Lectures 15-16

Gradient Descent and Convexity

DSC 40A, Fall 2025

Question 🤔

Answer at q.dsc40a.com

Remember, you can always ask questions at q.dsc40a.com!

If the direct link doesn't work, click the "🤔 Lecture Questions" link in the top right corner of dsc40a.com.

When is gradient descent guaranteed to work?

Convex functions

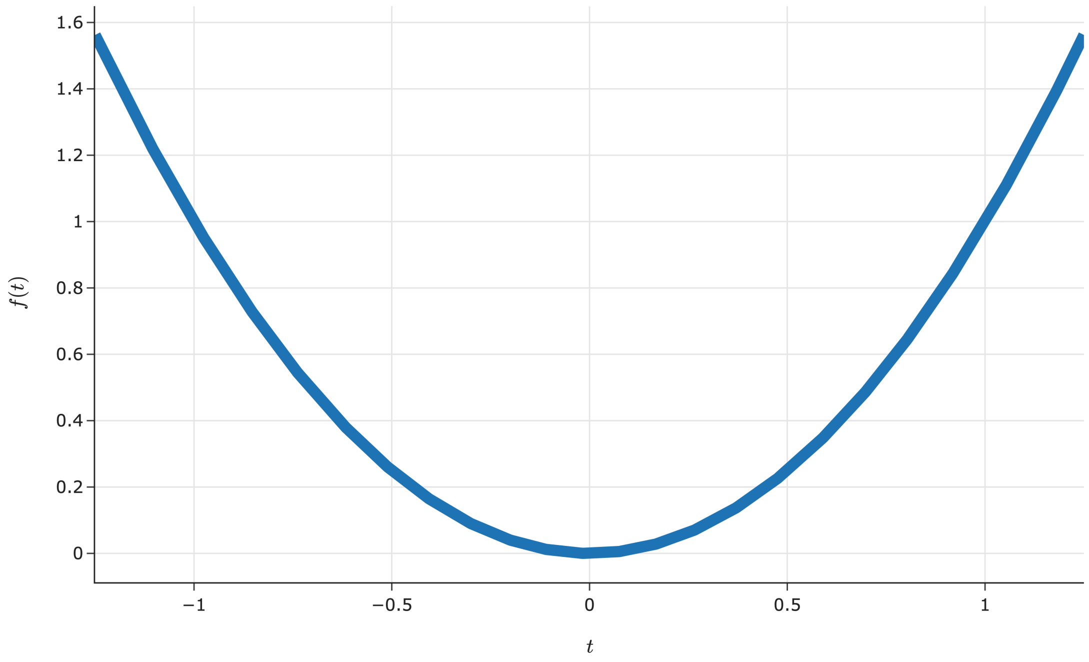

A **convex** function ✅

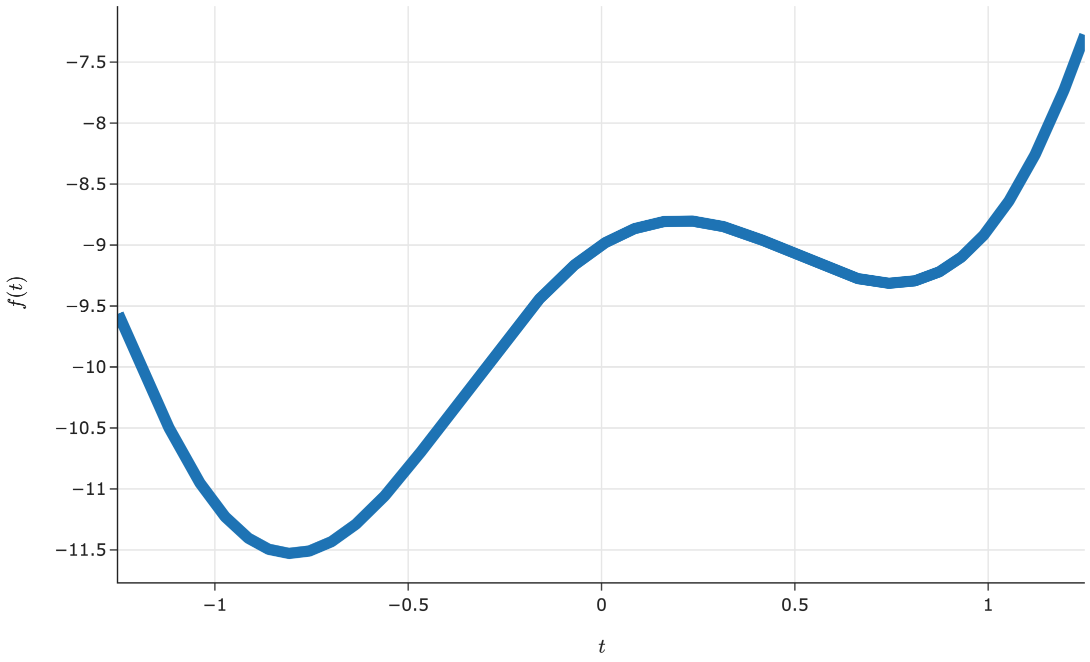

A **non-convex** function ❌

Convexity

- A function f is **convex** if, for **every** a, b in the domain of f, the line segment between:

$$(a, f(a)) \text{ and } (b, f(b))$$

does not go below the plot of f.

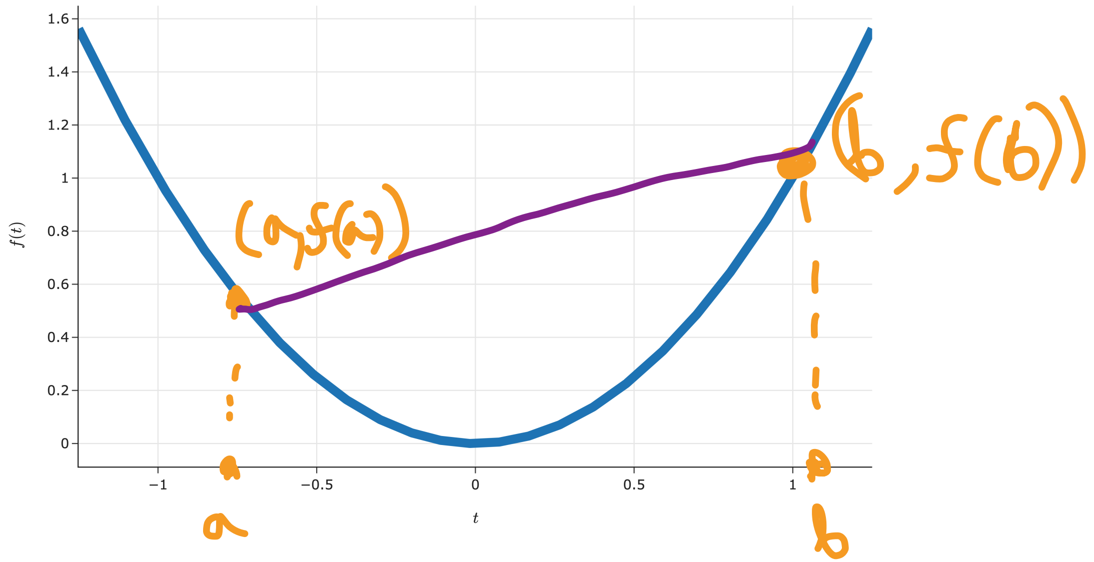

A **convex** function ✅

Question 🤔

Answer at q.dsc40a.com

Which of these functions are **not** convex?

- A. $f(x) = |x - 4|$.
- B. $f(x) = e^x$.
- C. $f(x) = \sqrt{x - 1}$.
- D. $f(x) = (x - 3)^{24}$.
- E. More than one of the above are non-convex.

Convex vs. concave

|x-4|

Convex

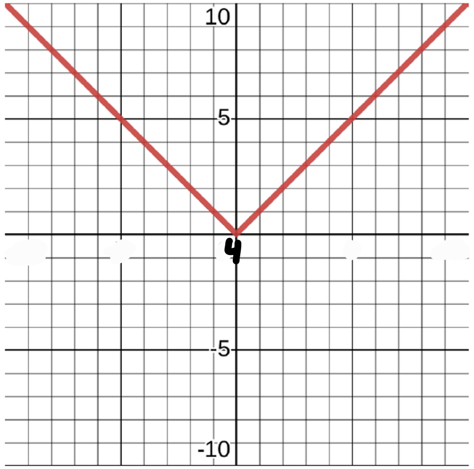

$f(x) = |x|$

Convex

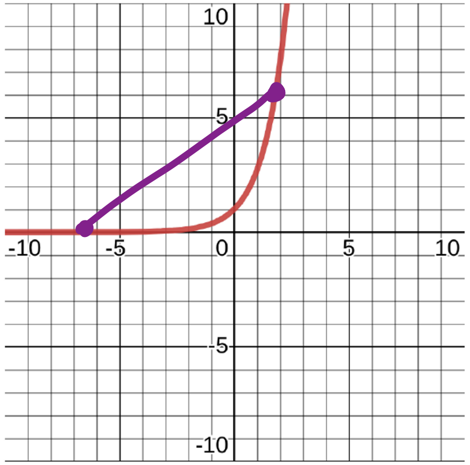

$f(x) = e^x$

Convex

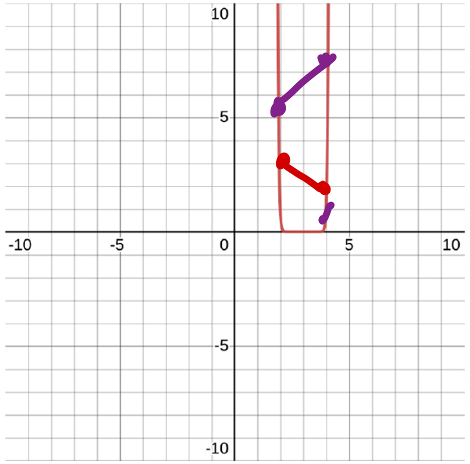

$f(x) = (x - 3)^{24}$ → even

Concave

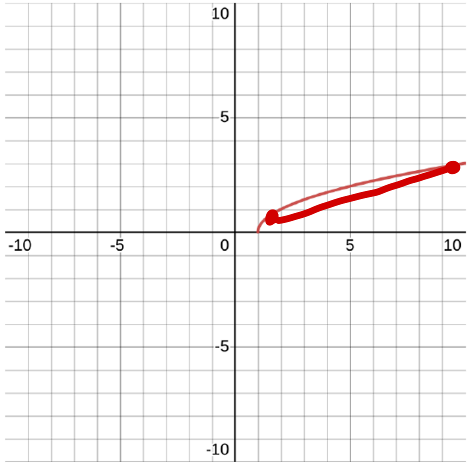

$f(x) = \sqrt{x - 1}$

Nonconvex functions and gradient descent

- We say a function is **nonconvex** if it does not meet the criteria for convexity.
- Nonconvex functions are (relatively) difficult to minimize.
- Gradient descent **might** still work, but it's not guaranteed to find a global minimum.
    - We saw this at the start of the lecture, when trying to minimize $f(t) = 5t^4 - t^3 - 5t^2 + 2t - 9$.

Choosing a step size in practice

- In practice, choosing a step size involves a lot of trial-and-error.
- In this class, we've only touched on "constant" step sizes, i.e. where α is a constant.

$$t_{i+1} = t_i - \alpha \frac{df}{dt}(t_i)$$

- **Remember**: α is the "step size", but the amount that our guess for t changes is $\alpha \frac{df}{dt}(t_i)$, not just α.
- In future courses, you'll learn about "decaying" step sizes, where the value of α decreases as the number of iterations increases.
  - Intuition: take much bigger steps at the start, and smaller steps as you progress, as you're likely getting closer to the minimum.

More examples

$H(x) = h$

Example: Huber loss and the constant model

- First, we learned about squared loss,

$$L_{\text{sq}}(y_i, H(x_i)) = (y_i - H(x_i))^2.$$

pro: differentiable, easy to minimize

con: sensitive to outliers

- Then, we learned about absolute loss,

$$L_{\text{abs}}(y_i, H(x_i)) = |y_i - H(x_i)|.$$

pro: robust to outliers

con: not differentiable

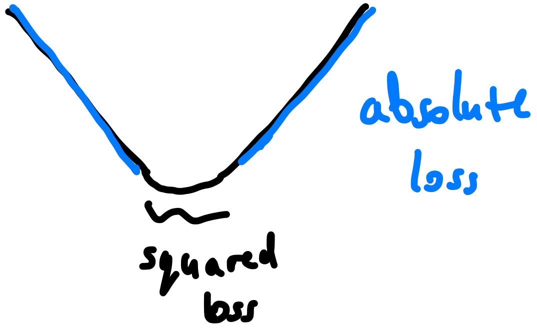

- Let's look at a new loss function, **Huber loss**:

$$L_{\text{huber}}(y_i, H(x_i)) = \begin{cases} \frac{1}{2}(y_i - H(x_i))^2 & \text{if } |y_i - H(x_i)| \leq \delta \\ \delta \cdot \left(|y_i - H(x_i)| - \frac{1}{2}\delta\right) & \text{otherwise} \end{cases}$$

δ → hyperparam

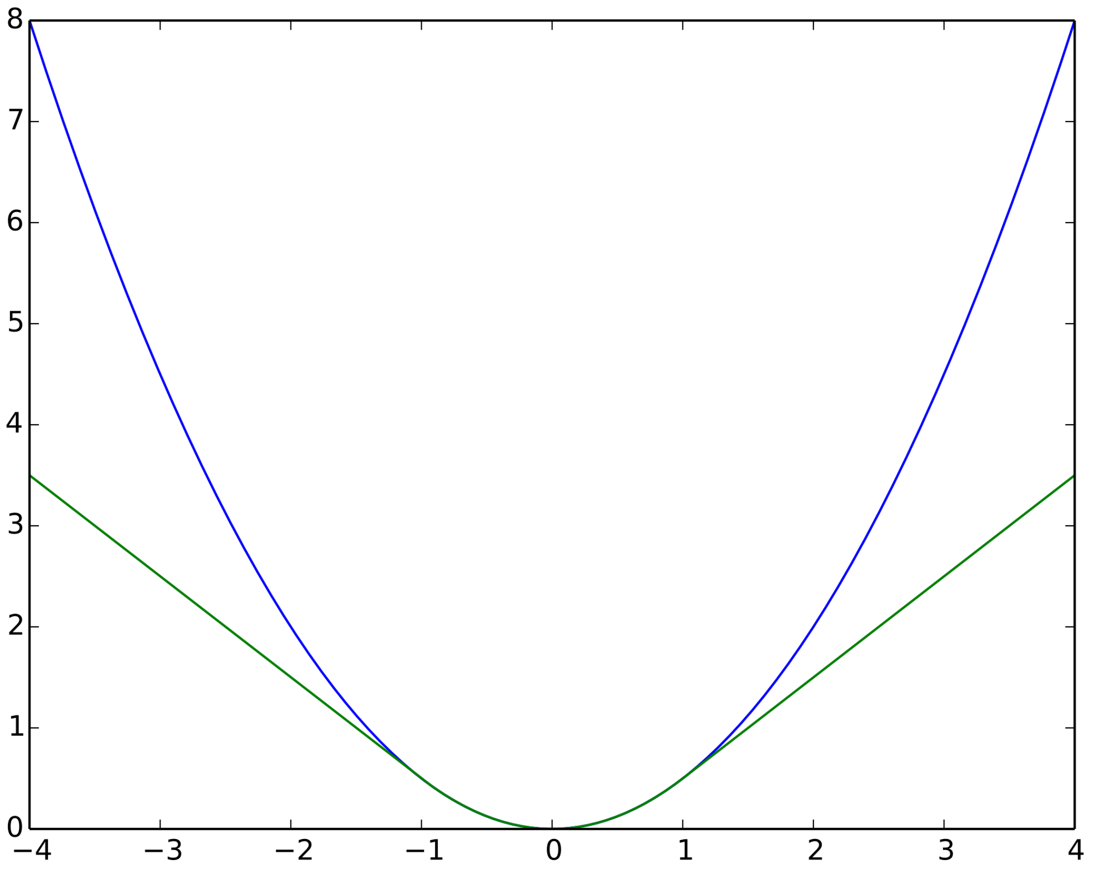

Squared loss in blue, **Huber** loss in green.

Note that both loss functions are convex!

Minimizing average Huber loss for the constant model

- For the constant model, $H(x) = h$:

$$L_{\text{huber}}(y_i, h) = \begin{cases} \frac{1}{2}(y_i - h)^2 & \text{if } |y_i - h| \leq \delta \\ \delta \cdot (|y_i - h| - \frac{1}{2}\delta) & \text{otherwise} \end{cases}$$

$$\implies \frac{\partial L}{\partial h}(h) = \begin{cases} -(y_i - h) & \text{if } |y_i - h| \leq \delta \\ -\delta \cdot \text{sign}(y_i - h) & \text{otherwise} \end{cases}$$

$$\text{sign}(x) = \begin{cases} +1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0 \end{cases}$$

- So, the **derivative** of empirical risk is:

$$\frac{dR_{\text{huber}}}{dh}(h) = \frac{1}{n}\sum_{i=1}^{n} \begin{cases} -(y_i - h) & \text{if } |y_i - h| \leq \delta \\ -\delta \cdot \text{sign}(y_i - h) & \text{otherwise} \end{cases}$$

- It's **impossible** to set $\frac{dR_{\text{huber}}}{dh}(h) = 0$ and solve by hand: we need gradient descent!

$$R_{sq} = \frac{1}{n}\sum_i \ell_{huber}(y_i, h)$$

Let's try this out in practice! Follow along in this notebook.