Bayesian Methods in Machine Learning
Seminar: 8

Evgenii Egorov, evgenyy.egorov@ya.ru

Skoltech

September 18, 2020

Stochastic Optimization 1D

Consider problem:

$$p(x) = \mathcal{N}(x; \mu, 1),$$
$$\min_{\mu} \langle x^2 \rangle_{p(x;\mu)}.$$

It is easy to see analytic solution: $\mu = 0$.
However, we consider the problem as a toy stochastic optimization problem, to compare variance of gradients:

1. Reinforce
2. Re-parametrization

Stochastic Optimization 1D: REINFORCE

Recall Reinforce:

$$\nabla_\theta \langle f(x) \rangle_{p(x;\theta)} = \langle f(x) \nabla \log p(x;\theta) \rangle_{p(x;\theta)}.$$

$$p(x) = \mathcal{N}(x; \mu, 1),$$
$$\min_\mu \langle x^2 \rangle_{p(x;\mu)}.$$

Problem:

1. Obtain 1-sample estimator of gradient

Stochastic Optimization 1D: REINFORCE

Estimator:

$$\nabla_\mu \langle x^2 \rangle_{p(x;\mu)} \approx (x_k - \mu) x_k^2, x_k \sim \mathcal{N}(x; \mu, 1).$$

Problem:

1. Find the variance of the gradient estimator
2. Optimize the variance of estimator over the constant baseline λ

Estimator with baseline:

$$(x_k - \mu)(x_k^2 + \lambda).$$

Useful:

$$x \sim \mathcal{N}(x|\vec{0}, \sigma^2),$$
$$\mathbb{E}x^p = \begin{cases} 0 & \text{if p is odd,} \\ \sigma^p (p-1)!! & \text{if p is even, } 6!! = 1 \cdot 3 \cdot 5 = 15. \end{cases}$$

Stochastic Optimization 1D: Re-Parametrization

Recall Re-Parametrization:

$$\nabla_\theta \langle f(x) \rangle_{p(x;\theta)} = \langle \nabla f(g(\varepsilon, \theta)) \rangle_{p_b(\varepsilon)}.$$

So, we can re-write our model as:

$$\min_\mu \langle (\varepsilon + \mu)^2 \rangle_{\mathcal{N}(\varepsilon;0,1)}.$$

Problem:

1. Obtain estimator of the gradient
2. Evaluate its variance

REINFORCE: Dependence over Dimension

Consider problem:

$$x \sim \mathcal{N}(x; \mu, I_d), \ \ \mu \in \mathbb{R}^d,$$

$$\max_{\mu} \left\langle \sum_{i=1}^{d} f_i(x_i) \right\rangle_{p(x;\mu)} .$$

Also, let's introduce several assumptions on the $f_i(x)$, $\forall i$:

- $\langle f_i(x) \rangle_{p(x;\mu)} = 0, \ \ \forall i,$
- $a \leq \mathbb{D} f_i(x) \leq b.$

Problem:

1. Obtain the straight forward estimator of gradient by REINFORCE
2. Low-bound its variance
3. Obtain estimator of gradient with taking all possible expectations before taking the gradient