HKU Musketeers Foundation
Institute of Data Science
香港大學同心基金數據科學研究院

Deep Learning
Methods for Sampling

Jianrui Wu
jrwu@connect.hku.hk

RPg Student Journal Club
Apr 25, 2025

The Problem

To Sample from...

- $\pi(x) = \frac{\gamma(x)}{Z}, Z = \int \gamma(x) \mathrm{d}x$
  - $\gamma(x)$ is unnormalized; Z is intractable
  - Called a Boltzmann distribution when $\gamma(x) = e^{-f(x)}$

Is It Useful?

Applications

- Bayesian inference
  - $\pi(\theta|x) \propto \pi(\theta)p(x|\theta)$
- Optimization problems
- Physics simulations

A Bit of History

Monte Carlo (MC)

- Markov chain Monte Carlo (MCMC)
  - Unadjusted Langevin Algorithm (ULA), variant of Metropolis-Hastings that uses gradient
    - $x_{t+1} = x_t - \eta_t \nabla f(x_t) + \sqrt{2\eta_t}\varepsilon_t, \varepsilon_t \sim \mathcal{N}(0, I)$
- Annealed Importance Sampling (AIS), special case of Sequential Monte Carlo (SMC)
  - Forward kernel F_k invariant w.r.t. $\pi_k \propto \gamma_k = \pi_0^{1-\beta_k}\gamma^{\beta_k}$; backward kernel B_k inverts F_{k+1}
  - Sample from π_0 to π_k as the extended proposal, and then back to π_0 again as the extended target
  - $w(x_{0:K}) = \frac{\Gamma(x_{0:K})}{Q(x_{0:K})}$ is an unbiased estimator of Z

Variational Inference (VI)

- Approximating p with q by minimizing reverse KL
  - Equivalent to maximizing ELBO
- Often extended by introducing an auxiliary variable, z
  - $q(x,z) = q(x|z)q(z), D_{\text{KL}}(q(x)\|\pi(x)) \leq D_{\text{KL}}(q(x,z)\|p(z|x)\pi(x)) = -\mathbb{E}_{q(x,z)}\left[\log \frac{p(z|x)\gamma(x)}{q(x,z)}\right] + \log Z$
  - Gaussian Mixture Models (GMM) commonly used for $q(x, z)$

Evaluation

Estimates of Normalizing Constant

- Importance weight $w(x) = \frac{\gamma(x)}{q(x)}$
- ELBO $\mathbb{E}_q\left[\log w\right]$, lower bound

Mode Collapse

- Reverse KL is zero-forcing
  - Distributions in the wild are often highly non-convex and multi-modal
  - Forward KL is not accessible in training
- Forward criteria for evaluation
  - Benchmark the methods on known, tractable distributions
  - $D_{\mathrm{KL}}(\pi \parallel q) = \mathbb{E}_\pi[\log w] - \log Z$

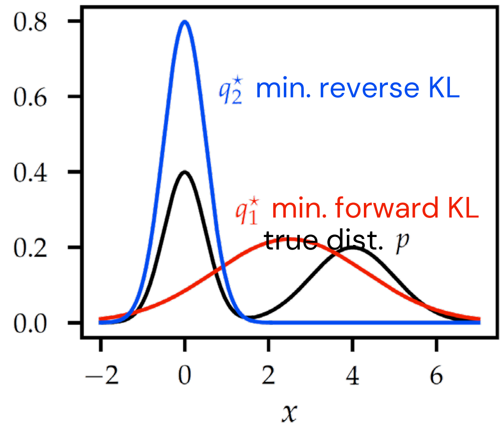

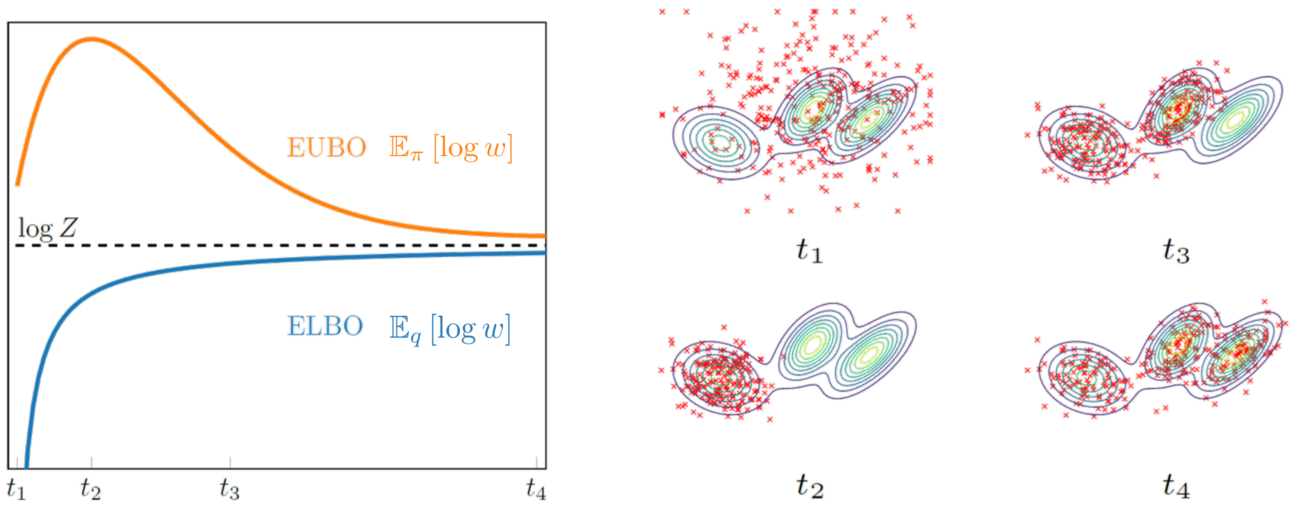

Blessing et al., "Beyond ELBOs."
Figure comparing forward/reverse KL: Krause and Hübotter, "Probabilistic Artificial Intelligence."

Recent Developments

Combining MC and VI

- AIS's proposal & target paradigm
  - Replace backward/forward kernels with DL components
- VI's loss function, reverse KL

Recent Developments

Monte Carlo Diffusion (MCD)

Motivation

- AIS is suboptimal in terms of variance; yet the optimal B_k is not directly available
  - $B_k^{\text{opt}}(x_k \mid x_{k+1}) = \frac{q_k(x_k)F_{k+1}(x_{k+1}|x_k)}{q_{k+1}(x_{k+1})}$

Method

- Choose F_k to be ULA's (not learnable)
  - ULA is just discretized diffusion
- Score matching applicable for B_k
- Minimize $D_{\text{KL}}(Q||P_\vartheta)$ w.r.t. ϑ

Algorithm 1 Unadjusted Langevin AIS/MCD – red instructions for AIS and blue for MCD

Require: Unnormalized target $\gamma(x)$, initial state proposal $\pi_0(x)$, number steps K, stepsize δ, annealing schedule $\{\beta_k\}_{k=0}^K$, score model $s_\theta(k,x)$
Sample $x_0 \sim \pi_0(x_0)$
Set $\log w = -\log \pi_0(x_0)$
for $k = 1$ to K **do**
  Define $\log \gamma_k(\cdot) = \beta_k \log \gamma(\cdot) + (1-\beta_k)\log \pi_0(\cdot)$
  Define $F_k(x_k|x_{k-1}) = \mathcal{N}(x_k; x_{k-1} + \delta\nabla \log \gamma_k(x_{k-1}), 2\delta I)$
  Sample $x_k \sim F_k(\cdot|x_{k-1})$
  Define $B_{k-1}(x_{k-1}|x_k) = F_k(x_{k-1}|x_k)$ ▷ AIS
  Define $B_{k-1}(x_{k-1}|x_k) = \mathcal{N}(x_{k-1}; x_k - \delta\nabla \log \gamma_k(x_k) + 2\delta s_\theta(k, x_k), 2\delta I)$ ▷ MCD
  Set $\log w = \log w + \log B_{k-1}(x_{k-1}|x_k) - \log F_k(x_k|x_{k-1})$
end for
Set $\log w = \log w + \log \gamma(x_K)$

$$X_0 \sim \pi_0 \longrightarrow \mathrm{d}X_t = \nabla \log \pi_t(X_t)\mathrm{d}t + \sqrt{2}\mathrm{d}B_t \longrightarrow X_T \approx \pi$$

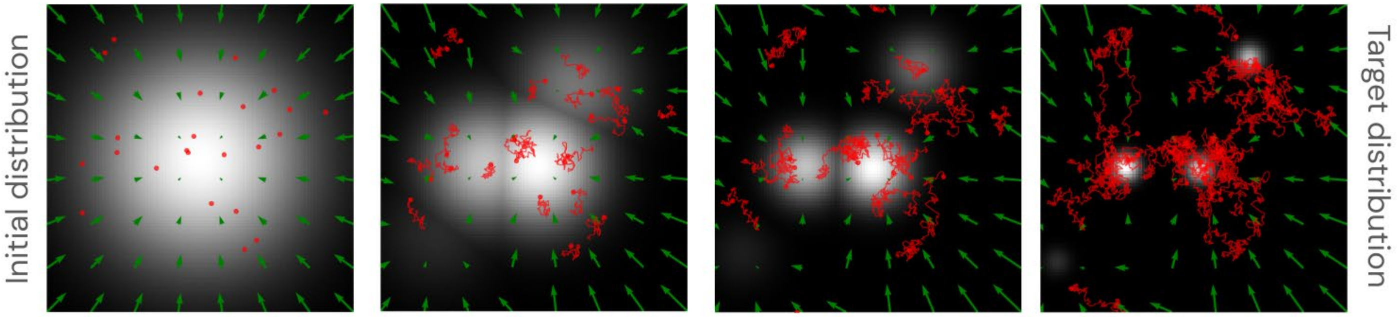

$$\bar{X}_T \approx \pi_0 \longleftarrow \mathrm{d}\bar{X}_t = \{-\nabla \log \pi_{T-t}(\bar{X}_t) + 2\nabla \log q_{T-t}(\bar{X}_t)\}\,\mathrm{d}t + \sqrt{2}\mathrm{d}\bar{B}_t \longleftarrow \bar{X}_0 \sim \pi$$

Doucet et al., "Score-Based Diffusion Meets Annealed Importance Sampling."

Sequential Controlled Langevin Diffusions (SCLD)

HKU Musketeers Foundation Institute of Data Science 香港大學同心基金數據科學研究院

Motivation
- Model forward process with DL
- Btw: KL has high variance

Method
- Continuous time formulization
  - $w = \log \frac{d\mathbb{P}}{d\mathbb{Q}}$ computed with RND
  - Can be refined with MCMC any time
- Log-variance divergence as loss
  - $D_{\mathrm{LV}}^{\widetilde{\mathbb{P}}}(\mathbb{P}, \mathbb{Q}) = \mathrm{Var}_{\widetilde{\mathbb{P}}}[w]$
  - Reference measure can use arbitrary control, allowing "off-policy training"

Algorithm 1 Sequential Controlled Langevin Diffusion (SCLD). ▷ See Algorithm 3 for details.

Require: Annealing path π, learned control u, time grid $0 = t_0 < \cdots < t_N = T$

1: *Initialize:* $\overline{X}_0 := X_0^{(1:K)} \sim p_{\mathrm{prior}}$ and $\overline{w}_0 := w_0^{(1:K)} = 1$
2: **for** $n = 1$ to $n = N$ **do**
3: *Transport:* $\overline{X}_{[t_{n-1},t_n]} = \mathtt{simulate\_SDE}(\overline{X}_{t_{n-1}}, u)$ ▷ See (6) and (19)
4: *Compute RNDs:* $\overline{w}_{[t_{n-1},t_n]} = \frac{\mathrm{d}\vec{\mathbb{P}}_{[t_{n-1},t_n]}}{\mathrm{d}\overleftarrow{\mathbb{P}}_{[t_{n-1},t_n]}}(\overline{X}_{[t_{n-1},t_n]})$ ▷ See (12) and (30)
5: *Update weights:* $\overline{w}_n = \overline{w}_{n-1}\overline{w}_{[t_{n-1},t_n]}$ ▷ See (13)
6: *Resample:* $\overline{X}_{t_n}, \overline{w}_n = \mathtt{resample}(\overline{X}_{t_n}, \overline{w}_n)$ ▷ See Algorithm 5
7: **return** Samples $\overline{X}_T := X_T^{(1:K)}$ approximately from p_{target}

"control"; u is parameterized with NN

$$\mathrm{d}X_t^u = u(X_t^u, t)\mathrm{d}t + \sigma(t)\vec{\mathrm{d}}W_t, \quad X_0^u \sim p_{\mathrm{prior}}$$

$$\mathrm{d}Y_t^v = v(Y_t^v, t)\mathrm{d}t + \sigma(t)\overleftarrow{\mathrm{d}}W_t, \quad Y_T^v \sim p_{\mathrm{target}}$$

$p_{X^u} = p_{Y^v} = \pi$

$(u - \sigma^2 \nabla \log \pi)(Y_t^u, t)\,\mathrm{d}t$ pins down unique u^*

Chen et al., "Sequential Controlled Langevin Diffusions."
Log-variance divergence as loss: Richter and Berner, "Improved Sampling via Learned Diffusions."

Further Reading

Reviews

- Blessing et al., "Beyond ELBOs."
- He et al., "No Trick, No Treat."
- A lot more papers: https://github.com/J-zin/Awesome-Diffuison-Flow-Samplers

Not Covered Today

- GFlowNet: https://milayb.notion.site/The-GFlowNet-Tutorial-95434ef0e2d94c24aab90e69b30be9b3
- VAN: Wu, Wang, and Zhang, "Solving Statistical Mechanics Using Variational Autoregressive Networks."