

Chapter 4

Optimization

Ernest K. Ryu
Seoul National University

Mathematical Machine Learning Theory
Spring 2024

"...in fact, the great watershed in optimization isn't between linearity and nonlinearity, but convexity and nonconvexity."

— R. Tyrrell Rockafellar, *SIAM Review*, 1993 —

Risk decomposition

Let

$$\theta^* \in \operatorname{argmin}_{\theta \in \mathbb{R}^p} \mathcal{R}[f_\theta], \quad \hat{\theta}^* \in \operatorname{argmin}_{\hat{\theta} \in \mathbb{R}^p} \hat{\mathcal{R}}[f_{\hat{\theta}}].$$

Then,

$$\begin{aligned} \mathcal{R}[f_{\hat{\theta}}] - \mathcal{R}^* &= \underbrace{(\mathcal{R}[f_{\hat{\theta}}] - \hat{\mathcal{R}}[f_{\hat{\theta}}])}_{=\text{Estimation error}} + \underbrace{(\hat{\mathcal{R}}[f_{\theta^*}] - \mathcal{R}[f_{\theta^*}])}_{=\text{Estimation error}} \\ &\quad \underbrace{(\mathcal{R}[f_{\theta^*}] - \mathcal{R}^*)}_{=\text{Approximation error}} + \underbrace{(\hat{\mathcal{R}}[f_{\hat{\theta}}] - \hat{\mathcal{R}}[f_{\hat{\theta}^*}])}_{=\text{Optimization error}} \\ &\leq (\mathcal{R}[f_{\hat{\theta}}] - \hat{\mathcal{R}}[f_{\hat{\theta}}]) + (\hat{\mathcal{R}}[f_{\theta^*}] - \mathcal{R}[f_{\theta^*}]) \\ &\quad (\mathcal{R}[f_{\theta^*}] - \mathcal{R}^*) + \underbrace{(\hat{\mathcal{R}}[f_{\hat{\theta}}] - \hat{\mathcal{R}}[f_{\hat{\theta}^*}])}_{=\text{Optimization error}} \geq 0 \end{aligned}$$

We now discuss algorithms for solving

$$\operatorname{minimize}_{\hat{\theta} \in \mathbb{R}^d} \hat{\mathcal{R}}[f_{\hat{\theta}}].$$

Gradient descent

Consider the optimization problem

$$\underset{\theta \in \mathbb{R}^d}{\text{minimize}} \quad F(\theta)$$

where $F: \mathbb{R}^d \rightarrow \mathbb{R}$ is differentiable.

We consider *gradient descent*

$$\theta^{k+1} = \theta^k - \alpha_k \nabla F(\theta^k)$$

where $\theta^0 \in \mathbb{R}^d$ is a starting point and $\alpha_0, \alpha_1, \dots \in \mathbb{R}$ is a positive sequence of stepsizes.

Outline

Quadratic optimization

Convex optimization

Least-squares to standard quadratic form

Consider the least-squares problem

$$\underset{\theta \in \mathbb{R}^d}{\text{minimize}} \quad \frac{1}{2} \|\Phi\theta - Y\|^2,$$

where $\Phi \in \mathbb{R}^{N \times d}$ and $Y \in \mathbb{R}^d$. Let $\theta^* = \Phi^\dagger Y$. Then,

$$\begin{aligned} \frac{1}{2} \|\Phi\theta - Y\|^2 &= \frac{1}{2} \|\Phi(\theta - \theta^*) + \Phi\theta^* - Y\|^2 \\ &= \frac{1}{2} \|\Phi(\theta - \theta^*)\|^2 + (\theta - \theta^*)^\top \underbrace{\Phi^\top (\Phi\Phi^\dagger - I) Y}_{=0} + \underbrace{\frac{1}{2} \|\Phi\theta^* - Y\|^2}_{\stackrel{\text{def}}{=} c} \\ &= \frac{1}{2} (\theta - \theta^*)^\top \underbrace{\Phi^\top \Phi}_{\stackrel{\text{def}}{=} H} (\theta - \theta^*) + c. \end{aligned}$$

Note that $H \in \mathbb{R}^{d \times d}$ is symmetric positive semidefinite.

Convex quadratic optimization

More generally, consider the optimization problem

$$\underset{\theta \in \mathbb{R}^d}{\text{minimize}} \quad F(\theta) = \frac{1}{2}\theta^\top H\theta + b^\top\theta + c,$$

where $H \in \mathbb{R}^{d \times d}$ is symmetric positive semidefinite, $b \in \mathbb{R}^d$, and $c \in \mathbb{R}$.
Then

$$\nabla F(\theta) = H\theta + b.$$

Note that if $\theta^\top H\theta + b^\top\theta + c$ had an asymmetric $H \in \mathbb{R}^{d \times d}$, then the function is equal to

$$\frac{1}{2}\theta^\top (H + H^\top)\theta + b^\top\theta + c.$$

So there is no loss of generality in assuming $H \in \mathbb{R}^{d \times d}$ is symmetric. This loss function is convex if and only if $H \succeq 0$. (To be proved in homework.)

Convex quadratic optimization in standard form

For

$$F(\theta) = \frac{1}{2}\theta^\top H\theta + b^\top\theta + c,$$

there exists some $\theta^* \in \mathbb{R}^d$ and $c' \in \mathbb{R}$ such that

$$F(\theta) = \frac{1}{2}(\theta - \theta^*)^\top H(\theta - \theta^*) + c'.$$

(To be proved in homework.) Of course, the c' is irrelevant in the optimization.

Therefore, W.L.O.G., consider

$$\underset{\theta \in \mathbb{R}^d}{\text{minimize}} \quad F(\theta) = \frac{1}{2}(\theta - \theta^*)^\top H(\theta - \theta^*),$$

where $H \in \mathbb{R}^{d \times d}$ is symmetric positive semidefinite. Then,

$$\nabla F(\theta) = H(\theta - \theta^*).$$

Of course, we won't have access to θ^* , and the actual code for computing the gradient will be something like $\nabla F(\theta) = H\theta + b$. We are simply stating the mathematical fact $\nabla F(\theta) = H\theta + b = H(\theta - \theta^*)$.

GD on strongly convex quadratic: Convergence rate

Theorem

Consider gradient descent applied to convex quadratic optimization with constant stepsize $\alpha_k = \alpha$. Let $0 < \mu \leq L < \infty$. Assume $\mu I \preceq H \preceq LI$. Then,

$$\|\theta^k - \theta^*\|^2 \leq \left(\max\{|1 - \alpha\mu|, |1 - \alpha L|\} \right)^{2k} \|\theta^0 - \theta^*\|^2, \quad \text{for } k = 0, 1, \dots$$

If $\alpha = 2/(\mu + L)$, then,

$$\|\theta^k - \theta^*\|^2 \leq \left(1 - \frac{2}{\kappa + 1} \right)^{2k} \|\theta^0 - \theta^*\|^2 \leq \exp\left(-\frac{4k}{\kappa + 1} \right) \|\theta^0 - \theta^*\|^2,$$

where $\kappa = L/\mu$, for $k = 0, 1, \dots$

(When $\kappa = \infty$, the bound merely guarantees $\|\theta^k - \theta^*\|^2 \leq \|\theta^0 - \theta^*\|^2$.)

GD on strongly convex quadratic: Convergence rate

Proof. Note

$$\begin{aligned}\theta^k - \theta^* &= \theta^{k-1} - \alpha \nabla F(\theta^{k-1}) - \theta^* \\ &= (I - \alpha H)(\theta^{k-1} - \theta^*) = (I - \alpha H)^k(\theta^0 - \theta^*),\end{aligned}$$

and, with $\Lambda(H)$ being the set of eigenvalues of H ,

$$\min_{\lambda \in [\mu, L]} \{1 - \alpha\lambda\}I \preceq \min_{\lambda \in \Lambda(H)} \{1 - \alpha\lambda\}I \preceq I - \alpha H \preceq \max_{\lambda \in \Lambda(H)} \{1 - \alpha\lambda\}I \preceq \max_{\lambda \in [\mu, L]} \{1 - \alpha\lambda\}I.$$

This implies

$$(I - \alpha H)^2 \preceq \left(\max_{\lambda \in [\mu, L]} |1 - \alpha\lambda| \right)^2 I.$$

Therefore,

$$\begin{aligned}\|\theta^k - \theta^*\|^2 &= (\theta^0 - \theta^*)^\top (I - \alpha H)^{2k} (\theta^0 - \theta^*) \\ &\leq \left(\max_{\lambda \in [\mu, L]} |1 - \alpha\lambda| \right)^{2k} \|\theta^0 - \theta^*\|^2 \\ &\leq \left(\max\{|1 - \alpha\mu|, |1 - \alpha L|\} \right)^{2k} \|\theta^0 - \theta^*\|.\end{aligned}$$

□

GD on strongly convex quadratic: Iteration complexity

Corollary

Consider gradient descent applied to convex quadratic optimization such that $\mu I \preceq H \preceq LI$. Consider a constant stepsize $\alpha_k = 2/(\mu + L)$. If

$$K \geq \frac{\kappa + 1}{4} (\log(1/\varepsilon) + 2 \log \|\theta^0 - \theta^*\|),$$

then,

$$\|\theta^K - \theta^*\|^2 \leq \varepsilon.$$

In optimization and numerical linear algebra, $\kappa = L/\mu \geq 1$ is called the *condition number* of the problem, and it characterizes the difficulty of the problem. In algorithm design, we want efficiency for harder problem instances since easy problems are easy anyway. Therefore, we are primarily interested in the algorithm's performance when $\kappa \gg 1$.

The iteration complexity is $\mathcal{O}(\kappa)$, when all other dependences are ignored. We will soon see that $\mathcal{O}(\sqrt{\kappa})$ can be obtained via “acceleration” and that this complexity is optimal.

GD on cvx quadratics: Function values

Theorem

Consider gradient descent applied to convex quadratic optimization with constant stepsize $\alpha_k = 1/L$. Assume $\mu I \preceq H \preceq LI$. Then,

$$F(\theta^k) - F(\theta^*) \leq \left(1 - \frac{1}{\kappa}\right)^{2k} (F(\theta^0) - F(\theta^*)) \leq \exp\left(-\frac{2k}{\kappa}\right) (F(\theta^0) - F(\theta^*))$$

for $k = 0, 1, \dots$.

Proof. We first note that

$$F(\theta^k) - F(\theta^*) = \frac{1}{2}(\theta^0 - \theta^*)^\top (I - \alpha H)^{2k} H (\theta^0 - \theta^*)$$

and

$$F(\theta^0) - F(\theta^*) = \frac{1}{2}(\theta^0 - \theta^*)^\top H (\theta^0 - \theta^*).$$

We conclude the statement with the same line of argument as before. \square

(When $\kappa = \infty$, bound merely guarantees $F(\theta^k) - F(\theta^*) \leq F(\theta^0) - F(\theta^*)$.)

GD on cvx quadratics: Sublinear convergence

Theorem

Consider gradient descent applied to convex quadratic optimization with constant stepsize $\alpha_k = 1/L$. Assume $0 \preceq H \preceq LI$. Then,

$$F(\theta^k) - F(\theta^*) \leq \frac{L}{8k} \|\theta^0 - \theta^*\|^2, \quad \text{for } k = 0, 1, \dots$$

Proof. Again, note that

$$F(\theta^k) - F(\theta^*) = \frac{1}{2} (\theta^0 - \theta^*)^\top (I - \alpha H)^{2k} H (\theta^0 - \theta^*).$$

By a similar argument as before,

$$(I - \alpha H)^{2k} H \preceq \left(\max_{\lambda \in [0, L]} |\lambda(1 - \alpha\lambda)^{2k}| \right) I.$$

For $\lambda \geq 0$ and $0 \leq \alpha \leq 1/L$,

$$\begin{aligned} |\lambda(1 - \alpha\lambda)^{2k}| &\leq \lambda \exp(-2k\alpha\lambda) = \frac{1}{2k\alpha} 2k\alpha\lambda \exp(-2k\alpha\lambda) \\ &\leq \frac{1}{2k\alpha} \sup_{\tau \geq 0} \{\tau \exp(-\tau)\} = \frac{1}{2ek\alpha} \leq \frac{1}{4k\alpha}. \end{aligned}$$

□

GD on cvx quadratics: Iteration complexity

Corollary

Consider gradient descent applied to convex quadratic optimization with constant stepsize $\alpha_k = 1/L$. Assume $0 \preceq H \preceq LI$. If

$$K \geq \frac{L}{8\varepsilon} \|\theta^0 - \theta^*\|^2,$$

then

$$F(\theta^k) - F(\theta^*) \leq \varepsilon.$$

We attain a convergence rate of $\mathcal{O}(1/k)$ and iteration complexity $\mathcal{O}(1/\varepsilon)$. We will soon see that accelerated methods achieve $\mathcal{O}(1/k^2)$ and $\mathcal{O}(1/\sqrt{\varepsilon})$ and that these are optimal.

Linear vs. sublinear convergence

In optimization, a rate of the form

$$\text{performance measure} \leq \exp(-k/C),$$

where $C > 0$, is (confusingly) referred to be both an exponential rate and a linear rate.¹

In contrast, a rate of the form

$$\text{performance measure} \leq \frac{C}{k^\gamma},$$

where $\gamma > 0$, is said to be a sublinear rate.

¹This is not a Q-linear rate but rather an R-linear rate. We will not worry about this distinction.

Why gradient methods?

When F is quadratic, one can use direct methods (linear algebraic) to solve the linear system $\nabla F(\theta) = 0$. Why consider gradient descent?

- ▶ When $\theta \in \mathbb{R}^d$, direct solves of $\nabla F(\theta) = 0$ requires $\mathcal{O}(d^3)$ cost, whereas $\theta^{k+1} = \theta^k - \alpha_k \nabla F(\theta^k)$ requires $\mathcal{O}(d^2)$ cost or less. (Cost per iteration is dominated by the cost of evaluating ∇F .) If GD converges to acceptable accuracy in fewer than $\mathcal{O}(d)$ iterations, GD is worthwhile.
- ▶ The guarantees of convex quadratic optimization serve as a conceptual baseline later when we try to get the same types of guarantees for convex non-quadratic optimization.

Outline

Quadratic optimization

Convex optimization

Convex non-quadratic optimization

In convex optimization, we solve

$$\underset{\theta \in \mathbb{R}^d}{\text{minimize}} \quad F(\theta),$$

where $F: \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex function that is not necessarily a quadratic.

Since F is non-quadratic, we can no longer use the linear algebraic tools. Nevertheless, can we get the same rates as in the quadratic case?

We start by assuming F is differentiable and we later consider the case where F is non-differentiable.

GD on smooth strongly convex F : Convergence rate

We first analyze the GD with a contraction argument.

Theorem

Let $0 < \mu \leq L < \infty$. Let $F: \mathbb{R}^d \rightarrow \mathbb{R}$ be L -smooth and μ -strongly convex. Consider GD with $\alpha_k = 1/L$. Then, for $k = 0, 1, \dots$,

$$\|\theta^k - \theta^*\|^2 \leq (1 - 1/\kappa)^k \|\theta^0 - \theta^*\|^2 \leq \exp(-k/\kappa) \|\theta^0 - \theta^*\|^2.$$

Proof.

$$\begin{aligned} \|\theta^{k+1} - \theta^*\|^2 &= \|\theta^k - \theta^*\|^2 - 2\alpha_k \langle \nabla F(\theta^k), \theta^k - \theta^* \rangle + \alpha_k^2 \|\nabla F(\theta^k)\|^2 \\ &\leq \|\theta^k - \theta^*\|^2 - \alpha_k (2 - \alpha_k L) \langle \nabla F(\theta^k), \theta^k - \theta^* \rangle \\ &\leq (1 - \mu\alpha_k(2 - \alpha_k L)) \|\theta^k - \theta^*\|^2 = (1 - 1/\kappa) \|\theta^k - \theta^*\|^2, \end{aligned}$$

where the first and second inequalities follows from

$$\begin{aligned} \langle \nabla F(\theta) - \nabla F(\eta), \theta - \eta \rangle &\geq \frac{1}{L} \|\nabla F(\theta) - \nabla F(\eta)\|^2, \quad \forall \theta, \eta \in \mathbb{R}^d \\ \langle \nabla F(\theta) - \nabla F(\eta), \theta - \eta \rangle &\geq \mu \|\theta - \eta\|^2, \quad \forall \theta, \eta \in \mathbb{R}^d \end{aligned}$$

by L -smoothness and μ -strong convexity, respectively. □

GD on smooth strongly convex F : Iteration complexity

Corollary

Let $F: \mathbb{R}^d \rightarrow \mathbb{R}$ be L -smooth and μ -strongly convex. Consider gradient descent with constant stepsize $\alpha_k = 1/L$. Then, if

$$K \geq \kappa \log(1/\varepsilon) + 2 \log \|\theta^0 - \theta^*\| = \mathcal{O}(\kappa \log(1/\varepsilon)),$$

then

$$\|\theta^k - \theta^*\|^2 \leq \varepsilon.$$

Polyak–Łojasiewicz inequality

Lemma

Let $F: \mathbb{R}^d \rightarrow \mathbb{R}$ be μ -strongly convex, differentiable, and with a minimizer θ^* . Then,

$$\|\nabla F(\theta)\|^2 \geq 2\mu(F(\theta) - F(\theta^*)), \quad \forall \theta \in \mathbb{R}^d.$$

Proof. By μ -s.c.,

$$\begin{aligned} F(\eta) &\geq F(\theta) + \langle \nabla F(\theta), \eta - \theta \rangle + \frac{\mu}{2} \|\eta - \theta\|^2 \\ &\geq \inf_{\eta \in \mathbb{R}^d} \left\{ F(\theta) + \langle \nabla F(\theta), \eta - \theta \rangle + \frac{\mu}{2} \|\eta - \theta\|^2 \right\} = F(\theta) - \frac{1}{2\mu} \|\nabla F(\theta)\|^2. \end{aligned}$$

(Infimum is attained at $\eta = \theta - \frac{1}{\mu} \nabla F(\theta)$.) Plugging $\eta = \theta^*$ into the LHS, we arrive at the conclusion. \square

This is called the Polyak–Łojasiewicz (PL) inequality. Strong convexity implies PL. However, the converse is not true.

GD on smooth strongly convex F : Convergence rate

Using the Polyak–Łojasiewicz inequality, we can obtain a rate on F .

Theorem

Let $F: \mathbb{R}^d \rightarrow \mathbb{R}$ be L -smooth, convex, and μ -P.L. Consider gradient descent with constant stepsize $\alpha_k = 1/L$. Then, for $k = 0, 1, \dots$,

$$F(\theta^k) - F(\theta^*) \leq (1 - 1/\kappa)^k (F(\theta^0) - F(\theta^*)) \leq \exp(-k/\kappa) (F(\theta^0) - F(\theta^*)).$$

Proof.

$$\begin{aligned} F(\theta^{k+1}) - F(\theta^*) &= F(\theta^k - \alpha \nabla F(\theta^k)) - F(\theta^*) \\ &\leq F(\theta^k) - \alpha \langle \nabla F(\theta^k), \nabla F(\theta^k) \rangle + \frac{\alpha^2 L}{2} \|\nabla F(\theta^k)\|^2 - F(\theta^*) \\ &= F(\theta^k) - F(\theta^*) - \frac{1}{2L} \|\nabla F(\theta^k)\|^2 \\ &\leq (1 - 1/\kappa) (F(\theta^k) - F(\theta^*)), \end{aligned}$$

where the first inequality follows from L -smoothness and the third follows from PL. □

Specifically, the following inequality was used in the previous proof:

$$\begin{aligned} F(\theta^{k+1}) &\leq F(\theta^k) + \langle \nabla F(\theta^k), \theta^{k+1} - \theta^k \rangle + \frac{L}{2} \|\theta^{k+1} - \theta^k\|^2 \\ &= F(\theta^k) - \frac{1}{2L} \|\nabla F(\theta^k)\|^2 \end{aligned}$$

Complexity lower bound

Later, we will establish the following complexity lower bounds

Theorem

Consider first-order algorithms minimizing $F: \mathbb{R}^d \rightarrow \mathbb{R}$ by only accessing $(F(\theta^k), \nabla F(\theta^k))$ for $k = 0, \dots, K - 1$, where θ^0 is given and $\theta^1, \dots, \theta^{K-1}$ are chosen by the algorithm. Denote the output of the algorithm as θ^K . Then, there are constants $C_1 > 0$ and $C_2 > 0$ such that the following hold: for any K and sufficiently large d , there is an L -smooth and μ -strongly convex F such that

$$F(\theta^K) - F(\theta^*) \geq C_1 \exp(-C_2 K / \sqrt{\kappa}) \|\theta^0 - \theta^*\|^2.$$

We will discuss the precise conditions of this lower bound and prove it later in this course. For now, understand that the $\mathcal{O}(\kappa)$ iteration complexity of GD is suboptimal, and it can be accelerated.

Complexity lower bound

Later, we will establish the following complexity lower bounds

Theorem

Consider first-order algorithms minimizing $F: \mathbb{R}^d \rightarrow \mathbb{R}$ by only accessing $(F(\theta^k), \nabla F(\theta^k))$ for $k = 0, \dots, K - 1$, where θ^0 is given and $\theta^1, \dots, \theta^{K-1}$ is chosen by the algorithm. Denote the output of the algorithm as θ^K . Then, there is a constant $C > 0$ such that the following hold: for any K and sufficiently large d , there is an L -smooth convex F such that

$$F(\theta^K) - F(\theta^*) \geq \frac{C}{k^2} \|\theta^0 - \theta^*\|^2.$$

Again, note that the $\mathcal{O}(1/k)$ rate of GD is suboptimal and that it can be accelerated.

GD on smooth convex F : Convergence rate

Next, we establish a sublinear rate on gradient descent.

Theorem

Let $F: \mathbb{R}^d \rightarrow \mathbb{R}$ be L -smooth convex. Assume F has a minimizer θ^* . Consider gradient descent with constant stepsize $\alpha_k = 1/L$. Then, for $k = 1, 2, \dots$,

$$F(\theta^k) - F(\theta^*) \leq \frac{L}{2k} \|\theta^0 - \theta^*\|^2.$$

Before we get to the direct analysis of GD, let us do a warm-up exercise by considering a continuous-time model of gradient descent. Since

$$\theta^{k+1} = \theta^k - \alpha \nabla F(\theta^k) \quad \Rightarrow \quad \frac{\theta^{k+1} - \theta^k}{\alpha} = -\nabla F(\theta^k),$$

for small enough $\alpha > 0$, we can approximate the discrete-time dynamics with the continuous-time gradient flow

$$\dot{\theta}(t) = -\nabla F(\theta(t)).$$

Continuous-time analysis of GD

Consider the continuous-time dynamics

$$\dot{\theta}(t) = -\nabla F(\theta(t))$$

Define the energy function (not an obvious choice)

$$\mathcal{E}(t) = t(F(\theta) - F(\theta^*)) + \frac{1}{2}\|\theta - \theta^*\|^2.$$

Then, we can show $\mathcal{E}(t)$ is dissipative:

$$\begin{aligned} \frac{d}{dt}\mathcal{E}(t) &= F(\theta) - F(\theta^*) + t\langle \nabla F(\theta), \dot{\theta} \rangle + \langle \theta - \theta^*, \dot{\theta} \rangle \\ &= \underbrace{F(\theta) - F(\theta^*) + \langle \theta^* - \theta, \nabla F(\theta) \rangle}_{\leq 0 \text{ by convexity}} - t\|\nabla F(\theta)\|^2 \leq 0. \end{aligned}$$

Therefore,

$$t(F(\theta) - F(\theta^*)) \leq \mathcal{E}(t) \leq \mathcal{E}(0) \leq \frac{1}{2}\|\theta(0) - \theta^*\|^2$$

and we conclude

$$F(\theta) - F(\theta^*) \leq \frac{1}{2t}\|\theta(0) - \theta^*\|^2.$$

Discrete-time analysis of GD

Proof. Recall

$$\theta^{k+1} = \theta^k - \alpha \nabla F(\theta^k).$$

Define the energy function

$$\mathcal{E}_k = k(F(\theta^k) - F(\theta^*)) + \frac{L}{2} \|\theta^k - \theta^*\|^2.$$

Then, we show $\{\mathcal{E}_k\}_{k=0}^\infty$ is dissipative

$$\begin{aligned} \mathcal{E}_{k+1} - \mathcal{E}_k &= (k+1)(F(\theta^{k+1}) - F(\theta^*)) - k(F(\theta^k) - F(\theta^*)) \\ &\quad - \alpha L \langle \nabla F(\theta^k), \theta^k - \theta^* \rangle + \alpha^2 L \|\nabla F(\theta^k)\|^2 \\ &\leq F(\theta^k) - F(\theta^*) - \frac{k+1}{2L} \|\nabla F(\theta^k)\|^2 - \langle \nabla F(\theta^k), \theta^k - \theta^* \rangle + \frac{1}{L} \|\nabla F(\theta^k)\|^2 \\ &\leq -\frac{1}{2L} \|\nabla F(\theta^k)\|^2 - \frac{k+1}{2L} \|\nabla F(\theta^k)\|^2 + \frac{1}{L} \|\nabla F(\theta^k)\|^2 \\ &= -\frac{k}{2L} \|\nabla F(\theta^k)\|^2 \leq 0, \end{aligned}$$

where the first and second inequalities follow from L -smoothness. The conclusion follows from

$$k(F(\theta^k) - F(\theta^*)) \leq \mathcal{E}_k \leq \mathcal{E}_0 = \frac{L}{2} \|\theta^0 - \theta^*\|^2. \quad \square$$

Discrete-time analysis of GD

Specifically, the following inequalities were used in the previous proof:

$$\begin{aligned} F(\theta^{k+1}) &\leq F(\theta^k) + \langle \nabla F(\theta^k), \theta^{k+1} - \theta^k \rangle + \frac{L}{2} \|\theta^{k+1} - \theta^k\|^2 \\ &= F(\theta^k) - \frac{1}{2L} \|\nabla F(\theta^k)\|^2 \end{aligned}$$

$$F(\theta^k) - F(\theta^*) - \langle \nabla F(\theta^k), \theta^k - \theta^* \rangle \leq -\frac{1}{2L} \|\nabla F(\theta^k)\|^2.$$

Subgradient descent

Consider the optimization problem

$$\underset{\theta \in \mathbb{R}^d}{\text{minimize}} \quad F(\theta)$$

where $F: \mathbb{R}^d \rightarrow \mathbb{R}$ is convex but not differentiable.

The *subgradient descent* method is

$$\begin{aligned} g^k &\in \partial F(\theta^k) \\ \theta^{k+1} &= \theta^k - \alpha_k g^k \end{aligned}$$

where $\theta^0 \in \mathbb{R}^d$ is a starting point and $\alpha_0, \alpha_1, \dots \in \mathbb{R}$ is a positive sequence of stepsizes. With $g^k \in \partial F(\theta^k)$, we assume that we are given a subgradient. (We cannot choose a particular subgradient.)

(Some say the name subgradient “descent” is a misnomer since there is no guarantee that the function value $F(\theta^k)$ is monotonically decreasing.)

Lipschitz continuity \Leftrightarrow bounded subgradients

Lemma

Let $F: \mathbb{R}^d \rightarrow \mathbb{R}$ be convex. Then, F is G -Lipschitz continuous, i.e.,

$$|F(x) - F(y)| \leq G\|x - y\|_2, \quad \forall x, y \in \mathbb{R}^d$$

if and only if

$$\|\partial F(x)\|_2 \leq G, \quad \forall x \in \mathbb{R}^d.$$

To clarify, L -smoothness concerns Lipschitz continuity of ∇F . Here, we are concerned with Lipschitz continuity of F .

To clarify, $\|\partial F(x)\|_2 \leq G$ means $\|g\|_2 \leq G$ for all $g \in \partial F(x)$.

If F is differentiable, then $\partial F = \nabla F$, and this lemma follows from standard calculus arguments. The proof of this lemma is beyond the scope of this course.

Convergence rate of subgradient descent

Theorem

Let $F: \mathbb{R}^d \rightarrow \mathbb{R}$ be a G -Lipschitz continuous convex function. Assume F has a minimizer θ^* . Let $\theta^0 \in \mathbb{R}^d$ be a starting point and write $R = \|\theta^0 - \theta^*\|_2$. Let $K > 0$ be the total iteration count. Then, subgradient descent with the constant stepsize

$$\alpha_k = \alpha = \frac{R}{G\sqrt{K+1}}$$

exhibits the rate

$$\min_{0 \leq k \leq K} F(\theta^k) - f(\theta^*) \leq \frac{GR}{\sqrt{K+1}}$$

and

$$f(\bar{\theta}^K) - f(\theta^*) \leq \frac{GR}{\sqrt{K+1}},$$

where

$$\bar{\theta}^K = \frac{1}{K+1} \sum_{k=0}^K \theta^k.$$

Proof. For $k = 0, 1, 2, \dots$,

$$\begin{aligned}\|\theta^{k+1} - \theta^*\|_2^2 &= \|\theta^k - \alpha g^k - \theta^*\|_2^2 \\ &= \|\theta^k - \theta^*\|_2^2 - 2\alpha \langle g^k, \theta^k - \theta^* \rangle + \alpha^2 \|g^k\|_2^2 \\ &\leq \|\theta^k - \theta^*\|_2^2 - 2\alpha (F(\theta^k) - F(\theta^*)) + \alpha^2 G^2.\end{aligned}$$

Therefore,

$$2\alpha (F(\theta^k) - F(\theta^*)) \leq \|\theta^k - \theta^*\|_2^2 - \|\theta^{k+1} - \theta^*\|_2^2 + \alpha^2 G^2.$$

With a telescoping sum argument, we get

$$\begin{aligned}2\alpha \sum_{k=0}^K (F(\theta^k) - F(\theta^*)) &\leq \|\theta^0 - \theta^*\|_2^2 - \|\theta^{K+1} - \theta^*\|_2^2 + \sum_{k=0}^K \alpha^2 G^2 \\ &\leq R^2 + (K+1)\alpha^2 G^2,\end{aligned}$$

and

$$\frac{1}{K+1} \sum_{k=0}^K F(\theta^k) - F(\theta^*) \leq \frac{R^2 + \alpha^2 G^2 (K+1)}{2\alpha (K+1)} = \frac{GR}{\sqrt{K+1}}.$$

Therefore,

$$\begin{aligned}\min_{0 \leq k \leq K} F(\theta^k) - F(\theta^*) &= \frac{1}{K+1} \sum_{k=0}^K \min_{0 \leq k \leq K} F(\theta^k) - F(\theta^*) \\ &\leq \frac{1}{K+1} \sum_{k=0}^K F(\theta^k) - F(\theta^*) \leq \frac{GR}{\sqrt{K+1}}.\end{aligned}$$

Likewise, using Jensen's inequality, we conclude

$$F(\bar{\theta}^K) - F(\theta^*) \leq \frac{1}{K+1} \sum_{k=0}^K F(\theta^k) - F(\theta^*) \leq \frac{GR}{\sqrt{K+1}}.$$

□

Complexity lower bound

Subgradient descent exhibits a $\mathcal{O}(1/\sqrt{K})$ rate, which is slower than gradient descent on L -smooth convex functions. This $\mathcal{O}(1/\sqrt{K})$ rate turns out to be optimal.

Theorem

Consider first-order algorithms minimizing $F: \mathbb{R}^d \rightarrow \mathbb{R}$ by only accessing $(F(\theta^k), g^k)$, where $g^k \in \partial F(\theta^k)$ is a subgradient, for $k = 0, \dots, K-1$, where θ^0 is given and $\theta^1, \dots, \theta^{K-1}$ is chosen by the algorithm. Denote the output of the algorithm as θ^K . Then, there is a constant $C > 0$ such that the following hold: for any K and sufficiently large d , there is an G -Lipschitz convex F such that

$$F(\theta^K) - F(\theta^*) \geq \frac{C}{\sqrt{K}} \|\theta^0 - \theta^*\|^2.$$