ell_t(x_{t-1}^{I_{t-1}}; x_t^{I_t}) > 0$ for $t \in \{2 \dots T\}$, which implies $k_t(x_t^i; x_{t-1}^{a_t^i}) > 0$ for $(i, t) \in \{(I_t, t) | t \in \{2 \dots T\}\}$. We also have $k_t(x_t^i; x_{t-1}^{a_t^i}) > 0$ for $(i, t) \in \{(i, t) | i \neq I_t, t \in \{2 \dots T\}\}$. Therefore $p(u, z) > 0$.

The weight $p(u, z)/q(u; z)$ is then defined for all $u, z \sim p(u, z)$ and $u|z \sim q(u; z)$, and is:

$$\begin{aligned} \frac{p(u, z)}{q(u; z)} &= \frac{\left[\prod_{i=1}^N k_1(x_1^i) \right] \left[\prod_{t=2}^T \prod_{i=1}^N \frac{w_{t-1}^{a_t^i}}{\sum_{j=1}^N w_{t-1}^j} k_t(x_t^i; x_{t-1}^{a_t^i}) \right] \left[\frac{w_T^{I_T}}{\sum_{j=1}^N w_T^j} k_{T+1}(z; x_T^{I_T}) \right]}{\left[\frac{1}{N^T} \right] \left[\ell_{T+1}(x_T^{I_T}; z) \prod_{t=2}^T \ell_t(x_{t-1}^{I_{t-1}}; x_t^{I_t}) \right] \left[\prod_{\substack{i=1 \\ i \neq I_1}}^N k_1(x_1^i) \right] \left[\prod_{t=2}^T \prod_{\substack{i=1 \\ i \neq I_t}}^N \frac{w_{t-1}^{a_t^i}}{\sum_{j=1}^N w_{t-1}^j} k_t(x_t^i; x_{t-1}^{a_t^i}) \right]} \\ &= \frac{k_1(x_1^{I_1}) \left[\prod_{t=2}^T \prod_{i=1}^N w_{t-1}^{a_t^i} k_t(x_t^i; x_{t-1}^{a_t^i}) \right] \left[w_T^{I_T} k_{T+1}(z; x_T^{I_T}) \right]}{\left[\prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j \right] \left[\ell_{T+1}(x_T^{I_T}; z) \prod_{t=2}^T \ell_t(x_{t-1}^{I_{t-1}}; x_t^{I_t}) \right] \left[\prod_{t=2}^T \prod_{\substack{i=1 \\ i \neq I_t}}^N w_{t-1}^{a_t^i} k_t(x_t^i; x_{t-1}^{a_t^i}) \right]} \\ &= \frac{k_1(x_1^{I_1}) \left[\prod_{t=2}^T w_{t-1}^{I_{t-1}} k_t(x_t^{I_t}; x_{t-1}^{I_{t-1}}) \right] \left[w_T^{I_T} k_{T+1}(z; x_T^{I_T}) \right]}{\left[\prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j \right] \left[\ell_{T+1}(x_T^{I_T}; z) \prod_{t=2}^T \ell_t(x_{t-1}^{I_{t-1}}; x_t^{I_t}) \right]} \\ &= \frac{k_1(x_1^{I_1}) \left[\prod_{t=2}^T \frac{k_t(x_t^{I_t}; x_{t-1}^{I_{t-1}})}{\ell_t(x_{t-1}^{I_{t-1}}; x_t^{I_t})} \right] \frac{k_{T+1}(z; x_T^{I_T})}{\ell_{T+1}(x_T^{I_T}; z)} \left[\frac{\tilde{p}_1(x_1^{I_1})}{k_1(x_1^{I_1})} \prod_{t=2}^T \frac{\tilde{p}_t(x_t^{I_t}) \ell_t(x_{t-1}^{I_{t-1}}; x_t^{I_t})}{\tilde{p}_{t-1}(x_{t-1}^{I_{t-1}}) k_t(x_t^{I_t}; x_{t-1}^{I_{t-1}})} \right]}{\left[\prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j \right]} \\ &= \frac{\tilde{p}_1(x_1^{I_1}) \left[\prod_{t=2}^T \frac{\tilde{p}_t(x_t^{I_t})}{\tilde{p}_{t-1}(x_{t-1}^{I_{t-1}})} \right] \frac{k_{T+1}(z; x_T^{I_T})}{\ell_{T+1}(x_T^{I_T}; z)}}{\prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j} = \frac{\tilde{p}_T(x_T^{I_T}) \frac{k_{T+1}(z; x_T^{I_T})}{\ell_{T+1}(x_T^{I_T}; z)}}{\prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j} = \frac{1}{w_{T+1}^1 \prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j} \end{aligned}$$

Appendix B: SMC with sequential observation and detailed balance kernels

In the experiments, we use SMC programs defined as follows. Let \mathcal{X}_0 be a hypothesis space, corresponding to ‘global’ latent variables. Let \mathcal{E}_t for $t \in \{1 \dots T\}$ be additional hypothesis space extensions, corresponding for ‘local’ latent variables for each of T observations $y_t \in \mathcal{Y}_t$ for $t \in \{1 \dots T\}$. Define $\mathcal{X}_t := \mathcal{X}_{t-1} \times \mathcal{E}_t$ for $t \in \{1 \dots T\}$. We use indexing notation where $a_{s:t} := (a_s, \dots, a_t)$ for any a, s, t . Let $p(\theta, e_{1:T}, y_{1:T})$ denote the model’s joint probability of $(\theta, e_{1:T}, y_{1:T}) \in \mathcal{X}_T \times \mathcal{Y}_1 \times \dots \times \mathcal{Y}_T$, where $\theta \in \mathcal{X}_0$ and $e_t \in \mathcal{E}_t$ for $t \in \{1 \dots T\}$. Assume for the given observation set $y_{1:T}$ that $p(\theta, e_{1:T}, y_{1:T}) > 0$ for all $(\theta, e_{1:T}) \in \mathcal{X}_T$. The target distribution of the SMC algorithm is the conditional distribution $p(\theta, e_{1:T} | y_{1:T}) \propto p(\theta, e_{1:T}, y_{1:T})$. Define the intermediate target distributions by $p_t(\theta, e_{1:t}) := p(\theta, e_{1:t} | y_{1:t})$, and the unnormalized target probability functions by $\tilde{p}_t(\theta, e_{1:t}) := p(\theta, e_{1:t}, y_{1:t})$ for $t \in \{1 \dots T\}$. Define the initialization kernel as: $k_1(\theta, e_1) := p(\theta, e_1)$. This kernel samples from the model’s prior distribution over the global latents and the local latents for the first observation. Suppose there exist ‘detailed balance kernels’ $d_t(\theta', e'_{1:t}; \theta, e_{1:t})$ for $t \in \{1, \dots, T\}$. Kernel d_t is a collection of distributions over elements of \mathcal{X}_t , indexed by elements of \mathcal{X}_t . Each detailed balance kernel d_t must satisfy the detailed balance property with respect to the intermediate target distribution p_t :

$$d_t(\theta', e'_{1:t}; \theta, e_{1:t}) p_t(\theta, e_{1:t}) = d_t(\theta, e_{1:t}; \theta', e'_{1:t}) p_t(\theta', e'_{1:t}) \text{ for all } (\theta, e_{1:t}), (\theta', e'_{1:t}) \in \mathcal{X}_t \quad (5)$$

Equivalently:

$$d_t(\theta', e'_{1:t}; \theta, e_{1:t}) \tilde{p}_t(\theta, e_{1:t}) = d_t(\theta, e_{1:t}; \theta', e'_{1:t}) \tilde{p}_t(\theta', e'_{1:t}) \text{ for all } (\theta, e_{1:t}), (\theta', e'_{1:t}) \in \mathcal{X}_t \quad (6)$$

Define $k_t(\theta', e'_{1:t}; \theta, e_{1:t-1}) := d_{t-1}(\theta', e'_{1:t-1}; \theta, e_{1:t-1}) p(e'_t | \theta', e'_{1:t-1}, y_{1:t-1})$ for $t \in \{2 \dots T\}$. Each k_t for $t \in \{2 \dots T\}$ is a collection of distributions over \mathcal{X}_t , indexed by elements of \mathcal{X}_{t-1} . It is possible to sample from k_t by sampling from the detailed balance kernel d_{t-1} and then sampling from the model prior distribution over the new local latents $p(e'_t | \theta', e'_{1:t-1}, y_{1:t-1})$. Intuitively, the kernel k_t first performs inference d_{t-1} targeting $p(\theta, e_{1:t-1} | y_{1:t-1})$, then extends the hypothesis space to include values of the local latent variables for observation t by sampling from the prior. Define $k_{T+1} := d_T$. Intuitively, kernel k_{T+1} performs inference d_T targeting the final target distribution $p(\theta, e_{1:T} | y_{1:T})$. Define the ‘backward kernels’ by $\ell_t(\theta', e'_{1:t-1}; \theta, e_{1:t}) := d_{t-1}(\theta', e'_{1:t-1}; \theta, e_{1:t-1})$ for $t \in \{2 \dots T\}$. Each kernel ℓ_t for $t \in \{2 \dots T\}$ is a collection of distributions over \mathcal{X}_{t-1} , indexed by elements of \mathcal{X}_t . To sample from ℓ_t , we simply sample from the detailed balance kernel d_{t-1} . Finally, define $\ell_{T+1} := d_T$. First, we show that $k_t(\theta', e'_{1:t}; \theta, e_{1:t-1}) > 0 \iff \ell_t(\theta, e_{1:t-1}; \theta', e'_{1:t}) > 0$ for all $t \in \{2 \dots T\}$. This follows from the detailed balance requirement (Equation 5) and from the fact that $p_t(\theta, e_{1:t}) > 0$ for all $(\theta, e_{1:t}) \in \mathcal{X}_t$ and for all $t \in \{2 \dots T\}$. The same argument applies to k_{T+1} and ℓ_{T+1} . Note that we do not require the detailed balance kernels to be ergodic. For example, a given kernel d_t may only update one of the components of $(\theta, e_{1:t})$. Given these definitions, the weight functions become:

$$w_1(\theta, e_1) := \frac{\tilde{p}_1(\theta, e_1)}{k_1(\theta, e_1)} = \frac{p(\theta, e_1, y_1)}{p(\theta, e_1)} = p(y_1 | \theta, e_1) \quad (7)$$

$$w_t((\theta, e_{1:t-1}), (\theta', e'_{1:t})) := \frac{\tilde{p}_t(\theta', e'_{1:t})}{\tilde{p}_{t-1}(\theta, e_{1:t-1})} \frac{\ell_t(\theta, e_{1:t-1}; \theta', e'_{1:t})}{k_t(\theta', e'_{1:t}; \theta, e_{1:t-1})} \quad (8)$$

$$= \frac{p(\theta', e'_{1:t}, y_{1:t})}{p(\theta, e_{1:t-1}, y_{1:t-1})} \frac{d_{t-1}(\theta, e_{1:t-1}; \theta', e'_{1:t-1})}{d_{t-1}(\theta', e'_{1:t-1}; \theta, e_{1:t-1}) p(e'_t | \theta', e'_{1:t-1}, y_{1:t-1})} \quad (9)$$

$$(10)$$

Then by detailed balance:

$$\frac{d_{t-1}(\theta, e_{1:t-1}; \theta', e'_{1:t-1})}{d_{t-1}(\theta', e'_{1:t-1}; \theta, e_{1:t-1})} = \frac{\tilde{p}_{t-1}(\theta, e_{1:t-1})}{\tilde{p}_{t-1}(\theta', e'_{1:t-1})} = \frac{p(\theta, e_{1:t-1}, y_{1:t-1})}{p(\theta', e'_{1:t-1}, y_{1:t-1})} \quad (11)$$

$$w_t((\theta, e_{1:t-1}), (\theta', e'_{1:t})) = \frac{p(\theta', e'_{1:t}, y_{1:t})}{p(\theta, e_{1:t-1}, y_{1:t-1})} \frac{p(\theta, e_{1:t-1}, y_{1:t-1})}{p(\theta', e'_{1:t-1}, y_{1:t-1}) p(e'_t | \theta', e'_{1:t-1}, y_{1:t-1})} \quad (12)$$

$$= \frac{p(\theta', e'_{1:t}, y_{1:t})}{p(\theta, e_{1:t-1}, y_{1:t-1})} \frac{p(\theta, e_{1:t-1}, y_{1:t-1})}{p(\theta', e'_{1:t}, y_{1:t-1})} = \frac{p(\theta', e'_{1:t}, y_{1:t})}{p(\theta', e'_{1:t}, y_{1:t-1})} \quad (13)$$

$$= p(y_t | \theta', e'_{1:t}, y_{1:t-1}) \quad (14)$$

Finally, $w_{T+1}((\theta, e_{1:T}), (\theta', e'_{1:T})) = 1/p(\theta', e'_{1:T}, y_{1:T})$. Algorithm 3 shows SIMULATE and REGENERATE specialized for sequential observation and detailed balance kernels, as used in the experiments. In Algorithm 3, parenthesized superscripts indicate the step t of the SMC algorithm, whereas subscripts indicate observation indices (e.g. $e_{t-1}^{i(t)}$ is the value of local latents for observation $t-1$ in particle i at step t of SMC).

Algorithm 3 SIMULATE and REGENERATE for SMC with sequential observation and detailed balance

```

procedure SIMULATE( )
  for  $i \leftarrow 1 \dots N$  do
     $\theta^{i(1)} \sim p(\theta)$  ▷ Sample global latents  $\theta$  from the prior
     $e_1^{i(1)} \sim p(e_1|\theta)$  ▷ Sample local latents for observation 1 by forward sampling in the model
     $w_1^i \leftarrow p(y_1|\theta^{i(1)}, e_1^{i(1)})$  ▷ Evaluate likelihood of  $y_1$ 
  end for
  for  $t \leftarrow 2 \dots T$  do
    for  $i \leftarrow 1 \dots N$  do
       $a_{t-1}^i \sim \text{Categorical}(\text{NORMALIZE}(\mathbf{w}_{t-1}))$  ▷ Sample the index of the parent particle
       $\theta^{i(t)}, e_{1:t-1}^{i(t)} \sim d_{t-1}(\cdot; \theta^{a_{t-1}^i(t-1)}, e_{1:t-1}^{a_{t-1}^i(t-1)})$  ▷ Detailed balance targeting  $p(\theta, e_{1:t-1}|y_{1:t-1})$ 
       $e_t^{i(t)} \sim p(e_t|\theta^{i(t)}, e_{1:t-1}^{i(t)})$  ▷ Sample local latents for observation  $t$  by forward sampling in the model
       $w_t^i \leftarrow p(y_t|\theta^{i(t)}, e_{1:t}^{i(t)})$  ▷ Evaluate likelihood of  $y_t$ 
    end for
  end for
   $I_T \sim \text{Categorical}(\text{NORMALIZE}(\mathbf{w}_T))$  ▷ Sample output particle index
   $\theta, e_{1:T} \sim d_T(\cdot; \theta^{I_T(T)}, e_{1:T}^{I_T(T)})$  ▷ Detailed balance targeting  $p(\theta, e_{1:T}|y_{1:T})$ 
  return  $\left( (\theta, e_{1:T}), \log \frac{p(\theta, e_{1:T}, y_{1:T})}{\prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j} \right)$  ▷ Return output sample and log-weight
end procedure

procedure REGENERATE( $(\theta, e_{1:T})$ )
   $(I_1, \dots, I_T) \sim \text{RAND-ANCESTRY}(N, T)$ 
   $\theta^{I_T(T)}, e_{1:T}^{I_T(T)} \sim d_T(\cdot; \theta, e_{1:T})$  ▷ Detailed balance targeting  $p(\theta, e_{1:T}|y_{1:T})$ 
  for  $t \leftarrow T-1 \dots 1$  do
     $\theta^{I_t(t)}, e_{1:t}^{I_t(t)} \sim d_t(\cdot; \theta^{I_{t+1}(t+1)}, e_{1:t}^{I_{t+1}(t+1)})$  ▷ Detailed balance targeting  $p(\theta, e_{1:t}|y_{1:t})$ 
  end for
  for  $i \leftarrow 1 \dots N$  do
    if  $i \neq I_1$  then
       $\theta^{i(1)} \sim p(\theta)$  ▷ Sample global latents  $\theta$  from the prior
       $e_1^{i(1)} \sim p(e_1|\theta)$  ▷ Sample local latents for observation 1 by forward sampling in the model
    end if
     $w_1^i \leftarrow p(y_1|\theta^{i(1)}, e_1^{i(1)})$  ▷ Evaluate likelihood of  $y_1$ 
  end for
  for  $t \leftarrow 2 \dots T$  do
    for  $i \leftarrow 1 \dots N$  do
      if  $i = I_t$  then
         $a_{t-1}^i \leftarrow I_{t-1}$ 
      else
         $a_{t-1}^i \sim \text{Categorical}(\text{NORMALIZE}(\mathbf{w}_{t-1}))$  ▷ Sample the index of the parent particle
         $\theta^{i(t)}, e_{1:t-1}^{i(t)} \sim d_{t-1}(\cdot; \theta^{a_{t-1}^i(t-1)}, e_{1:t-1}^{a_{t-1}^i(t-1)})$  ▷ Detailed balance targeting  $p(\theta, e_{1:t-1}|y_{1:t-1})$ 
         $e_t^{i(t)} \sim p(e_t|\theta^{i(t)}, e_{1:t-1}^{i(t)})$  ▷ Sample local latents for observation  $t$  by forward sampling in the model
      end if
       $w_t^i \leftarrow p(y_t|\theta^{i(t)}, e_{1:t}^{i(t)})$  ▷ Evaluate likelihood of  $y_t$ 
    end for
  end for
  return  $\log \frac{p(\theta, e_{1:T}, y_{1:T})}{\prod_{t=1}^T \frac{1}{N} \sum_{j=1}^N w_t^j}$  ▷ Return the log-weight
end procedure

```

Appendix C: Using cycles of detailed balance kernels

Recall that we did not require the detailed balance kernels d_t for $t \in \{1 \dots T\}$ to be ergodic. In particular, each d_t can update only a subset of the random variables in $(\theta, e_{1:t}) \in \mathcal{X}_t$. We now show that Algorithm 3 can be used without modification when the kernels k_t for $t \in \{2 \dots T-1\}$ utilize instead *cycles* of detailed balance kernels each targeting the same distribution $p_{t-1}(\theta, e_{1:t-1})$, provided the corresponding kernels ℓ_t sample from the same cycle *in reverse order*. Note that the cycle of detailed balance kernels may not itself satisfy detailed balance.

For some $1 < r \leq s < T$, suppose that $\mathcal{X}_{r-1} = \mathcal{X}_r = \dots = \mathcal{X}_s$, and $\tilde{p}_{r-1} = \tilde{p}_r = \dots = \tilde{p}_s$ (meaning the target distributions do not change from $t = r-1$ through $t = s$). Suppose $k_t(x'; x) = d_{t-1}(x'; x)$ and $\ell_t(x'; x) = d_{t-1}(x'; x)$ for $x, x' \in \mathcal{X}_{r-1}$ for $t \in \{r \dots s\}$ where d_{t-1} is a detailed balance kernel targeting \tilde{p}_{r-1} , for $r \leq t \leq s$. Then for $r \leq t \leq s$ the weights are:

$$w_t(x_{t-1}, x_t) = \frac{\tilde{p}_{r-1}(x_t)}{\tilde{p}_{r-1}(x_{t-1})} \frac{d_{t-1}(x_{t-1}; x_t)}{d_{t-1}(x_t; x_{t-1})} = \frac{\tilde{p}_{r-1}(x_t)}{\tilde{p}_{r-1}(x_{t-1})} \frac{\tilde{p}_{r-1}(x_{t-1})}{\tilde{p}_{r-1}(x_t)} = 1$$

Consider modifying the SIMULATE and REGENERATE procedures to replace $a_{t-1}^i \sim \text{Categorical}(\text{NORMALIZE}(\mathbf{w}_{t-1}))$ with $a_{t-1}^i \leftarrow i$ for $r+1 \leq t \leq s+1$, and modifying RAND-ANCESTRY (used by REGENERATE) to replace $I_t \sim \text{Uniform}(1 \dots N)$ with $I_t \leftarrow I_{t-1}$ for $r+1 \leq t \leq s+1$. The joint probability for SIMULATE is then divided by the probability of the excluded random choices, which is $1/N^{N(s-r+1)}$ since each weight w_t^i is deterministically 1 for $r \leq t \leq s$. The joint probability for REGENERATE is then divided by the probability of the excluded random choices, which is $1/N^{(N-1)(s-r+1)} \cdot 1/N^{s-r+1} = 1/N^{N(s-r+1)}$. The weight expression is therefore unchanged, but can be simplified (because weights w_t^i for $r \leq t \leq s$ are deterministically 1) to:

$$\frac{p(u, z)}{q(u; z)} = \frac{1}{w_{T+1}^1 \left(\prod_{t=1}^{r-1} \frac{1}{N} \sum_{j=1}^N w_t^j \right) \cdot \left(\prod_{t=s+1}^T \frac{1}{N} \sum_{j=1}^N w_t^j \right)}$$

In a concise implementation of these modified procedures, the steps r through $s+1$ are collapsed into one step, with the cycle of detailed balance kernels in $(k_r = d_{r-1}, \dots, k_s = d_{s-1}, k_{s+1})$ taking the role of a single k_t in SIMULATE and REGENERATE and the reverse cycle in $(\ell_{s+1}, \ell_s = d_{s-1}, \dots, \ell_r = d_{r-1})$ taking the role of the corresponding ℓ_t in REGENERATE.

To see that k_{T+1} and ℓ_{T+1} can also be replaced with cycles of detailed balance kernels, consider introducing new random variables $z_1, \dots, z_R \in \mathcal{X}^T$ for some $R \geq 1$ into both SIMULATE and REGENERATE as follows: In the joint probability expression $p(u, z)$ replace $k_{T+1}(z; x_T^I)$ with:

$$k_{T+1}(z_1; x_T^I) \prod_{r=2}^R k_{T+r}(z_r; z_{r-1}) k_{T+R+1}(z; z_R)$$

Let $k_{T+1}, \dots, k_{T+R+1}$ be detailed balance kernels targeting \tilde{p}_T . This corresponds to applying a sequence of detailed balance kernels immediately prior to returning the output z in SIMULATE. Similarly, in the joint probability expression $q(u; z)$, replace $\ell_{T+1}(x_T^I; z)$ with:

$$k_{T+1}(x_T^I; z_1) \prod_{r=2}^R k_{T+r}(z_{r-1}; z_r) k_{T+R+1}(z_R; z)$$

This corresponds to applying the same cycle of detailed balance kernels to the input of REGENERATE that were applied before the output of SIMULATE, but in the reverse order. The new log weight is then adjusted by a factor of:

$$\begin{aligned} & \frac{k_{T+1}(z_1; x_T^I) \prod_{r=2}^R k_{T+r}(z_r; z_{r-1}) k_{T+R+1}(z; z_R) \ell_{T+1}(x_T^I; z)}{k_{T+1}(x_T^I; z_1) \prod_{r=2}^R k_{T+r}(z_{r-1}; z_r) k_{T+R+1}(z_R; z) k_{T+1}(z; x_T^I)} \\ &= \frac{\tilde{p}_T(z_1)}{\tilde{p}_T(x_T^I)} \prod_{r=2}^R \frac{\tilde{p}_T(z_r)}{\tilde{p}_T(z_{r-1})} \frac{\tilde{p}_T(z)}{\tilde{p}_T(z_R)} \frac{\tilde{p}_T(x_T^I)}{\tilde{p}_T(z)} = 1 \end{aligned}$$

Therefore, k_{T+1} and ℓ_{T+1} can be replaced with a cycle of detailed balance kernels targeting p_T and the reversed cycle, respectively, without modifying the expression for the returned log-weight value, which is still computed using $w_{T+1}^1 = w_{T+1}(x_T^I, z)$, and does not depend on the intermediate values z_1, \dots, z_R .