KL/TV Resampling: Statistical Distance Based Offspring Selection in SMC Methods

Oskar Kviman, Hazal Koptagel, Harald Melin and Jens Lagergren
Correspondence: okviman@kth.se

Introduction

The research interest in combining the **variational inference (VI)** and sequential Monte Carlo (SMC) methodologies is rapidly increasing, especially in domains allowing for model parameter learning via amortized VI [1, 2, 3]. Yet, **utilizing VI in order to directly design resampling/offspring selection schemes** – one of the defining elements of particle filters (PFs) – has, to the best of our knowledge, received little attention so far.

In light of this, we propose two novel offspring selection schemes which multiply/discard particles in order to minimize the **Kullback-Leibler (KL) divergence** or the **total variation (TV) distance** with respect to the real-valued particle distribution (prior to resampling). The reference distribution can either be the rational-valued particle distribution (post resampling), or the model's target distribution. By **regarding offspring selection as a problem of minimizing statistical distances**, we further bridge the gap between optimization-based density estimation and SMC theory.

Our proposed methods **outperform or compare favourably** with the multinomial, systematic and stratified resampling schemes on common density-estimation benchmark datasets in terms of estimating the true sequence of latent variables via the resulting smoothing density.

Background

We are interested in approximating the intractable posterior $p(x_{1:n}|y_{1:n})$, where $x_{1:n}$ and $y_{1:n}$ are sequences of latent variables and observations of length n, respectively.

To do this, we employ a PF which provides an approximation of the posterior in the form of a particle distribution

$$q(x_{1:n}|y_{1:n}) = \sum_{s=1}^{S} w_n^s \delta_{x_{1:n}^s}(x_{1:n})$$

Above we let S denote the number of particles, $\delta_{x_{1:n}^s}(x_{1:n})$ as the Dirac distribution with mass in $x_{1:n}^s$ and w_n^s the normalized importance weight of particle s at time n. We obtain w_n^s by normalizing the unnormalized importance weight

$$\tilde{w}_n^s = \frac{p(y_n|x_n^s)p(x_n^s|x_{n-1}^{a_s})}{\pi(x_n^s|x_{n-1}^{a_s})}$$

over the S particles. In the above equation we note that the ratio between the model's target distribution and the proposal density is not multiplied by $\tilde{w}_n^s$ as the focus of our study is on the case where resampling occurs at every time step. Finally, a_s is the index of the sth particle's ancestor.

Problem Setting

Resampling, or offspring selection, is a methodology for discarding less promising particles in favour of particles which we wish to multiply. Designing resampling methods means modeling this selection process. Namely, we should decide how to choose the rational-valued particle distribution

$$\bar{q}(x_{1:n}|y_{1:n}) = \sum_s \frac{m_s}{S} \delta_{x_{1:n}^s}(x_{1:n})$$

where m_s is the multiplicity of particle s, given the real-valued particle distribution

$$\tilde{q}(x_{1:n}|y_{1:n}) = \sum_s \tilde{w}_n^s \delta_{x_{1:n}^s}(x_{1:n})$$

As an example, in multinomial resampling this is done by first normalizing $\tilde{q}(x_{1:n}|y_{1:n})$ and then sampling S offspring indices from the resulting categorical distribution with replacement.

References

[1] Naesseth, Christian, et al. "Variational sequential monte carlo." International conference on artificial intelligence and statistics. PMLR, 2018.
[2] Corenflos, Adrien, et al. "Differentiable particle filtering via entropy-regularized optimal transport." International Conference on Machine Learning. PMLR, 2021.
[3] Matthews, Alexander GDG, et al. "Continual Repeated Annealed Flow Transport Monte Carlo." arXiv preprint arXiv:2201.13117 (2022).

KL Divergence and ELBO

Let $q(x|y)$ be the variational distribution and $r(x|y)$ be the reference distribution we wish to approximate. Then the KL divergence is defined as

$$\mathrm{KL}(q(x|y)\|r(x|y)) = \mathbb{E}_{q(x|y)}\left[\log \frac{q(x|y)}{r(x|y)}\right]$$

The KL divergence is non-negative, and zero if and only if $q(x|y) = r(x|y)$. Although $r(x|y)$ might be intractable, minimizing $\mathrm{KL}(q(x|y)\|r(x|y))$ is apparently equivalent to maximizing the evidence lower bound (ELBO)

$$\mathcal{L}(q,r,y) = \mathbb{E}_{q(x|y)}\left[\log \frac{r(y|x)r(x)}{q(x|y)}\right]$$

where $r(x,y)$ is tractable.

TV Distance

Let Q and R be discrete conditional probability measures on $\mathcal{X}$, then the TV distance is defined as

$$\mathrm{TV}(Q,R) = \frac{1}{2}\sum_{x\in\mathcal{X}} |q(x|y) - r(x|y)|$$

SV Model

Here we consider the stochastic volatility (SV) model

$$\mu(x_1) = \mathcal{N}(0, \frac{\sigma^2}{1-\phi^2})$$
$$X_n = \phi X_{n-1} + \sigma V_n$$
$$Y_n = \beta \exp(X_n/2) R_n$$

where R_n and V_n are random variables following the standard normal distribution. All experiments are averaged over five random seeds and we fixed $S = 500$.

In order to estimate the true (generating) latent variable x_n, we use the expected value, the median and the mode of the smoothing density $q(x_n|y_{1:N})$.

KL Resampling

Let the variational distribution be the rational-valued particle distribution, then we get

$$\begin{aligned}
\mathcal{L}(\bar{q},p,y_{1:n}) &= \sum_{s=1}^{S} \bar{q}(x_{1:n}^s)\log\frac{p(x_{1:n}^s, y_{1:n})}{\bar{q}(x_{1:n}^s)} \\
&= \sum_{s=1}^{S} \frac{m_s}{S}\log\frac{p(x_{1:n}^s, y_{1:n})}{\frac{m_s}{S}} \\
&\propto \sum_{s=1}^{S} m_s \log\frac{p(x_{1:n}^s, y_{1:n})}{m_s}
\end{aligned}$$

The multiplicities that maximize the ELBO constitute the rational-valued particle distribution that we use post resampling. In Algorithm 1, $u^s = p(x_{1:n}^s, y_{1:n})$.

Algorithm 1 KL Resampling

Input: $\{u^s\}_{s=1}^S$
Result: $\{a_s\}_{s=1}^S$

order $u^1 \geq \cdots \geq u^S$
set $m_s = 0, \forall s$
define $f(m,s) = m_s \log \frac{u^s}{m_s}$
define $C^+(m,s) = f(m+1,s) - f(m,s)$

while $\sum_{s=1}^S m_s < S$ **do**
  compute $t = argmax_s\ C^+(m,s)$
  set $m_t = m_t + 1$
end
return $\{a_s\}_{s=1}^S$

TV Resampling

Here we let the reference distribution be the normalized real-valued particle distribution, and aim to choose $\bar{q}(x_{1:n}|y_{1:n})$ such that we minimize

$$\begin{aligned}
\mathrm{TV}(\bar{Q},Q) &= \frac{1}{2}\sum_{s=1}^{S} |\bar{q}(x_{1:n}^s|y_{1:n}) - q(x_{1:n}^s|y_{1:n})| \\
&= \frac{1}{2}\sum_s \left|\frac{m_s}{S} - w_n^s\right|
\end{aligned}$$

The multiplicities that minimize the TV distance form the rational-valued particle distribution that we use post resampling. In Algorithm 2, $u^s = w_n^s$.

Algorithm 2 TV Resampling

Input: $\{u^s\}_{s=1}^S$
Result: $\{a_s\}_{s=1}^S$

set $m_s = u_s S, \forall s$
set $\alpha^s = m_s - \lfloor m_s \rfloor, \forall s$
compute $\alpha = \sum_s \alpha^s$
order $\alpha^1 \geq \cdots \geq \alpha^S$
for $s = 1, ..., S$ **do**
  if $s \leq \alpha$
    set $m_s \leftarrow \lceil m_s \rceil$
  else $s \leq \alpha$
    set $m_s \leftarrow \lfloor m_s \rfloor$
  end
end
return $\{a_s\}_{s=1}^S$

Experimental Results