

11. Exercise Sheet

1. Baseline trick

Write down and proof the baseline gradient representation for infinite discounted MDPs.

2. PL-condition

a) Prove that μ -strong convexity implies the PL-condition (5.9.), i.e.

$$\|\nabla f(x)\|^2 \geq 2r(f(x) - f_*) \quad \forall x \in \mathbb{R}^d \quad (1)$$

for $r = \mu$ and $f_* = \min_{x \in \mathbb{R}^d} f(x) > -\infty$.

b) Show that $f(x) = x^2 + 3 \sin^2(x)$ satisfies the PL-condition (1) and prove that f is not convex. Plot the function to see why gradient descent converges. Hint: The plot can also help to find the parameter r of the PL-condition.

3. Stochastic gradient descent

In the lecture we proved convergence of SGD to stationary points if the function is L -smooth and bounded. Consider the setting from the theorem of the lecture and additionally assume μ -strong convexity. Prove that $\|X_n - x_*\| \rightarrow 0$ almost surely.