Chapter 0: Introduction and Preliminaries

Ernest K. Ryu

MATH 164: Optimization
University of California, Los Angeles
Department of Mathematics

Last edited: January 29, 2025

Goal of this class

This class is about (mathematical) optimization.[1]

- Many engineering problems: We need to make a choice, and we want to make the "best" choice.
- Many scientific problems are: Nature will equilibrate at the "minimum" energy configuration, and we wish to find this configuration.
- Many data science problems: We want to find the model configuration (parameter) that "best" explains the data.

Mathematical optimization is the underlying math problem that abstracts away the engineering/scientific context. (Calculus is used to model physical systems, but calculating derivatives and integrals is independent of the physical context the calculus problems originate from.)

[1] Mathematical optimization contrasts with, say, compiler optimization or code optimization

Unconstrained optimization

An unconstrained optimization problem has the form

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad f(x),$$

where $f\colon \mathbb{R}^n \to \mathbb{R}$ has appropriate assumptions.

We refer to x as the *optimization variable* or *decision variable* and f as the *objective function* or *loss function*.

We refer to $p_\star = \inf_{x \in \mathbb{R}^n} f(x) \in [-\infty, \infty)$ as the *optimal value* of theoptimization problem.

In this class, we assume x is a continuous variable and that f is continuous and (usually) differentiable. Problems with such structure are referred to as *continuous optimization* problems.

Problems with Boolean- or integer-values x are referred to as *combinatorial optimization* problems. (Not our focus.)

Local vs. global minima

$x_\star$ is a *local minimum* if $f(x) \geq f(x_\star)$ within a small neighborhood.[2]

$x_\star$ is a *global minimum* if $f(x) \geq f(x_\star)$ for all $x \in \mathbb{R}^n$

In the worst case, finding the global minimum of an optimization problem is difficult. (The class of non-convex optimization problems is NP-hard.)

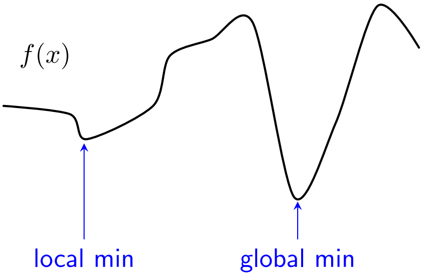

[2] if $\exists\, r > 0$ s.t. $\forall\, x$ s.t. $\|x - x_\star\| \leq r \Rightarrow f(x) \geq f(x_\star)$

Minimization vs. maximization

Why consider minimization problems? Why not maximize?

Minimization and maximization problems are equivalent since

$$\underset{x \in \mathbb{R}^n}{\text{maximize}} \quad f(x) \qquad \Leftrightarrow \qquad \underset{x \in \mathbb{R}^n}{\text{minimize}} \quad -f(x).$$

When maximizing, we refer to f as the *objective function*, *reward function*, and *merit function*.

Min vs. max: Choose what is more natural given the problem context.

The baseline convention is to minimize because of convexity.
More on this later.

Definition of solutions

For minimization problems, we define *solutions* to be global minimizers.

- A solution may or may not exist.
- A solution may or may not be unique.

Some refer to a local minimizer as a "local solution." We will not use this terminology.

Some refer to any (feasible) point as a " solution." (In a business context, if a company is selling you a "solution," this is an actionable plan that is hopefully decent, but there is no promise of the optimality of this plan.) We will not use this terminology.

In maximization problems, we define solutions to be global maximizers.

Solving unconstrained optimization with calculus

In calculus, you have actually seen some optimization.

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad f(x),$$

and assume $f \colon \mathbb{R}^n \to \mathbb{R}$ is differentiable. Then,

$$\operatorname{argmin} f \subseteq \{x \in \mathbb{R}^n \mid \nabla f(x) = 0\}.$$

($\min f$ is the minimum *value* of f while $\operatorname{argmin} f$ is the set of *input* x's minimizing f.)

In other words, the minimizers must have zero-gradient ($\nabla f(x) = 0$ is a necessary condition). However, this is not a sufficient condition, and you did things like the second derivative test.

Solving unconstrained optimization with calculus

Consider

$$\underset{x,y\in\mathbb{R}}{\text{maximize}} \quad f(x,y) = 2xy + 2x - x^2 - 2y^2.$$

Then,

$$\nabla f(x,y) = \begin{bmatrix} 2y + 2 - 2x \\ 2x - 4y \end{bmatrix}.$$

Solving for $\nabla f(x,y) = 0$ yields $(x,y) = (2,1)$. Next, carry out the second derivative test.

$$\nabla^2 f(2,1) = \begin{bmatrix} \frac{\partial^2 f}{\partial x^2}(2,1) & \frac{\partial^2 f}{\partial x \partial y}(2,1) \\ \frac{\partial^2 f}{\partial x \partial y}(2,1) & \frac{\partial^2 f}{\partial y^2}(2,1) \end{bmatrix} = \begin{bmatrix} -2 & +2 \\ +2 & -4 \end{bmatrix}$$

Then $\det\left(\nabla^2 f(x)\right) > 0$ and $f_{xx}(2,1) < 0$, so $(2,1)$ is a local maximum of f. (Alternatively, we can note that both eigenvalues of $\nabla^2 f(x,y)$ are negative.) With a little bit more work, we can show that $(2,1)$ is the global maximum.

Very nice. We can do this all the time?

Can't solve unconstrained optimization with calculus

Consider minimizing the Mishra's Bird function

$$\underset{x,y\in\mathbb{R}}{\text{minimize}} \quad f(x,y) = \sin(x)e^{(1-\cos(y))^2} + \cos(y)e^{(1-\sin(x))^2} + (x-y)^2.$$

(This is a commonly used non-convex test function to evaluate the performance of optimization algorithms.) Then,

$$\nabla f(x,y) = \begin{bmatrix} \cos(x)e^{(1-\cos(y))^2} - 2\cos(y)\cos(x)(1-\sin(x))e^{(1-\sin(x))^2} + 2(x-y) \\ 2\sin(x)\sin(y)(1-\cos(y))e^{(1-\cos(y))^2} - \sin(y)e^{(1-\sin(x))^2} - 2(x-y) \end{bmatrix}.$$

Solving for $\nabla f(x,y) = 0$ analytically is impossible.

Recommended solution 1: Plot the 2D function and eyeball the solution.

Recommended solution 2: Take the eyeballed solution and run GD to refine it to local optimality.

Can't solve unconstrained optimization with calculus

Consider the ℓ_2-regularized logistic regression problem

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \sum_{i=1}^{N} \log\left(1 + \exp(v_i^\intercal x)\right) + \frac{\lambda}{2}\|x\|^2,$$

for some $\lambda > 0$ and $v_1, \ldots, v_N \in \mathbb{R}$. (These arise in statistics and machine learning.)

Then,

$$\nabla f(x) = \lambda x + \sum_{i=1}^{N} \frac{1}{1 + \exp(-v_i^\intercal x)} v_i$$

Solving for $\nabla f(x, y) = 0$ analytically is impossible. When $d > 2$, plotting and eyeballing the solution is impossible.

We must use a numerical algorithm.

It turns out that there is a unique point satisfying $\nabla f(x) = 0$, and it can be computed reliably with GD.

Constrained optimization

Constrained optimization problems have the form

$$\begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & x \in C \end{array} \quad \text{or} \quad \begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & g(x) = 0 \end{array} \quad \text{or} \quad \begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & g(x) \le 0 \end{array}$$

where $C \subseteq \mathbb{R}^n$ is nonempty, $f \colon \mathbb{R}^n \to \mathbb{R}$, and $g \colon \mathbb{R}^n \to \mathbb{R}$. We can also have combinations of these kinds of constraints.

We call $x \in C$ a set constraint, $g(x) = 0$ an *equality constraint*, and $g(x) \le 0$ an *inequality constraint*.

We say an optimization variable $x \in \mathbb{R}^n$ is *feasible* if it satisfies the constraints. Otherwise, we say x is *infeasible*.

Local vs. global minima with constraints

$x_\star$ is a *local minimum* if $f(x) \geq f(x_\star)$ among the feasible points within a small neighborhood.[3]

$x_\star$ is a *global minimum* if $f(x) \geq f(x_\star)$ for all feasible x.

The optimal value $p_\star \in [-\infty, \infty)$ is

$$p_\star = \begin{cases} \inf\{f(x) \,|\, x \in C\} & \text{in case 1} \\ \inf\{f(x) \,|\, g(x) = 0,\ x \in \mathbb{R}^n\} & \text{in case 2} \\ \inf\{f(x) \,|\, g(x) \leq 0,\ x \in \mathbb{R}^n\} & \text{in case 3.} \end{cases}$$

I.e., $p_\star$ is the minimum (infimum) value attained among feasible points.

[3] if $\exists\, r > 0$ s.t. $\forall\, x$ s.t. $\|x - x_\star\| \leq r$ and x feasible $\Rightarrow f(x) \geq f(x_\star)$

Local vs. global minima with constraints

Show figure:

Transforming constraints

Constrained optimization problems in different forms can be transformed into one another.

Any inequality can be written as LHS ≤ 0:

$$\begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & g(x) \leq h(x) \end{array} \quad \Leftrightarrow \quad \begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & g(x) - h(x) \leq 0 \end{array}$$

Multiple constraints can be combined into one:

$$\begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & g_1(x) \leq 0 \\ & \vdots \\ & g_m(x) \leq 0 \end{array} \quad \Leftrightarrow \quad \begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & \max\limits_{i=1,\ldots,m} g_i(x) \leq 0 \end{array}$$

Multiple constraints can be combined into one:

$$\begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & x \in C_1 \\ & \vdots \\ & x \in C_m \end{array} \quad \Leftrightarrow \quad \begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & x \in \bigcap\limits_{i=1,\ldots,m} C_i \end{array}$$

Solving constrained optimization with calculus: Method of Lagrange multipliers

Consider

$$\begin{aligned} \underset{x,y\in\mathbb{R}}{\text{minimize}} \quad & f(x,y) = -8x^2 + 2y \\ \text{subject to} \quad & g(x,y) = x^2 + y^2 - 1 = 0. \end{aligned}$$

Extrema solve the system of equations

$$\nabla f(x,y) = \lambda \nabla g(x,y), \qquad g(x,y) = 0,$$

where a Lagrange multiplier $\lambda \in \mathbb{R}$ is introduced. For the problem at hand,

$$\begin{bmatrix} -16x \\ 2 \end{bmatrix} = \lambda \begin{bmatrix} 2x \\ 2y \end{bmatrix}, \qquad x^2 + y^2 = 1.$$

With calculations, we find that there are four solutions to the equations

$$(x, y, \lambda) = (0, -1, 1),\ (0, +1, -1),\ (-\tfrac{3\sqrt{7}}{8}, -\tfrac{1}{8}, 8)\ (+\tfrac{3\sqrt{7}}{8}, -\tfrac{1}{8}, 8).$$

Plugging (x, y) values into f, we get the objective values

$$-2,\ +2,\ -\tfrac{65}{8},\ -\tfrac{65}{8}$$

So, the solutions are $(x_\star, y_\star) = (\pm\frac{3\sqrt{7}}{8}, -\frac{1}{8})$.

Failure of the method of Lagrange multipliers

For most problems of practical interest, the method of Lagrange multipliers fails because:

(i) $\nabla f(x) = \lambda \nabla g(x)$ may not have an analytical solution or

(ii) the constraint function g may not be differentiable.

In most cases, we must use a numerical algorithm.

Reducing constrained opt. to unconstrained opt. using penalty functions

Consider the equality constrained optimization problem

$$\begin{array}{ll} \underset{x\in\mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & Ax = b. \end{array}$$

The unconstrained optimization problem

$$\underset{x\in\mathbb{R}^n}{\text{minimize}} \quad f(x) + \frac{\rho}{2}\|Ax - b\|^2$$

with sufficiently large $\rho > 0$, can be a good approximation to the constrained optimization problem.

How large should ρ be? Figure it out with trial and error.

Reducing constrained opt. to unconstrained opt. using penalty functions

Consider the inequality constrained optimization problem

$$\begin{array}{ll} \underset{x\in\mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & g(x) \leq 0. \end{array}$$

The unconstrained optimization problems

$$\underset{x\in\mathbb{R}^n}{\text{minimize}} \quad f(x) + \rho \max\{0, g(x)\} \quad \text{or} \quad \underset{x\in\mathbb{R}^n}{\text{minimize}} \quad f(x) + \rho\big(\max\{0, g(x)\}\big)^2$$

with sufficiently large $\rho > 0$, can be a good approximation to the constrained optimization problem.

Approximating a constrained optimization problem with an unconstrained one and using unconstrained optimization algorithms to solve it is a simple heuristic that is widely used. This is not always a bad idea.

However, there are better approaches that directly address constraints.