Continual Repeated Annealed Flow Transport Monte Carlo (CRAFT)

Alex Matthews

DeepMind

Michael Arbel

INRIA

Univ. Grenoble Alpes

Danilo J. Rezende

DeepMind

Arnaud Doucet

DeepMind

github.com/deepmind/annealed_flow_transport

Combining SMC samplers with normalizing flows

Normalizing flows → CRAFT ← SMC samplers

Fast
Incorporate symmetries

Parameter intensive
Topologically constrained

MCMC sampling
Annealing
Resampling

Importance sampling

Significantly improves on prior work

Annealed Flow Transport (AFT)

Arbel et al. 2021

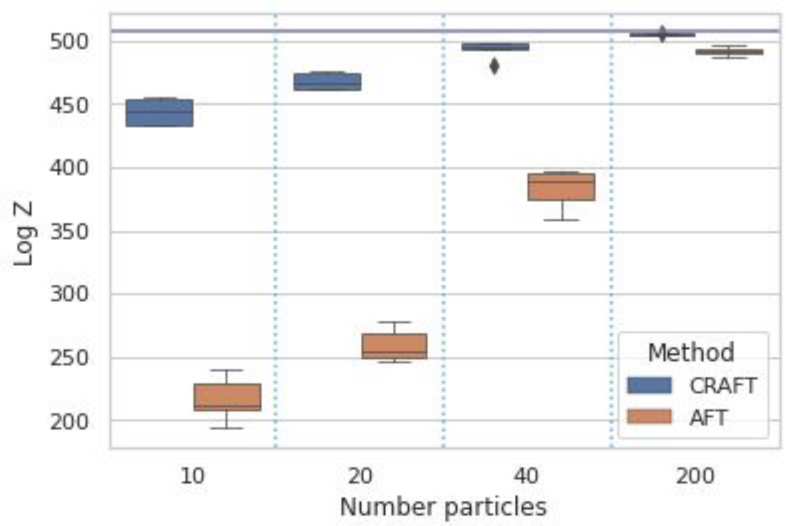

Stochastic Normalizing Flows (SNF)

Wu et al. 2020

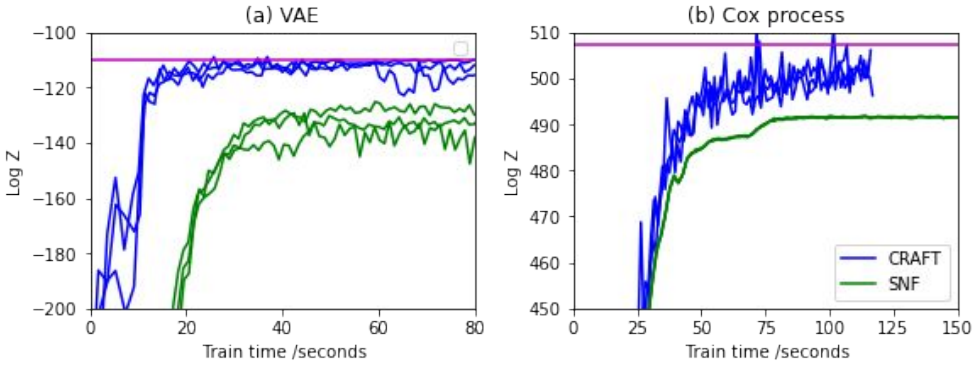

There are good conceptual reasons for improvements.

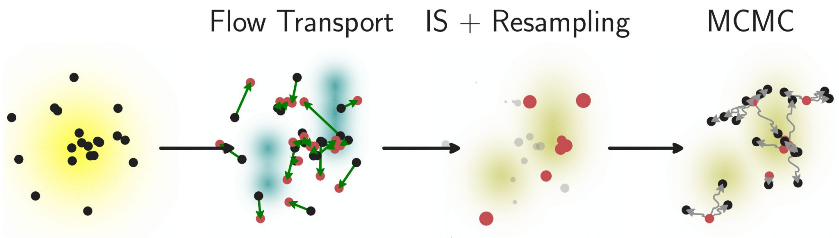

For fixed/learnt flows we have a sequence of steps between annealing distributions.

Algorithm 1 SMC-NF-step

Input: Approximations (π_{k-1}^N, Z_{k-1}^N) to (π_{k-1}, Z_{k-1}), normalizing flow T_k, unnormalized annealed targets γ_{k-1} and γ_k and resampling threshold $A \in [1/N, 1)$.
Output: Particles at iteration k: $\pi_k^N = (X_k^i, W_k^i)_{i=1}^N$, approximation Z_k^N to Z_k.
Transport particles: $Y_k^i = T_k(X_{k-1}^i)$.
Compute IS weights:
  $w_k^i \leftarrow W_{k-1}^i G_k(X_{k-1}^i)$ // unnormalized
  $W_k^i \leftarrow w_k^i / \sum_{j=1}^N w_k^j$ // normalized
Estimate normalizing constant Z_k:
  $Z_k^N \leftarrow Z_{k-1}^N (\sum_{i=1}^N w_k^i)$.
Compute effective sample size ESS_k^N.
if $\mathrm{ESS}_k^N \leq NA$ **then**
  Resample N particles denoted abusively also Y_k^i according to the weights W_k^i, then set $W_k^i = \frac{1}{N}$.
end if
Sample $X_k^i \sim \mathcal{K}_k(Y_k^i, \cdot)$. // MCMC
Return (π_k^N, Z_k^N).

Algorithm 3 CRAFT-deployment

Input: Fixed/trained NFs $\{T_k\}_{k=1}^K$, number of particles N, unnormalized annealed targets $\{\gamma_k\}_{k=0}^K$ with $\gamma_0 = \pi_0$, resampling threshold $A \in [1/K, 1)$.
Output: Approximations (π_K^N, Z_K^N) to (π, Z).

Sample $X_0^i \sim \pi_0$ and set $W_0^i = \frac{1}{N}$ and $Z_0^N = 1$.
for $k = 1, \ldots, K$ **do**
  $(\pi_k^N, Z_k^N) \leftarrow \texttt{SMC-NF-step}\,(\pi_{k-1}^N, Z_{k-1}^N, T_k)$
end for
Return (π_K^N, Z_K^N).

See also our AFT talk from ICML 2021.

CRAFT flow training objective

Previous distribution passed through a flow *T*

Current distribution

Zero if flow transport perfect.

$$H = \sum_{l=1}^{K} \mathcal{KL}[T_{\#}^{(l)} \pi_{l-1} || \pi_l]$$

Sum over transitions between temperatures.

Estimate objective and gradients using current importance sampling estimate.

Gradients local to transitions so no need to backprop through discrete steps.

Further analysis in paper.

Train particle MCMC samplers for hard physics problems

Use CRAFT to generate particle MCMC proposals.

Test on ϕ^4 -theory near criticality.

Can incorporate symmetries of the problem into the flow.

CRAFT outperforms strong baselines.