Gradient Descent

David Woodruff
CMU

Norms and Inner Products

The inner product between two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ is written as $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_i x_i y_i$. Recall that the Euclidean norm of $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ is given by

$$\|\mathbf{x}\| = \sqrt{\sum_{i=1}^{n} x_i^2} = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}.$$

For any $c \in \mathbb{R}$ and $\mathbf{x} \in \mathbb{R}^n$, we get $\|c\mathbf{x}\| = |c| \cdot \|\mathbf{x}\|$, and also $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$. Moreover,

$$\|\mathbf{x} + \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 + 2 \langle \mathbf{x}, \mathbf{y} \rangle$$

Convexity

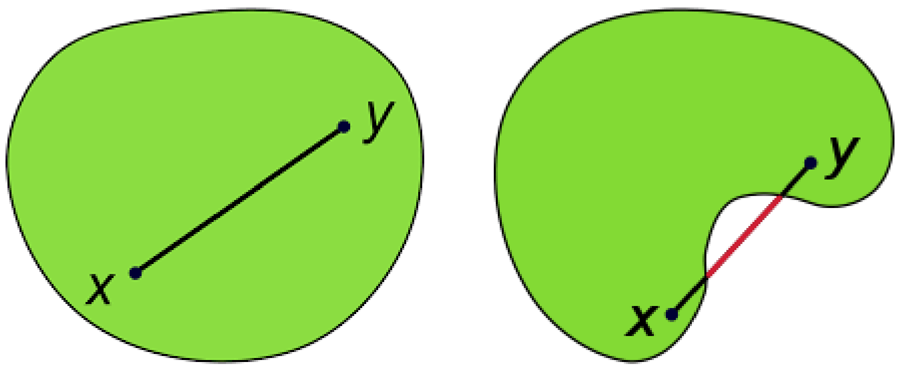

Definition 1 *A set $K \subseteq \mathbb{R}^n$ is said to be convex if*

$$(\lambda \mathbf{x} + (1-\lambda)\mathbf{y}) \in K \qquad \forall \mathbf{x}, \mathbf{y} \in K, \forall \lambda \in [0,1]$$

Definition 2 *For a convex set $K \subseteq \mathbb{R}^n$, a function $f : K \to \mathbb{R}$ is said to be convex over K iff*

$$f(\lambda \mathbf{x} + (1-\lambda)\mathbf{y}) \leq \lambda f(\mathbf{x}) + (1-\lambda) f(\mathbf{y}) \qquad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n, \forall \lambda \in [0,1]$$

Whenever K is not specified, assume $K = \mathbb{R}^n$.

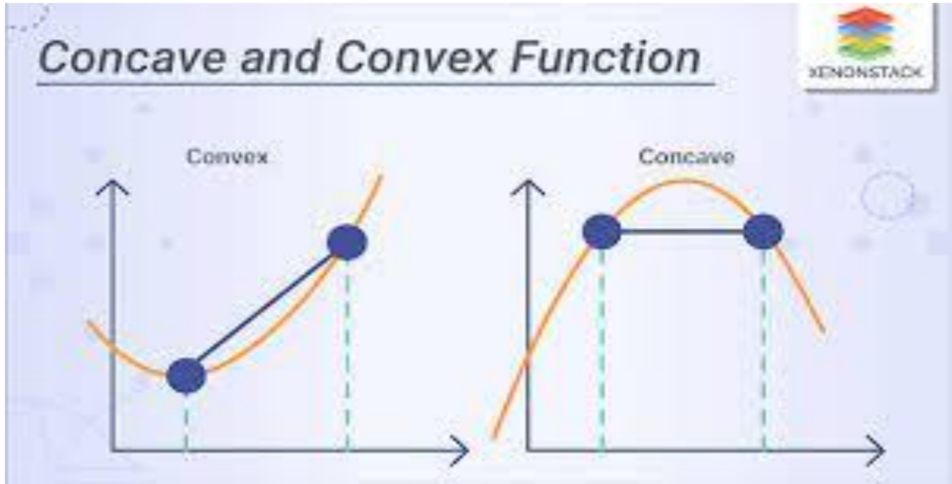

Optimization in the Real World

- How to solve: $\min \langle c, x \rangle$

  such that $Ax \leq b$

  $x \geq 0$

- Real life is not so pretty

- How to solve: $\min f(x)$

  such that $x \in K$

  f is a convex function and K is a convex set

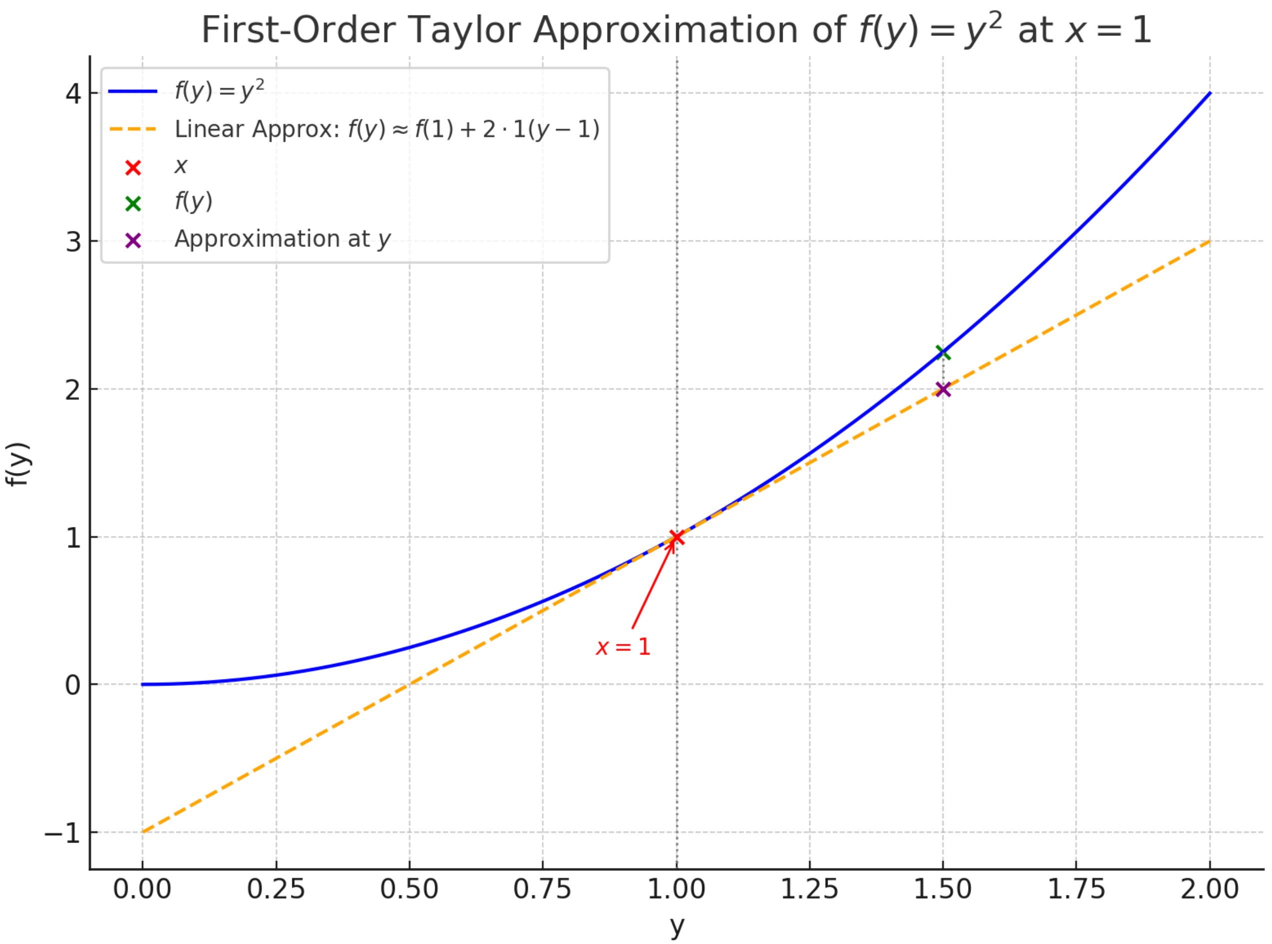

We can try to approximate f(y) around a value x as $f(y) = f((y-x) + x) \approx f(x) + f'(x)\cdot(y-x)$

If $f(x) = x^2$, then $f(y) \approx f(x) + 2x \cdot (y-x)$

Two Variable Calculus

- Now suppose $f(x_1, x_2)$ is a function of two variables, how to approximate $f(y_1, y_2)$?
- If $x_2 = y_2$, we could look at $\frac{\partial f}{\partial x_1}$ and $f(y_1, y_2) \approx f(x_1, x_2) + \frac{\partial f}{\partial x_1}(y_1 - x_1)$
- If $x_1 = y_1$, we could instead look at $\frac{\partial f}{\partial x_2}$ and $f(y_1, y_2) \approx f(x_1, x_2) + \frac{\partial f}{\partial x_2}(y_2 - x_2)$
- But both $x_2 \neq y_2$ and $x_1 \neq y_1$ might happen
- Use $f(y_1, y_2) \approx f(x_1, x_2) + \frac{\partial f}{\partial x_1}(y_1 - x_1) + \frac{\partial f}{\partial x_2}(y_2 - x_2)$

Gradients

In the context of this lecture, we will always assume that the function f is differentiable. The analog of the derivative in the multivariate case is the *gradient* ∇f, which is itself a function from $K \to \mathbb{R}^n$ defined as follows:

$$\nabla f(\mathbf{x}) = \left(\frac{\partial f}{\partial x_1}(\mathbf{x}), \ldots, \frac{\partial f}{\partial x_n}(\mathbf{x}) \right).$$

We assume that the gradient is well-defined at all points in K

Fact 3 *A differentiable function $f : K \to \mathbb{R}$ is convex iff $\forall \mathbf{x}, \mathbf{y} \in K$,*

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle.$$

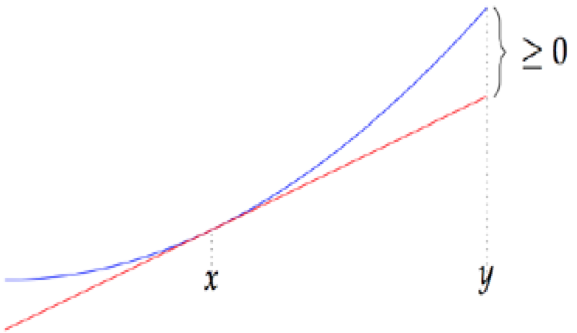

$f(x) + \langle \nabla f(x), y - x \rangle$ is the linear approximation of f(y) around x

Minimizing a Function

- To minimize a function, set its gradient to equal 0
  - Finds global minimum if f is a convex function
  - For non-convex function, still often finds a good solution, i.e., a local minimum
- Gradient is a very complicated expression, can’t solve by setting to 0
- Instead, update iteratively. This is called gradient descent

Gradient Descent

The basic gradient descent "framework" says:

Start with some point $\mathbf{x}_0$. At each step $t = 1, 2, \ldots, T-1$, set

$$\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t - \eta_t \cdot \nabla f(\mathbf{x}_t). \quad (1)$$

return $\hat{\mathbf{x}} = \frac{1}{T}\sum_{t=0}^{T-1} \mathbf{x}_t$.

- η_t is a "learning rate"
- We "move a little" in a direction (the negative gradient) that reduces loss function
  - Think of rolling ball down a hill
- Can just output x_t in practice, though often easier to prove statements about $\hat{x}$

$$f(x) = x^2$$

$$\frac{df}{dx} = 2x$$

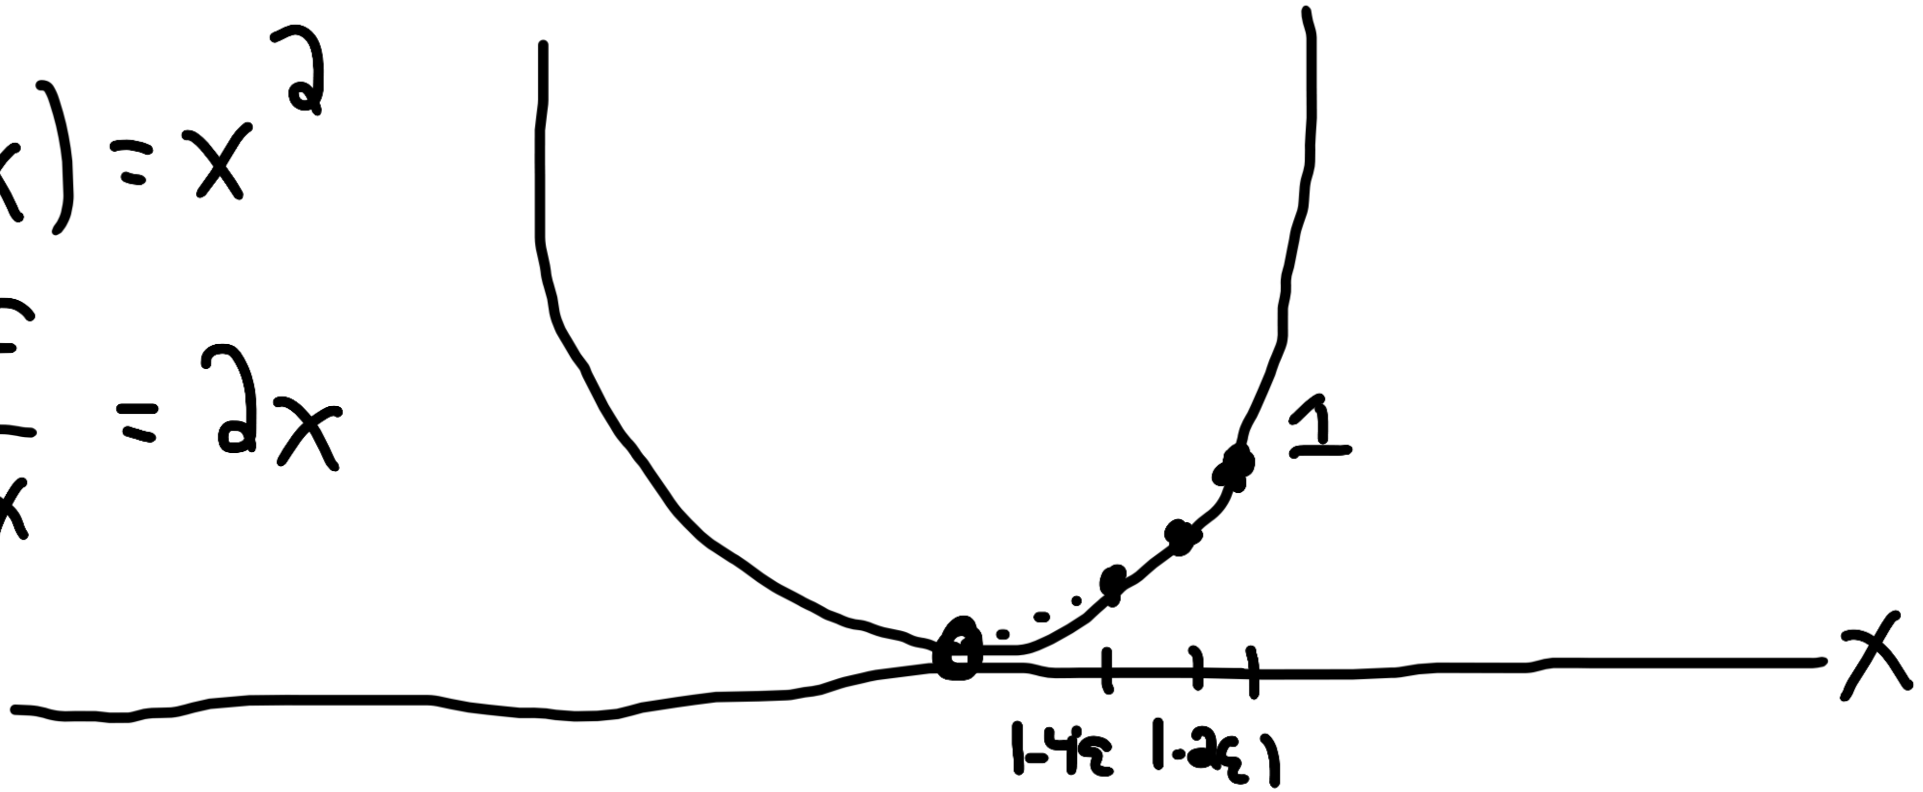

If $\eta = \epsilon$

If $x_0 = 1$, then $x_1 = x_0 - \epsilon \cdot 2 = 1 - 2\epsilon$

Also $x_2 = x_1 - \epsilon(2 - 4\epsilon) = 1 - 4\epsilon + O(\epsilon^2)$

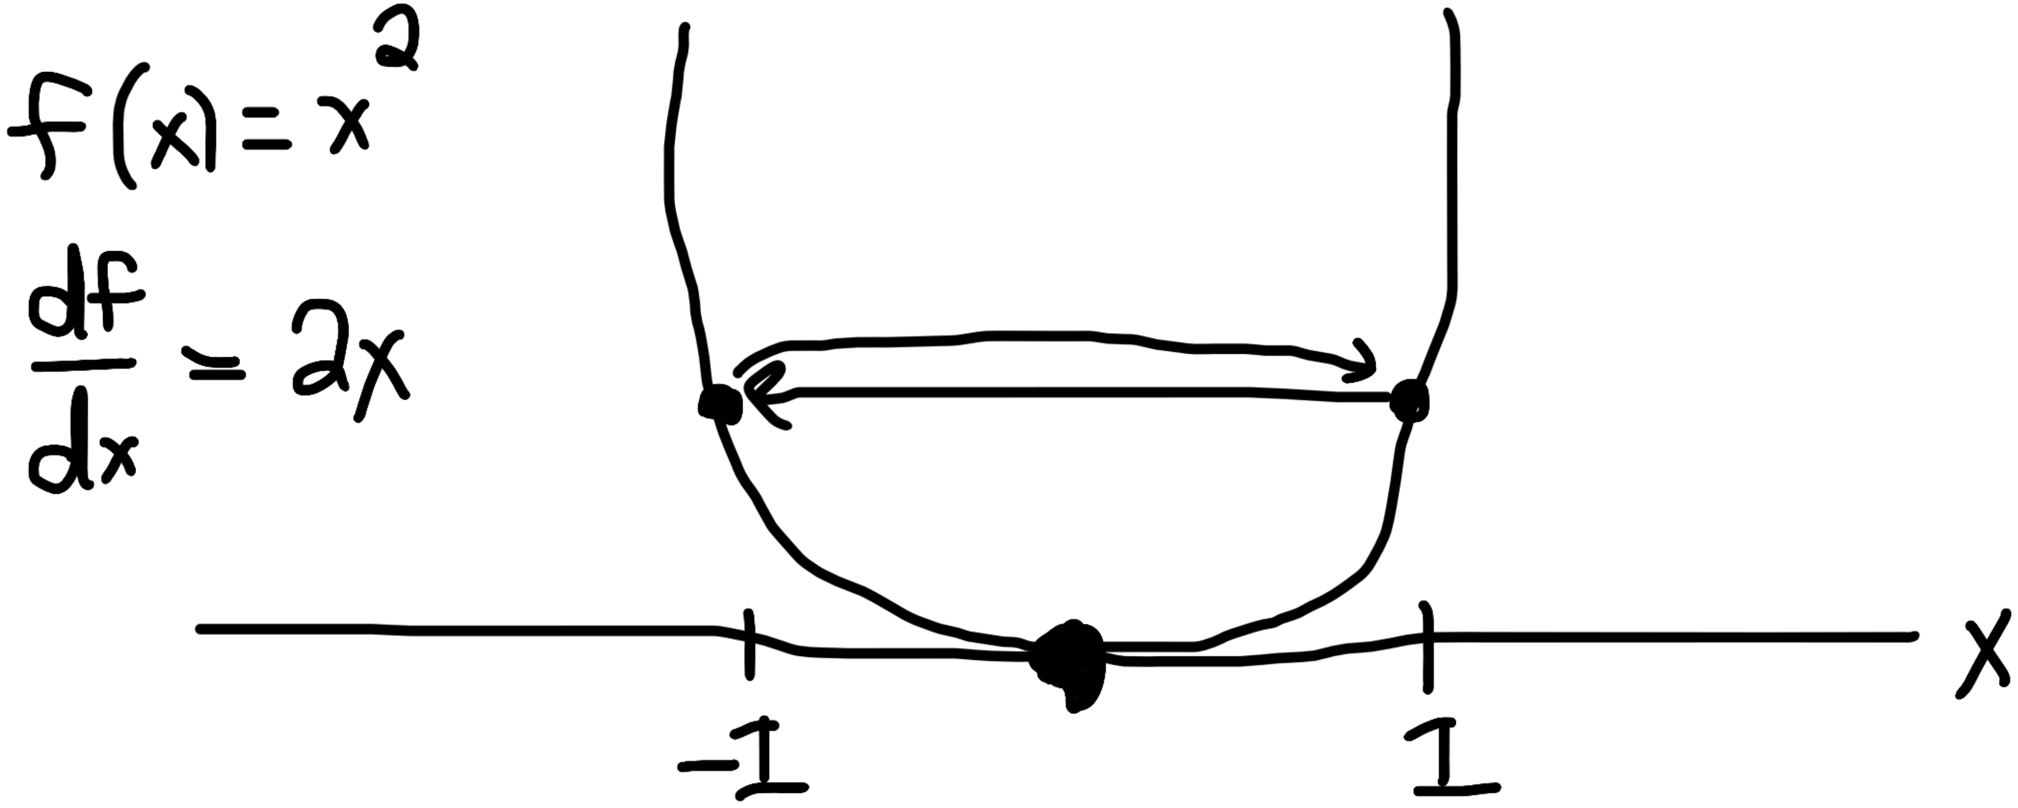

If $\eta = 1$

If $x_0 = 1$, then $x_1 = x_0 - 1 \cdot 2 = -1$

Also $x_2 = x_1 - 1 \cdot (-2) = 1$

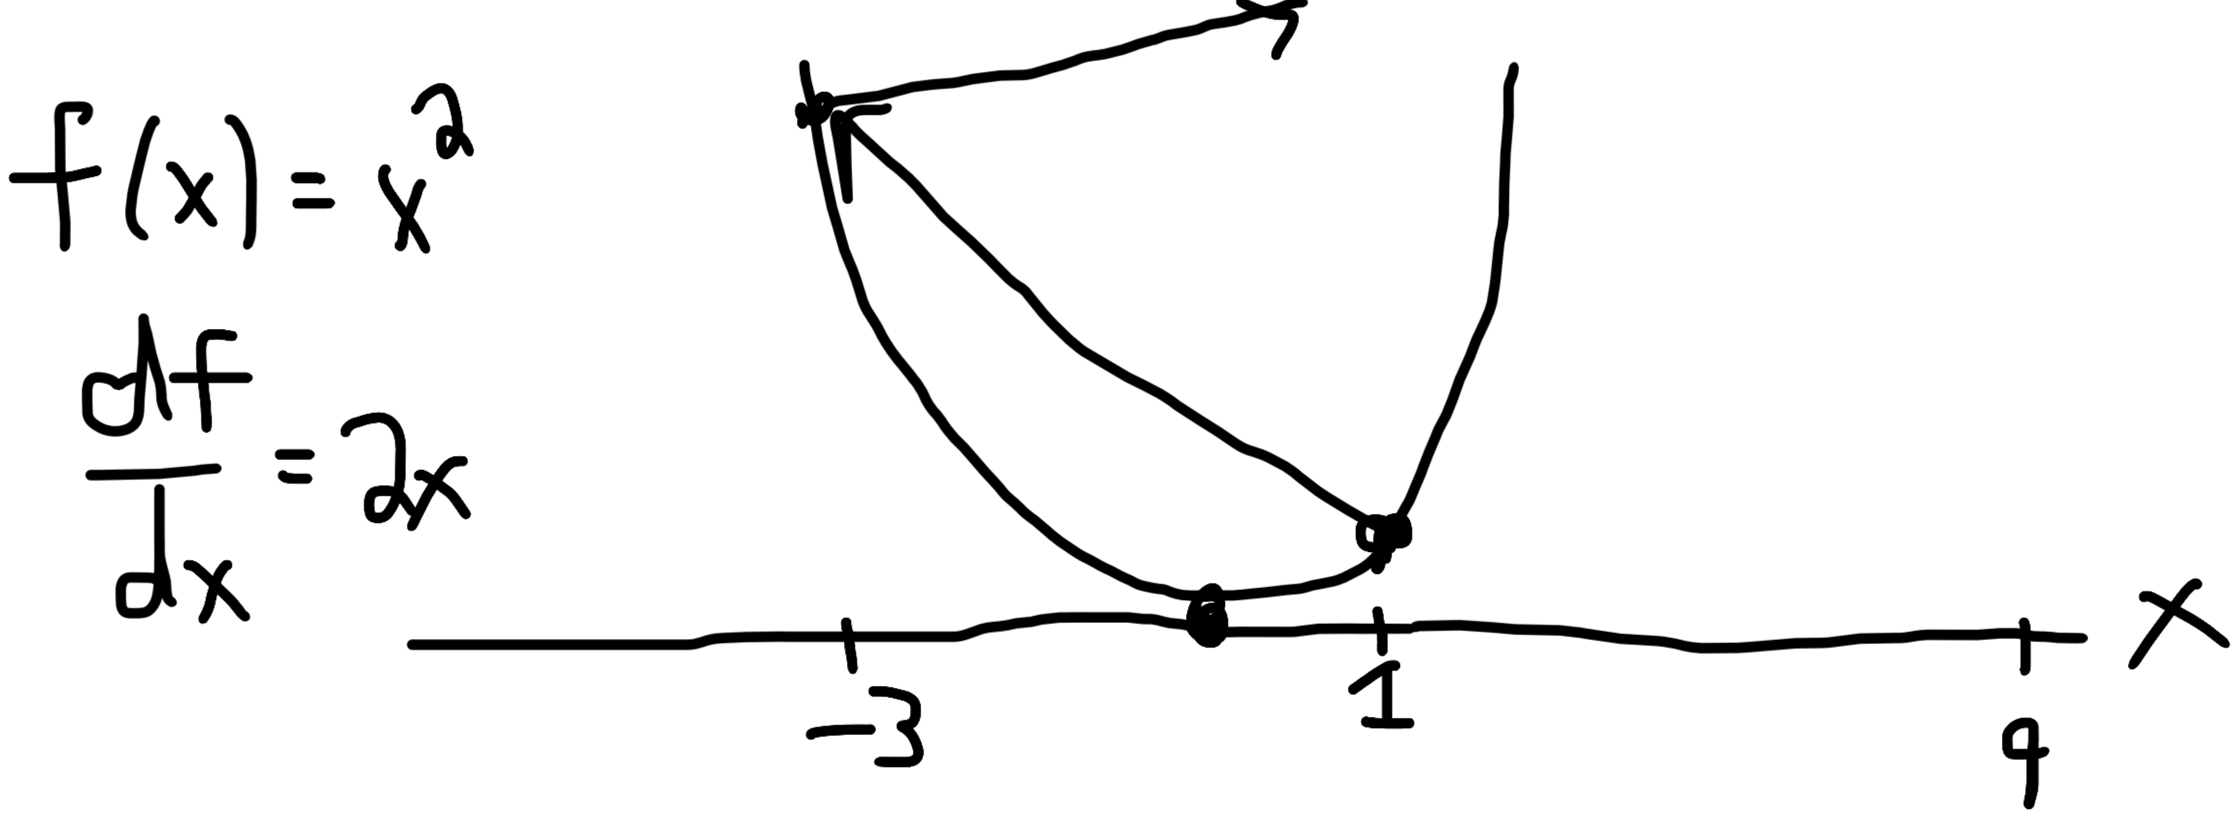

If $\eta = 2$

If $x_0 = 1$, then $x_1 = x_i - 2 \cdot 2 = -3$

Also $x_2 = x_1 - 2 \cdot (-6) = 9$

Gradient Descent Convergence

The analysis we give works for all convex functions. Its guarantee will depend on two things:

- The distance of the starting point $\mathbf{x}_0$ from the optimal point $\mathbf{x}^*$. Define $D := \|\mathbf{x}_0 - \mathbf{x}^*\|$.
- A bound G on the norm of the gradient at any point $\mathbf{x} \in \mathbb{R}^n$. Specifically, we want that $\|\nabla f(\mathbf{x})\| \leq G$ for all $\mathbf{x} \in \mathbb{R}^n$.[4]

A Stronger Statement – Online Gradient Descent

- Suppose we even allow the function f_t to change at each time step

Here's how to solve this problem. We can use almost the same update rule as (1), with one slight modification. The update rule is now taken with respect to gradient of the current function f_t.

$$\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t - \eta_t \cdot \nabla f_t(\mathbf{x}_t). \tag{2}$$

Theorem (Online Gradient Descent) *For any (differentiable) convex function* $f : \mathbb{R}^n \to \mathbb{R}$ *and any starting point* $\mathbf{x}_0$*, if we set* $\eta_t := \eta$*, then for any point* $\mathbf{x}^* \in \mathbb{R}^n$*,*

$$\sum_{t=0}^{T-1} f_t(\mathbf{x}_t) \leq \sum_{t=0}^{T-1} f_t(\mathbf{x}^*) + \frac{\eta}{2} G^2 T + \frac{1}{2\eta} D^2. \tag{3}$$

where G *is an upper bound on* $\max_t \|\nabla f_t\|$*, and* $D := \|\mathbf{x}_0 - \mathbf{x}^*\|$.

Theorem (Online Gradient Descent) *For any (differentiable) convex function $f : \mathbb{R}^n \to \mathbb{R}$ and any starting point $\mathbf{x}_0$, if we set $\eta_t := \eta$, then for any point $\mathbf{x}^* \in \mathbb{R}^n$,*

$$\sum_{t=0}^{T-1} f_t(\mathbf{x}_t) \leq \sum_{t=0}^{T-1} f_t(\mathbf{x}^*) + \frac{\eta}{2} G^2 T + \frac{1}{2\eta} D^2. \tag{3}$$

where G is an upper bound on $\max_t \|\nabla f_t\|$, and $D := \|\mathbf{x}_0 - \mathbf{x}^*\|$.

Proof: The proof is a short and sweet potential function argument. Define

$$\Phi_t := \frac{\|\mathbf{x}_t - \mathbf{x}^*\|^2}{2\eta}.$$

Note that $\Phi_0 = \frac{1}{2\eta} D^2$. We will show that

$$f_t(\mathbf{x}_t) + (\Phi_{t+1} - \Phi_t) \leq f_t(\mathbf{x}^*) + \frac{\eta}{2} G^2. \tag{4}$$

Summing this up over all times gives

$$\sum_{t=0}^{T-1} f_t(\mathbf{x}_t) + (\Phi_T - \Phi_0) \leq \sum_{t=0}^{T-1} f_t(\mathbf{x}^*) + \frac{\eta}{2} G^2 T.$$

Now using that $\Phi_T \geq 0$ and $\Phi_0 = D^2/(2\eta)$ completes the proof.

To prove (4), let's calculate

$$\begin{aligned}
\Phi_{t+1} - \Phi_t = \frac{1}{2\eta}\left(\|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 - \|\mathbf{x}_t - \mathbf{x}^*\|^2\right) &\overset{(F1)}{=} \frac{1}{2\eta}\left(\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 + 2\langle \mathbf{x}_{t+1} - \mathbf{x}_t, \mathbf{x}_t - \mathbf{x}^* \rangle\right) \\
&= \frac{1}{2\eta}\left(\eta^2 \|\nabla f_t(\mathbf{x}_t)\|^2 - 2\eta \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^* \rangle\right) \\
&\leq \frac{\eta}{2} G^2 - \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^* \rangle. \qquad (5)
\end{aligned}$$

Proof: The proof is a short and sweet potential function argument. Define

$$\Phi_t := \frac{\|\mathbf{x}_t - \mathbf{x}^*\|^2}{2\eta}.$$

Note that $\Phi_0 = \frac{1}{2\eta}D^2$. We will show that

$$f_t(\mathbf{x}_t) + (\Phi_{t+1} - \Phi_t) \leq f_t(\mathbf{x}^*) + \frac{\eta}{2}G^2. \tag{4}$$

Summing this up over all times gives

$$\sum_{t=0}^{T-1} f_t(\mathbf{x}_t) + (\Phi_T - \Phi_0) \leq \sum_{t=0}^{T-1} f_t(\mathbf{x}^*) + \frac{\eta}{2}G^2 T.$$

Now using that $\Phi_T \geq 0$ and $\Phi_0 = D^2/(2\eta)$ completes the proof.

To prove (4), let's calculate

$$\begin{aligned}
\Phi_{t+1} - \Phi_t = \frac{1}{2\eta}\left(\|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 - \|\mathbf{x}_t - \mathbf{x}^*\|^2\right) &\overset{(F1)}{=} \frac{1}{2\eta}\left(\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 + 2\langle \mathbf{x}_{t+1} - \mathbf{x}_t, \mathbf{x}_t - \mathbf{x}^*\rangle\right) \\
&= \frac{1}{2\eta}\left(\eta^2\|\nabla f_t(\mathbf{x}_t)\|^2 - 2\eta\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^*\rangle\right) \\
&\leq \frac{\eta}{2}G^2 - \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^*\rangle. \qquad (5)
\end{aligned}$$

Next we use the convexity of f (via Fact 3) to bound the difference

$$f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) \leq \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^*\rangle \tag{6}$$

Summing up (5) and (6) means the inner-product term cancels, and gives us the amortized cost bound (4), and hence proves Theorem ■

Constrained Minimization

Having done the analysis for the unconstrained case, we get the constrained case almost for free. The main difference is that the update step may take us outside K. So we just "project back into K". The algorithm is almost the same, let the blue parts highlight the changes.

Start with some point $\mathbf{x}_0$. At each step $t = 1, 2, \ldots, T-1$, set

$$\mathbf{y}_{t+1} \leftarrow \mathbf{x}_t - \eta_t \cdot \nabla f(\mathbf{x}_t).$$

Let $\mathbf{x}_{t+1}$ be the point in K closest to $\mathbf{y}_{t+1}$.
return $\widehat{\mathbf{x}} = \frac{1}{T}\sum_{t=0}^{T-1} \mathbf{x}_t$.

Now if we satisfy that $\|\nabla f(x)\| \leq G$ for all $x \in K$, the online optimization theorem remains exactly the same.

Theorem (Constrained Online Gradient Descent) *For any convex body $K \subseteq \mathbb{R}^n$, and sequence of (differentiable) convex functions $f_t : K \to \mathbb{R}$ and any starting point $\mathbf{x}_0$, if we set $\eta_t := \eta$, then for any point $\mathbf{x}^* \in \mathbb{R}^n$,*

$$\sum_{t=0}^{T-1} f_t(\mathbf{x}_t) \leq \sum_{t=0}^{T-1} f_t(\mathbf{x}^*) + \frac{\eta}{2} G^2 T + \frac{1}{2\eta} D^2. \tag{7}$$

where G is an upper bound on $\max_t \max_{\mathbf{x} \in K} \|\nabla f_t(\mathbf{x})\|$, and $D := \|\mathbf{x}_0 - \mathbf{x}^*\|$.

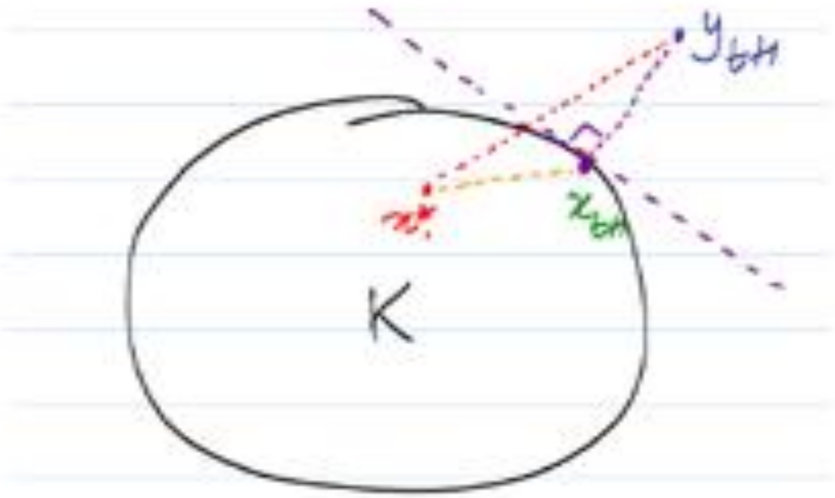

Figure 3: The projection ensures that $\mathbf{x}^*$ lies on the other side of the tangent hyperplane at $\mathbf{x}_{t+1}$, so the angle is obtuse. This means the squared length of the "hypotenuse" is larger than the squared length of either of the sides.

But now we claim that

$$\|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 \le \|\mathbf{y}_{t+1} - \mathbf{x}^*\|^2.$$

Indeed, since $\mathbf{x}_{t+1}$ is the "projection" of $\mathbf{y}_{t+1}$ onto the convex set K. The proof is essentially by picture (see Figure 3):

This means the changes die out almost immediately. Indeed,

$$\begin{aligned}\Phi_{t+1} - \Phi_t &\le \frac{1}{2\eta}\left(\|\mathbf{y}_{t+1} - \mathbf{x}^*\|^2 - \|\mathbf{x}_t - \mathbf{x}^*\|^2\right) = \frac{1}{2\eta}\left(\|\mathbf{y}_{t+1} - \mathbf{x}_t\|^2 + 2\langle \mathbf{y}_{t+1} - \mathbf{x}_t, \mathbf{x}_t - \mathbf{x}^*\rangle\right) \\ &= \frac{1}{2\eta}\left(\eta^2\|\nabla f_t(\mathbf{x}_t)\|^2 - 2\eta\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^*\rangle\right) \\ &\le \frac{\eta}{2}G^2 - \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^*\rangle.\end{aligned}$$