

MS&E 313 / CS 269G : Chapter 4 - Acceleration*

By Aaron Sidford (sidford@stanford.edu)

April 20, 2017

In the last chapter we proved the following result about gradient descent for minimizing L -smooth μ -strongly convex functions.

Theorem 1. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a L -smooth μ -strongly convex function for $\mu \geq 0$. Let $x_0 \in \mathbb{R}^n$ and $x_* \in X_*(f)$ be arbitrary and let $x_{k+1} = x_k - \frac{1}{L} \nabla f(x_k)$ for all $k \geq 0$. Then*

$$f(x_k) - f_* \leq \min \left\{ \left(1 - \frac{\mu}{L}\right)^k [f(x_0) - f_*], \frac{L \cdot \|x_0 - x_*\|_2^2}{k + 4} \right\}.$$

Consequently we can compute an ϵ -suboptimal point with $O(\lceil \min\{\frac{L}{\mu} \log(\frac{f(x_0) - f_}{\epsilon}), \frac{L\|x_0 - x_*\|_2^2}{\epsilon}\} \rceil)$ oracle calls.*

A natural question to ask, is this optimal? If all we have is a gradient oracle, can we do better? Here we address this question, showing how better running times can be achieved through a technique typically referred to as *acceleration*.

1 Acceleration

So how should we accelerate? The idea here we use is pretty simple. If f is L -smooth μ -strongly convex we know that for all x the functions

$$L_x(y) = f(x) + \nabla f(x)^\top (y - x) + \frac{\mu}{2} \|y - x\|_2^2 \text{ and } U_x(y) = f(x) + \nabla f(x)^\top (y - x) + \frac{L}{2} \|y - x\|_2^2$$

Lower and upper bound f , i.e.

$$L_x(y) \leq f(y) \leq U_x(y)$$

for all $y \in \mathbb{R}^n$. Now before our algorithm worked simply by greedily decreasing our function value using the upper bound. Here we try to do better by using the lower bound as well. Whereas the gradient descent algorithm we analyzed before used a fixed step size that only depended upon the smoothness of the function here we use strong convexity in designing our steps as well.

2 Acceleration Approach

The algorithm we use for our analysis is fairly straightforward. In every iteration k we maintain some $x_k \in \mathbb{R}^n$ and some lower bound function $L_k : \mathbb{R}^n \rightarrow \mathbb{R}$ such that for all $x \in \mathbb{R}^n$ we have $L_k(x) \leq f(x)$. We

*These notes are a work in progress. They are not necessarily a subset or superset of the in-class material and there may also be occasional *TODO* comments which demarcate material I am thinking of adding in the future. These notes will converge to a superset of the class material that is *TODO*-free. Your feedback is welcome and highly encouraged. If anything is unclear, you find a bug or typo, or if you would find it particularly helpful for anything to be expanded upon, please do not hesitate to post a question on the discussion board or contact me directly at sidford@stanford.edu.

store $L_k(x) = \psi_k + \frac{\mu}{2}\|x - v_k\|_2^2$ for some $\psi_k \in \mathbb{R}$ and some $v_k \in \mathbb{R}^n$. Now, the idea is that $\min_x L_k(x) = \psi_k$ and thus

$$f(x_k) - \psi_k \geq f(x_k) - \min_x f(x) = f(x_k) - f_*.$$

Consequently, it suffices to show we can decrease $f(x_k) - \psi_k$ at a fast rate.

The way we do this is simple. We let

$$y_k = \alpha \cdot x_k + (1 - \alpha) \cdot v_k$$

and use this point y_k to improve both our lower bound and upper bound. We let $x_{k+1} = y_k - \frac{1}{L} \nabla f(x_k)$ and we let

$$L_{k+1} = \beta \cdot L_k + (1 - \beta) \cdot U_{y_k}(y)$$

That is it, that is the entire algorithm. We pick our lower bounds $L_k(x) = \psi_k + \frac{\mu}{2}\|x - v_k\|_2^2$ as the above update rule keeps the L_k of this form and we can store these L_k compactly.

What is tricky about this algorithm and the analysis of it is reasoning about exactly what happens when we combine lower bounds. In the next section we analyze this through a self contained helper lemma.

3 Quadratics and Combining Lower Bounds

Here we give a self contained lemma about the effect of taking convex combinations of quadratics. Proving these will be done in homework.

Lemma 2. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a twice differentiable function where $\nabla^2 f(x) = \mathbf{A} \in \mathbb{R}^{n \times n}$ for all $x \in \mathbb{R}^n$ then for all x, y we have that*

$$f(x) = f(y) + \nabla f(x)^\top (y - x) + \frac{1}{2}(y - x)^\top \mathbf{A}(y - x).$$

Proof. Homework. □

This lemma shows that we can re-write our lower bounds as quadratics centered around a particular point.

Corollary 3. *If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a differentiable μ -strongly convex function then for all $x, y \in \mathbb{R}^n$ we have*

$$f(y) \geq f(x) - \frac{1}{2\mu} \|\nabla f(x)\|_2^2 + \frac{\mu}{2} \left\| y - \left[x - \frac{1}{\mu} \nabla f(x) \right] \right\|_2^2.$$

Proof. We have already seen that

$$f(y) \geq L_x(y) \stackrel{\text{def}}{=} f(x) + \nabla f(x)^\top (y - x) + \frac{\mu}{2} \|y - x\|_2^2.$$

However, we know that

$$L_x \left(x - \frac{1}{\mu} \nabla f(x) \right) = f(x) - \frac{1}{2\mu} \|\nabla f(x)\|_2^2$$

and

$$\nabla L_x \left(x - \frac{1}{\mu} \nabla f(x) \right) = \nabla f(x) - \frac{\mu}{\mu} \nabla f(x) = \vec{0}.$$

Since $\nabla^2 L_x(y) = \mathbf{I}$ for all y the result follows from Lemma 2. □

Next, using Lemma 2 we show how to combine quadratic lower bounds.

Lemma 4. Let $f_1, f_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ be defined for all $x \in \mathbb{R}^n$ by

$$f_1(x) \stackrel{\text{def}}{=} \psi_1 + \frac{\mu}{2} \|x - v_1\|_2^2 \quad \text{and} \quad f_2(x) = \psi_2 + \frac{\mu}{2} \|x - v_2\|_2^2$$

For $\psi_1, \psi_2 \in \mathbb{R}$, $v_1, v_2 \in \mathbb{R}^n$, and $\mu \geq 0$. Then for all $\alpha \in [0, 1]$ we have

$$\alpha \cdot f_1(x) + (1 - \alpha) \cdot f_2(x) = \psi_\alpha + \frac{\mu}{2} \|x - v_\alpha\|_2^2$$

where

$$v_\alpha = \alpha v_1 + (1 - \alpha) v_2 \quad \text{and} \quad \psi_\alpha = \alpha \psi_1 + (1 - \alpha) \psi_2 + \frac{\mu}{2} \alpha (1 - \alpha) \|v_1 - v_2\|_2^2.$$

Proof. Homework. □

Intuitively, the above equation says that the farther away the centers are, the more that combining them increases the lower bound.

4 Building An Accelerated Gradient Step

Using the analysis in the previous section yields the following bound for improving our lower bounds.

Lemma 5. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable μ -strongly convex function and let $L_k(x) = \psi_k + \frac{\mu}{2} \|x - v_k\|_2^2$ by such that $f(x) \geq L_k(x)$ for all x . Then for all $\beta \in [0, 1]$ and $y_k \in \mathbb{R}^n$ we have that

$$L_{k+1}(x) = \beta \cdot L_k(x) + (1 - \beta) \cdot L_{y_k}(x)$$

where $L_{y_k}(x) = f(y_k) + \nabla f(y_k)^\top (x - y_k) + \frac{\mu}{2} \|x - y_k\|_2^2$ satisfies $f(x) \geq L_{k+1}(x)$ for all $x \in \mathbb{R}^n$ and

$$L_{k+1}(x) = \psi_{k+1} + \frac{\mu}{2} \|x - v_{k+1}\|_2^2$$

where

$$v_{k+1} = \beta \cdot v_k + (1 - \beta) \cdot \left[y_k - \frac{1}{\mu} \nabla f(y_k) \right]$$

and

$$\psi_{k+1} = \beta \cdot \psi_k + (1 - \beta) \cdot \left[f(y_k) - \frac{1}{2\mu} \|\nabla f(y_k)\|_2^2 \right] + \frac{\mu}{2} \beta (1 - \beta) \cdot \left\| v_k - \left[y_k - \frac{1}{\mu} \nabla f(y_k) \right] \right\|_2^2.$$

Proof. First, note that $\nabla U_y(y_k - \frac{1}{\mu} \nabla f(y_k)) = \vec{0}$ and $f(y_k - \frac{1}{\mu} \nabla f(y_k)) = f(y_k) - \frac{1}{2\mu} \|\nabla f(y_k)\|_2^2$ and consequently, by Lemma 2 we have that

$$L_{y_k}(x) = f(y_k) - \frac{1}{2\mu} \|\nabla f(y_k)\|_2^2 + \frac{\mu}{2} \|x - (y_k - \frac{1}{\mu} \nabla f(y_k))\|_2^2$$

The result then follows by applying Lemma 4. □

Using this we can analyze a gradient descent step.

Lemma 6. Under the same assumptions of Lemma 5 if $y_k = \alpha \cdot x_k + (1 - \alpha) \cdot v_k$ for some $\alpha \in (0, 1)$ and $x_{k+1} = y_k - \frac{1}{L} \nabla f(y_k)$ then we have

$$f(x_{k+1}) - \psi_{k+1} \leq \beta [f(x_k) - \psi_k] + \beta \cdot \left[1 - \alpha \cdot \frac{1 - \beta}{1 - \alpha} \right] (f(y_k) - f(x_k)) + \left[\frac{(1 - \beta)^2}{2\mu} - \frac{1}{2L} \right] \|\nabla f(y_k)\|_2^2.$$

Consequently if $\kappa = \frac{L}{\mu}$, $\beta = 1 - \sqrt{\frac{1}{\kappa}}$, and $\alpha = \frac{\sqrt{\kappa}}{1+\sqrt{\kappa}}$ then

$$f(x_{k+1}) - \psi_{k+1} \leq \left(1 - \frac{1}{\sqrt{\kappa}}\right) [f(x_k) - \psi_k]$$

Proof. Now by our assumption on y_k we have

$$v_k - y_k = \frac{1}{1-\alpha} [y_k - \alpha \cdot x_k] - y_k = \frac{\alpha}{1-\alpha} [y_k - x_k] .$$

Consequently, since $f(x_k) \geq f(y_k) + \nabla f(y_k)^\top (x_k - y_k)$ by convexity we have

$$\begin{aligned} \frac{\mu}{2} \left\| v_k - \left[y_k - \frac{1}{\mu} \nabla f(y_k) \right] \right\|_2^2 &= \frac{\mu}{2} \left[\|v_k - y_k\|_2^2 + \frac{2}{\mu} \nabla f(y_k)^\top (v_k - y_k) + \frac{1}{\mu^2} \|\nabla f(y_k)\|_2^2 \right] \\ &\geq \frac{\alpha}{1-\alpha} \cdot \nabla f(y_k)^\top (y_k - x_k) + \frac{1}{2\mu} \|\nabla f(y_k)\|_2^2 \\ &\geq \frac{\alpha}{1-\alpha} \cdot [f(y_k) - f(x_k)] + \frac{1}{2\mu} \|\nabla f(y_k)\|_2^2 . \end{aligned}$$

Consequently, by Lemma 5 we have

$$\psi_{k+1} \geq \beta \cdot \psi_k + (1-\beta) \cdot \left[f(y_k) - \frac{1}{2\mu} \|\nabla f(y_k)\|_2^2 \right] + \beta(1-\beta) \cdot \left[\frac{\alpha}{1-\alpha} \cdot [f(y_k) - f(x_k)] + \frac{1}{2\mu} \|\nabla f(y_k)\|_2^2 \right]$$

Combining this with the fact that $f(x_{k+1}) \leq f(y_k) - \frac{1}{2L} \|\nabla f(y_k)\|_2^2$ yields

$$\begin{aligned} f(x_{k+1}) - \psi_{k+1} &\leq \beta\alpha \cdot \frac{1-\beta}{1-\alpha} \cdot f(x_k) - \beta\psi_k + \left[1 - (1-\beta) - \alpha\beta \cdot \frac{1-\beta}{1-\alpha} \right] f(y_k) \\ &\quad + \left[\frac{1-\beta}{2\mu} - \beta \cdot (1-\beta) \cdot \frac{1}{2\mu} - \frac{1}{2L} \right] \|\nabla f(y_k)\|_2 \end{aligned}$$

This yields the first formula. The values for β and α were chosen by solving for $(1-\beta)^2 = \frac{\mu}{L} = \frac{1}{\sqrt{\kappa}}$ yielding the first formula and then solving for $\alpha \cdot \frac{1-\beta}{1-\alpha} = 1$ which yields that $\frac{\alpha}{\sqrt{\kappa}} = 1 - \alpha$ which then yields $\alpha = \frac{1}{1+\frac{1}{\sqrt{\kappa}}} = \frac{\sqrt{\kappa}}{1+\sqrt{\kappa}}$. \square

5 Accelerated Gradient Descent Guarantees

In the last section we showed how to construct a step that decreased an upper bound on $f(x_k) - f_*$ by a multiplicative $1 - \sqrt{\frac{\mu}{L}}$ in every iteration. To turn this into a full algorithm, all that remains is to show how to bound the initial error, i.e. how to get an initial quadratic lower bound on our function. However, by Lemma 3 we already know that how to get a lower bound, so all that remains is to analyze the initial error with this lower bound.

Lemma 7. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a L -smooth μ -strongly convex function then for any x_0 we have that for*

$$\psi_0 = f(x_0) - \frac{1}{2\mu} \|\nabla f(x_0)\|_2^2 \text{ and } v_0 = x_0 - \frac{1}{\mu} \nabla f(x_0)$$

it is the case that $f(x) \geq L_0(x) \stackrel{\text{def}}{=} \psi_0 + \frac{\mu}{2} \|x - v_0\|_2^2$ and $f(x_0) - \psi_0 \leq \frac{L}{\mu} \cdot [f(x_0) - f_]$.*

Proof. The fact that $f(x) \geq L_0(x)$ is immediate from Lemma 3. Since $\|\nabla f(x_0)\|_2^2 \leq 2L \cdot [f(x_0) - f_*]$ we obtain the desired upper bound on $f(x_0) - \psi_0$. \square

Putting this all together yields the following.

Theorem 8 (Accelerated Gradient Descent). *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a L -smooth μ -strongly convex function and let $\kappa = \frac{L}{\mu}$. For arbitrary $x_0 \in \mathbb{R}^n$ compute $v_0 = x_0 - \frac{1}{\mu} \nabla f(x_0)$ and for all $k \geq 0$ let*

- $y_{k+1} = \alpha \cdot x_k + (1 - \alpha) \cdot v_k$ for $\alpha = \frac{\sqrt{\kappa}}{1 + \sqrt{\kappa}}$
- $v_{k+1} = \beta \cdot v_k + (1 - \beta) \cdot \left[y_k - \frac{1}{\mu} \nabla f(y_k) \right]$ for $\beta = 1 - \frac{1}{\sqrt{\kappa}}$
- $x_{k+1} = y_k - \frac{1}{L} \nabla f(y_k)$

Then we have that $f(x_k) - f_ \leq \left(1 - \frac{1}{\sqrt{\kappa}}\right)^k \cdot \kappa \cdot [f(x_0) - f_*]$ and consequently we can compute an ϵ -suboptimal point for f with $1 + \lceil \sqrt{\kappa} \log(\kappa \cdot [f(x_0) - f_*]) \rceil$ queries to a gradient oracle.*

Proof. For all k if we let $L_k(x) = \psi_k + \frac{\mu}{2} \|x - v_0\|_2^2$ then we have by previous lemmas that there is a way to chose the ψ_k such that $f(x) \geq L_k(x)$ for all x and therefore $f(x_k) - f_* \leq f(x_k) - \psi_k$. We have also proven that this can be done so that $f(x_{k+1}) - \psi_{k+1} \leq \left(1 - \frac{1}{\sqrt{\kappa}}\right) [f(x_k) - \psi_k]$ and $f(x_0) - \psi_0 \leq \kappa \cdot [f(x_0) - f_*]$ yielding the result. \square

6 Improving the Analysis

A natural question to ask is can the accelerated gradient descent algorithm be further improved? It can be shown that the dependence on κ in the asymptotic rate cannot, in general, be improved if the function can be accessed only through a gradient oracle.

However, the κ in the logarithmic term can be improved just by slightly improving the analysis. Rather than tracking $f(x_k) - \psi_k$ there is another natural potential function can be used, the sum of the function error of $f(x_k)$ and the appropriately scaled distance of v_k to the optimal point, i.e. $\frac{\mu}{2} \|x_* - v_k\|_2^2$. In the homework, you will show that the same elements of the above proof yield the following.

Lemma 9. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a L -smooth μ -strongly convex function and let $\kappa = \frac{L}{\mu}$. Show that if for some x_k and v_k we let*

- $y_{k+1} = \alpha \cdot x_k + (1 - \alpha) \cdot v_k$ for $\alpha = \frac{\sqrt{\kappa}}{1 + \sqrt{\kappa}}$
- $v_{k+1} = \beta \cdot v_k + (1 - \beta) \cdot \left[y_k - \frac{1}{\mu} \nabla f(y_k) \right]$ for $\beta = 1 - \frac{1}{\sqrt{\kappa}}$
- $x_{k+1} = y_k - \frac{1}{L} \nabla f(y_k)$

Then if x_ is the unique minimizer of f and we let $\epsilon_k \stackrel{\text{def}}{=} f(x_k) - f_*$ and $r_k \stackrel{\text{def}}{=} \frac{\mu}{2} \|v_k - x_*\|_2^2$ then*

$$\epsilon_{k+1} + r_{k+1} \leq \left(1 - \frac{1}{\sqrt{\kappa}}\right) [\epsilon_k + r_k].$$

This gives the following improved analysis of accelerated gradient descent

Theorem 10 (Accelerated Gradient Descent (Improved)). *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a L -smooth μ -strongly convex function and let $\kappa = \frac{L}{\mu}$. For arbitrary $x_0 \in \mathbb{R}^n$ let $v_0 = x_0$ and for all $k \geq 0$ let*

- $y_k = \alpha \cdot x_k + (1 - \alpha) \cdot v_k$ for $\alpha = \frac{\sqrt{\kappa}}{1 + \sqrt{\kappa}}$

- $v_{k+1} = \beta \cdot v_k + (1 - \beta) \cdot \left[y_k - \frac{1}{\mu} \nabla f(y_k) \right]$ for $\beta = 1 - \frac{1}{\sqrt{\kappa}}$
- $x_{k+1} = y_k - \frac{1}{L} \nabla f(y_k)$

Then we have that $f(x_k) - f_* \leq \left(1 - \frac{1}{\sqrt{\kappa}}\right)^k \cdot 2 \cdot [f(x_0) - f_*]$ and consequently we can compute an ϵ -suboptimal point for f with $\lceil \sqrt{\kappa} \log(2 \cdot [f(x_0) - f_*]) \rceil$ queries to a gradient oracle.

Proof. By the previous theorem we have that for $x_* \in X_*(f)$ it is the case that

$$f(x_k) - f_* + \frac{\mu}{2} \|v_k - x_*\|_2^2 \leq \left(1 - \frac{1}{\sqrt{\kappa}}\right)^k \cdot \left[f(x_0) - f_* + \frac{\mu}{2} \|v_0 - x_*\|_2^2 \right].$$

Since $v_0 = x_0$ and by strong convexity we have that $\frac{\mu}{2} \|x_0 - x_*\|_2^2 \leq f(x_0) - f_*$ the result follows. \square

7 Momentum

Another popular viewpoint or perspective on acceleration is that it can be viewed as gaining momentum in some sense, i.e. once you move in the direction of the gradient you keep moving in that direction for some time afterwards. This view can be confirmed by a rearranging of the variables in the method we derived. This gives another popular statement of the accelerated gradient descent algorithm.

Theorem 11. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a L -smooth μ -strongly convex function and let $\kappa = \frac{L}{\mu}$. For arbitrary $x_0 \in \mathbb{R}^n$ let $x_1 = x_0 - \frac{1}{L} \nabla f(x_0)$ and for all $k \geq 1$ let

- $y_k = x_k + \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right) (x_k - x_{k-1})$
- $x_{k+1} = y_k - \frac{1}{L} \nabla f(y_k)$

Then we have that $f(x_k) - f_* \leq \left(1 - \frac{1}{\sqrt{\kappa}}\right)^k \cdot 2 \cdot [f(x_0) - f_*]$ and consequently we can compute an ϵ -suboptimal point for f with $\lceil \sqrt{\kappa} \log(2 \cdot [f(x_0) - f_*]) \rceil$ queries to a gradient oracle.

Proof. We obtain this by massaging the previous algorithm. Note that in Theorem 10 we have that

$$v_{k-1} = \frac{1}{1-\alpha} [y_{k-1} - \alpha \cdot x_{k-1}] \text{ and } \nabla f(y_{k-1}) = L(y_{k-1} - x_k)$$

and therefore

$$\begin{aligned} y_k &= \alpha \cdot x_k + (1 - \alpha) \cdot v_k \\ &= \alpha \cdot x_k + (1 - \alpha) \cdot \left[\beta \cdot v_{k-1} + (1 - \beta) \cdot \left[y_{k-1} - \frac{1}{\mu} \nabla f(y_{k-1}) \right] \right] \\ &= \alpha \cdot x_k + \beta \cdot [y_{k-1} - \alpha \cdot x_{k-1}] + (1 - \alpha)(1 - \beta) \cdot [y_{k-1} - \kappa \cdot [y_{k-1} - x_k]] \\ &= [\alpha + \kappa \cdot (1 - \alpha) \cdot (1 - \beta)] \cdot x_k - \alpha \cdot \beta \cdot x_{k-1} - [\beta + (1 - \alpha)(1 - \beta)(1 - \kappa)] \cdot y_{k-1} \end{aligned}$$

Now since $\alpha = \frac{\sqrt{\kappa}}{1+\sqrt{\kappa}}$ and $\beta = 1 - \frac{1}{\sqrt{\kappa}}$ we have

$$(1 - \alpha) \cdot (1 - \beta) \cdot (1 + \kappa) = \frac{1}{1 + \sqrt{\kappa}} \cdot \frac{1}{\sqrt{\kappa}} \cdot (1 - \kappa) = \frac{1 - \sqrt{\kappa}}{\sqrt{\kappa}} = -\beta$$

and we have that

$$\begin{aligned}\alpha - \kappa \cdot (1 - \alpha) \cdot (1 - \beta) &= \frac{\sqrt{\kappa}}{1 + \sqrt{\kappa}} + \kappa \cdot \left(1 - \frac{\sqrt{\kappa}}{1 + \sqrt{\kappa}}\right) \cdot \left(1 - \left(1 - \frac{1}{\sqrt{\kappa}}\right)\right) \\ &= \frac{\sqrt{\kappa}}{1 + \sqrt{\kappa}} + \sqrt{\kappa} \cdot \frac{1}{1 + \sqrt{\kappa}} = \frac{2\sqrt{\kappa}}{1 + \sqrt{\kappa}} = 1 + \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}\end{aligned}$$

and

$$\alpha \cdot \beta = \frac{\sqrt{\kappa}}{1 + \sqrt{\kappa}} \cdot \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa}} = \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}$$

and thus the result follows by Theorem 10. \square

8 Non-strongly Convex Functions

How can we use the above result to minimize non-strongly convex functions? There is a fairly general trick to reduce non-strongly convex function minimization to strongly convex function minimization and that is *regularization*. This is a fairly overloaded term with all sorts of applications and interpretations, particularly in machine learning. When we use this term in the class though, we will simply use it to refer to the idea of adding a simple function we understand to improve the behavior of our iterative methods.

The idea we use here is simple. Instead of minimizing $f(x)$ directly, given some point we just minimize $g(x) = f(x) + \frac{\mu}{2}\|x - x_0\|_2^2$. Clearly this function is μ strongly convex and thus we can apply accelerated gradient descent as analyzed above to it. Below we analyze the performance of this scheme.

Lemma 12. *If f is a L -smooth convex function then given any $x_0 \in \mathbb{R}^n$ we can compute an ϵ -suboptimal point with*

$$\left\lceil \sqrt{1 + \frac{L \cdot \|x_0 - x_*\|_2^2}{\epsilon}} \log \left(\frac{L \cdot \|x_0 - x_*\|_2^2}{\epsilon} \right) \right\rceil$$

queries for any $x_* \in X_*(f)$.

Proof. Given $x_0 \in \mathbb{R}^n$ we run accelerated gradient descent to minimize $g(x) = f(x) + \frac{\mu}{2}\|x - x_0\|_2^2$ for a value of μ we pick later. Since f is μ -strongly convex and $L + \mu$ smooth we have that we can compute an ϵ -sub-optimal point x_ϵ for g with $\sqrt{\frac{L+\mu}{\mu}} \log(2 \cdot [g(x_0) - g_*]/\epsilon)$. Now, since $g(x) \geq f(x)$ for all x we have $g_* \geq f_*$ and $g(x_0) - g_* \leq f(x_0) - f_* \leq \frac{L}{2} \cdot \|x_0 - x_*\|_2^2$. Furthermore, this implies that

$$f(x_\epsilon) \leq \epsilon + \min_{x \in \mathbb{R}^n} f(x) + \frac{\mu}{2}\|x - x_*\|_2^2 \leq \epsilon + f_* + \frac{\mu}{2}\|x_0 - x_*\|_2^2.$$

Consequently, so long as we compute $x_{\frac{\epsilon}{2}}$ and $\mu = \frac{\epsilon}{\|x_0 - x_*\|_2^2}$ then we have the desired result. \square

There are two natural ways to remove the log factor in the above analysis. The first is to change the accelerated gradient descent algorithm itself to decay the value of μ used, the second is to minimize $f(x) + \frac{\mu}{2}\|x - x_0\|_2^2$ in phases changing perhaps what the regularization is with respect to. Both can be used to remove the logarithmic factors and they are similar in some sense. The first has the advantage of perhaps being a more natural way of running the algorithm, but the second has the virtue of being a fairly general reduction. We will talk more on this type of reduction in the next chapter.