Stochastic Convex Optimization

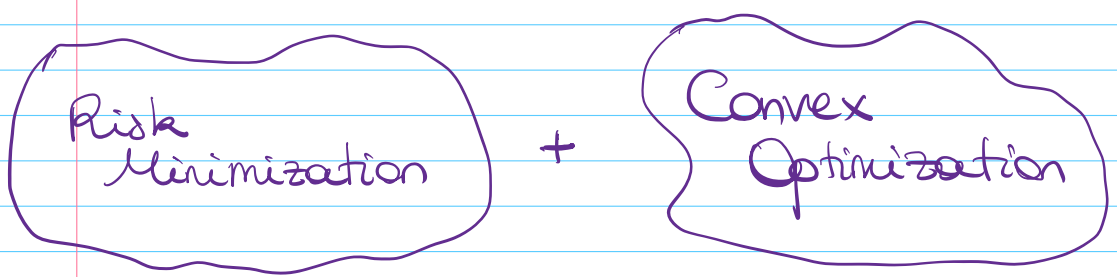

Back to the basics :

$(x_1, y_1) \dots (x_n, y_n) \sim P$

- $\mathcal{X}$ set of inputs
- $\mathcal{Y}$ set of outputs

$\ell(y, y')$ loss function.

Goal: Find a predictor $f: \mathcal{X} \to \mathcal{Y}$ that minimizes the P-risk ← test risk

$$\boxed{R_P(f) = \mathbb{E}_P[\ell(f(X), Y)]}$$

To do so: introduce a class of predictors

$$\mathcal{F} = \{ f_\theta : \mathcal{X} \to \mathcal{Y} : \quad \theta \in \mathbb{R}^d \}$$

$\Rightarrow$ Consider $\theta^* = \operatorname{argmin}_\theta R_P(f_\theta)$

How we did so far:

1. Approximate $R_P(f_\theta)$ by $R_n(f_\theta)$
2. Minimize $R_n(f_\theta)$ using GD. (or closed form in certain cases)

$$R_n(f_\theta) = \frac{1}{n} \sum_{i=1}^{n} \ell_\theta(f_\theta(X_i), Y_i)$$

$$\nabla R_n(f_\theta) = \frac{1}{n} \sum_{i=1}^{n} \underbrace{\nabla_\theta \ell_\theta(f_\theta(X_i), Y_i)}_{n \text{ gradients to compute}}$$

GD: $\theta^{t+1} = \theta^t - s \nabla R_n(f_{\theta^t})$

$\Rightarrow$ T steps of GD $= O(nT)$

Can we do better ?

Idea: Apply GD directly on

$$\theta \mapsto R_P(f_\theta) = \mathbb{E}_P[\ell(f_\theta(X), Y)]$$

$$\nabla R_P(f_\theta) = \mathbb{E}_P[\nabla_\theta \ell(f_\theta(X), Y)]$$

... But we do not have access to $\nabla R_P(f_\theta)$

$\Rightarrow$ Estimate it by $\nabla_\theta \ell(f_\theta(x_i), y_i)$

$$\theta_{t+1} = \theta_t - s \nabla_\theta \ell(f_{\theta_t}(x_i), y_i)$$

STOCHASTIC GRADIENT DESCENT

$\rightsquigarrow$ 1 step of SGD $= O(1)$

① Stochastic Gradient Descent

A more general setting

$F: \mathbb{R}^d \longrightarrow \mathbb{R}$

SGD : • θ^1 initialization

For $t = 1 \dots T-1$:

• V^t random vector such that

$$\mathbb{E}[V^t | \theta^t] = \nabla F(\theta^t)$$

Unbiased estimate of the gradient

• $\theta^{t+1} = \theta^t - \delta_t V^t$

Output : average $\bar{\theta} = \frac{1}{T} \sum_{t=1}^{T} \theta_t$.

Toy Example : $F(\theta_1, \theta_2) = a\theta_1^2 + b\theta_2^2$

$V^t = \nabla F(\theta_1^t, \theta_2^t) + \varepsilon$

ε ← random noise

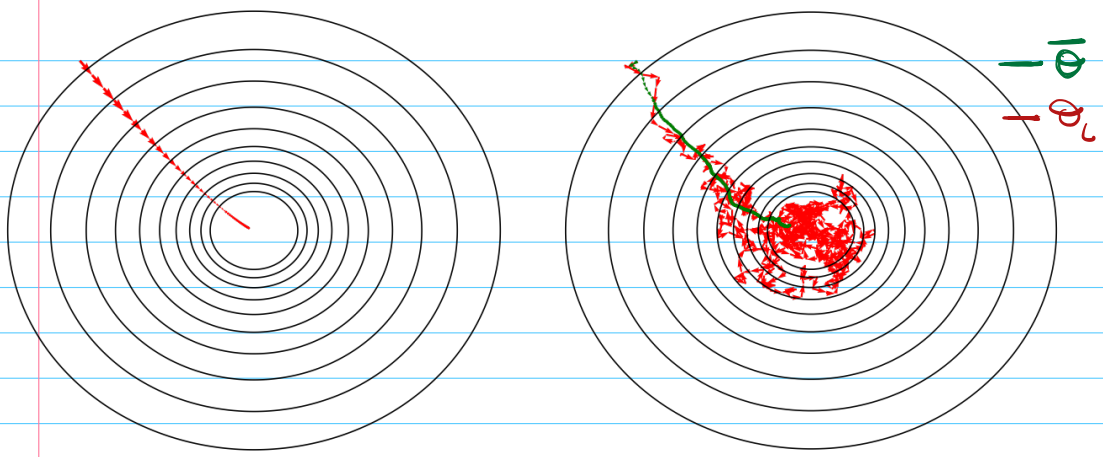

GRADIENT DESCENT

STOCHASTIC GRADIENT DESCENT

→ For risk minimization:

$$F(\theta) = R_P(f_\theta)$$

$$T = n \quad \text{(number of samples)}$$

$$V^t = \nabla \ell(f_\theta(x_i), y_i)$$

$$\mathbb{E}[V^t \mid \theta^t] = \mathbb{E}[V^t] = \nabla F(\theta)$$

only depends on $(x_j, y_j)_{j<i}$

THEOREM :

$F : \mathbb{R}^d \to \mathbb{R}$ convex + differentiable

minimizer $\theta^* \in B(0; R)$

$\mathbb{E}[\|v^t\|^2] \leq \rho^2$

Step size $\eta = \sqrt{\frac{4R^2}{\rho^2 T}}$.

Initialization $\theta^0 \in B(0; R)$

$$\Rightarrow \mathbb{E}[F(\bar{\theta})] - F(\theta^*) \leq 2\frac{R\rho}{\sqrt{T}}$$

proof 1 : special case $v^t = \nabla F(\theta^t)$

proof 2 : General case

Remark: In practice, a good idea is to discard the first iterates when computing $\bar{\theta}$: Forget about initialization

$$\bar{\theta} = \frac{1}{T - T_0} \sum_{t=T_0}^{T} \theta_t$$

Another Example: → arbitrary functions

$$F(\theta) = \frac{1}{n} \sum_{i=1}^{n} F_i(\theta)$$

Let $I \sim \text{Unif}\{1 \dots n\}$

$$\Rightarrow F(\theta) = \mathbb{E}[F_I(\theta)]$$

To get estimates of $\nabla F(\theta)$:

$$V^t = \nabla F_{I_t}(\theta) \qquad I_1 \dots I_T \text{ iid samples.}$$

Comparison with GD:

Recall:

(1) F β-smooth + α-strongly convex

T steps of GD ; $s = 1/\beta$

$$F(\theta^T) - F(\theta^*) \leq \exp\left(-\frac{\alpha}{\beta} T\right)$$

$T = O(\log(1/\varepsilon))$ steps to reach precision ε.

(2) F β-smooth + cvx

T steps of GD ; $s = 1/\beta$

$$F(\theta^T) - F(\theta^*) \leq \frac{\beta}{T}$$

$T = O(1/\varepsilon)$ steps to reach precision ε.

(3) F cvx + diff

T steps of SGD

$$E[F(\bar{\theta})] - F(\theta^*) \lesssim \frac{1}{\sqrt{T}}$$

Question : Can we improve the rate for SGD with more assumptions?

→ β - smoothness ? ⇝ no

→ α - strongly cvx ? ⇝ yes ! 

$$\text{proj}_R(\theta) = \begin{cases} \theta & \text{if } \|\theta\| \le R \\ \frac{R}{\|\theta\|}\theta & \text{else.} \end{cases}$$

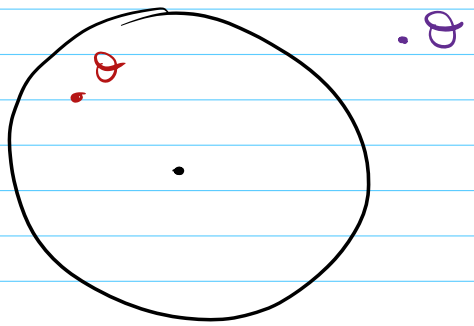

SGD with projection :

For $t = 1 \dots T-1$:

- V^t random vector such that

$$\mathbb{E}[V^t | \theta^t] = \nabla F(\theta^t)$$

- $\theta^{t+1} = \text{proj}_R(\theta^t - \delta_t V^t)$

Output : average $\bar{\theta} = \frac{1}{T} \sum_{t=1}^{T} \theta_t$.

THEOREM :

$F : \mathbb{R}^d \to \mathbb{R}$ α-strongly convex

minimizer $\theta^* \in B(0; R)$

$\mathbb{E}[\|V^t\|^2] \leq \rho^2$

Step size $\delta_t = 1/\alpha t$

Initialization $\theta^0 \in B(0; R)$

$$\Rightarrow \mathbb{E}[F(\bar{\theta})] - F(\theta^*) \leq \frac{\rho^2}{2\alpha T}(1 + \log T)$$

proof :

② Application to Risk Minimization

Back to risk minimization

$$F(\theta) = R_P(f_\theta) \qquad \theta \in \mathbb{R}^k$$

minimizer θ^*

Sample $(X_1, Y_1) \dots (X_n, Y_n)$

Two methods to approximate θ^*:

(SGD) $\hat{\theta}_{SGD}$: n steps of SGD with gradients $\nabla_\theta \ell(f_\theta(X_i), Y_i)$

(GD-ER) $\hat{\theta}_T$: Gradient Descent for T steps applied to

$$R_n(f_\theta) = \frac{1}{n} \sum_{i=1}^{n} \ell(f_\theta(X_i), Y_i)$$

$$\mathbb{E}[R_P(f_{\hat{\theta}}) - R_P^*] = \underbrace{\mathbb{E}[R_P(f_{\hat{\theta}}) - R_P(f_{\theta^*})]}_{\text{optimization error}} + \underbrace{R_P(f_{\theta^*}) - R_P^*}_{\text{approximation error}}$$

Two Questions :

1. What is the minimal number n of samples required to get an optimization error smaller than ε ?
2. What is the associated time complexity ?

$\Rightarrow$ We answer (1) and (2) in the setting where : $\forall x \in \mathcal{X}, y \in \mathcal{Y}$

$\theta \mapsto \ell(f_\theta(x), y)$ is $\begin{cases} \alpha\text{- strongly convex} \\ \beta\text{- smooth} \end{cases}$

(SGD) $v^i = \nabla \ell(f_{\theta^i}(x_i), y_i)$

$\rightsquigarrow \|v^i\| \leq \beta \|\theta^i - \theta^-\| \leq 2\beta R$

$\Rightarrow \mathbb{E}[R_p(f_{\theta_{SGD}}) - R_p(f_{\theta^*})] \lesssim \frac{1}{n}$ $\quad \|\varepsilon$

(1) $n = \tilde{O}(1/\varepsilon)$ $\quad \rightarrow$ up to log factors

(2) Time complexity : $\tilde{O}(kn)$

(GD-ER) Let $\hat{\theta}_\infty = \arg\min_\theta R_n(f_\theta)$.

$$R_\phi(f_{\theta_\tau}) - R_\phi(f_{\theta^*}) \leq$$

$$\leq 2 \cdot \sup_\theta |R_n(f_\theta) - R_\phi(f_\theta)| + R_n(f_{\theta_\tau}) - R_n(f_{\theta_\infty})$$

estimation error

$\lesssim \frac{1}{\sqrt{n}}$ $\qquad$ $\lesssim \exp(-\frac{\alpha}{\beta}\tau)$

To get

$$R_\phi(f_{\theta_\tau}) - R_\phi(f_{\theta^*}) \leq \varepsilon$$

1. $n = O(1/\varepsilon^2)$ samples
2. $\tau = O(\log(\varepsilon^{-1}))$

Complexity $O(k\tau n) = \tilde{O}\left(\frac{k}{\varepsilon^2}\right)$

SUMMARY: in the { α-strongly convex, β-smooth } setting

Algo	Num. of samples	Complexity
SGD	$n = \tilde{O}(1/\varepsilon)$	$\tilde{O}(k/\varepsilon)$
GD	$n = O(1/\varepsilon^2)$	$\tilde{O}(k/\varepsilon^2)$

To go further...

Variance Reduction technique:

SVRG / SAGA

↳ see lecture notes for references.