L01 Introduction

CS 169 / CS 268
Introduction to Optimization

Ioannis Panageas

Fall 2026

Course information

CS 169 and CS 268 meet together.
Assessments are distinct, with the same grading weights.

Tuesdays and Thursdays, 3:30–4:50 p.m.
ICS 174. All times are Pacific Time.

Instructor: Ioannis Panageas
ipanagea@ics.uci.edu
Office hours by appointment.

TA: Parnian Shahkar
shahkarp@uci.edu
TA: Aidan Sciortino
asciorti@uci.edu
Office hours to be announced on **EdStem**.

Course material

Announcements and questions: **EdStem**
The course-specific link will be announced.

Slides, syllabus, and course resources:

[CS 169 / CS 268 course website]

Recommended textbook

An Introduction to Optimization, fifth edition
Edwin K. P. Chong, Wu-Sheng Lu,
and Stanislaw H. Żak

Grading

Three midterms: 75% total
25% for each midterm

Three homeworks: 15% total
5% for each homework. Individual work.

Take-home final project: 10%
The project is the final assessment.
There is no in-person final exam.

Total: 100%
Deadlines and submission instructions on **EdStem**.

Important dates

Midterm 1: Thursday, October 15

Midterm 2: Thursday, November 5

Midterm 3: Thursday, December 3

Exams take place during class in ICS 174.
One hour for the exam, plus administration time.

Thursday, November 17: **No class**
Thursday, November 26: **Thanksgiving No class**

Missed-midterm policy

A missed midterm requires **instructor approval**.

Email the instructor as soon as you know about a conflict. Give advance notice whenever the conflict is known in advance.

If the reason is approved, the missing midterm's points are allocated equally among the other grading components. There is no makeup exam.

Academic integrity and access

Submit your own work. Follow the permitted resource and collaboration rules for each assessment.

For homework and the project, AI may improve writing, spelling, or grammar. Other uses require **express permission**. **Acknowledge any use.**

Academic dishonesty ordinarily results in an F. Suspected violations are reported to OAISC.

For accommodations, contact **UCI DSC**.
For an access barrier, email the instructor.

What is optimization?

Choose a feasible value of x that makes $f(x)$ as small as possible.

$$\min_{x \in S} f(x), \qquad S \subseteq \mathbb{R}$$

Decision variable: x
Objective function: f
Feasible set: S

Unconstrained: S is the real line.
Constrained: S restricts the permitted values of x.

The lemonade example

A manager wants to choose the price p.
Keep temperature and staffing fixed.

Toy model: demand is $10 - p$, with $0 \leq p \leq 10$.

$$R(p) = p(10 - p)$$

Maximize revenue, or minimize its negative:

$$-R(p) = p^2 - 10p = (p-5)^2 - 25$$

Best price: $p = 5$. This model has one variable.

Why study optimization?

Computer science and machine learning

Fit a predictor by minimizing prediction error.
Neural-network training also minimizes a loss.

Engineering

Choose one design parameter to reduce cost.

Operations research

Choose a departure time to reduce travel time.

Economics

Choose a price to maximize revenue.

Learning one parameter

Predict the same value θ for two observations, whose measured values are 2 and 4.

Choose θ to minimize the average squared error:

$$L(\theta) = \frac{(\theta - 2)^2 + (\theta - 4)^2}{2}$$

$$L(\theta) = (\theta - 3)^2 + 1$$

The best prediction is $\theta = 3$. Minimum loss: 1.

Convex sets in one variable

Definition (Convex set).
A set S of real numbers is convex if every point between any two points of S also belongs to S.

$$x, y \in S, \quad 0 \leq t \leq 1$$

$$(1-t)x + ty \in S$$

The weight t moves from x to y.
At $t = 1/2$, the weighted average is the midpoint.

Convex and nonconvex sets

Convex examples

$$[-2,2], \quad (0,1), \quad [0,\infty), \quad \mathbb{R}$$

A set with a gap is not convex:

$$S = [-2,-1] \cup [1,2]$$

Both -1 and 1 belong to S.
Their midpoint 0 does not belong to S.

Convex

$[0,2]$

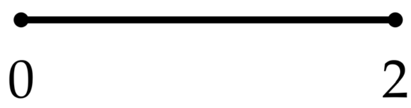

Nonconvex

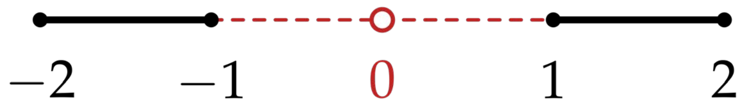

Convex functions

Definition (Convex function).
Let f be a real-valued function on a convex set S.
The function is convex if, for every x, y in S
and every t between 0 and 1,

$$f((1-t)x + ty) \leq (1-t)f(x) + tf(y)$$

The graph lies on or below every chord
joining two points on the graph.

A chord is the straight segment between two points.

Convexity in pictures

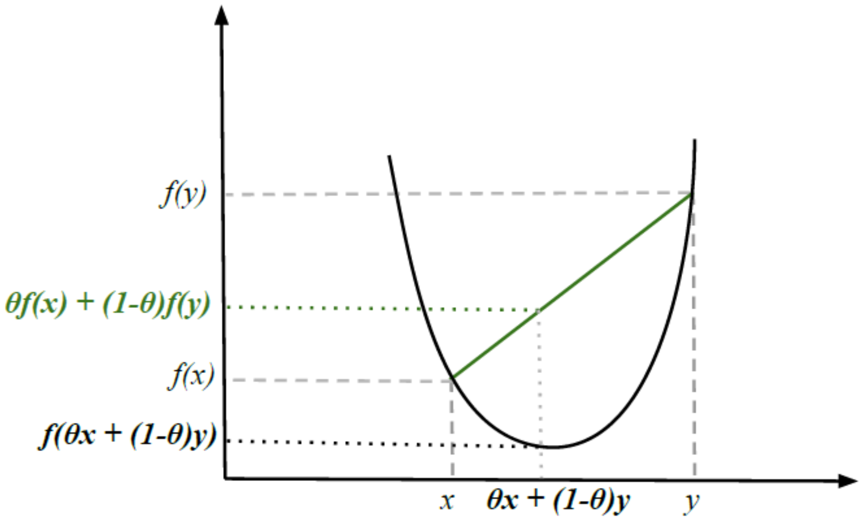

Convex

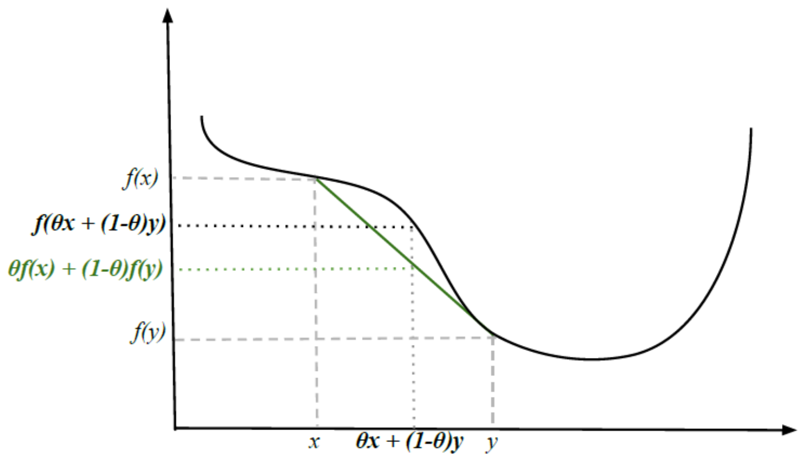

Nonconvex

Example: a quadratic function

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

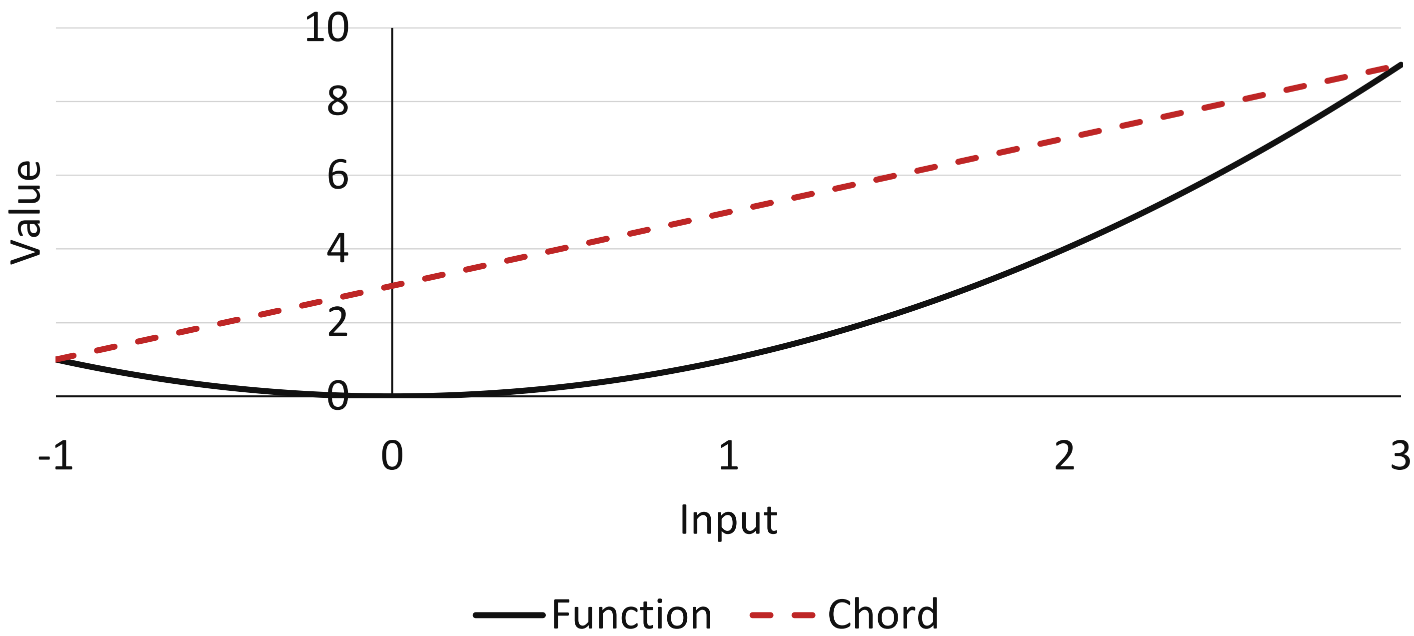

$$f(1) = 1 \leq \frac{f(-1) + f(3)}{2} = 5$$

Why the quadratic is convex

For any real x, y and any $0 \le t \le 1$:

$$(1-t)x^2 + ty^2 - ((1-t)x + ty)^2$$

$$= t(1-t)(x-y)^2 \ge 0$$

Every factor on the right is nonnegative.
Rearranging gives the convexity inequality.

More convex functions

$$f(x) = |x|$$

Convex, although the derivative does not exist at 0.

$$f(x) = ax + b$$

Affine functions satisfy the convexity inequality with equality. Here a and b are fixed constants.

Convexity does not require a unique minimizer.

A nonconvex function

$$f(x) = -x^2, \qquad x \in [-1, 1]$$

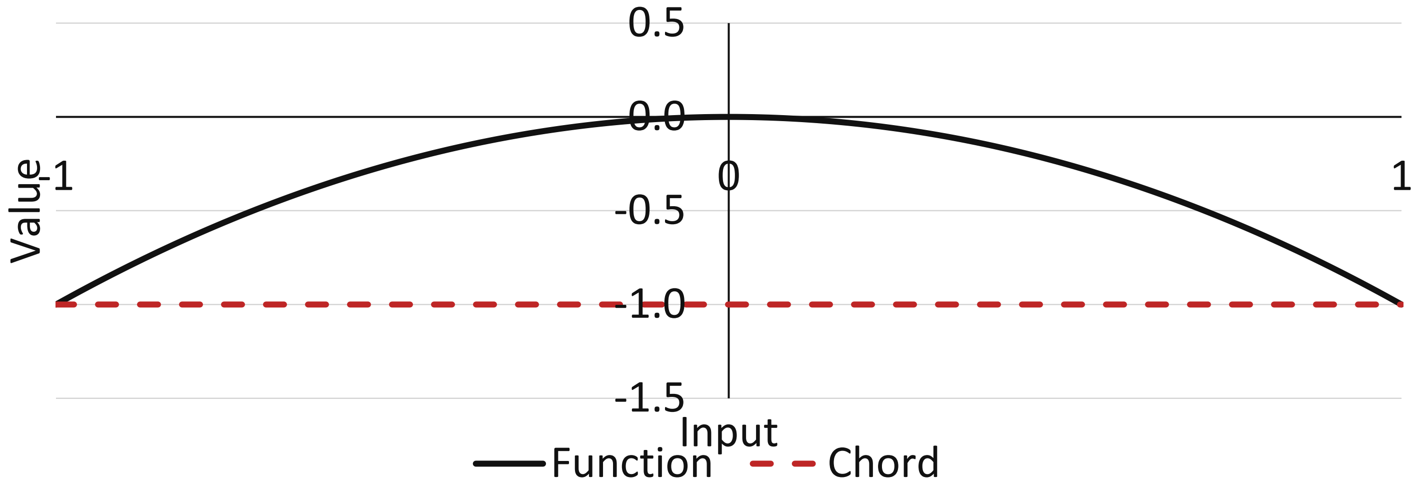

$$f(0) = 0 > \frac{f(-1) + f(1)}{2} = -1$$

The derivative and direction

The derivative $f'(x)$ is the local slope.

Positive derivative: a sufficiently small move left reduces the function value.

Negative derivative: a sufficiently small move right reduces the function value.

For a differentiable convex function on the real line, a zero derivative identifies a global minimizer.

Gradient descent in one variable

In one variable, the gradient is the derivative.
Start with an initial value x_0.

$$x_{k+1} = x_k - \eta f'(x_k)$$

x_k: the current iterate
$f'(x_k)$: the derivative at the current iterate
$\eta > 0$: the step size
$k = 0, 1, 2, \ldots$: the iteration number

Gradient descent for the quadratic

$$f(x) = x^2, \qquad f'(x) = 2x$$

Substitute the derivative into the update:

$$x_{k+1} = x_k - 2\eta x_k$$

$$x_{k+1} = (1 - 2\eta)x_k$$

The step size determines the multiplier.

A worked example

$$x_0 = 4, \quad \eta = \frac{1}{4}, \quad x_{k+1} = \frac{1}{2} x_k$$

Iteration	**Iterate**	**Objective value**
0	4	16
1	2	4
2	1	1
3	$\frac{1}{2}$	$\frac{1}{4}$
4	$\frac{1}{4}$	$\frac{1}{16}$

Each iterate halves. Each objective value quarters.

The iterates on the graph

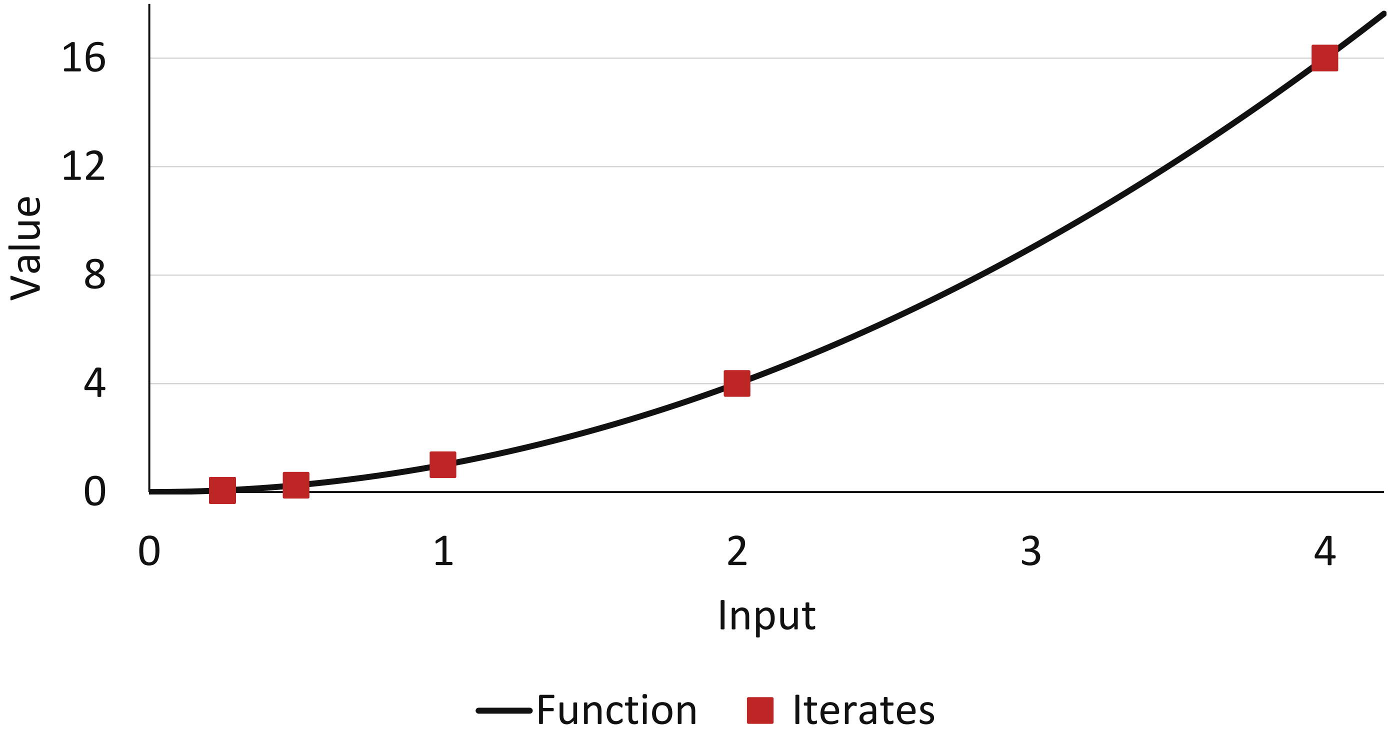

$$x_k = 4\left(\frac{1}{2}\right)^k \longrightarrow 0$$

The step size matters

$$x_{k+1} = (1 - 2\eta)x_k$$

Step size	Multiplier	Behavior
$\frac{1}{4}$	$\frac{1}{2}$	Converges
$\frac{1}{2}$	0	One step to zero
$\frac{3}{4}$	$-\frac{1}{2}$	Alternates, converges
1	-1	Oscillates
> 1	< -1	Diverges

For x^2 and a nonzero starting point:
convergence to 0 occurs exactly when $0 < \eta < 1$.

Questions to check understanding

Is the set $\{-1, 1\}$ convex?
Which two points show your answer?

Is every convex function differentiable?
Give an example.

For $f(x) = x^2$, start at $x_0 = 4$ and use $\eta = 3/4$.
What are the next two iterates?
Does the sequence converge?