Prof. Dr. Leif Döring
Sara Klein, Benedikt Wille

Reinforcement Learning

8. Exercise Sheet

1. Convergence of Stochastic Gradient Descent

The goal of this exercise is to prove the convergence of the stochastic version of the gradient descent method. Let $F : \mathbb{R}^d \to \mathbb{R}$ be a function of the form $F(x) = \mathbb{E}[f(x, Z)]$ for some $Z \sim \mu$, whose minimum we want to find but whose gradient we cannot exactly compute. The idea is to approximate the gradient of F by $\nabla_x f(x, Z_i)$ with independent realisations $Z_i \sim \mu$ in each step, leading to the following algorithm:

Algorithm 1: Stochastic Gradient Descent

Data: Realisation of initial random variable X_0, stepsizes α_k
Result: Approximation X of a stationary point of F
Set $k = 0$
while *not converged* **do**
  simulate $Z_{k+1} \sim \mu$ independently
  approximate the gradient $\nabla_x F(X_k)$ through
  $G_k = \nabla_x f(X_k, Z_{k+1})$
  set $X_{k+1} = X_k - \alpha_k G_k$
  set $k = k + 1$
end
return $X := X_k$

Assume the following:

- Let $(\Omega, \mathcal{F}, (\mathcal{F}_k)_{k\in\mathbb{N}}, \mathbb{P})$ be a filtered probability space, where the filtration is defined by

$$\mathcal{F}_k := \sigma(X_0, Z_m, m \le k) \text{ for } Z_k \sim_{\text{i.i.d}} \mu,$$

- let $F : \mathbb{R}^d \to \mathbb{R}$, $x \mapsto \mathbb{E}[f(x, Z)]$ for $Z \sim \mu$ be an L-smooth function for some $L < 1$, i.e.

$$\|\nabla F(x) - \nabla F(y)\| \le L\|x - y\| \quad \forall x, y \in \mathbb{R}^d$$

  and let $F_* := \inf_{x\in\mathbb{R}^d} F(x) > -\infty$,
- let $\nabla_x F(x) = \mathbb{E}[\nabla_x f(x, Z)]$ and $\mathbb{E}[\|\nabla_x f(x, Z)\|^2] \le c$ for some $c > 0$ and all $x \in \mathbb{R}^d$,
- let $(\alpha_k)_{k\in\mathbb{N}}$ be a sequence of $\mathcal{F}_k$-adapted and strictly positive random variables, where

$$\sum_{k=1}^{\infty} \alpha_k = \infty \text{ and } \sum_{k=1}^{\infty} \alpha_k^2 < \infty$$

- let X_0 be such that $\mathbb{E}[F(X_0)] < \infty$, and
- let $(X_k)_{k\in\mathbb{N}}$ be the random variables generated by applying Stochastic Gradient Descent.

a) For all L-smooth functions $f : \mathbb{R}^d \to \mathbb{R}$ it holds that

$$f(x+y) \leq f(x) + y^T \nabla f(x) + \frac{L}{2}\|y\|^2 \quad \forall x, y \in \mathbb{R}^d.$$

b) Define $M_{k+1} := \nabla_x F(X_k) - \nabla_x f(X_k, Z_{k+1})$ and show that

$$\mathbb{E}[M_{k+1}|\mathcal{F}_k] = 0 \text{ and } \mathbb{E}[\|M_{k+1}\|^2|\mathcal{F}_k] \leq c - \|\nabla_x F(X_k)\|^2 \quad \forall k \in \mathbb{N}.$$

c) Schow that $\lim_{k\to\infty} F(X_k) = F_\infty$ almost surely for some almost surely finite random variable.

d) Show that $\lim_{k\to\infty} \|\nabla_x F(X_k)\|^2 = 0$ almost surely.